



DVS Enterprise

Reference Architecture for



XenDesktop/ XenApp

12/5/2013
Phase 5
Version 1.17

THIS DOCUMENT IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

Copyright © 2013 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	1
1.1 Purpose.....	1
1.2 Scope.....	1
1.3 What's New in This Release.....	1
2 Solution Architecture Overview.....	2
2.1 Introduction.....	2
2.1.1 Physical Architecture Overview.....	2
2.1.2 DVS Enterprise – Solution Layers	3
2.2 Local Tier 1	4
2.2.1 Local Tier 1 – 90 User Combined Pilot.....	4
2.2.2 Local Tier 1 – 90 User Scale Ready Pilot.....	4
2.2.3 Local Tier 1 (iSCSI)	5
2.3 Shared Tier 1 Rack.....	7
2.3.1 Shared Tier 1 – Rack – 500 Users (iSCSI – EQL)	7
2.3.2 Shared Tier 1 – Rack (iSCSI – EQL).....	7
2.3.3 Shared Tier 1 – Rack – 1000 Users (FC – CML)	9
2.4 Shared Tier 1 Blade	11
2.4.1 Shared Tier 1 – Blade – 500 users (iSCSI – EQL)	11
2.4.2 Shared Tier 1 – Blade (iSCSI – EQL).....	12
2.4.3 Shared Tier 1 – Blade (FC – CML).....	14
3 Hardware Components	17
3.1 Network.....	17
3.1.1 Force10 S55 (ToR Switch).....	17
3.1.2 Force10 S60 (1Gb ToR Switch).....	18
3.1.3 Force10 S4810 (10Gb ToR Switch)	19
3.1.4 Brocade 6510 (FC ToR Switch)	20
3.1.5 PowerEdge M I/O Aggregator (10Gb Blade Interconnect)	21
3.1.6 PowerConnect M6348 (1Gb Blade Interconnect).....	22
3.1.7 Brocade M5424 (FC Blade Interconnect)	22
3.2 Servers.....	23
3.3 Storage.....	24
3.3.1 EqualLogic Tier 1 Storage (iSCSI)	24
3.3.2 EqualLogic Tier 2 Storage (iSCSI)	25
3.3.3 Compellent Storage (FC).....	30
3.3.4 NAS	33
3.4 Dell Wyse Cloud Clients	34
3.4.1 Dell Wyse T10.....	34
3.4.2 Dell Wyse D10D.....	35
3.4.3 Wyse Xenith 2.....	35
3.4.4 Xenith Pro 2.....	35
4 Software Components	36
4.1 Citrix XenDesktop.....	36
4.1.1 Provisioning Services (PVS).....	37
4.1.2 Machine Creation Services (MCS).....	37
4.1.3 Citrix Personal vDisk Technology	38
4.1.4 Citrix Profile Manager	39
4.2 Desktop and Application Delivery with Citrix XenApp.....	39



4.2.1 XenDesktop with XenApp and PvDisk Integration.....	41
4.2.2 PVS Integration with XenApp	41
4.2.3 XenApp Integration into DVS Enterprise Architecture.....	41
4.3 VDI Hypervisor Platforms.....	42
4.3.1 VMware vSphere 5.....	42
4.3.2 Microsoft Windows Server 2012 Hyper-V	42
5 Solution Architecture for XenDesktop 7.....	44
5.1 Compute Server Infrastructure	44
5.1.1 Local Tier 1 Rack.....	44
5.1.2 Shared Tier 1 Rack (iSCSI).....	44
5.1.3 Shared Tier 1 Rack (FC)	45
5.1.4 Shared Tier 1 Blade.....	46
5.2 Management Server Infrastructure	48
5.2.1 SQL Databases	48
5.2.2 DNS	49
5.3 Scaling Guidance.....	49
5.3.1 Windows 8 – vSphere	50
5.3.2 Windows 8 – Hyper-V.....	51
5.4 Storage Architecture Overview	51
5.4.1 Local Tier 1 Storage.....	51
5.4.2 Shared Tier 1 Storage.....	51
5.4.3 Shared Tier 2 Storage	52
5.4.4 Storage Networking - EqualLogic iSCSI.....	52
5.4.5 Storage Networking - Compellent Fiber Channel (FC).....	53
5.5 Virtual Networking	54
5.5.1 Local Tier 1 – Rack – iSCSI	54
5.5.2 Shared Tier 1 – Rack – iSCSI	58
5.5.3 Shared Tier 1 – Rack – FC	61
5.5.4 Shared Tier 1 – Blade – iSCSI.....	64
5.5.5 Shared Tier 1 – Blade – FC.....	67
5.6 Solution High Availability	70
5.6.1 Compute Layer HA (Local Tier 1)	71
5.6.2 vSphere HA (Shared Tier 1)	72
5.6.3 Hyper-V HA (Shared Tier 1).....	72
5.6.4 Management Server High Availability	72
5.6.5 XenApp Server High Availability	73
5.6.6 Provisioning Services High Availability	73
5.6.7 Windows File Services High Availability	74
5.6.8 SQL Server High Availability	74
5.6.9 Load Balancing	75
5.7 XenDesktop Communication Flow.....	76
6 Customer Provided Stack Components	77
6.1 Customer Provided Storage Requirements.....	77
6.2 Customer Provided Switching Requirements.....	77
7 End-user Workload Characterization	78
7.1 Workload Characterization Overview	78
7.1.1 Basic Workload Characterization	78
7.1.2 Standard Workload Characterization	78
7.1.3 Premium Workload Characterization.....	79
7.1.4 Workload Characterization Testing Details.....	79



8 Solution Performance and Testing	81
8.1 Load Generation and Monitoring.....	81
8.1.1 Login VSI – Login Consultants.....	81
8.1.2 Liquidware Labs Stratusphere UX.....	81
8.1.3 EqualLogic SAN HQ.....	82
8.1.4 VMware vCenter.....	82
8.1.5 Microsoft Perfmon.....	82
8.2 Testing and Validation.....	82
8.2.1 Testing Process.....	82
8.3 XenDesktop Test Results.....	83
8.3.1 Windows Server 2012 Hyper-V.....	85
8.3.2 vSphere 5.1 Update 1.....	89
8.3.3 XenDesktop with Personal vDisk Enabled.....	93
8.3.4 EqualLogic PS6110XS Hybrid Array Testing Summary.....	93
8.3.5 Dell Compellent Storage Testing Summary.....	96
8.4 XenApp Testing Results.....	100
8.4.1 Test Configuration.....	100
8.4.2 Test Methodology.....	100
8.4.3 Test Results Summary.....	100
8.4.4 XenApp Physical Host Server Scaling (Sandy Bridge).....	102
Appendix A – Branch Office WAN Acceleration	103
Appendix B – Load Balancing and Disaster Recovery	104
Acknowledgements	105
About the Authors	106



1 Introduction

1.1 Purpose

This document describes:

1. Dell DVS Reference Architecture for Citrix XenDesktop scaling from 50 to 50K+ VDI users.
2. A VDI Experience Proof of Concept (POC) or pilot Solution, an entry level configuration supporting up to 90 VDI users.
3. 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 on Windows Server Hyper-V 2012 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

- Intel Ivy Bridge CPU support for all servers (E5-2600 v2) with increased user densities
- Dell IO Aggregator added to extended SCL replacing M8024-K modules
- Update to Dell Compellent Tier 1 scaling and configuration (SCOS 6.3)
- Dell Wyse Cloud Client updates
- Citrix XenMobile description added to appendix A
- Dell PowerEdge VRTX offering for ROBO: [LINK](#)
- Shared graphics (vSGA) on NVIDIA Grid cards: [LINK](#)
- Microsoft Lync 2013 test results and sizing: [LINK](#)



2 Solution Architecture Overview

2.1 Introduction


The DVS Enterprise Solution leverages a core set of hardware and software components consisting of 4 primary layers:

- Networking Layer
- Compute Server Layer
- Management Server Layer
- Storage Layer

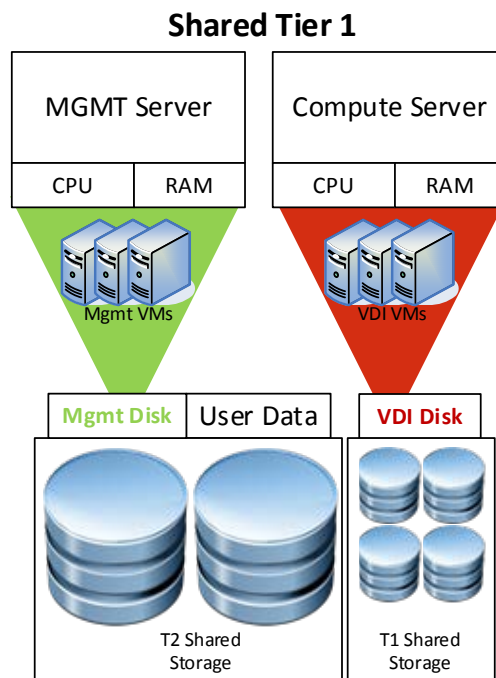
These components have been integrated and tested to provide the optimal balance of high performance and lowest cost per user. Additionally, the DVS Enterprise Solution includes an approved extended list of optional components in the same categories. These components give IT departments the flexibility to custom tailor the solution for environments with unique VDI feature, scale or performance needs. The DVS Enterprise 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 DVS Enterprise architecture consists of two models: Local Tier1 and Shared Tier1. "Tier 1" in the DVS context defines from which disk source the VDI sessions execute. Local Tier1 includes rack servers only 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. DVS Enterprise 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 DVS Enterprise – Solution Layers

Only a single high performance Force10 48-port switch is required to get started in the network layer. This switch will host all solution traffic consisting of 1Gb iSCSI and LAN sources for smaller stacks. Above 1000 users we recommend that LAN and iSCSI traffic be separated into discrete switching fabrics. Additional switches can be 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 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.



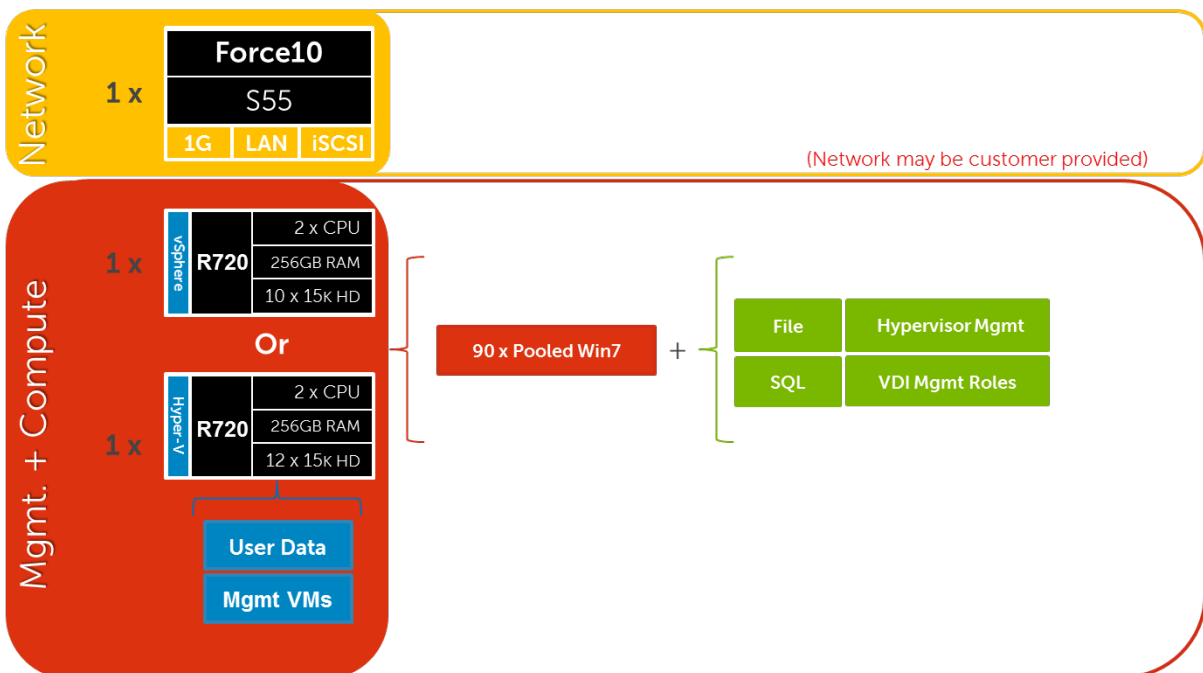
The storage layer consists of options provided by EqualLogic for iSCSI and Compellent arrays for Fiber Channel to suit your Tier 1 and Tier 2 scaling and capacity needs.



2.2 Local Tier 1

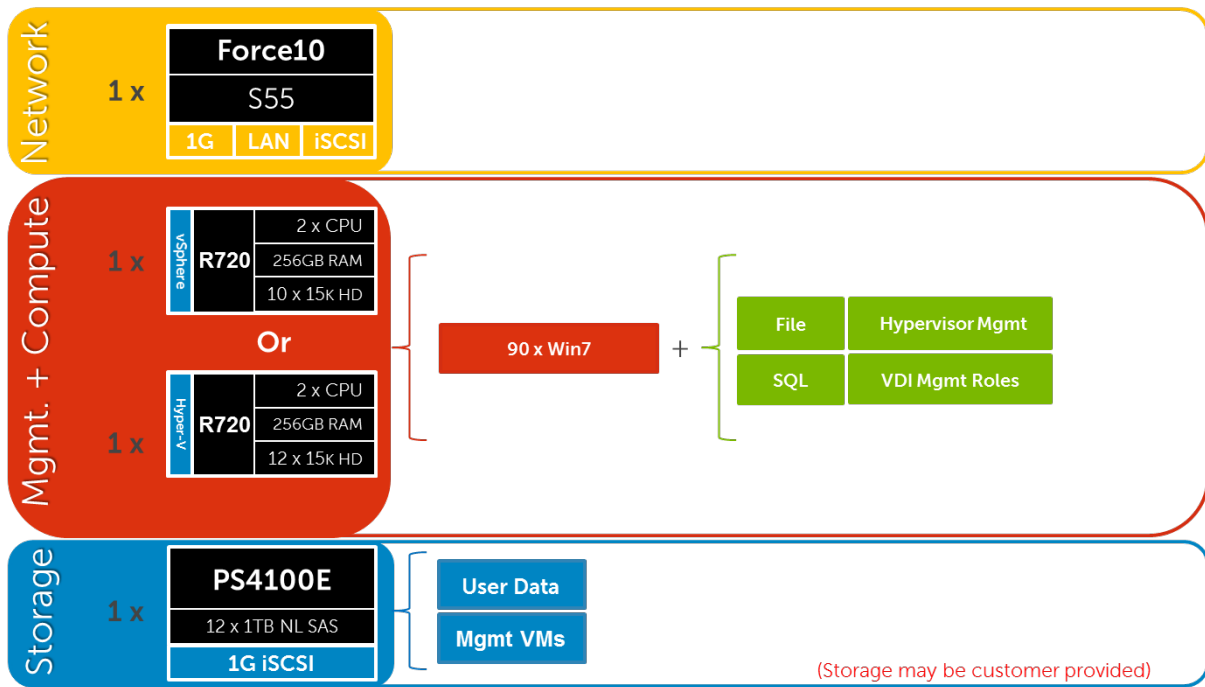
2.2.1 Local Tier 1 – 90 User Combined Pilot

For a very small deployment or pilot effort to familiarize yourself with the solution architecture, we offer a 90 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 MCS.



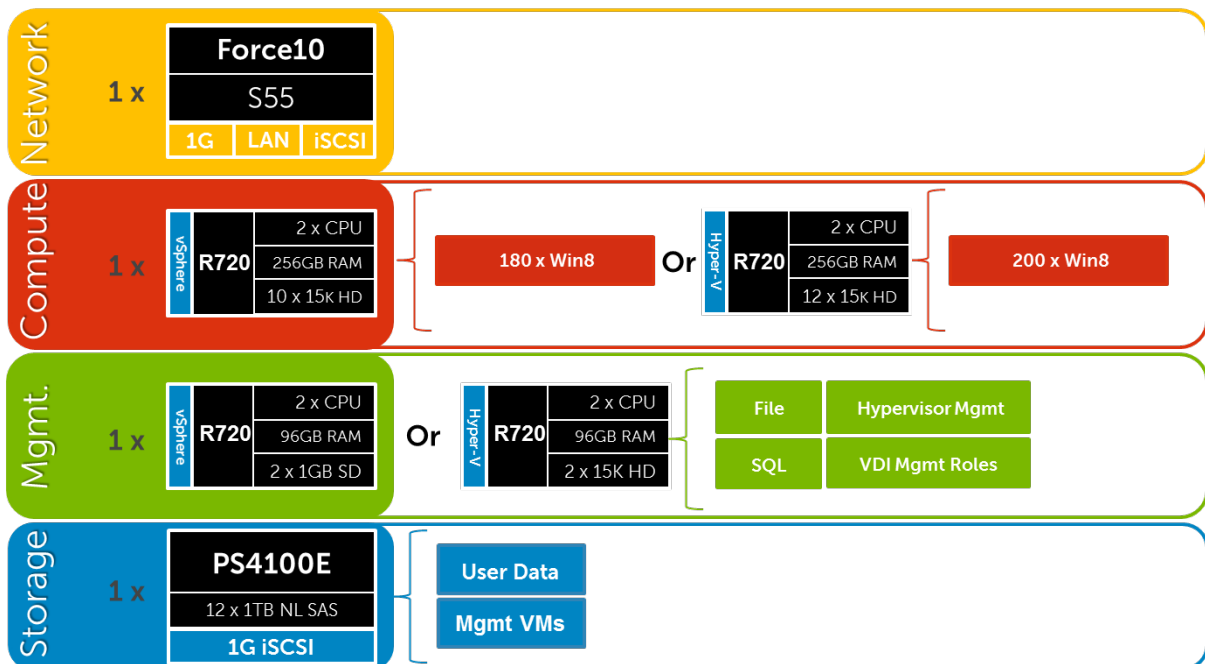
2.2.2 Local Tier 1 – 90 User Scale Ready Pilot

In addition to the 90 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 will benefit by building out into Tier 2 initially.



2.2.3 Local Tier 1 (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.




2.2.3.1 Local Tier 1 – Network Architecture (iSCSI)

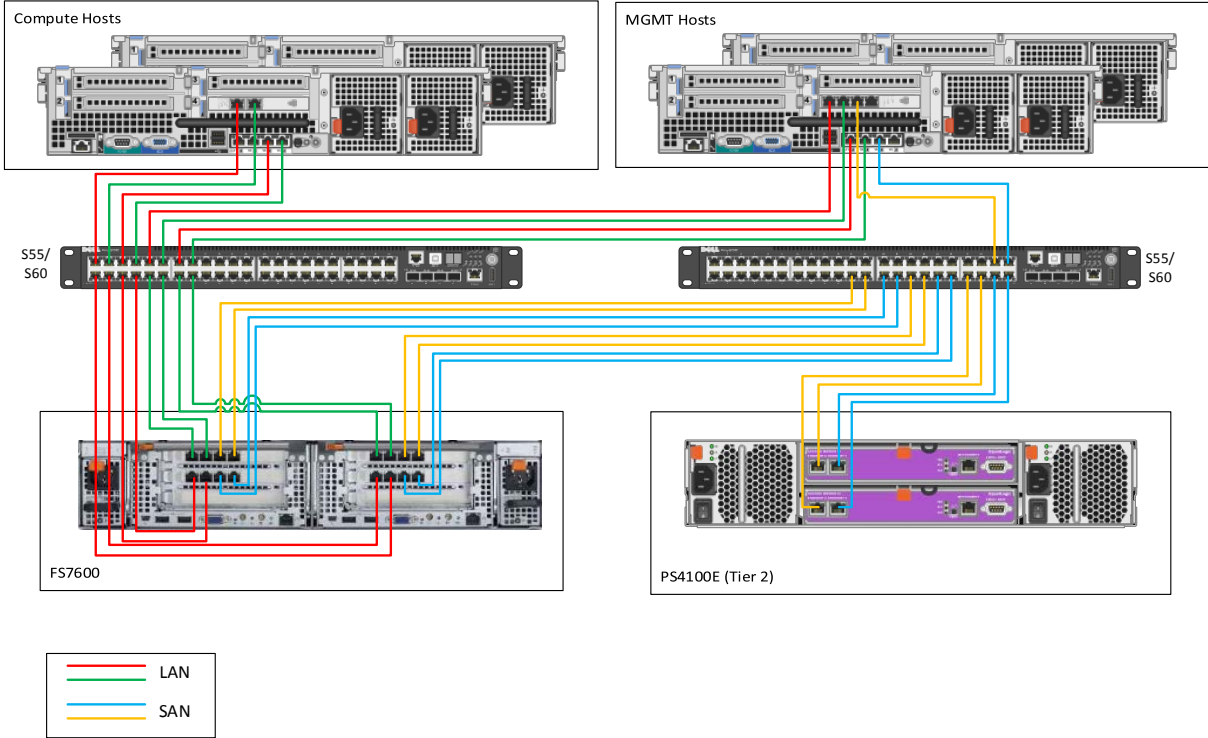
In the local tier 1 architecture, a single Force10 switch can be shared among all network connections for both management and compute, up to 1000 users. Over 1000 users DVS recommends separating the network fabrics to isolate iSCSI and LAN traffic as well as making each



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/ switched locally, with all layer 3/ routable VLANs trunked from a core or distribution switch. The following diagrams illustrate the logical data flow in relation to the core switch.



2.2.3.2 Local Tier 1 Cabling (Rack – HA)



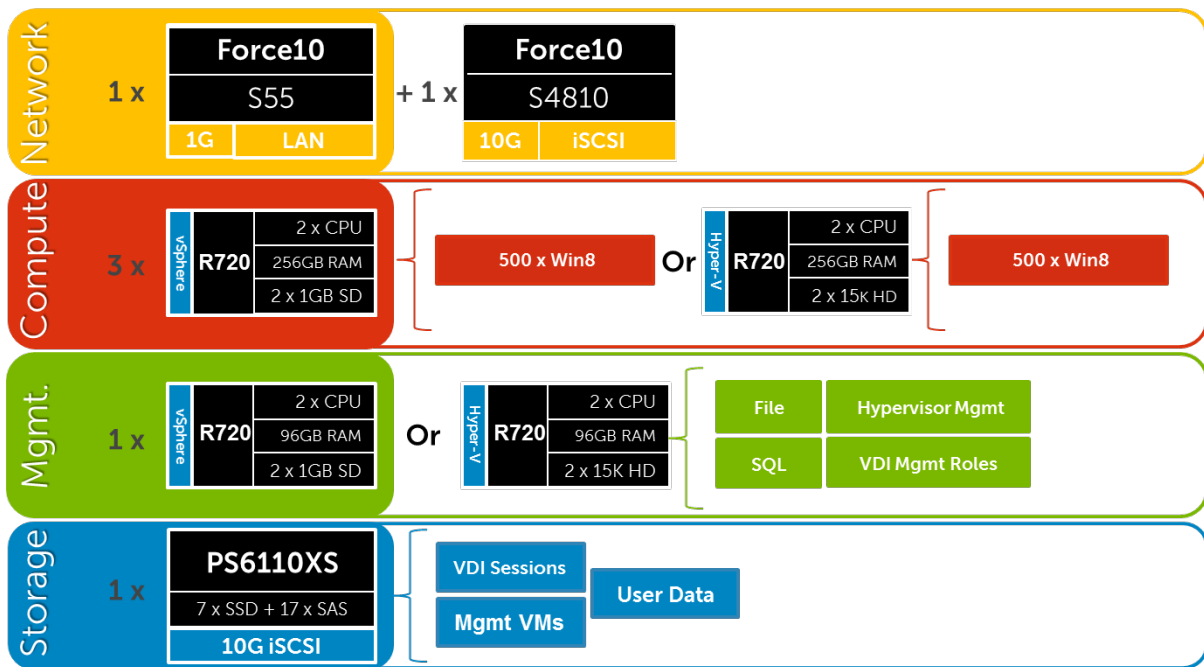
2.2.3.3 Local Tier 1 Rack Scaling Guidance (iSCSI)

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

2.3 Shared Tier 1 Rack

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

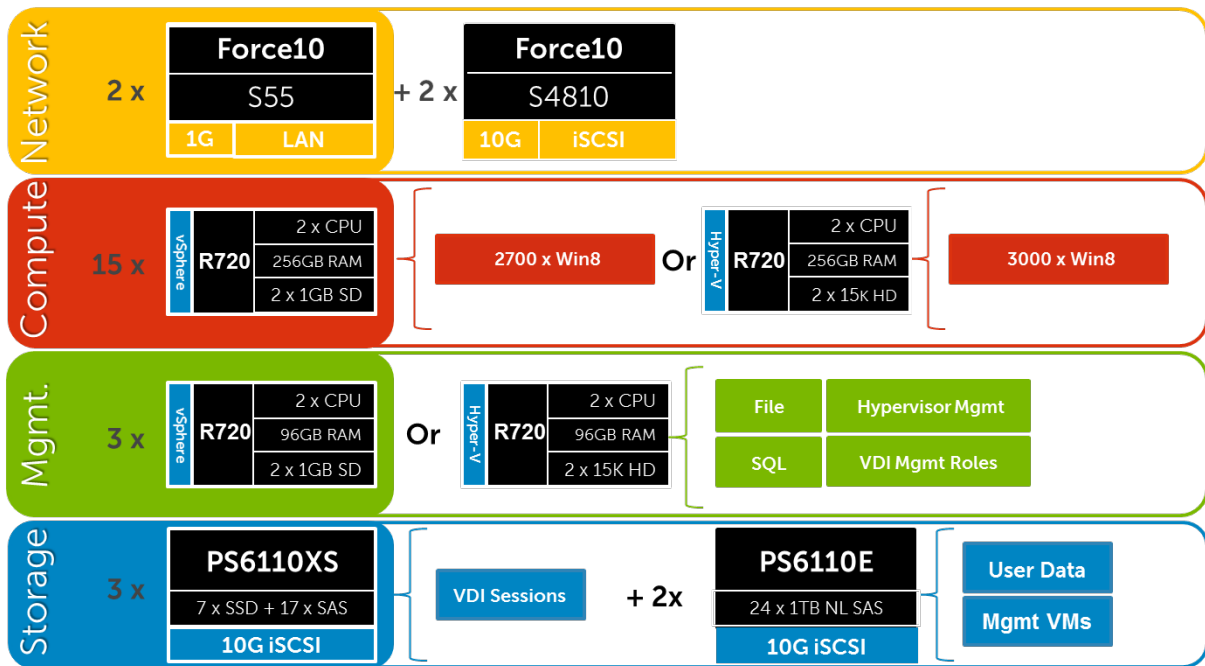
For POCs or small deployments, Tier1 and Tier2 can be combined on a single 6110XS storage array. Above 500 users, a separate array needs to be used for Tier 2.



2.3.2 Shared Tier 1 – Rack (iSCSI – EQL)

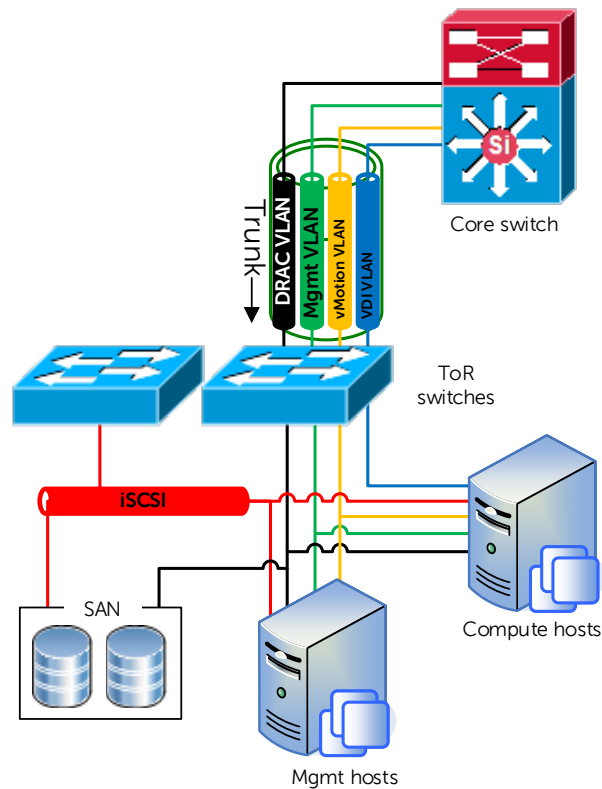
For 500 or more users on EqualLogic, the Storage layers are separated into discrete arrays. The drawing below depicts a 3000 user build where the network fabrics are separated for LAN and iSCSI traffic. Additional 6110XS arrays are added for Tier 1 as the user count scales, just as the Tier 2 array models change also based on scale. The 4110E, 6110E, 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 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.



2.3.2.2 Shared Tier 1 Cabling – (Rack – EQL)



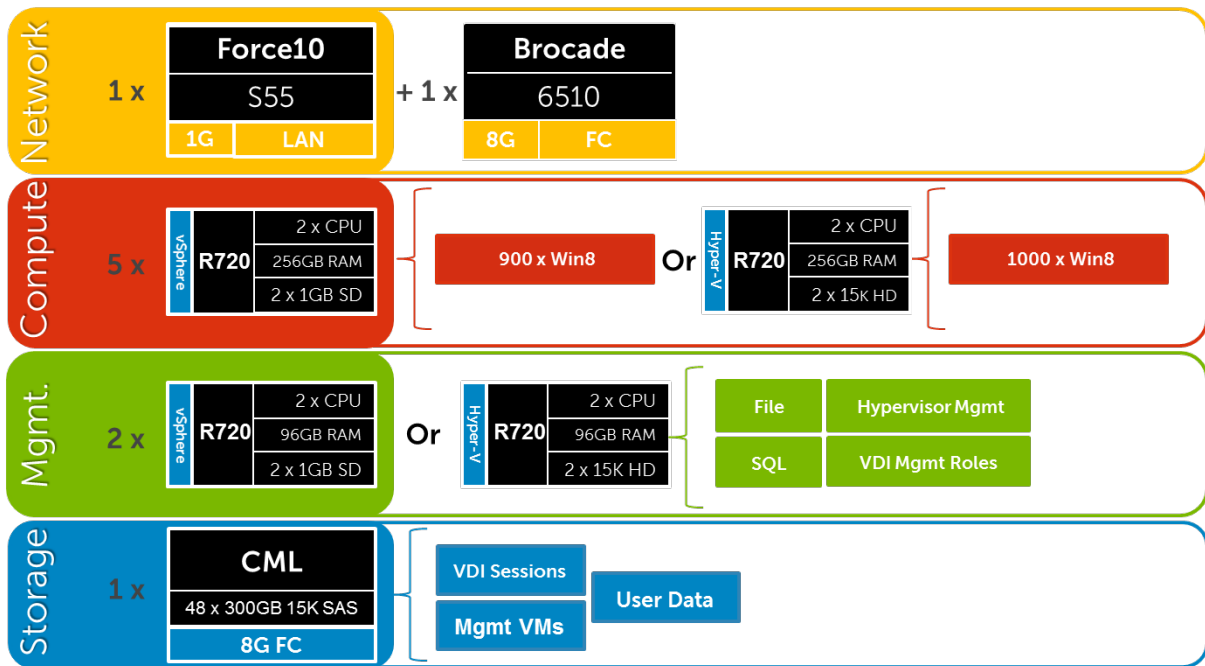
2.3.2.3 Shared Tier 1 Rack Scaling Guidance (iSCSI)

Shared Tier 1 HW scaling (Rack – iSCSI)					
User Scale	ToR LAN	ToR 10Gb iSCSI	EQL T1	EQL T2	EQL NAS
0-500	S55	S4810	6110XS	-	-
500-1000	S55	S4810	6110XS	4110E	-
0-1000 (HA)	S55	S4810	6110XS	4110E	NX3300
0-3000	S55	S4810	6110XS	6110E	NX3300
3000-6000	S55	S4810	6110XS	6510E	NX3300
6000+	S60	S4810	6110XS	6510E	NX3300

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


Utilizing Compellent storage for Shared Tier 1 provides a fiber channel solution where Tier 1 and Tier 2 are functionally combined in a single array. Tier 2 functions (user data + Mgmt VMs) can be removed from the array if the customer has another solution in place. Doing this will net an additional 30% resource capability per Compellent array for Tier 1 user desktop sessions based on our test results. Scaling this solution is very linear by predictably adding Compellent arrays for every 1000 users, on average.






2.3.3.1 Shared Tier 1 Rack – Network Architecture (FC)

In the Shared Tier 1 architecture for rack servers using fiber channel, a separate switching infrastructure is required for FC. Management and compute servers will 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)

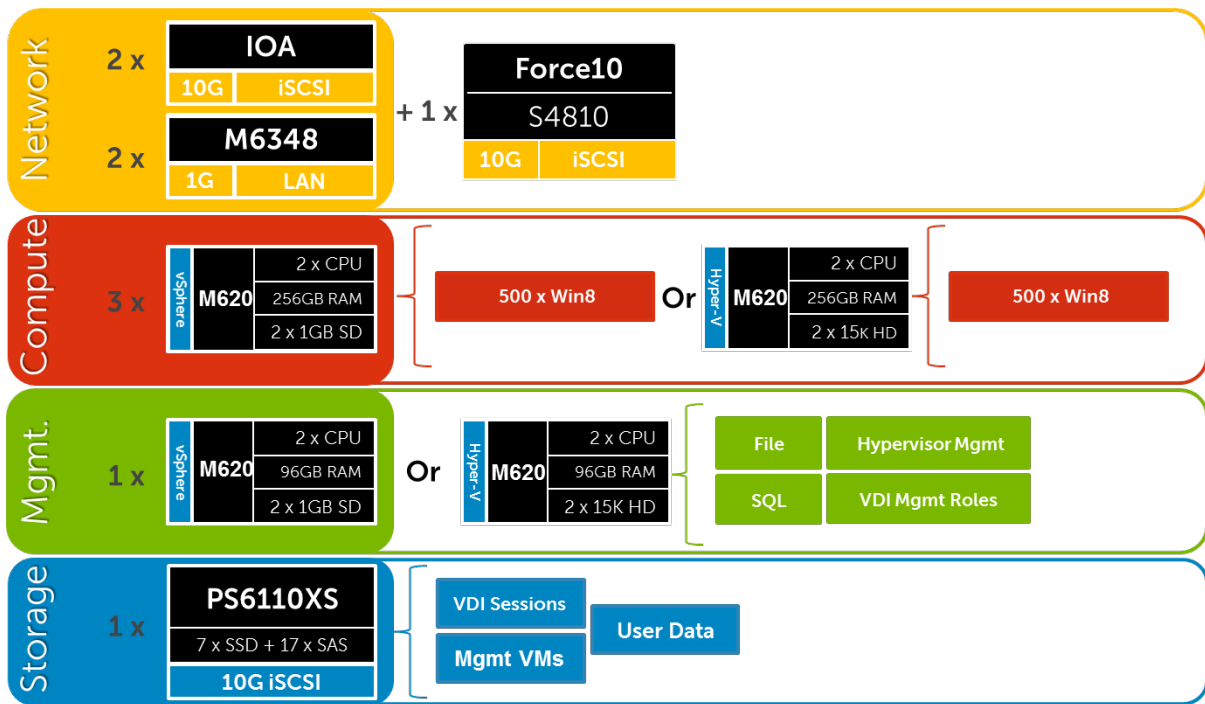
Shared Tier 1 HW scaling (Rack - FC)					
User Scale	LAN Network	FC Network	CML T1	CML T2	CML NAS
0-1000	S55	6510	15K SAS	-	-
0-1000 (HA)	S55	6510	15K SAS	NL SAS	FS8600
1000-6000	S55	6510	15K SAS	NL SAS	FS8600
6000+	S60	6510	15K SAS	NL SAS	FS8600

2.4 Shared Tier 1 Blade

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

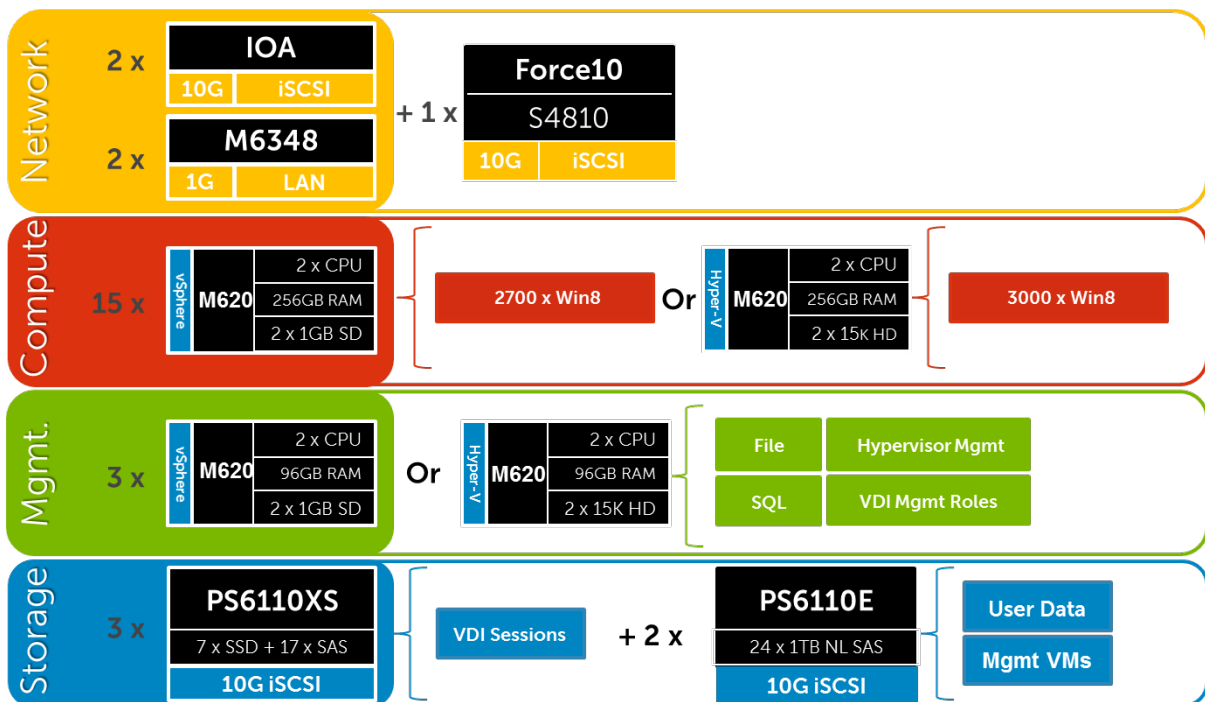
As is the case in the Shared Tier 1 model using rack servers, blades can also be used in a 500 user bundle by combining Tier 1 and Tier 2 on a single 6110XS array. Above 500 users, separate Tier 1 and Tier 2 storage into discrete arrays.






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 6110XS for VDI sessions. At this scale we also separate LAN from iSCSI switching. Load balancing and NAS can be added optionally for HA. The drawing below depicts a 3000 user solution:



2.4.2.2 Shared Tier 1 Blade Cabling (Blade – EQL)



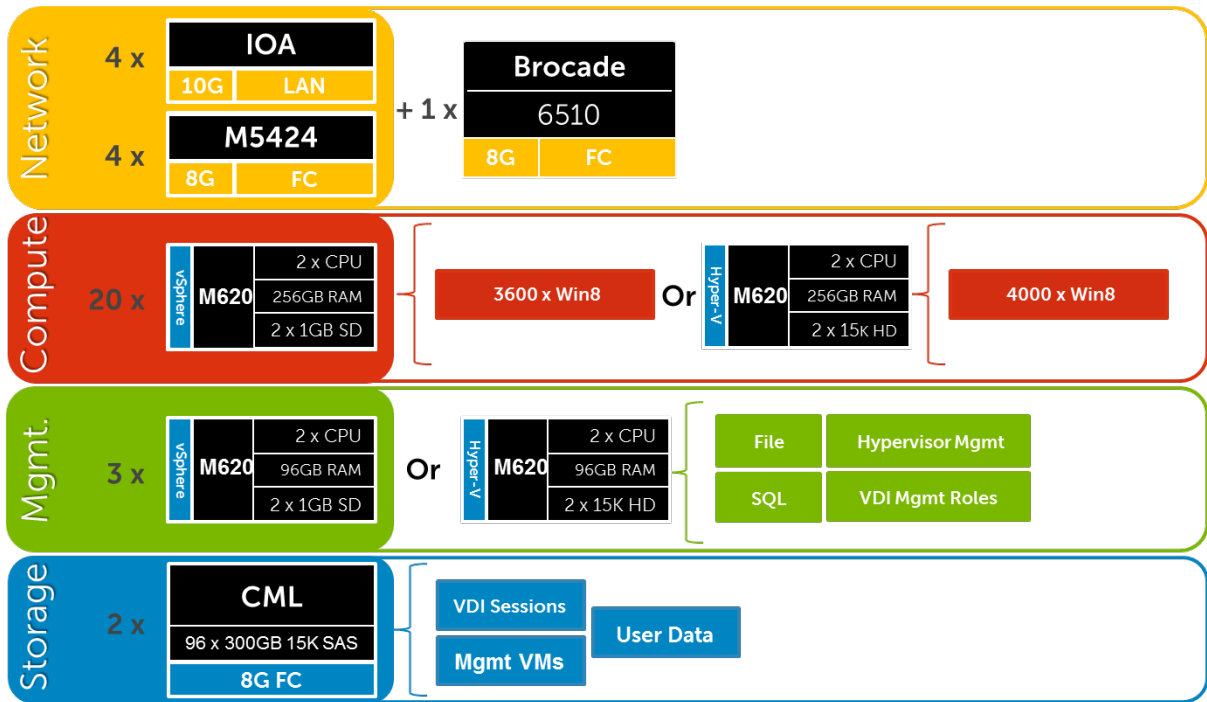
2.4.2.3 Shared Tier 1 Blade Scaling Guidance (iSCSI)

Shared Tier 1 HW scaling (Blade - iSCSI)						
User Scale	Blade LAN	Blade iSCSI	ToR 10Gb iSCSI	EQL T1	EQL T2	EQL NAS
0-500	M6348	IOA	S4810	6110XS	-	-
500-1000	M6348	IOA	S4810	6110XS	4110E	-
0-1000 (HA)	M6348	IOA	S4810	6110XS	4110E	NX3300
0-3000	M6348	IOA	S4810	6110XS	6110E	NX3300
3000-6000	M6348	IOA	S4810	6110XS	6510E	NX3300
6000+	M6348	IOA	S4810	6110XS	6510E	NX3300

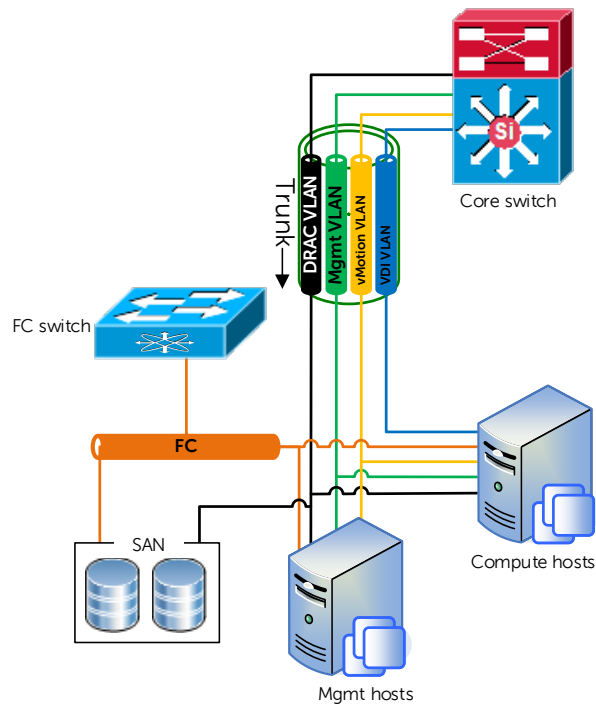
2.4.3 Shared Tier 1 – Blade (FC – CML)

Fiber channel is again an option in Shared Tier 1 using 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.






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	Blade FC	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 DVS Enterprise 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 1/10 GbE 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 DVS Enterprise 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 48 GbE access ports with up to four modular 10 GbE 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 can be 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 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	ToR switch for LAN and iSCSI in Local Tier 1 solution



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 can be used.
- The front 4 SFP ports can support copper cabling and can be 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 can be 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 1/10 GbE access switch optimized for lowering operational costs at the network edge and is recommended for DVS Enterprise 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 GbE ports and up to four optional 10 GbE 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 is 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 can be used.
- The front 4 SFP ports can support copper cabling and can be 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 switch can be 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 10/40 GbE 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 1/10 GbE (SFP+) ports as well as four 40 GbE QSFP+ uplinks to conserve valuable rack space and simplify the migration to 40 Gbps in the data center core (Each 40 GbE QSFP+ uplink can support four 10 GbE 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/ multi-mode 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 can be 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 can be 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 switch can be 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 16 Gbps Fibre Channel 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 16 Gbps



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 switch for all solutions. Optional for blades.



Guidance:

- The 6510 FC switch can be 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 can be 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
		4-port SFP+ 10Gb module	
		4-port 10GBASE-T copper module (one per IOA)	
		Stacking available only with Active System Manager	






Guidance:

- 10Gb uplinks to a ToR switch are the preferred design choice using TwinAx or optical cabling for longer runs.
- If copper-based uplinks are necessary, an additional FlexIO modules can be used.

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 12 switches	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 can be 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 switch and the Dell™ PowerEdge™ M1000e blade enclosure provide robust solutions for Fibre Channel 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 fibre channel (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 can be 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

The rack server platform for the DVS Enterprise 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.





The blade server platform for the DVS Enterprise 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.



3.3 Storage


3.3.1 EqualLogic Tier 1 Storage (iSCSI)

3.3.1.1 PS6110XS

Implement both high-speed, low-latency solid-state disk (SSD) technology and high-capacity HDDs from a single chassis. The PS6110XS 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.

Model	Features	Options	Uses
EqualLogic PS6110XS	24 drive hybrid array (SSD + 10K SAS), dual HA controllers, Snaps/Clones, Async replication, SAN HQ, 10Gb	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 can be 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.


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 PS6110E


Model	Features	Options	Uses
EqualLogic PS6110E	24 drive bays (NL-SAS/ 7200k RPM), dual HA controllers, Snaps/Clones, Async replication, SAN HQ, 10Gb, 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 (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 (10Gb – iSCSI)
		96TB – 48 x 2TB SATA	
		144TB – 48 x 3TB NL SAS	





3.3.2.7 EqualLogic Configuration


Each tier of EqualLogic storage is 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.



3.3.3 Compellent Storage (FC)


Dell DVS recommends that all Compellent storage arrays be implemented using 2 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.



3.3.3.1 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 can be 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. SSDs can be 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:



3.3.3.2 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 and arrays built for this purpose will need to be custom. 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 should be 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 will be 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:

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:



3.3.4 NAS

3.3.4.1 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.

3.3.4.2 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).
-------------------	--	---	--



3.3.4.3 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 end/ zero clients are the recommended choices for this solution.

3.4.1 Dell Wyse T10

The T10 handles everyday tasks with ease and also provides multimedia acceleration for task workers who need video. Users will enjoy integrated graphics processing and additional WMV & H264 video decoding capabilities from the Marvell



ARMADA™ PXA 510 v7 1.0 GHz System-on-Chip (SoC). In addition, the T10 is one of the only affordable thin clients to support dual monitors with monitor rotation, enabling increased productivity by providing an extensive view of task work. Designing smooth playback of high bit-rate HD video and graphics in such a small box hasn't been at the expense of energy consumption and heat emissions either. Using just 7 watts of electricity earns this device an Energy Star V5.0 rating. In addition, the T10's small size enables discrete mounting options: under desks, to walls, and behind monitors, creating cool workspaces in every respect.

3.4.2 Dell Wyse D10D



The Dell Wyse D10D is a high-performance and secure ThinOS 8 thin client that is absolutely virus and malware immune. The D10D features an advanced dual-core AMD processor that handles demanding multimedia apps with ease and delivers brilliant graphics. Powerful, compact and extremely energy efficient, the D10D is a great VDI end point for organizations that need high-end performance but face potential budget limitations.

3.4.3 Wyse 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 on your desk, on the wall or behind your display. Using about 7 Watts of power in full operation, the Xenith 2 creates very little heat for a greener, more comfortable working environment.

3.4.4 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 processor, 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. Completely 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](#)

4 Software Components

4.1 Citrix XenDesktop

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

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 and help desk 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.
- **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.
- **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.
- **Farm Database**
 - A Microsoft SQL database that hosts configuration and session information and as a result is hosted on Microsoft SQL Server 2012 SP1 in a SQL mirror configuration with a witness (HA).



- **Virtual Delivery Agent (VDA)**
 - The Virtual Desktop Agent is a transparent plugin that is installed on every virtual desktop and enables the direct connection between the virtual desktop and users' endpoint devices.
- **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.

4.1.1 Provisioning Services (PVS)

The default method of desktop image provisioning and delivery as well as delivering virtual XenApp server images within the DVS Enterprise Solution is by leveraging Citrix Provisioning Services. Provisioning Services enables real-time provisioning and re-provisioning which enable administrators to completely eliminate the need to manage and patch individual systems.

Instead, all image management is done on the master image. This greatly reduces the amount of storage required compared to other methods of creating virtual desktops.



Using Provisioning Services, vDisk images are configured in Standard Image mode. A vDisk in Standard Image mode allows many desktops to boot from it simultaneously; greatly reducing the number of images that must be maintained and the amount of storage that would be required. The Provisioning Server runs on a virtual instance of Windows Server 2012 on the Management Server(s).

4.1.1.1 PVS Write Cache

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

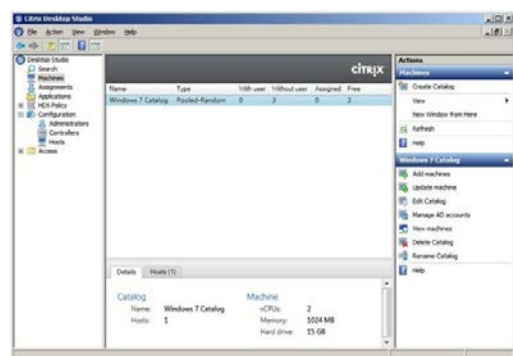
While there are 4 possible locations for storage of the write cache in PVS, the DVS Enterprise Solution places the PVS write cache on Tier 1 storage (local or shared) and configured as follows:

- VM Memory x 1.5 (up to 4092MB) + 1024MB (temporary session data)

4.1.2 Machine Creation Services (MCS)

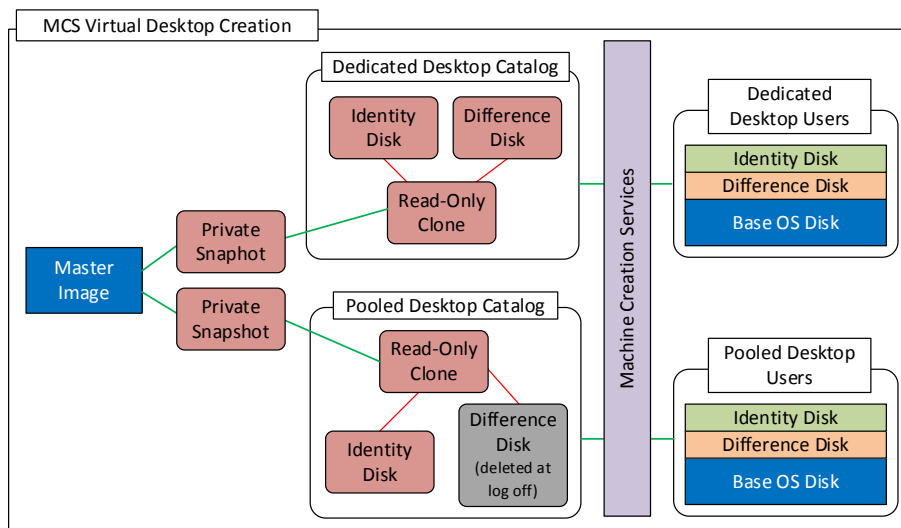
Citrix Machine Creation Services is an alternative mechanism within Citrix XenDesktop for desktop image creation and management. Machine Creation Services uses the hypervisor APIs to create, start, stop, and delete virtual desktop images. Desktop images are organized in a desktop catalog. Within a catalog there are three possible types of desktops to create and deploy access to:

- **Pooled-Random:** Desktops are assigned randomly. When they logoff, the desktop is reset to its original state and then free for another user login and use. When rebooted, any changes made are destroyed.



- **Pooled-Static:** Desktops are permanently assigned to a single user. When a user logs off, only that user can use the desktop, regardless if the desktop is rebooted. During reboots, any changes made are destroyed.
- **Dedicated:** Desktops are permanently assigned to a single user. When a user logs off, only that user can use the desktop, regardless if the desktop is rebooted. During reboots, any changes made will persist across subsequent start-ups.

All the desktops in a pooled or dedicated catalog will be based off a master desktop template which is selected when the catalog is first created. 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.



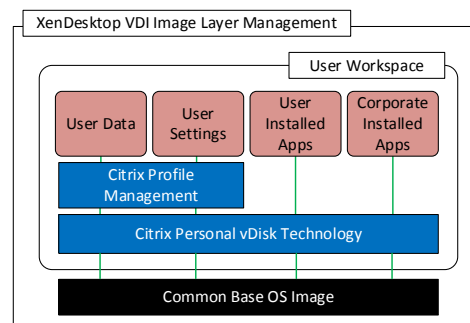
While typically intended for small to medium sized XenDesktop VDI deployments MCS can bring along with it some substantial Tier 1 storage cost savings because of the snapshot/identity/difference disk relationship. The Tier 1 disk space requirements of the identity and difference disks when layered on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

This disk space savings coupled with simplified management architecture makes Machine Creation Services an excellent option for VDI desktop image creation and management for 50-2000 user deployments.

4.1.3 Citrix Personal vDisk Technology

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

With Citrix Personal vDisk, each user receives personal storage in the form of a layered vDisk, which enables them to personalize and persist their desktop environment.



Additionally, this vDisk stores any user or departmental apps as well as any data or settings the VDI administrator chooses to store. Personal vDisk provides the following benefits to XenDesktop:

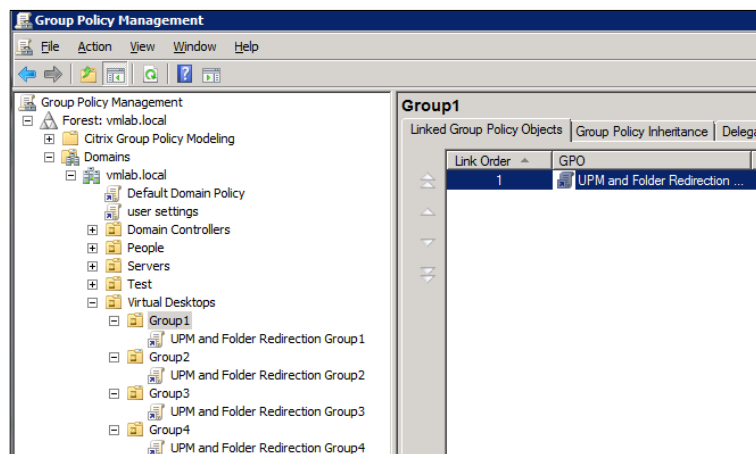
- Persistent personalization of user profiles, settings and data.
- Enables deployment and management of user installed and entitlement based applications
- Fully compatible with Application delivery solutions such as Microsoft SCCM, App-V and Citrix XenApp.
- 100% persistence with VDI pooled Storage 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 which is installed on any device where the user profiles will be managed by the toolset, which will be the virtual desktops. The second component is a Group Policy Administrative Template, which is imported to a group policy which is assigned to an organizational unit within active directory which contains the devices upon which the user profiles will be managed.

In order to further optimize the profile management folders within the user profile that can be used to store data will be redirected the users' home drive. The folder redirection will be managed via group policy objects within Active Directory. The following folders will be redirected:

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



4.2 Desktop and Application Delivery with Citrix XenApp


The DVS Enterprise Solution has been expanded to include integration with Citrix XenApp. XenApp, formerly known as WinFrame, MetaFrame and then Presentation Server, has been the cornerstone application virtualization product in the Citrix portfolio since the 1990's. Today, XenApp's proven architecture and virtualization technologies enable customers to instantly deliver any Windows-based application to users anywhere on any device.

XenApp perfectly complements a XenDesktop-based VDI deployment by enabling the delivery of applications within a user's virtual desktop. This gives the user a customized application set with a "locally-installed" application experience even though the applications are centrally installed and managed on XenApp servers. This can dramatically simplify their XenDesktop environment by

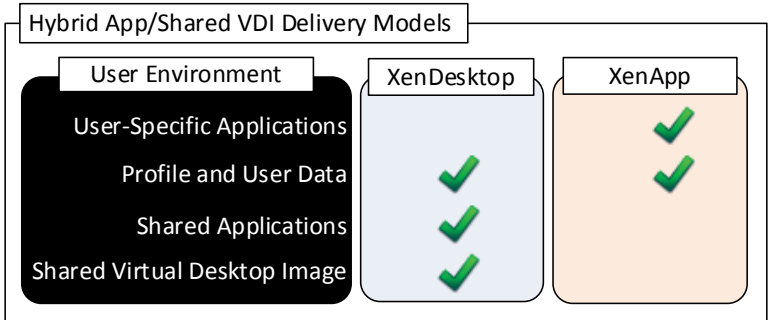


leveraging a widely shared virtual desktop image, while at the same extending the scalability of XenDesktop by alleviating the desktop compute servers from virtually all application loads by only having to run an instance of Citrix Receiver. This two-tiered approach to desktop and application delivery brings management simplification, a much quicker return on investment and the absolute best end-user experience. Synergies are also created between XenDesktop and XenApp:

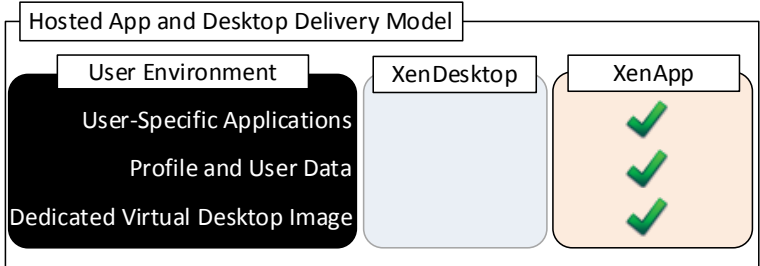
- Management of applications (single instance)
- PVS to stream XenApp servers as well as user desktops
- Scalability of XD = high level (CPU, IOPS reduction)
- Shared storage scalability less IOPS = more room to grow = save 3 IOPS per user 2.6 vs 7.0



The DVS Enterprise Solution 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 of which with commonly shared applications installed within the golden image. Then a user-specific application set is deployed and made accessible via the XenApp infrastructure, from within the virtual desktop. This deployment model is common when rolling out XenApp-based applications within an existing VDI deployment. Another example is in environments where application owners vary across IT such that the roles and responsibilities differ from those of the VDI administrator.



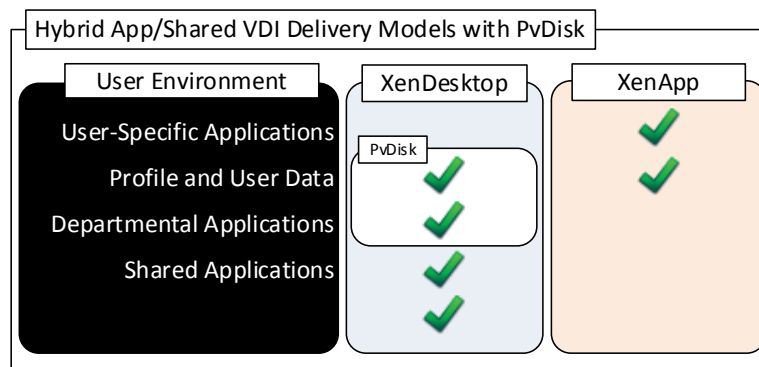
Alternatively, XenApp provides a platform for delivering a Windows server-based desktop to users who may not need a full VDI solution with XenDesktop. XenApp increases infrastructure resource utilization while reducing complexity as all hosted applications and desktops are managed from one central console. XenApp simplifies and automates the process of delivering these resources with a speedy return on investment.



4.2.1 XenDesktop with XenApp and PvDisk Integration

In a XenApp implementation, applications and desktops execute from a centralized Windows-based server and then are accessed via the Citrix ICA protocol and Citrix Receiver client plug-in. There are some instances, however, where certain departmental or custom applications cannot run hosted on a Windows server. 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 PvDisk technology is the 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 not only personalize their desktop, but provides 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.



4.2.2 PVS Integration with XenApp

One of the many benefits of PVS is the ability to quickly scale a XenApp farm, however when called upon to deliver large image volumes, PVS servers can have significant network requirements. 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.

4.2.3 XenApp Integration into DVS Enterprise Architecture

The XenApp server exists as a virtualized instance of Windows Server 2008 R2. A minimum of two (2), up to a maximum of eight (8) virtual servers can be installed per physical XenApp compute host. Since XenApp is being added to an existing DVS Enterprise stack, the only additional components required are:

- Two or more XenApp Virtual Server instances
- Creation of a XenApp data store on the existing SQL server
- Creation of an additional Storefront services web site
- Configuration of Provisioning Services for the virtual XenApp server image



- The installation of the Citrix Receiver within the golden image (for XenDesktop integration)

The number of virtual XenApp servers will vary between 2 and 8 per physical, with a minimum of 2 existing on the initial base physical server. The total number of virtual XenApp servers will be 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.

Additional resiliency can be attained by integrating the XenApp virtual servers with the virtualized management components of the XenDesktop solution, thus spreading a small farm across multiple physical hosts.

4.3 VDI Hypervisor Platforms

4.3.1 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.

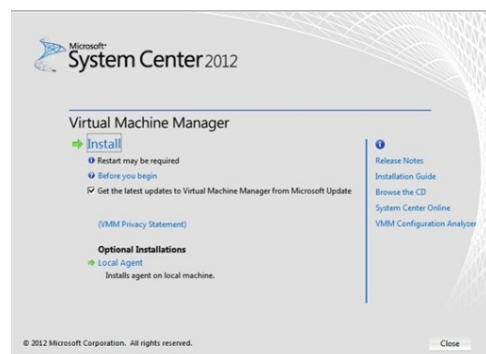
Throughout the DVS Enterprise solution, all VMware and Microsoft best practices and prerequisites for core services are adhered to (NTP, DNS, Active Directory, etc.). The vCenter 5 VM used in the solution will be a single Windows Server 2008 R2 VM, residing on a host in the management tier. SQL server is a core component of vCenter and will be hosted on another VM also residing in the management tier. All additional XenDesktop components need to be installed in a distributed architecture, 1 role per VM.

4.3.2 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 can be 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.



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. Due to the local disk requirement in the Compute layer, this model supports rack servers only. vSphere or Hyper-V can be used as the solution hypervisor. 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.

Local 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 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	OR 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)	OR 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 Shared Tier 1 Rack (iSCSI)

In the Shared Tier 1 model, VDI sessions execute on shared storage so there is no need for local disk on each server. 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. 2 additional DP NICs (2 x 1Gb + 2 x 10Gb) provide slot and connection redundancy for both network fabrics. All configuration options are identical except for CPU and RAM which are reduced on the Management host.



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 1GB 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 1GB 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 Shared Tier 1 Rack (FC)

Fiber Channel can be 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 DVS-supported using vSphere.

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 1GB 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)	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 1GB 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 DVS Enterprise Solution can support the following user counts per server (PVS).

Local/ Shared Tier 1 Rack Densities		
Workload	vSphere	Hyper-V
Basic (Win8)	180	200
Standard (Win8)	115	150
Premium (Win8)	100	125

- XenApp Compute Host

The XenApp Virtual Server compute hosts share the same PowerEdge R720 hardware configuration as the XenDesktop compute hosts for both VMware vSphere and Microsoft Hyper-V with one exception. The XenApp compute hosts ship with 96GB of RAM instead of the standard 256GB of RAM.

XenApp Virtual Server Compute Host – PowerEdge R720	
2 x Intel Xeon E5-2690v2 Processor (3GHz)	2 x Intel Xeon E5-2690v2 Processor (3GHz)
96GB Memory (6 x 16GB DIMMs @ 1600Mhz)	96GB Memory (6 x 16GB DIMMs @ 1600Mhz)
VMware vSphere on 2 x 1GB 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	OR 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	

5.1.4 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 1GB 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 serdes mezz (LAN)		1 x Intel i350 1Gb QP serdes 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)	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 1GB 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 serdes mezz (LAN)		1 x Intel i350 1Gb QP serdes mezz (LAN)
iDRAC7 Enterprise w/ vFlash, 8GB SD		iDRAC7 Enterprise w/ vFlash, 8GB SD

5.1.4.2 FC

Fiber Channel can be 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)	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 1GB 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)	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 1GB 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 DVS Enterprise Solutions can support the following single server user densities:



Shared Tier 1 Blade Densities		
Workload	vSphere	Hyper-V
Basic (Win8)	180	200
Standard (Win8)	115	150
Premium (Win8)	100	125

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 will be 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 will be 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 will be 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 would be 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 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 can be 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 XenApp farm.	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.

5.3.1 Windows 8 – vSphere

Rack or Blade, Win8, vSphere, PVS										
Basic User Count	Standard User Count	Premium 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
180	115	100	1	1	1	1	1	1	1	1
1000	690	600	1	6	1	1	1	1	1	1
2000	1380	1200	2	12	1	2	1	1	1	1
3000	1955	1700	3	17	2	3	1	1	1	1
4000	2645	2300	3	23	2	4	1	1	1	1
5000	3220	2800	4	28	2	5	1	1	1	1
6000	3910	3400	5	34	3	5	1	1	1	1
7000	4485	3900	5	39	3	6	1	1	1	1
8000	5175	4500	6	45	4	7	1	1	1	1
9000	5750	5000	7	50	4	8	1	1	1	1
10000	6440	5600	7	56	4	9	1	1	1	1

Rack or Blade, Win8, vSphere, MCS										
Basic User Count	Standard User Count	Premium User Count	Physical Mgmt Servers	Physical Host Servers	M1000e Blade Chassis	Virtual Delivery Controllers	Virtual License Servers	Virtual StoreFront Servers	Virtual vCenter Server	
180	115	100	1	1	1	1	1	1	1	
1000	690	600	2	6	1	1	1	1	1	
2000	1380	1200	2	12	1	1	1	1	1	
3000	1955	1700	2	17	2	1	1	1	1	
4000	2645	2300	2	23	2	1	1	1	1	
5000	3220	2800	3	28	2	1	1	1	1	



5.3.2 Windows 8 – Hyper-V

Rack or Blade, Win8, Hyper-V, PVS										
Basic User Count	Standard User Count	Premium 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
200	150	125	1	1	1	1	1	1	1	1
1000	750	625	1	5	1	1	1	1	1	1
2000	1500	1250	2	10	1	2	1	1	1	1
3000	2250	1875	3	15	2	3	1	1	1	1
4000	3000	2500	3	20	2	4	1	1	1	1
5000	3750	3125	4	25	2	5	1	1	1	1
6000	4500	3750	5	30	3	5	1	1	1	1
7000	5250	4375	5	35	3	6	1	1	1	1
8000	6000	5000	6	40	3	7	1	1	1	1
9000	6750	5625	7	45	4	8	1	1	1	1
10000	7500	6250	7	50	4	9	1	1	1	1

Rack or Blade, Win8, Hyper-V, MCS										
Basic User Count	Standard User Count	Premium User Count	Physical Mgmt Servers	Physical Host Servers	M1000e Blade Chassis	Virtual Delivery Controllers	Virtual License Servers	Virtual StoreFront Servers	Virtual SCVMM Server	
190	145	125	1	1	1	1	1	1	1	
1000	870	750	2	6	1	1	1	1	1	
2000	1595	1375	2	11	1	1	1	1	1	
3000	2320	2000	2	16	2	1	1	1	1	
4000	3190	2750	2	22	2	1	1	1	1	
5000	3915	3375	3	27	2	1	1	1	1	

5.4 Storage Architecture Overview

The DVS Enterprise solution has greatly expanded tier 1 and tier 2 storage strategy and flexibility over prior releases. 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 DVS Enterprise solution will consistently meet or outperform user needs and expectations.

5.4.1 Local Tier 1 Storage

Choosing the local tier 1 storage option means that the virtualization host servers use ten (10) locally installed 300GB 15k drives to house the user desktop vDisk images. In this model, tier 1 storage exists as local hard disks on the Compute hosts themselves. To achieve the required performance level, RAID 10 must be used 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.

5.4.2 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 will be 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
---------	-----------	---------------	---------	-------------



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

5.4.3 Shared Tier 2 Storage

Tier 2 is shared iSCSI storage used to host the Management server VMs and user data. EqualLogic 4100 series 1Gb arrays can be used for smaller scale deployments (Local Tier 1 only) or the 61x0 or 65x0 series for larger deployments (up to 16 in a group). The 10Gb 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 can be chosen to meet the capacity needs of the customer. The user data can be 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 can be 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 can be 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

5.4.4 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 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 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.

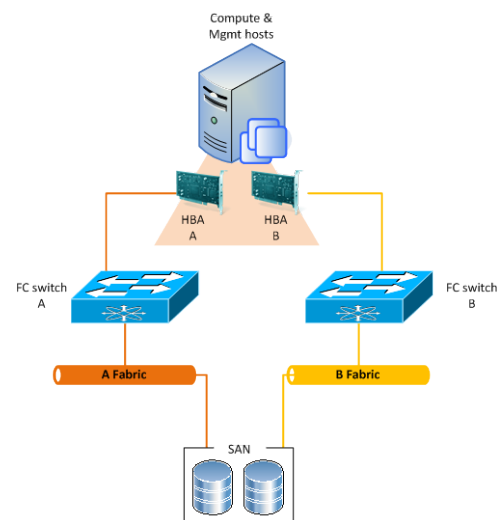
5.4.5 Storage Networking - Compellent Fiber Channel (FC)

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 can be configured to begin with to reduce costs, as a best practice recommendation, the environment needs to be configured with 2 fabrics to provide multi-pathing and end-to-end redundancy.




Using QLogic HBAs the following BIOS settings were 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 can be set to 255 because the ESXi VMkernel driver module and DSNRO can more conveniently control the queue depth

5.4.5.1 FC Zoning

Zone at least 1 port from each server HBA to communicate with a single Compellent fault domain. The result of this will be 2 distinct FC fabrics and 4 redundant paths per server. Round Robin or Fixed Paths are supported. 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.





5.5 Virtual Networking

5.5.1 Local Tier 1 – Rack – iSCSI

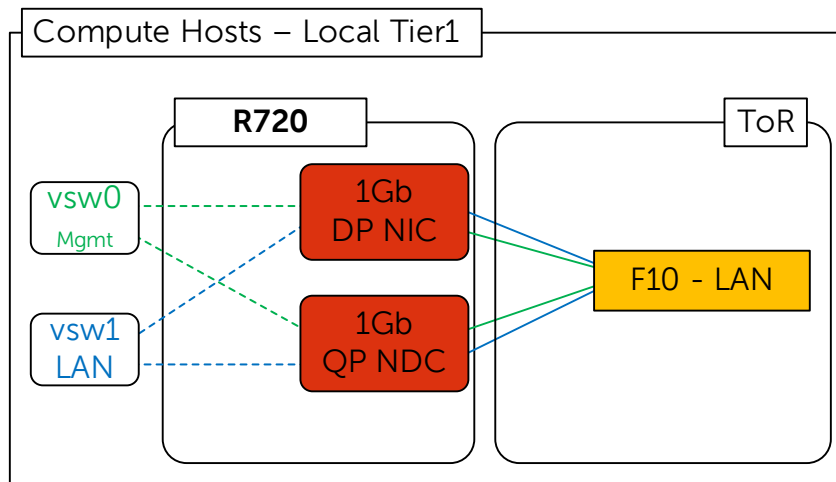
The network configuration in this model will vary between the Compute and Management hosts. The Compute hosts will not need access to iSCSI storage since they are hosting VDI sessions locally. Since the Management VMs will be hosted on shared storage, they can take advantage of Live Migration. The following 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
 - 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 can be configured for all hardware management traffic – L3 routed via core switch

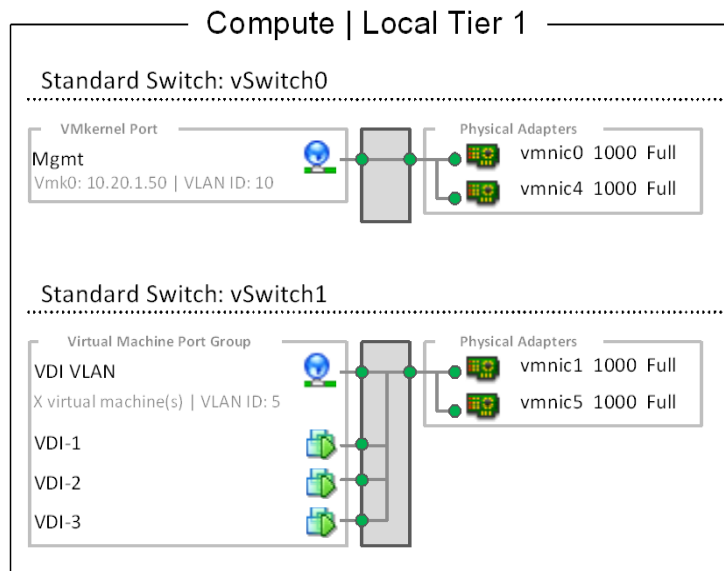
Following best practices, LAN and block storage traffic will be separated in solutions >1000 users. This traffic can be combined within a single switch in smaller stacks to minimize buy-in costs. Each Local Tier 1 Compute host will have a quad port NDC as well as a 1Gb dual port NIC. Configure the LAN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



5.5.1.1 vSphere

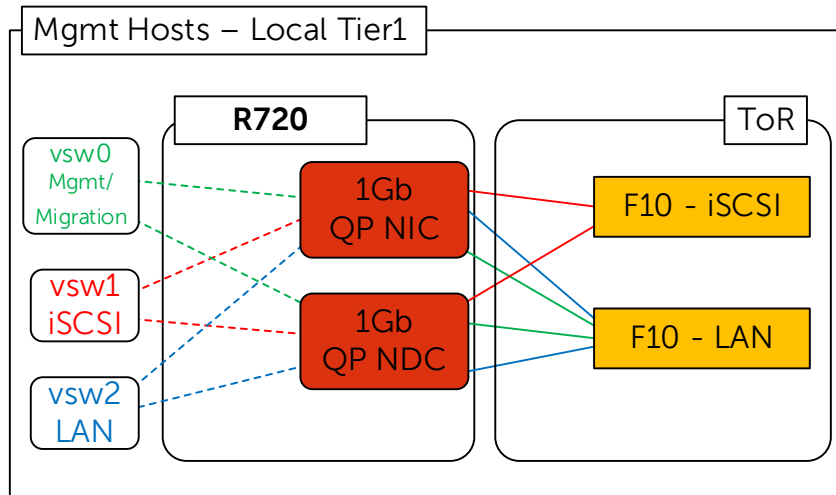


The Compute host will require 2 vSwitches, one for VDI LAN traffic, and another for the ESXi Management. Configure both vSwitches so that each is physically connected to both the onboard NIC as well as the add-on NIC. Set all NICs and switch ports to auto negotiate.

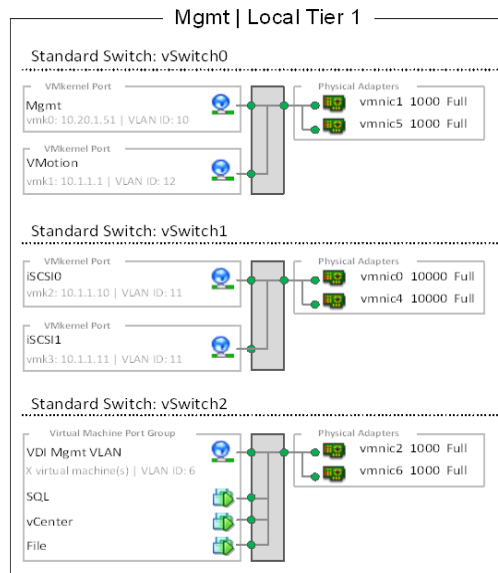


The Management hosts have a slightly different configuration since they will additionally access iSCSI storage. The add-on NIC for the Management hosts will be a 1Gb quad port NIC. 3 ports of both the NDC and add-on NIC will be used for the required connections. Isolate iSCSI onto its own vSwitch with redundant ports and connections from all 3 vSwitches. Connections should pass through both the NDC and add-on NIC per the diagram below. Configure the LAN traffic from the server to the ToR switch as a LAG.





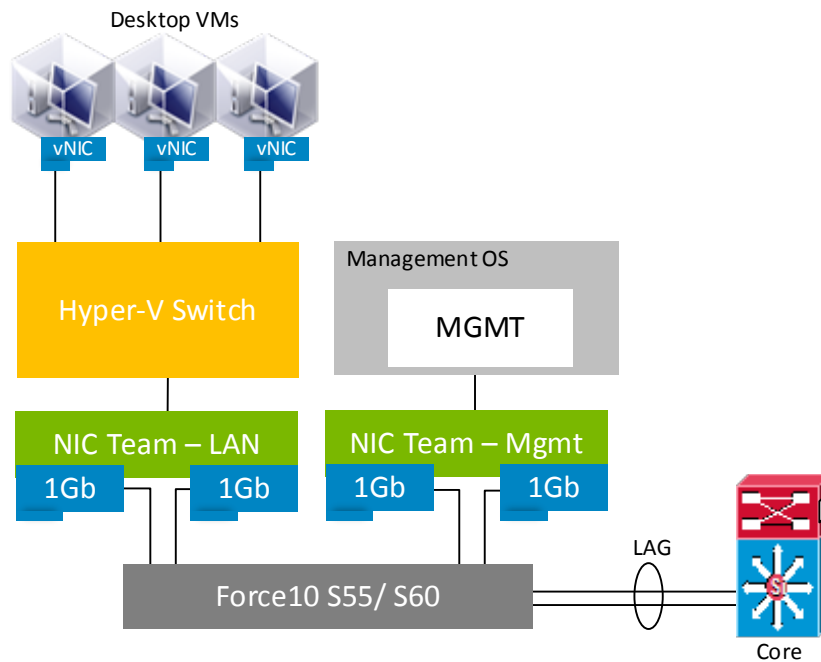
vSwitch0 carries traffic for both Management and Live Migration which needs be VLAN-tagged so that either NIC can serve traffic for either VLAN. The Management VLAN will be L3 routable while the Live Migration VLAN will be L2 non-routable.



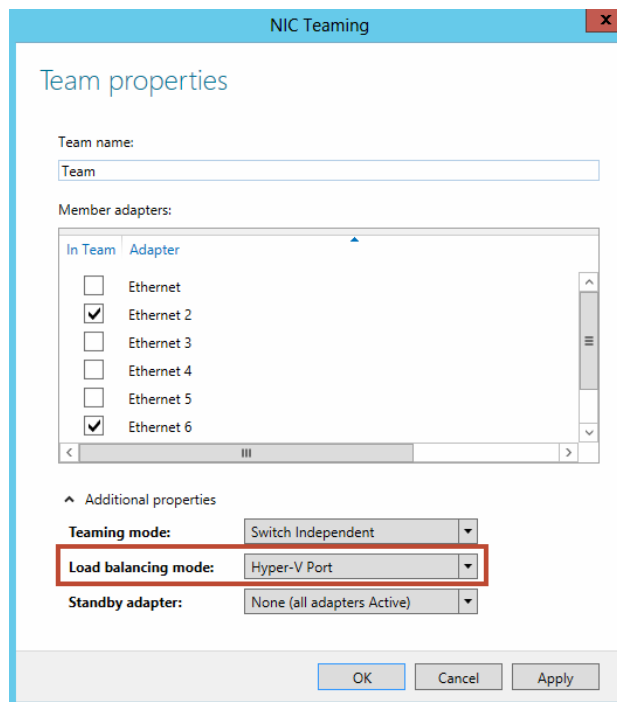
5.5.1.2 Hyper-V

The Hyper-V configuration, while identical in core requirements and hardware, is executed a bit differently. Native Windows Server 2012 NIC Teaming is utilized to load balance and provide resiliency for network connections. For the compute host in this scenario, 2 NIC teams should be configured, one for the Hyper-V switch and one for the management OS. 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.



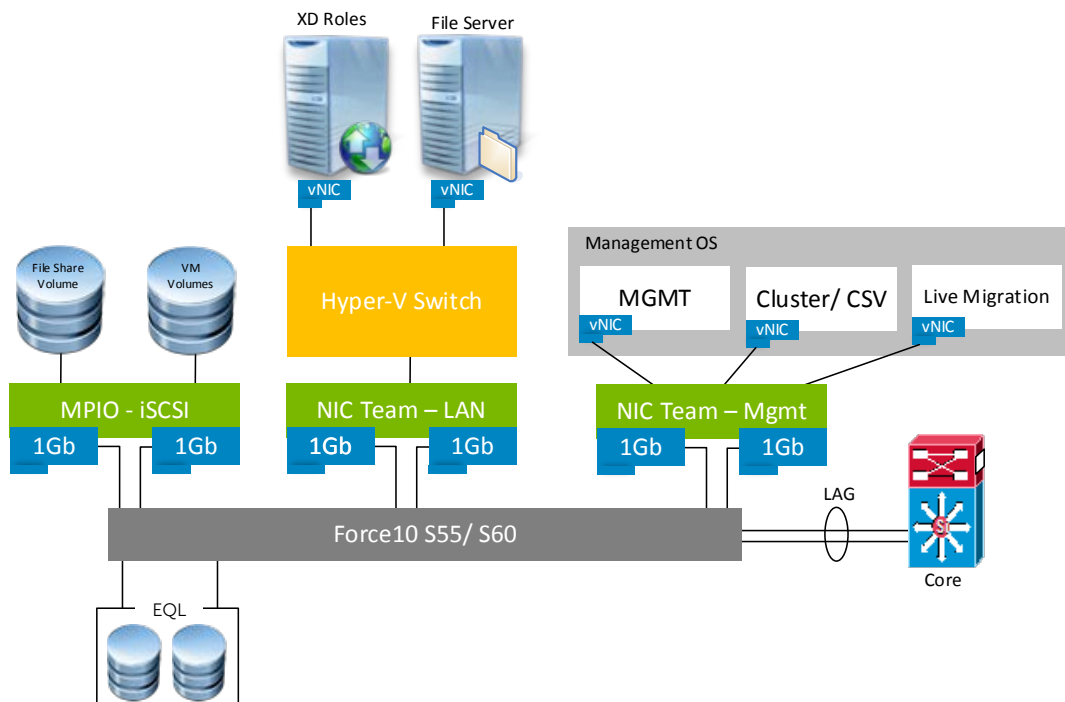


The NIC team for the Hyper-V switch should be configured as switch independent, Hyper-V port for the load balancing mode with all adapters set to active. This team will be used exclusively by Hyper-V.



The management hosts in this scenario connect to shared storage so 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 2 NIC teams, one for Hyper-V and the other for the management OS. If the management host is to be clustered or utilize live migration, additional vNICs (Team Interfaces) with dedicated VLANs should be configured on the Mgmt NIC team to support these functions.





5.5.2 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:

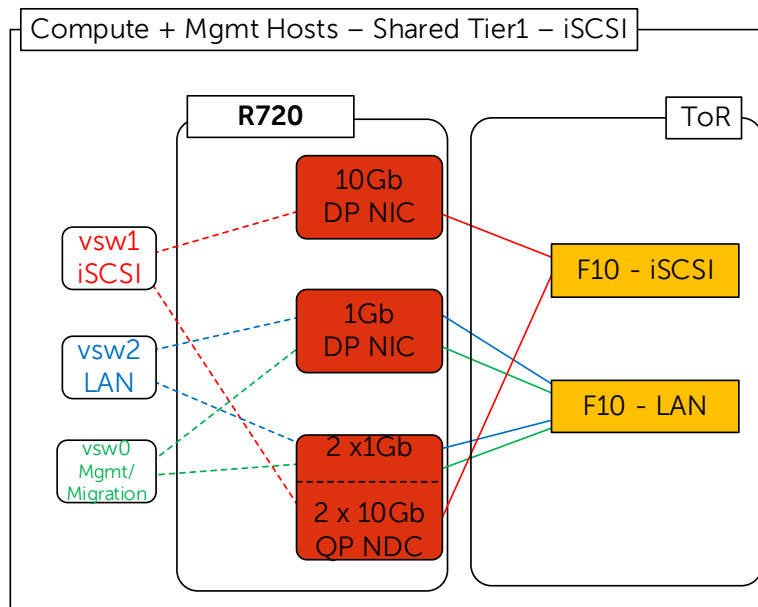
- Compute hosts (Shared Tier 1)
 - Management VLAN: Configured for hypervisor Management traffic – L3 routed via core switch
 - Live Migration VLAN: Configured for Live Migration traffic – L2 switched only, trunked from Core
 - iSCSI VLAN: Configured for iSCSI traffic – L2 switched only via ToR switch
 - 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 can be configured for all hardware management traffic – L3 routed via core switch

Following best practices, iSCSI and LAN traffic will be 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

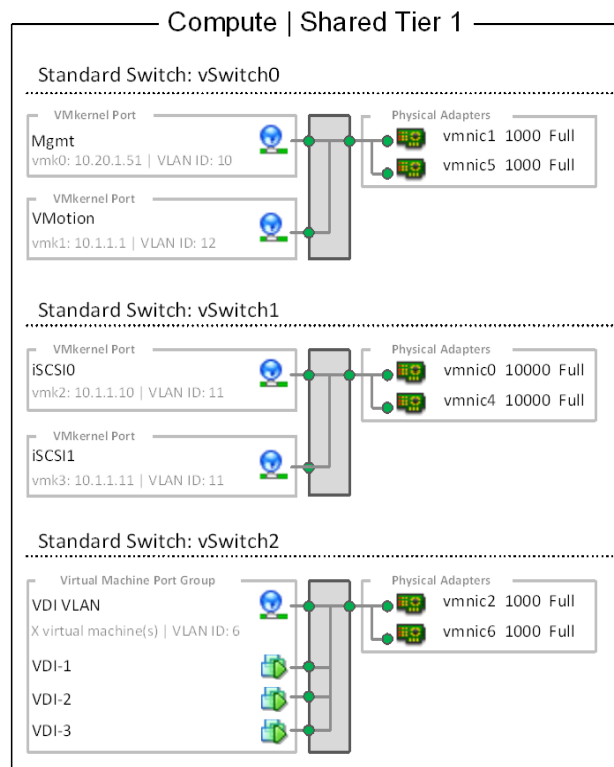


redundant ports. Connections from all 3 vSwitches should pass through both the NDC and add-on NICs per the diagram below. Configure the LAN traffic from the server to the ToR switch as a LAG.

5.5.2.1 vSphere



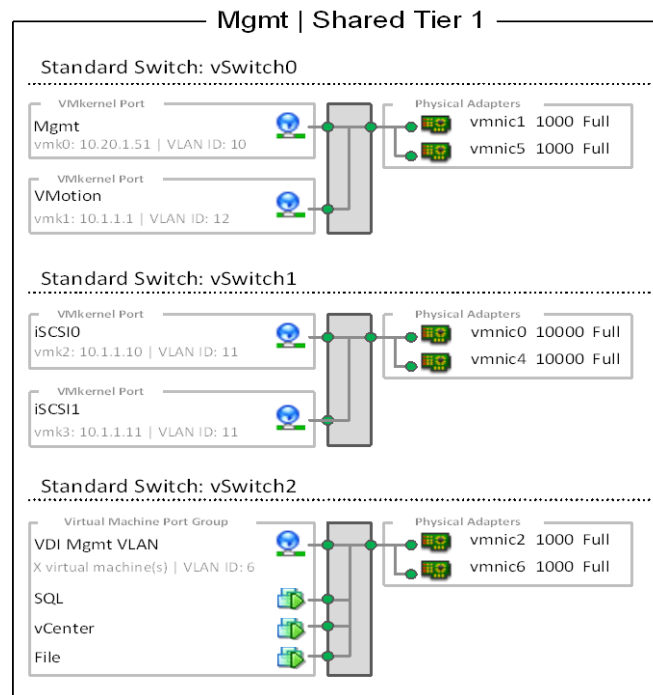
vSwitch0 carries traffic for both Management and Live Migration which needs to be VLAN-tagged so that either NIC can serve traffic for either VLAN. The Management VLAN will be L3 routable while the Live Migration VLAN will be L2 non-routable.



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

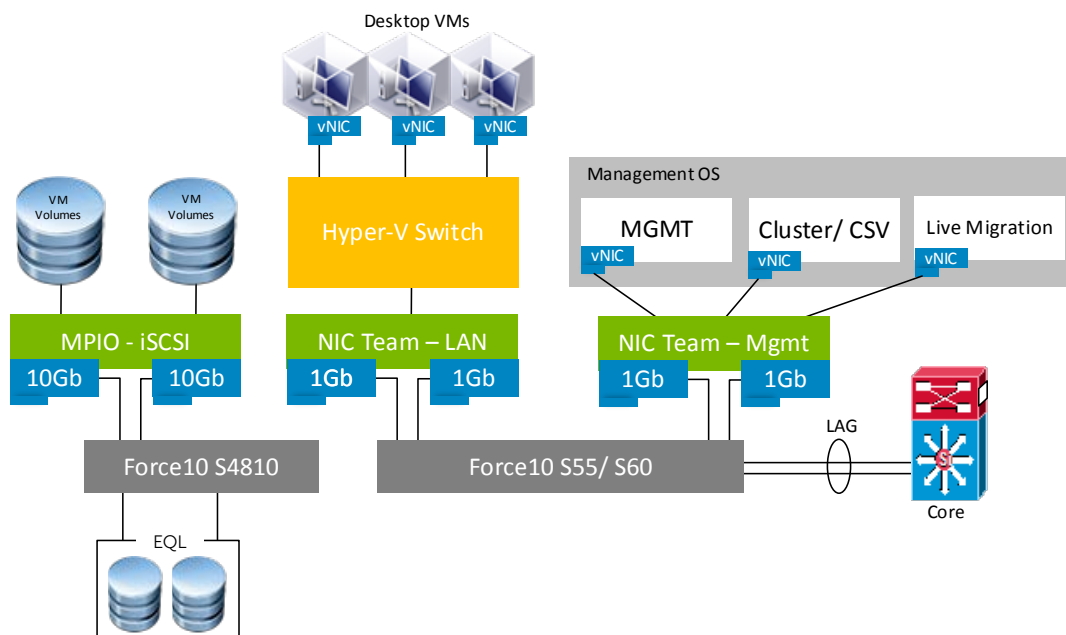


be taken to ensure that all vSwitches are assigned redundant NICs that are NOT from the same PCIe device.

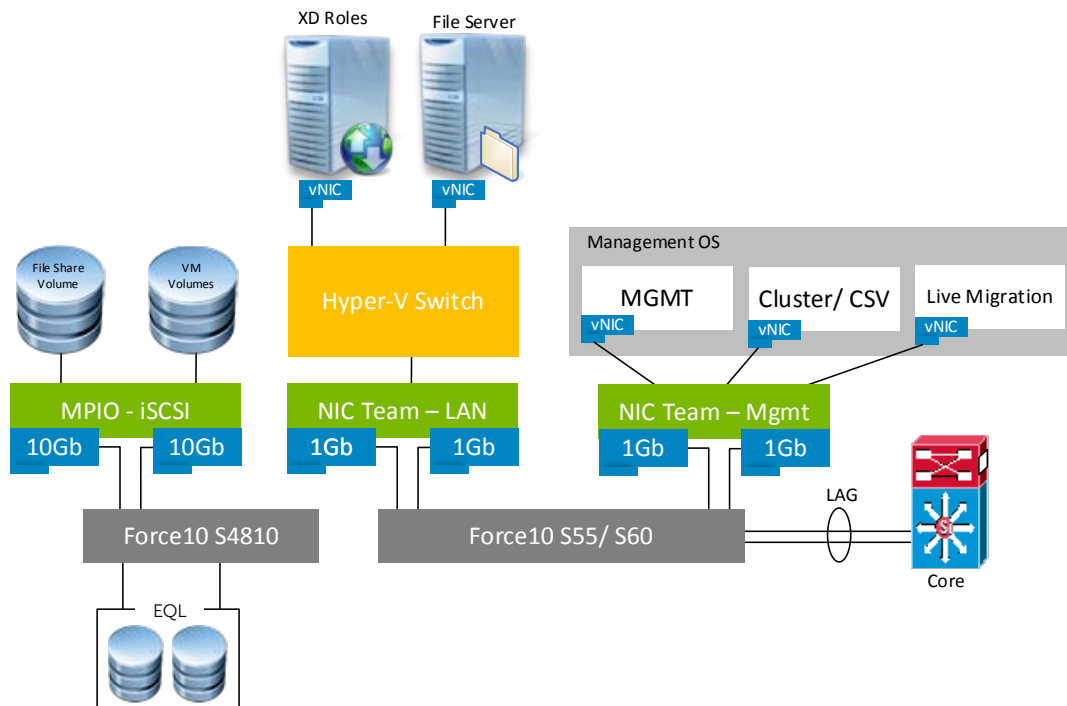


5.5.2.2 Hyper-V

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 should be 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.



5.5.3 Shared Tier 1 – Rack – FC

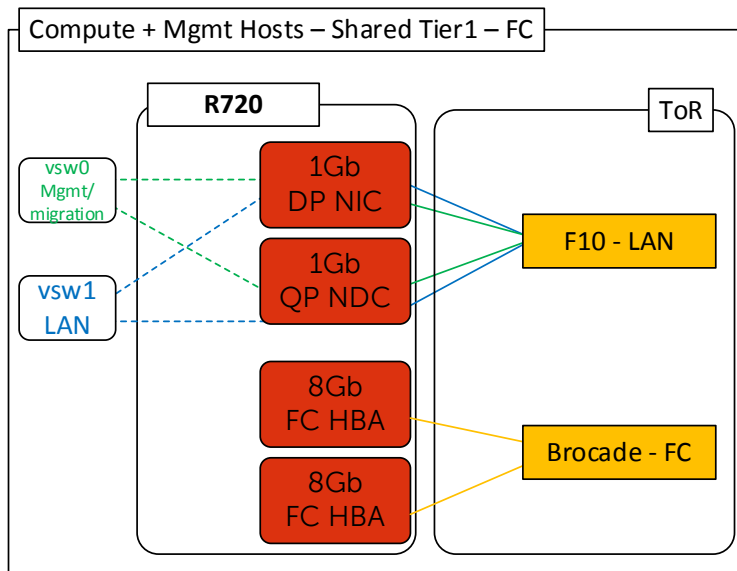
Using Fiber Channel based storage eliminates the need to build iSCSI into the network stack but requires additional fabrics to be built out. The network configuration in this model is identical between the Compute and Management hosts. Both need access to FC storage since they are hosting VDI sessions from shared storage and both can leverage Live Migration as a result 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
 - 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 can be configured for all hardware management traffic – L3 routed via core switch

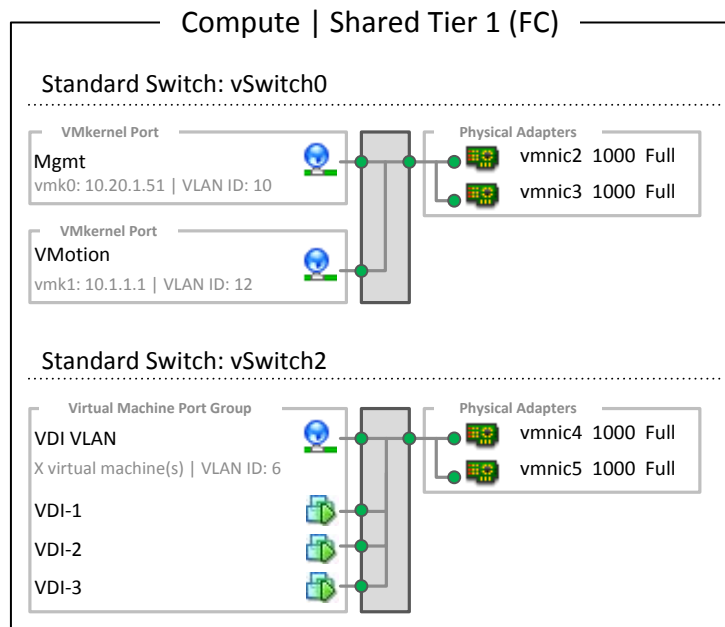


FC and LAN traffic are physically separated into discrete switching fabrics. Each Shared Tier 1 Compute and Management host will have a quad port NDC (4 x 1Gb), a 1Gb dual port NIC, as well as 2 x 8Gb dual port FC HBAs. Connections from both vSwitches should pass through both the NDC and add-on NICs per the diagram below. Configure the LAN traffic from the server to the ToR switch as a LAG.

5.5.3.1 vSphere

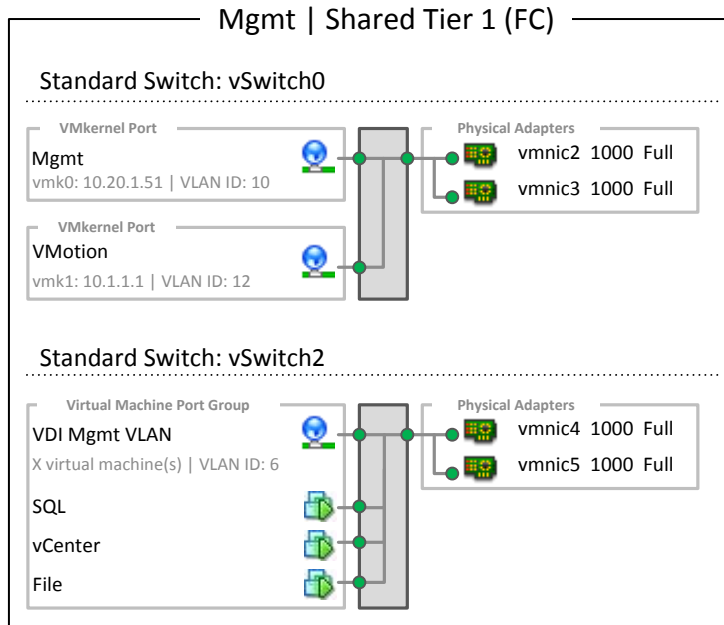


vSwitch0 carries traffic for both Management and Live Migration which needs to be VLAN-tagged so that either NIC can serve traffic for either VLAN. The Management VLAN will be L3 routable while the Live Migration VLAN will be L2 non-routable.



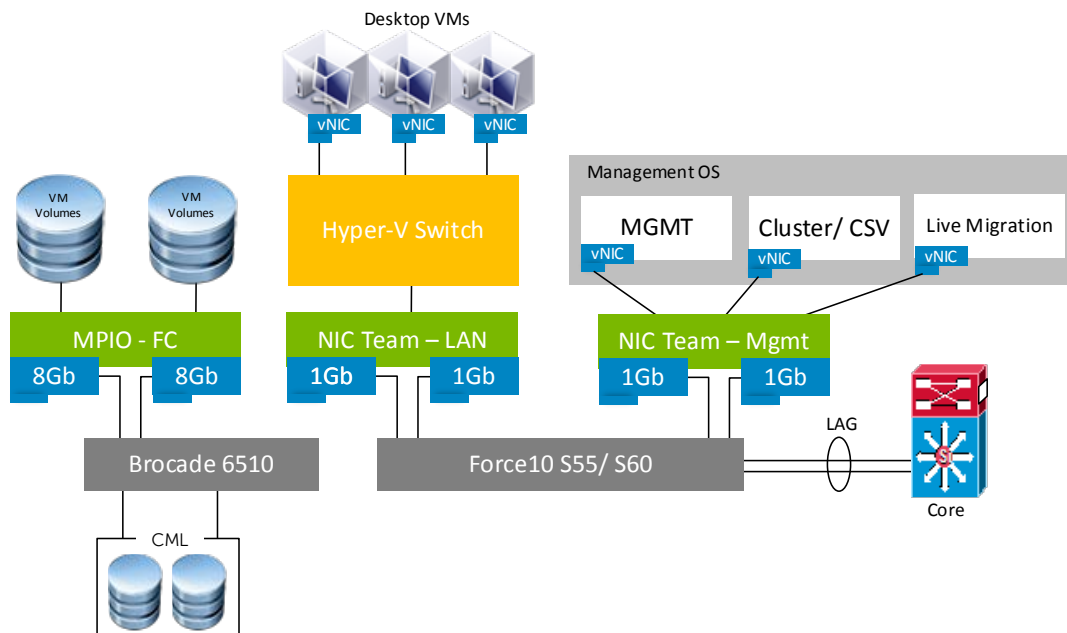
The Management server is configured identically except for the VDI Management VLAN which is fully routed but should be separated from the VDI VLAN used on the Compute host.





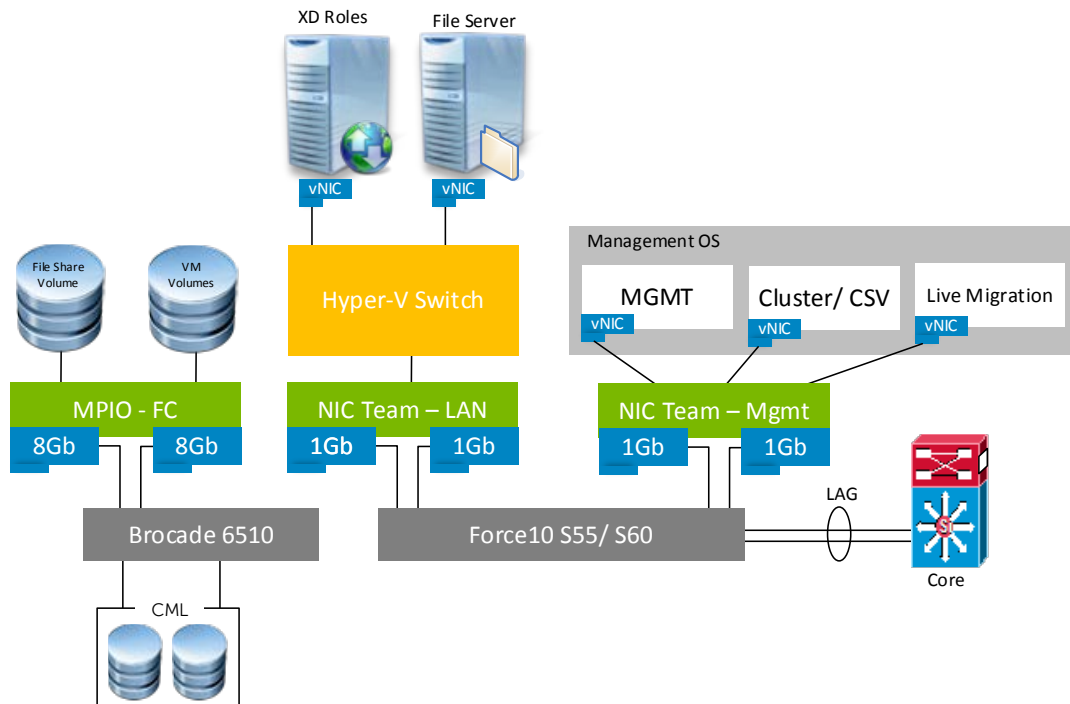
5.5.3.2 Hyper-V

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 should be 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.





5.5.4 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:

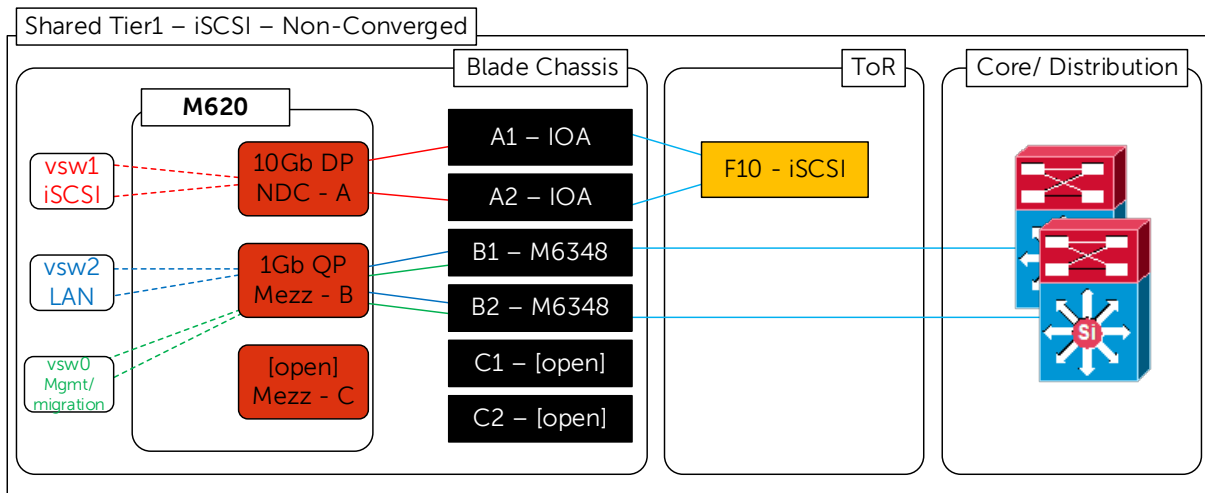
- Compute hosts (Shared Tier 1)
 - Management VLAN: Configured for hypervisor Management traffic – L3 routed via core switch
 - Live Migration VLAN: Configured for Live Migration traffic – L2 switched only, trunked from Core
 - iSCSI VLAN: Configured for iSCSI traffic – L2 switched only via ToR switch
 - 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 can be configured for all hardware management traffic – L3 routed via core switch

Following best practices, iSCSI and LAN traffic will be 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 will be left open for future expansion. Connections from 10Gb and 1Gb

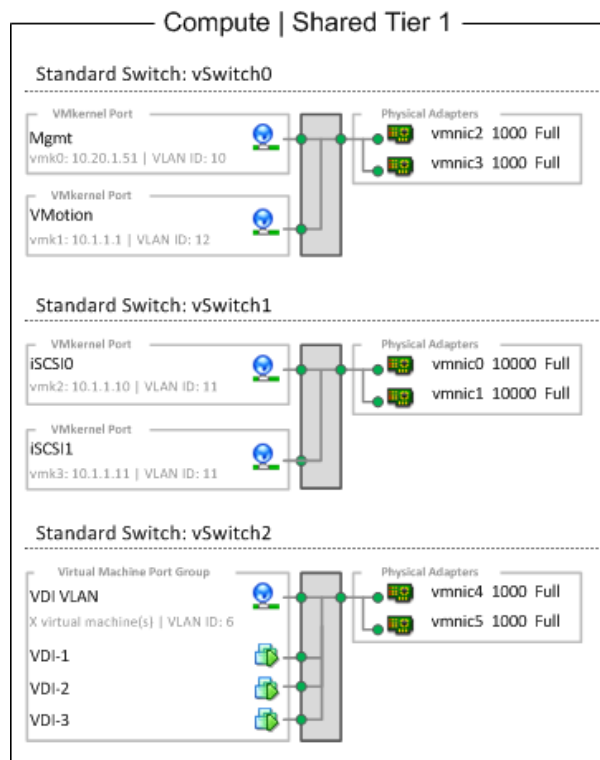


traffic vSwitches should pass through the blade mezzanines and interconnects per the diagram below. Configure the LAN traffic from the server to the ToR switch as a LAG if possible.

5.5.4.1 vSphere

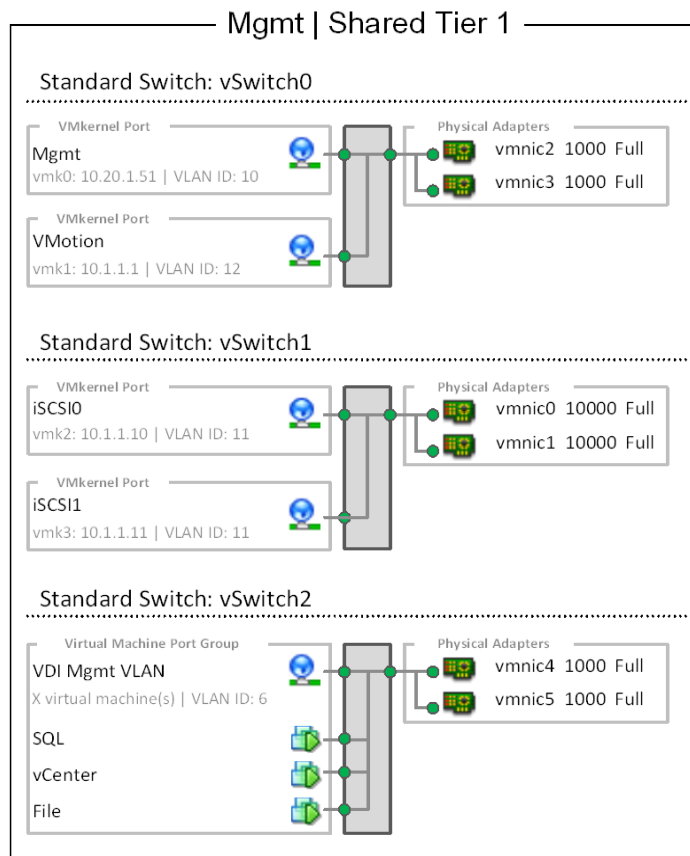


vSwitch0 carries traffic for both Management and Live Migration which needs to be VLAN-tagged so that either NIC can serve traffic for either VLAN. The Management VLAN will be L3 routable while the Live Migration VLAN will be L2 non-routable.



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

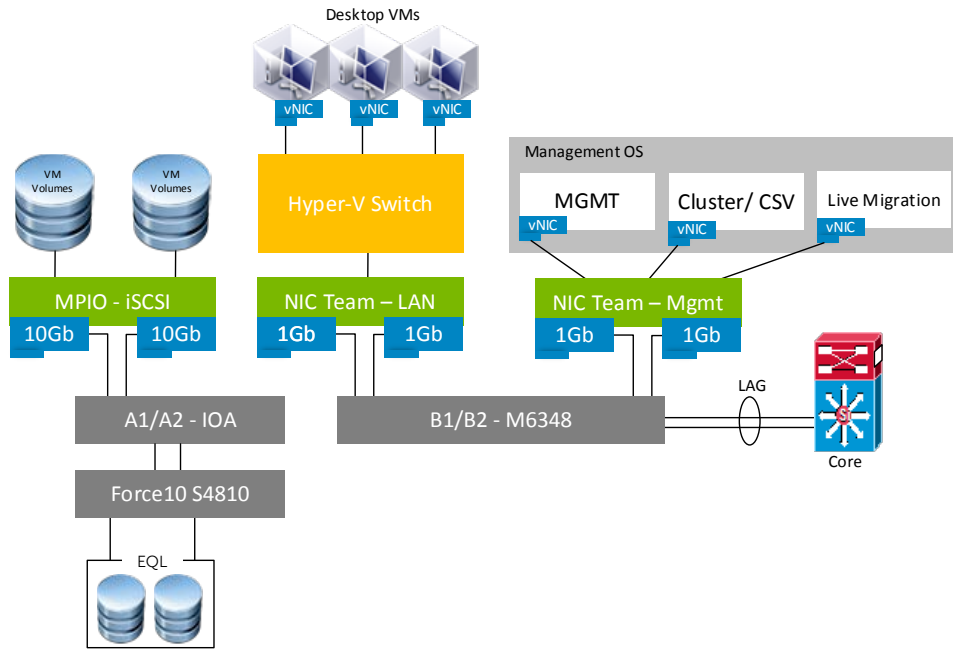




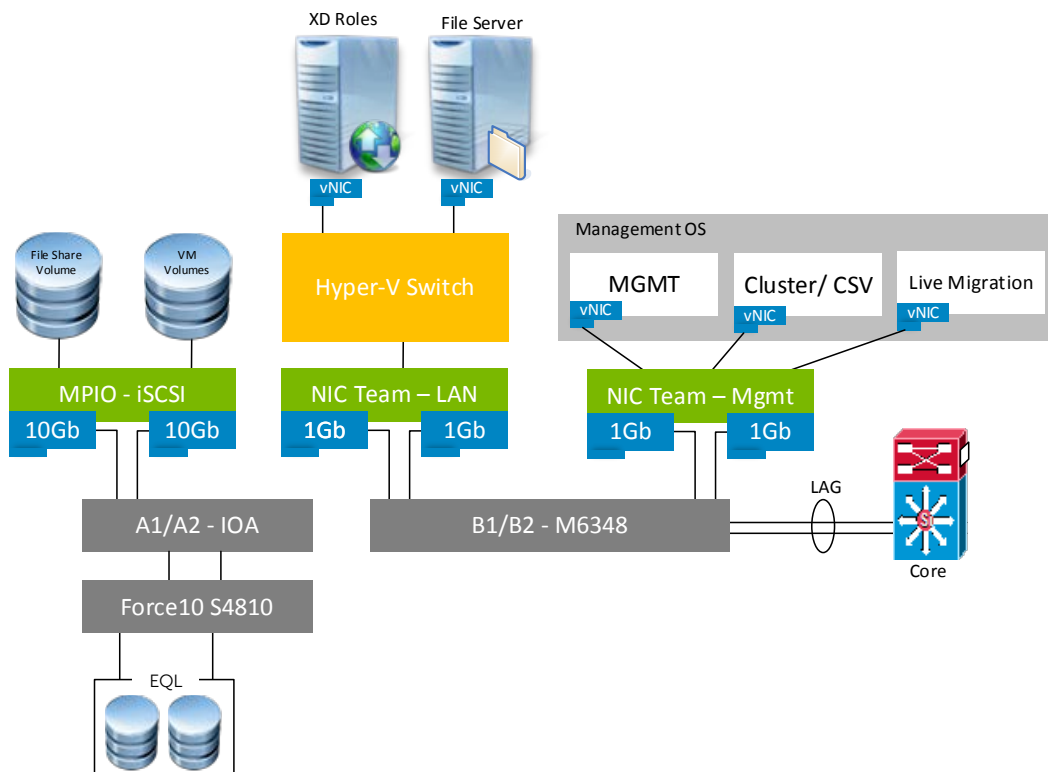
5.5.4.2 Hyper-V

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 should be 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.



5.5.5 Shared Tier 1 – Blade – FC

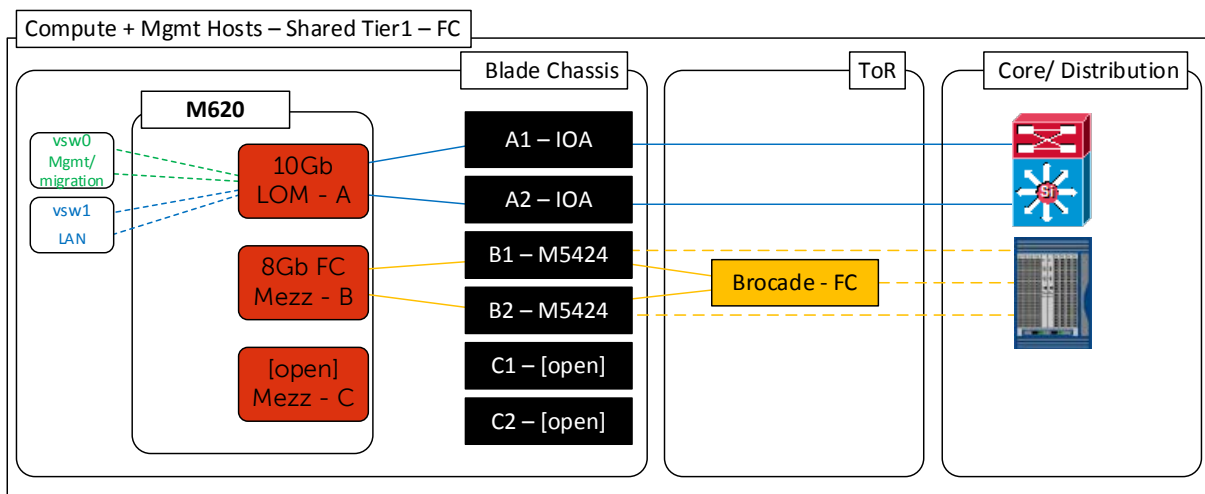
Using Fiber Channel based storage eliminates the need to build iSCSI into the network stack but requires additional fabrics to be built out. The network configuration in this model is identical between the Compute and Management hosts. The following outlines the VLAN requirements for the Compute and 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 can be 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 will be 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.

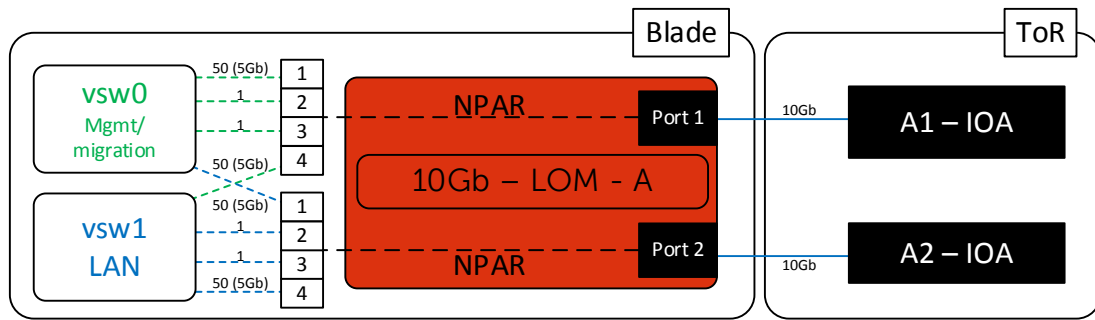
5.5.5.1 vSphere



5.5.5.1.1 Shared Tier 1 – Blade – Network Partitioning

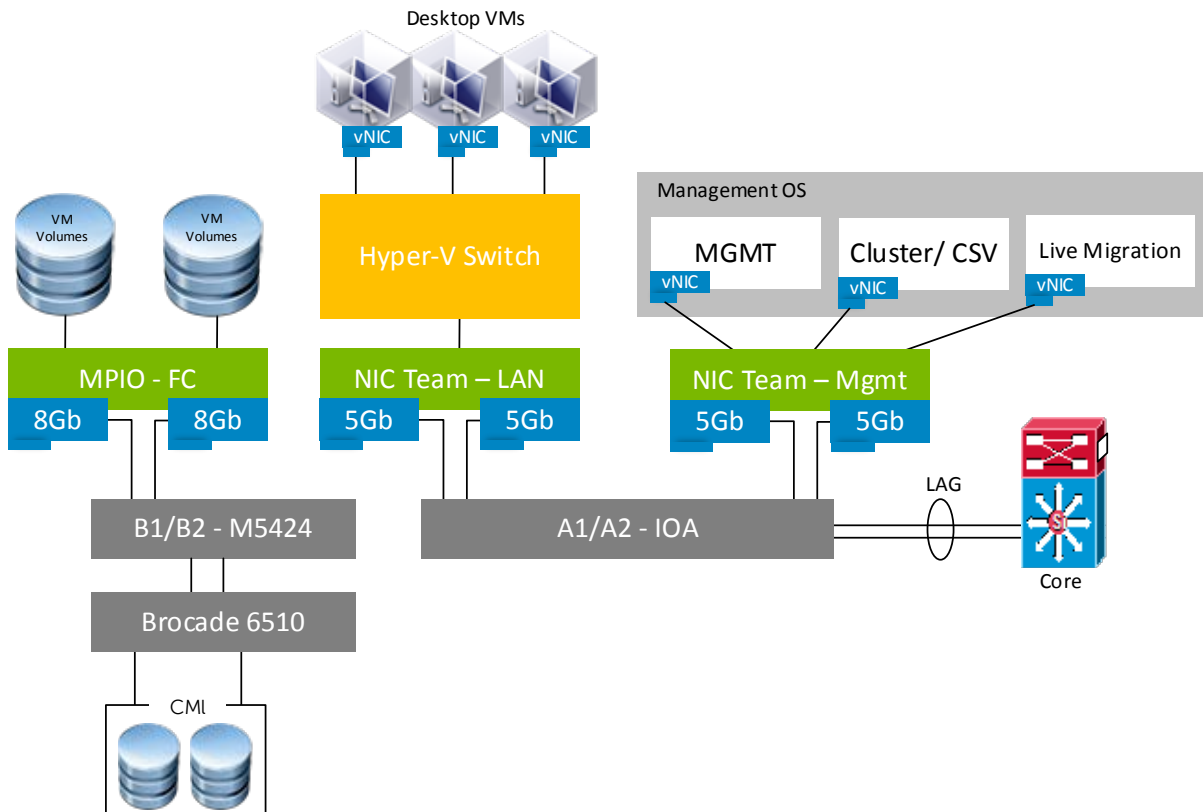
Network partitioning (NPAR) takes place within the UEFI of the 10GB LOMs of each blade in the A fabric. Partitioning allows a 10Gb NIC to be split into multiple logical NICs that can be assigned differing amounts of bandwidth. 4 partitions are defined per NIC with the amounts specified below. We only require 2 partitions per NIC port so the unused partitions receive a bandwidth of 1. Partitions can be oversubscribed, but not the reverse. We will be 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.





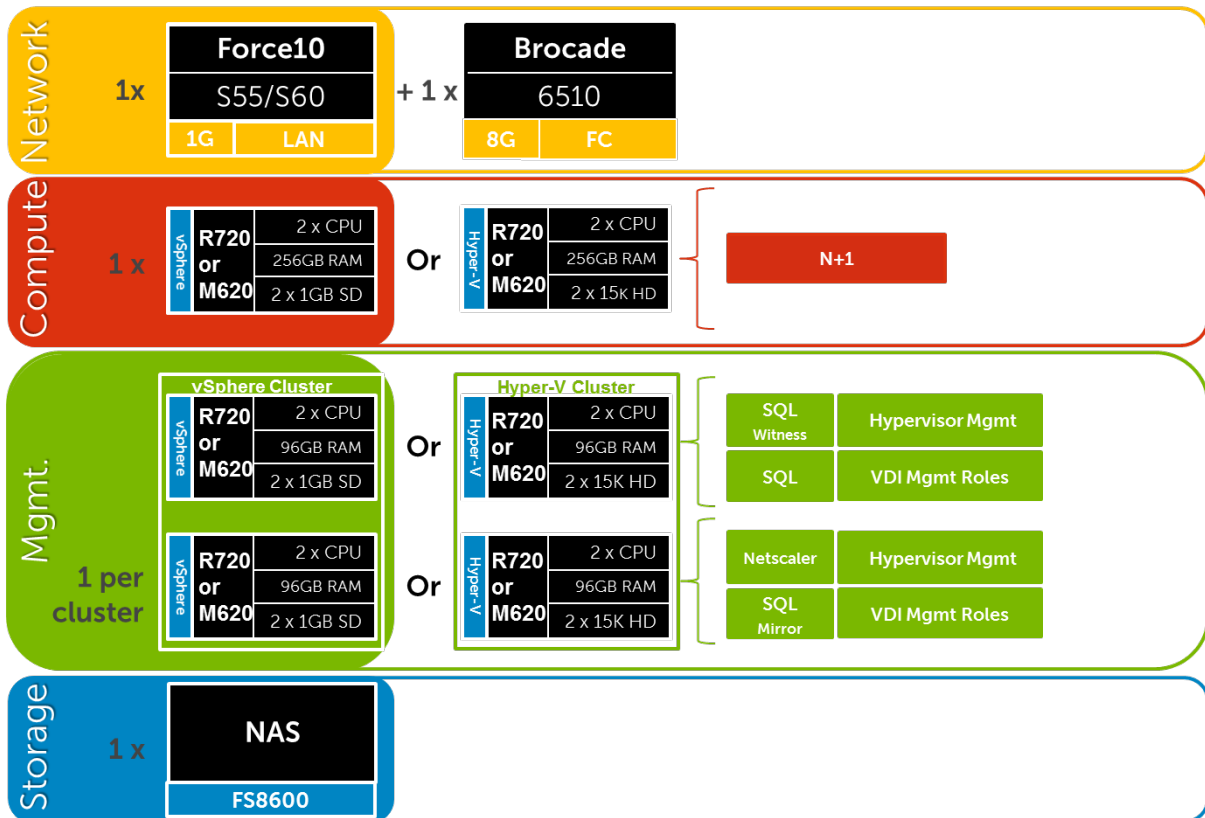
5.5.5.2 Hyper-V

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 should be 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 will be 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.



5.6 Solution High Availability

High availability (HA) is offered to protect each layers of the solution architecture, 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 mgmt 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 can be used to host file shares. Storage protocol switch stacks and NAS selection will vary based on chosen solution architecture.

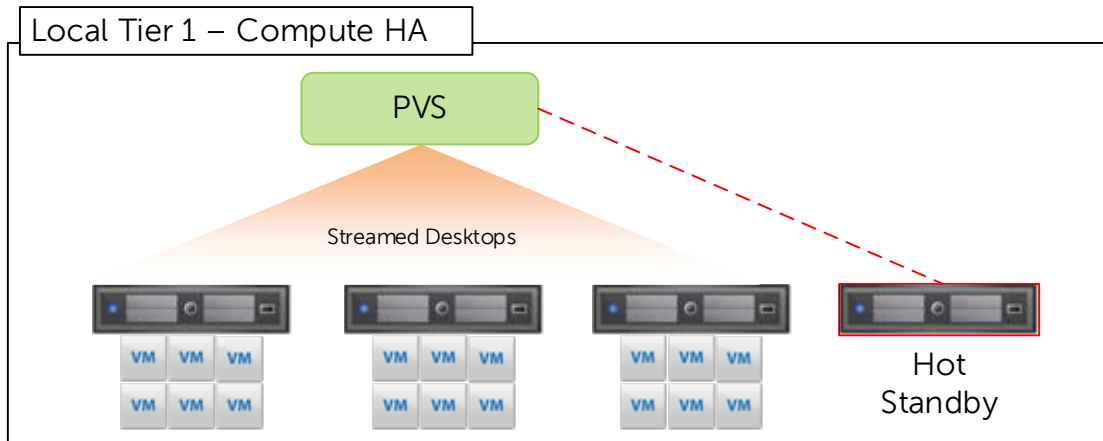


The HA options provides redundancy for all critical components in the stack while improving the performance and efficiency of the solution as a whole.

- An additional switch is added at the network tier which will be configured with the original as a stack and equally spreading each host's network connections across both.
- 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 will be no vSphere HA cluster in the compute tier as VMs that run here 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 will then be configured in an HA cluster. All applicable Citrix server roles can then be duplicated on the new host where connections to each will be load balanced via the addition of a virtual NetScaler appliance. SQL will also receive greater protection through the addition and configuration of a SQL mirror with a witness.

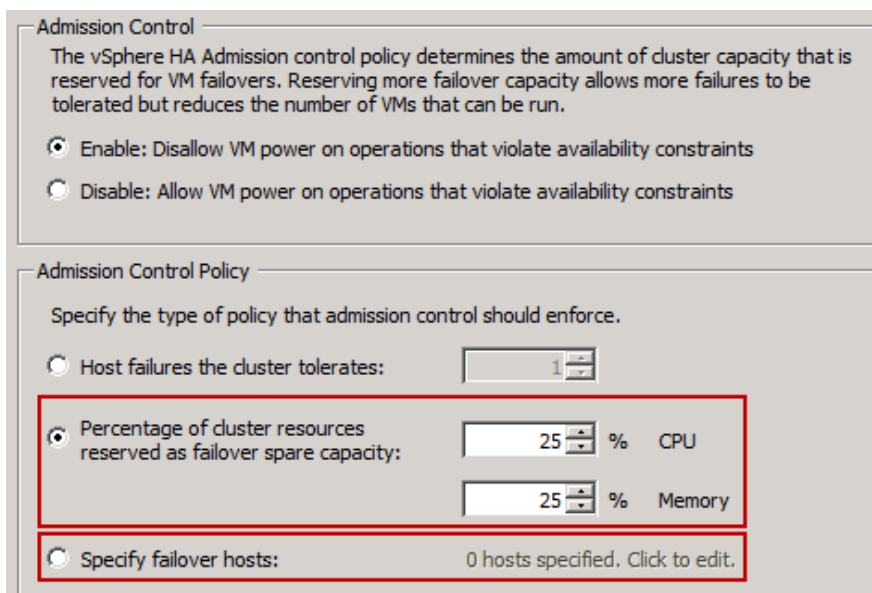
5.6.1 Compute Layer HA (Local Tier 1)

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 will be 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 can be configured in an HA cluster with or without the HA bundle. An extra server in the Management layer will provide protection should a host fail.

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



5.6.2 vSphere HA (Shared Tier 1)


Both Compute and Management hosts are identically configured, within their respective tiers, and leverage shared storage so can make full use of vSphere HA. The Compute hosts can be configured in an HA cluster following the boundaries of vCenter with respect to limits imposed by VMware (3000 VMs per vCenter). This will result in multiple HA clusters managed by multiple vCenter servers.



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

5.6.3 Hyper-V HA (Shared Tier 1)

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



5.6.4 Management Server High Availability

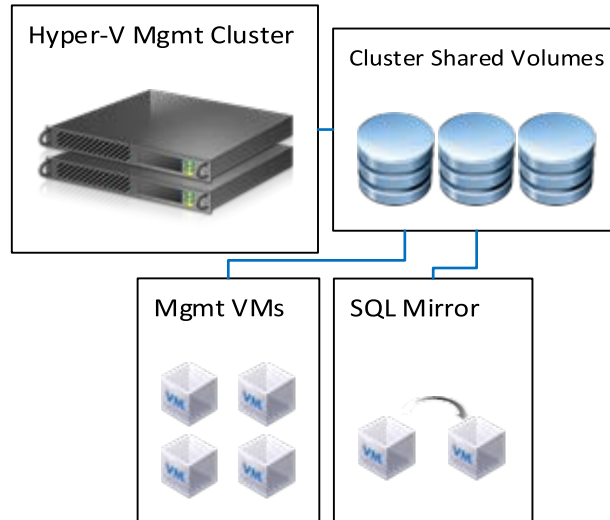
The applicable core Citrix roles will be load balanced via DNS by default. In environments requiring HA, NetScaler VPX will be introduced to manage load-balancing efforts. XenDesktop, PVS, and vCenter configurations are stored in SQL which will be protected via the SQL mirror.

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



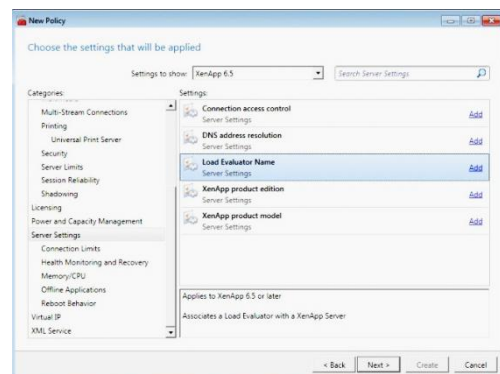
The following will protect each of the critical infrastructure components in the solution:

- The Management hosts will be configured in a Hyper-V cluster (Node and Disk Majority).
- The storage volume that hosts the Management VMs will be 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 can be 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 will be 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 will be 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 can be 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 will be 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 will be responsible for the Write cache of each desktop that it is hosting. This Write cache will be 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 will be unaware of the failure.

Streaming Services run under a user account with Service account credentials. The Service account credentials (user account name and password) will be 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 will be provided by the Dell FS7600, FS8600 or PowerVault NX3300 clustered NAS devices. To ensure proper redundancy, distribute the NAS cabling between ToR switches.

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

The DVS recommendation is to configure the original file server VM to use RDMs or PTDs to access the storage LUNs, therefore migration to the NAS will be 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 will be provided via a 3-server synchronous mirror configuration that includes a witness (High safety with automatic failover). This configuration will protect all critical data stored within the database from physical server as well as virtual server problems. DNS will be used to control access to the active SQL server, please refer to section 5.7.1 for more details. 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](#)

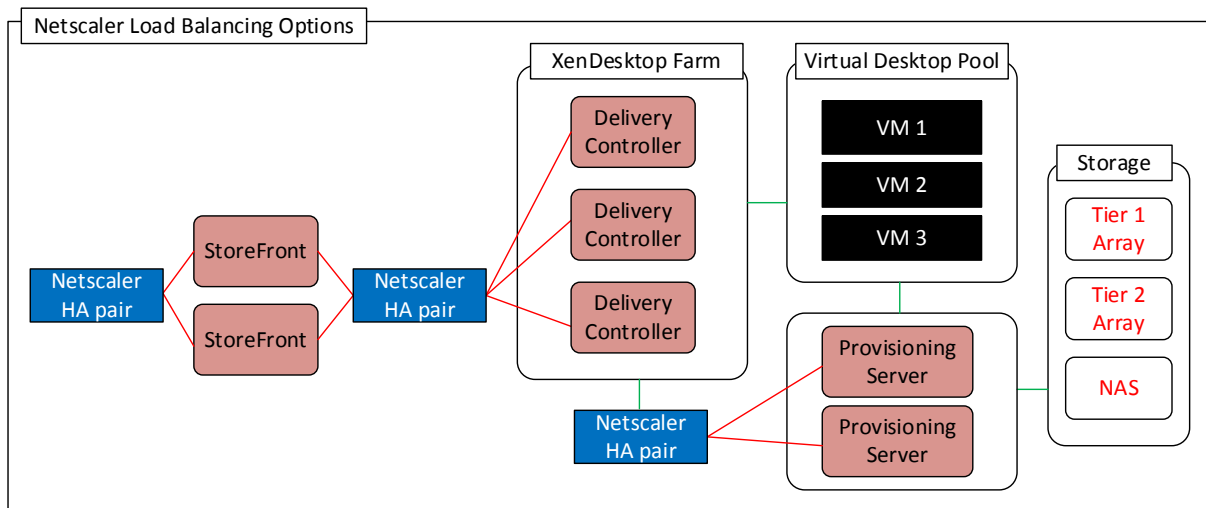
Additional resources can be found in TechNet: [LINK1](#) and [LINK2](#)

5.6.9 Load Balancing

Depending on which management components are to be made highly available, the use of a load balancer may be required. 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

Dell recommends the Citrix Netscaler for load balancing the DVS Enterprise Citrix solution.



5.6.9.1 DNS for Load Balancing

When considering DNS for non SQL-based components such as Citrix Storefronts or file servers, where a load balancing behavior is desired, invoke the native DNS round robin feature. To invoke round robin, enter the resource records for a service into DNS as A records with the same name.

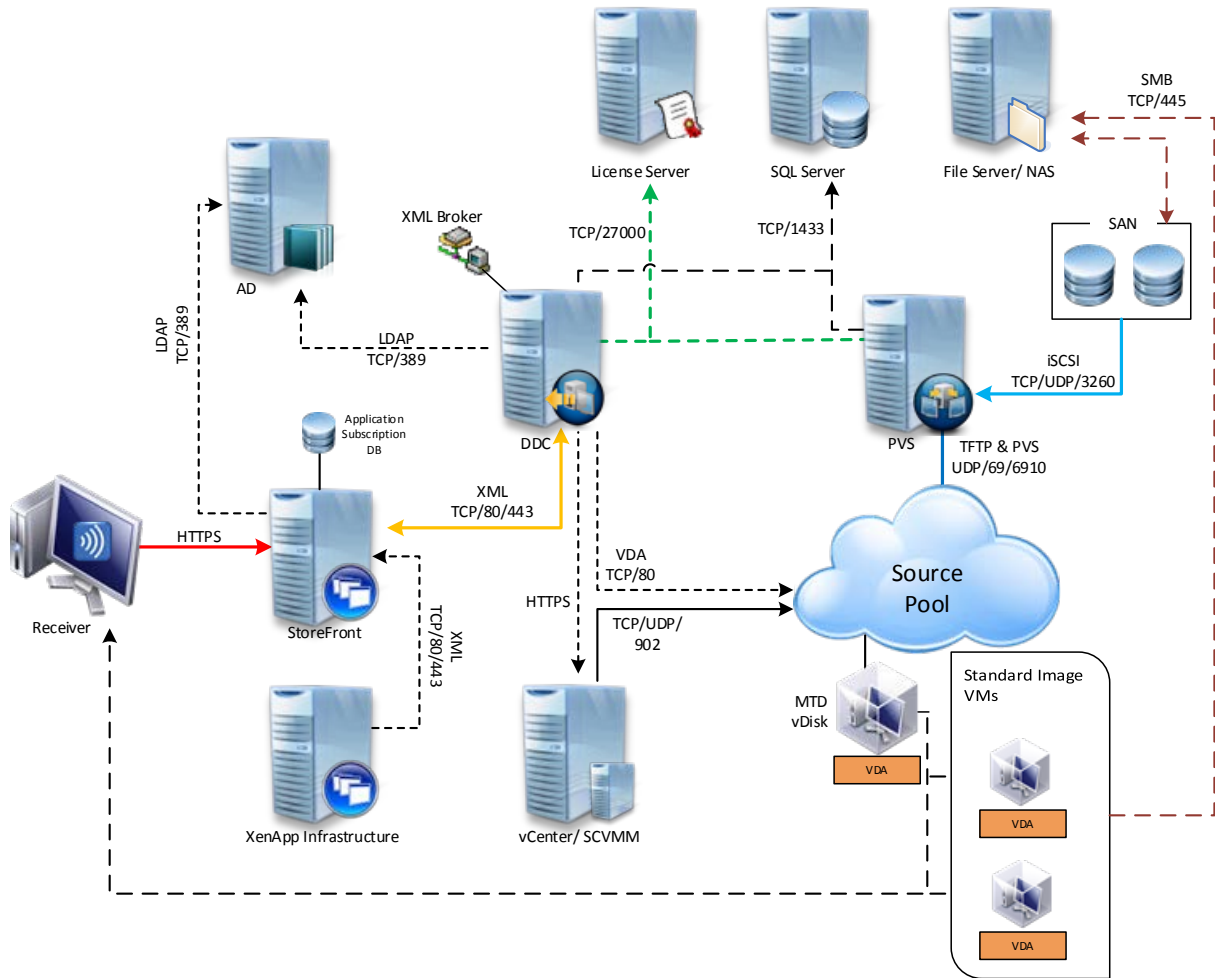
For example, in the base configuration the single Storefront server will have its own hostname registered in DNS as an A record. Create a new A record to be used should additional WI's come online or be retired for whatever reason. This creates machine portability at the DNS layer to remove the importance of actual server hostnames. The name of this new A record is unimportant but must be used as the primary name record to gain access to the resource, not the server's host name! In this case I have created three new A records called "WebInterface", all presumably pointing to three different servers.

When a client requests the name Web Interface, DNS will direct them to the 3 hosts in round robin fashion. The following resolutions were performed from 2 different clients. Repeat this method of creating an identical but load-balanced namespace for all applicable components of the architecture stack.



Name	Type	Data
WebInterface	Host (A)	10.1.1.1
WebInterface	Host (A)	10.1.1.2
WebInterface	Host (A)	10.1.1.3

5.7 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 DVS Enterprise solution, the following minimum hardware requirements must be met.

Feature	Minimum Requirement	Notes
Total Tier 2 Storage Space	User count and workload dependent	
Tier 2 Drive Support	7200rpm NL SAS	
Tier 1 IOPS Requirement	(Total Users) x 10 IOPS	
Tier 2 IOPS Requirement	(Total Users) x 1 IOPS	
Data Networking	1GbE Ethernet for LAN 10GbE Ethernet for T1 iSCSI	Up to 6 1Gb ports per host are required for Shared T1 scenarios.
Shared Array Controllers	1 with >4GB cache	4GB of cache minimum per controller is recommended for optimal performance and data protection.
RAID Support	10, 6	RAID 10 is leveraged for local Compute host storage and high performance shared arrays. Protect T2 storage via RAID 6.

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 DVS Enterprise solution, the following minimum hardware requirements must be met.

Feature	Minimum Requirement	Notes
Switching Capacity	Line rate switch	
10Gbps Ports	Uplink to Core Shared T1 iSCSI/ Converged	DVS Enterprise leverages 10Gb for Shared Tier 1 solution models. 10Gb uplinks to the core network are also recommended.
1Gbps Ports	5x per Management server 5x per Compute Server 6x per Storage Array	DVS Enterprise 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 End-user Workload Characterization

It's important to understand the user workloads when designing a Desktop Virtualization Solution. The Dell Desktop Virtualization Solution methodology includes a Blueprint process to assess and categorize a customer's environment according to the workloads defined in this section. In the Dell Desktop Virtualization solution this will map directly to the SLA levels we offer in our Integrated Stack. There are three levels, each of which is bound by specific metrics and capabilities.

7.1 Workload Characterization Overview

7.1.1 Basic Workload Characterization

The Basic User workload profile consists of simple task worker workloads. Typically a repetitive application use profile with a non-personalized virtual desktop image. Sample use cases may be a kiosk or call-center use cases which do not require a personalized desktop environment and the application stack is static. In a virtual desktop environment the image is dynamically created from a template for each user and returned to the desktop pool for reuse by other users. The workload requirements for a basic user is the lowest in terms of CPU, memory, network and Disk I/O requirements and will allow the greatest density and scalability of the infrastructure.

User Workload	VM vCPU	VM Memory Allocation	VM Memory Reservation	Approx. IOPS	VDI Session Disk Space	OS Image Notes
Basic	1	2GB	1GB	7-8	3GB	This user workload leverages a shared desktop image emulates a task worker. Only two apps are open simultaneously and session idle time is approximately one hour and forty-five minutes.

7.1.2 Standard Workload Characterization

The Standard User workload profile consists of email, typical office productivity applications and web browsing for research/training. There is minimal image personalization required in a standard user workload profile. The workload requirement for a Standard User is moderate and most closely matches the majority of office worker profiles in terms of CPU, memory, network and Disk I/O. This will allow moderate density and scalability of the infrastructure.

User Workload	VM vCPU	VM Memory Allocation	VM Memory Reservation	Approx. IOPS	VDI Session Disk Space	OS Image Notes
Standard	2	3GB	1.5GB	9-10	3.75GB	This user workload leverages a shared desktop image emulates a medium knowledge worker. Five applications are open simultaneously and session idle time is approximately 45 seconds.



7.1.3 Premium Workload Characterization

The Premium User workload is an advanced knowledge worker. All office applications are configured and utilized. The user has moderate-to-large file size (access, save, transfer requirements). There is some graphics creation or editing done for presentations or content creation tasks. Web browsing use is typically research/training driven, similar to Standard Users. The Premium User requires extensive image personalization, for shortcuts, macros, menu layouts and so on. The workload requirements for a Premium User are heavier than typical office workers in terms of CPU, memory, Network and Disk I/O. This will limit density and scalability of the infrastructure.

User Workload	VM vCPU	VM Memory Allocation	VM Memory Reservation	Approx. IOPS	VDI Session Disk Space	OS Image Notes
Premium	2	4GB	2GB	10-12	6GB	This user workload leverages a shared desktop image emulates a high level knowledge worker. Eight applications are open simultaneously and session idle time is approximately two minutes.

7.1.4 Workload Characterization Testing Details

User Workload	VM Memory	OS Image	Workload Description
Basic	2GB	Shared	<p>This workload emulates a task worker.</p> <ul style="list-style-type: none"> The light workload is very light in comparison to medium. Only 2 apps are open simultaneously. Only apps used are IE, Word and Outlook. Idle time total is about 1:45 minutes
Standard	3GB	Shared	<p>This workload emulates a medium knowledge working using Office, IE and PDF.</p> <ul style="list-style-type: none"> Once a session has been started the medium workload will repeat every 12 minutes. During each loop the response time is measured every 2 minutes. The medium workload opens up to 5 apps simultaneously. The type rate is 160 ms for each character. Approximately 2 minutes of idle time is included to simulate real-world users. <p>Each loop will open and use:</p> <ul style="list-style-type: none"> Outlook 2007, browse 10 messages. Internet Explorer, one instance is left open (BBC.co.uk), one instance is browsed to Wired.com, Lonelyplanet.com and heavy flash app gettheglass.com. Word 2007, one instance to measure response time, one instance to review and edit document. Bullzip PDF Printer & Acrobat Reader, the word document is printed and reviewed to PDF. Excel 2007, a very large randomized sheet is opened.



			<ul style="list-style-type: none"> • PowerPoint 2007, a presentation is reviewed and edited. • 7-zip: using the command line version the output of the session is zipped.
Premium	4GB	Shared plus Profile Virt, or, Private	<p>The heavy workload is based on the standard workload; the differences in comparison to the standard workload are:</p> <ul style="list-style-type: none"> • Type rate is 130 ms per character. • Idle time total is only 40 seconds. • The heavy workload opens up to 8 apps simultaneously



8 Solution Performance and Testing

8.1 Load Generation and Monitoring

8.1.1 Login VSI – 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.

8.1.2 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.) ?	20	0 <= 15	<= 60	<= unbounded	
Application Load Time : Avg. startup time for applications (sec.) ?	20	0 <= 10	<= 30	<= unbounded	
CPU Queue Length : Length of CPU queue at inspection time ?	20	0 <= 3	<= 6	<= unbounded	
Page Faults : Number of page faults during inspection interval ?	20	0 <= 2,000	<= 10,000	<= unbounded	
Non-Responding Applications : Number of unresponsive applications at inspection time ?	20	0 <= 2	<= 3	<= unbounded	

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



8.1.3 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 will detail 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)

8.1.4 VMware vCenter

VMware vCenter has been 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.

8.1.5 Microsoft Perfmon

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

8.2 Testing and Validation

8.2.1 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 4 runs. A pilot run to validate that the infrastructure is functioning and valid data can be captured and 3 subsequent runs allowing correlation of data. Summary of the test results will be 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 will be subjective, it will help provide a better understanding of the end user experience of the desktop sessions, particularly under high load, and ensure that the data gathered is reliable.

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 6 desktop hosts were conducted using the Login VSI "Parallel Launch" mode, 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.

8.3 XenDesktop Test Results

This validation was designed to evaluate the capabilities of the Ivy Bridge processors when used in a XenDesktop 7 environment with Windows 8. The E5-2690V2v2 processor was used in both rack and blade servers. The XenDesktop solutions were deployed on ESXi 5.1 U1 Hypervisor and Microsoft Hyper-V 2012.

This validation was performed on single server compute host solutions, running on Dell R720 and Dell M620 hosts. The compute hosts had 256GB of RAM, and dual E5-2690V2 v2 3.0GHz 10 core processors. Using Ivy Bridge processors the M620 can support the same processors as the R720 which was not true for the previous generation Sandy Bridge processors.

The Hyper-V results presented following were all gathered from PVS deployed XenDesktop solution. The ESXi results were gathered from MCS deployed solution. All the results presented below were performed on Shared tier1 storage. In the case of Hyper-v the shared storage platform was the Dell EqualLogic PS6110XS iSCSI array. In the ESXi solution the storage platform was the Dell Compellent SC8000 array connected via Fiber Channel. During the validation effort the latency on the storage was always observed to be less than 5 ms for both arrays.

Validation was performed using DVS standard testing methodology using LoginVSI load generation tool for VDI benchmarking that simulates production user workloads. The Windows 8 VMs were configured for memory and CPU as follows. It was noted that setting a value lower than 1GB for start-up memory on Hyper-V caused many of the Citrix Receiver sessions to fail.

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

User Workload	vCPUs	Hyper-V Start up Memory	Hyper-V Minimum Memory	Hyper-V Max Memory	ESXi Memory Reservation	ESXi memory configured
Basic User	1	1GB	1GB	2GB	1GB	2GB
Standard User	2	1.5GB	1GB	3GB	1.5GB	3GB
Premium User	2	1.5GB	1GB	4GB	1.5GB	4GB

Because of poor results shown for CPU queuing in previous tests with single CPU, 2 x vCPUs were assigned for the Premium and Standard workloads.

As a result of the testing, the following density numbers can be applied to the individual solutions. In all cases CPU percentage used was the limiting factor. Memory usage, IOPs and network usage were not strained.



The following table summarizes the user workload resources and densities as tested for **ESXi** and PVS on Windows 8 desktops:

User Workload	Max Users	Max CPU	IOPs per user
Basic User	190	92%	10
Standard User	115	84%	12
Premium User	105	81%	14

The following table summarizes the user workload resources and densities as tested for **Hyper-V** and PVS on Windows 8 desktops:

User Workload	Max Users	Max CPU	IOPS per user
Basic User	200	83%	9
Standard User	150	85%	13
Premium User	125	82%	16

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

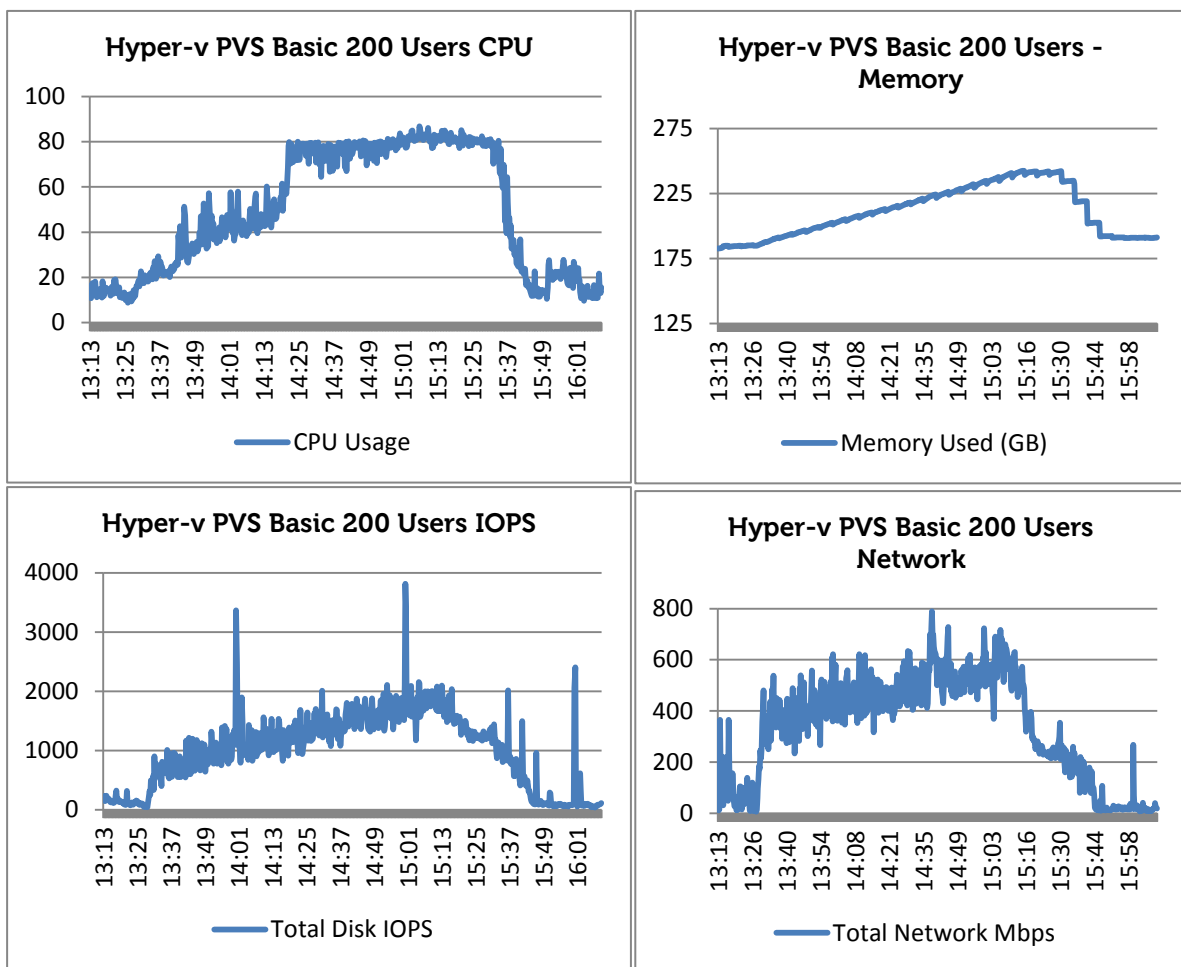


8.3.1 Windows Server 2012 Hyper-V

Recent VDI density testing using Windows 8 on Windows Server 2012 Hyper-V has generally shown a higher number of desktops than other hypervisors on the same hardware configurations. This trend has continued with the Ivy Bridge testing. However the density numbers obtained with ESXi are much closer to Hyper-V than in previous testing. Also, although Hyper-V results show very favorable CPU usage it was seen that the max CPU observed in many tests was reached long before all the VDI users were logged in during the test. In addition the Stratusphere UX results for Hyper-V (even with low CPU use) are not exceptional.

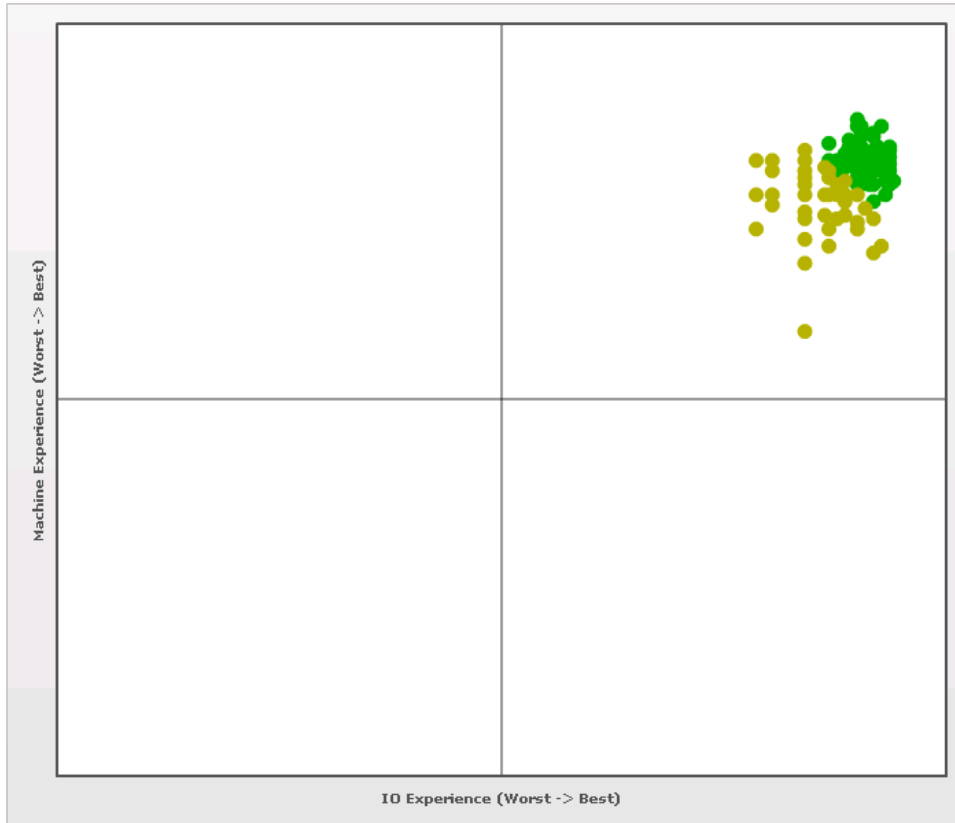
1.1.1 Basic User Workload (200 Users)

These graphs show CPU, memory, local disk IOPS, network and VDI UX scatter plot results. In addition, for this workload only the graph showing the Hyper-V logical processor and Virtual processor side by side is shown. A high logical processor % runtime combined with a low virtual processor % runtime is typical of an environment where there are more processors allocated to VMs than are physically available on the compute host, which is the case for this VDI environment.


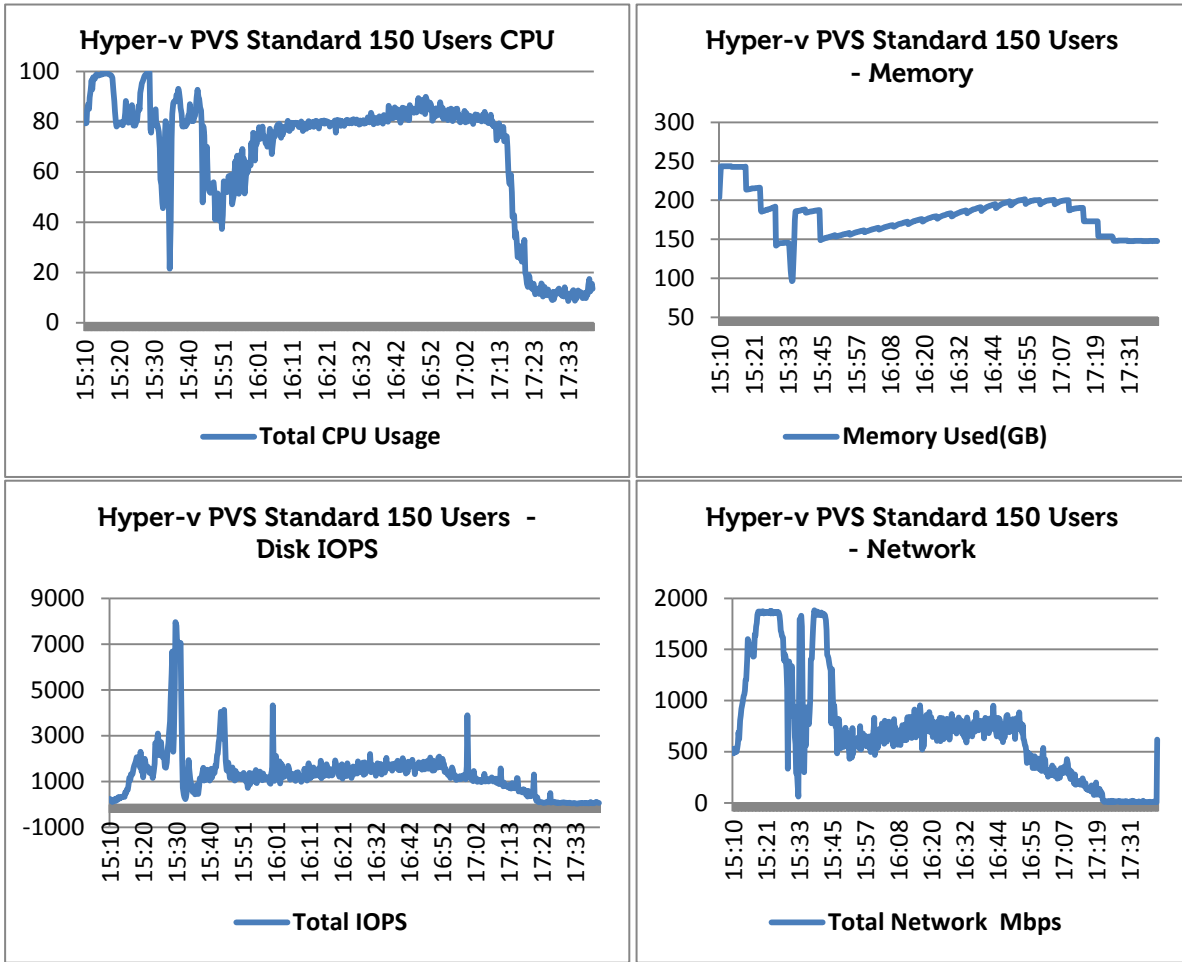


The Login VSI tests normally continue for 960 seconds before the sessions start to logoff. In the CPU chart above it can be seen that the 80% CPU was reached before all the sessions logged in but did not rise further as more sessions logged in. This pattern is repeated in other Hyper-V tests for standard and Premium workloads.







1.1.1 Standard User Workload (150 Users)



1.1.2 Premium User Workload (125 Users)




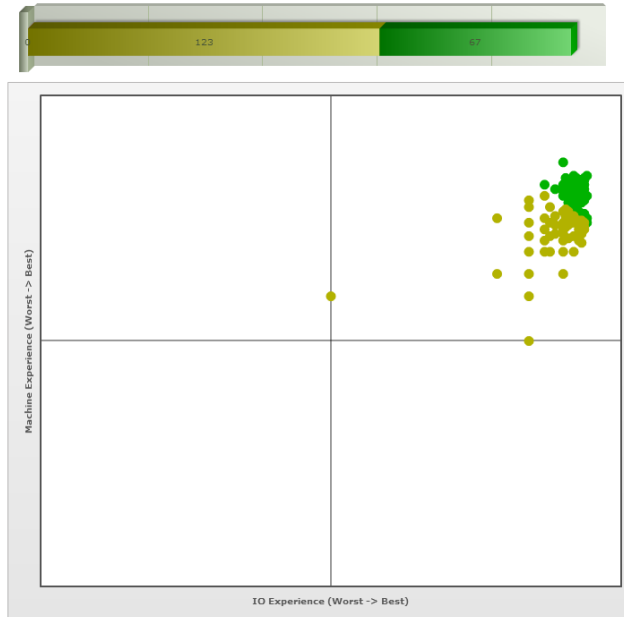
8.3.2 vSphere 5.1 Update 1

The R720 servers were connected to Compellent fiber channel storage which was used tier1 and tier2 storage. The density numbers achieved on ESXi were not as good as those obtained on Hyper-V. However, Stratusphere results and subjective tests produced similar results.

1.1.1 Basic User Workload (190 Users)


In these results CPU reaches 92% as all the users log on but falls off quickly. From the Stratusphere it can be seen that many of the users suffered long logon delays. This was traced to the fact that no domain users were logged into the Golden image prior to sealing the image and deploying the Desktop group. Subsequent runs for Standard and Premium workloads show improved Stratusphere results.


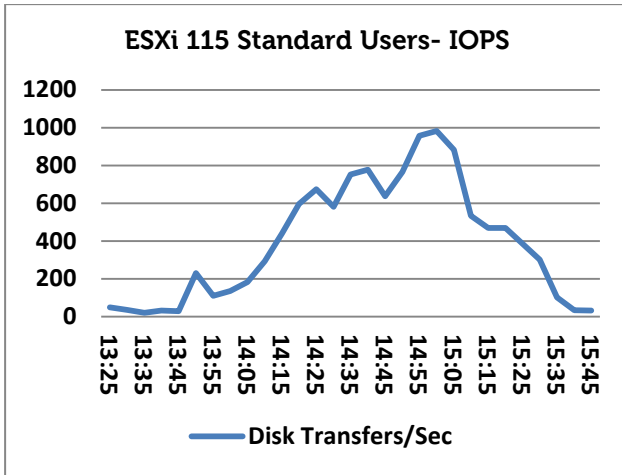




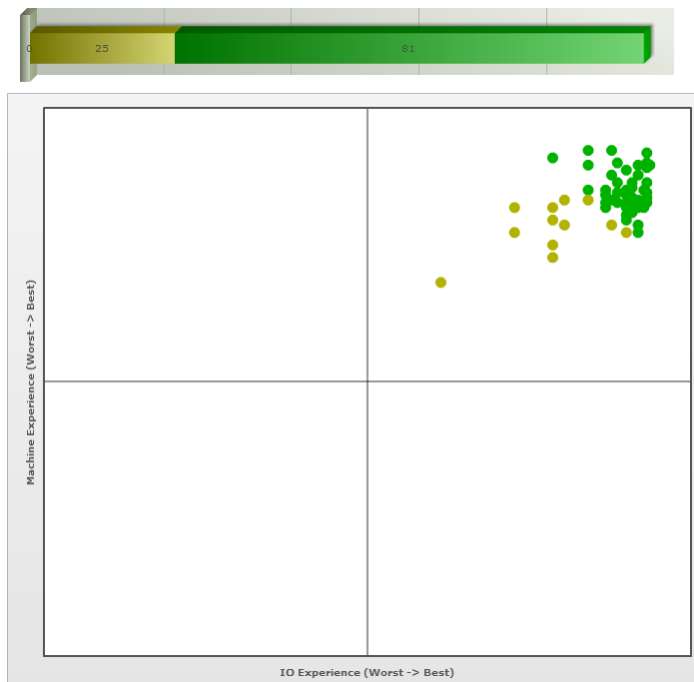
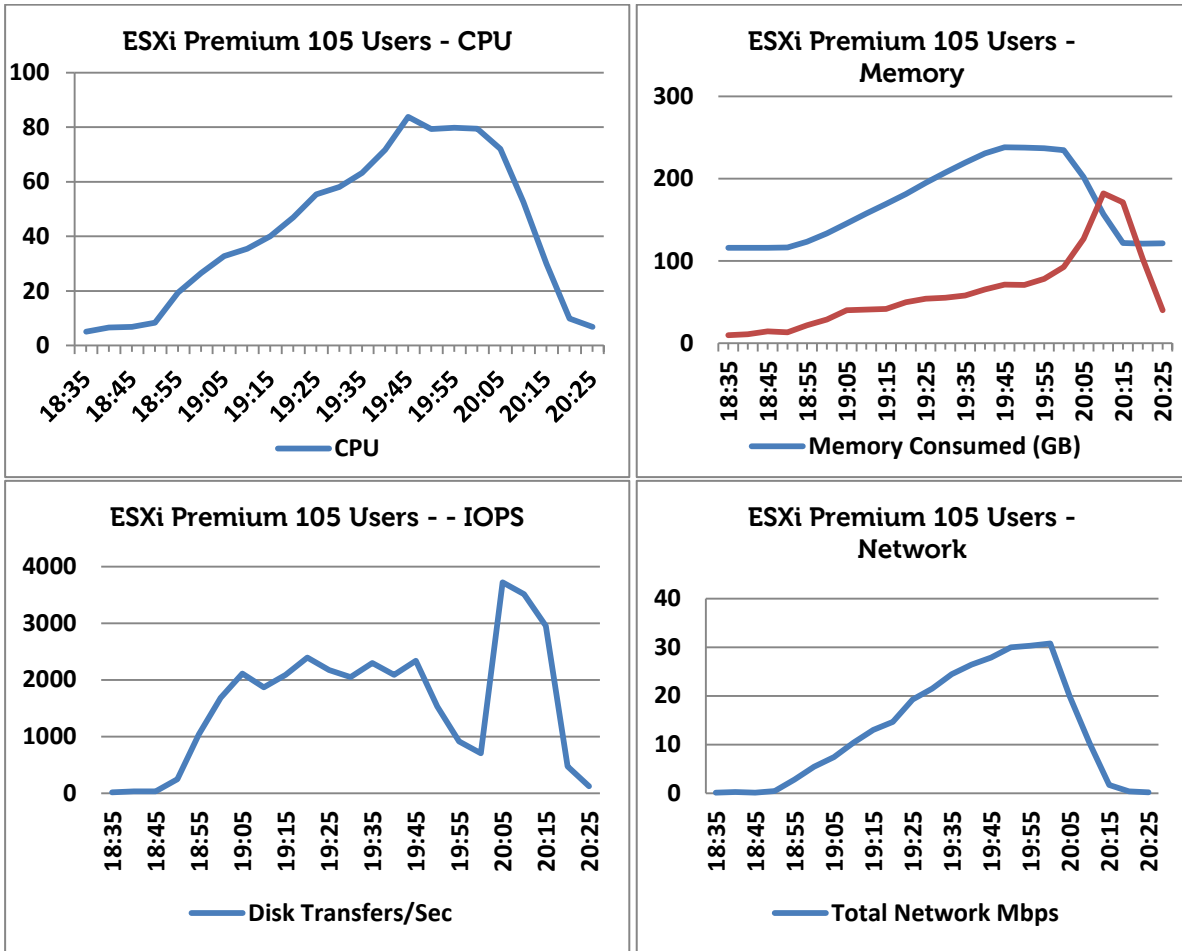
1.1.2 Standard User Workload (115 Users)

Presented below are 2 sets of results for a Standard workload on ESXi. The first set of results (120 users) is gathered when the test is run immediately after a host reboot and the VMs are started. The second set of results was gathered when the VMs were simply rebooted after a previous test. Note that the IOPs value in the first test is much higher than the second test. This is presumably due to caching. The memory and network usage were similar for each test.





1.1.3 Premium User Workload (105 Users)



8.3.3 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 can be supported, 116 basic user sessions can be supported with PvDisk enabled. Also it was determined that on an M620 blade server configuration where 135 basic users can be supported, 108 basic user sessions can be supported with PvDisk enabled.

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


<i>Basic Users</i>	<i>Basic Users w/PvDisk</i>	<i>Standard Users</i>	<i>Standard Users w/PvDisk</i>	<i>Premium Users</i>	<i>Premium Users w/PvDisk</i>	<i>Physical Mgmt Servers</i>	<i>Physical Host Servers</i>
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

8.3.4 EqualLogic PS6110XS Hybrid Array Testing Summary


This user story was a full bundle standard LoginVSI user validation which included a total of 720 Standard Workload image virtual desktops on the EqualLogic 6110XS Hybrid Array. This workload was chosen as it uses the highest amount of IOPS and is the best suited test to stress the array.

The following test was a 100% pre-booted test using DVS standard testing methodology using VSI to manage the images and Stratusphere UX to deliver performance results. Boot storm represents the peak IOPS on the array where Multiple VM's are booted at the same time and compete for System resources. For the testing used here the full load booted was 720 Virtual Machines giving a total of 12360 IOPS during the boot storm. No latency was experienced during the boot storm.






Below is the CPU data for the 720 User test including management Hosts. CPU performed excellently throughout this validation work. CPU never exceeded 66% during steady state.




Below is the active memory data for all hosts used for this validation work. Memory utilization across all hosts was consistent with the active management host only registered a small % of the compute node utilization.



Network data is also consistent across all hosts in use however the management host registered higher activity as it is handling network traffic at the management layer and ISCSI activity.





8.3.4.1 EqualLogic PS6110XS Test Results Summary Tables

Array	Average IOPS	Average Read IOPS	Average Write IOPS	Average Latency (ms)	Average Read Latency (ms)	Average Write Latency (ms)
M610-XS-01	2045.53	660.02	1385.50	1.91	4.15	1.19
Averages	2045.53	660.02	1385.50	1.91	4.15	1.19

The above table shows the average IOPS and latency figures for the duration of the validation workload. These figures are taken from the total duration of the test run indicating low latency on the array. Below are the Peak IOPS figures. It is important to note that these figures are based on the peak average IOPS during the 5minute polling period. The latency dropped off again significantly at the next polling period, this represents a random spike in activity at that time and could contribute to some sessions being represented in yellow on the StratusphereUX plotter graph.

Array	Peak IOPS	Peak Read IOPS	Peak Write IOPS	Peak Latency (ms)	Peak Read Latency (ms)	Peak Write Latency (ms)
M610-XS-01	12360	10371	5436	20.4	13.5	6.9
Averages	12360	10371	5436	20.4	13.5	6.9

In conclusion further synthetic testing has shown that the storage controller is performing well during steady state at 9500 IOPS indicating that there is sufficient headroom for additional sessions. It should be pointed out that the Peak IOPS figure shown above is the peak generated during the boot process, the peak during steady state was just below 7000 IOPS, so the 9500 IOPS seen during the extended synthetic testing shows that we are well within the threshold of total IOPS.

User Workload	Citrix XenDesktop Sessions
Basic User	1079
Standard User	1000



8.3.5 Dell Compellent Storage Testing Summary

The DVS 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 Gbps	16	Front end host connectivity
	SAS - 6 Gbps	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 3 – 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 4 – Steady State



Storage Scaling

Due to the nature of spinning drives, sizing a solution can be 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 will be 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 will be required to handle the desktop's performance.

Summary

In conclusion, a 2000 desktop Citrix XenDesktop based VDI solution can be 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.

8.3.5.1 Dell Compellent Storage Test Result Summary Tables

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

8.4 XenApp Testing Results

8.4.1 Test Configuration

The test environment included these core components:

- One DVS XenApp Virtual Server Compute Host based off a Dell R720 (Sandy Bridge).
 - See section 2.8.4 for hardware configuration details.
- Microsoft Windows Server 2008 R2 SP1 with Hyper-V.
- Citrix XenApp v6.5 configured with the default HDX and connection settings.
- Login VSI 3.5 (www.loginvsi.com).

8.4.2 Test Methodology

Each test run followed the following sequence of steps, until an optimum result and configuration was determined. While the dozens of test runs and their results are not included in this document, core data points for what proved to be the optimal configuration are

- The environment was configured with 12 Login VSI launchers, and after verifying first that they were ready for testing, a script was invoked the started PerfMon scripts to capture comprehensive system performance metrics, and then initiated the workload simulation portion of the test in which Login VSI launched new user connections at 20-second intervals.
- Once all users' sessions were connected, the steady state portion of the test began. During ramp-up and steady state, Login VSI tracked user experience statistics, looping through specific operations and measuring response times at regular intervals. Response times were used to determine Login VSIMax, the maximum number of users the test environment can support before performance degrades consistently.
- After a specified amount of elapsed steady state time, Login VSI started to log off the user sessions. After all sessions were logged off, the performance monitoring scripts were stopped and data gathered and analyzed.

8.4.3 Test Results Summary

This validation was performed for eight (8) virtual XenApp servers, running on an R720 host, 96GB of RAM and dual 2.9 GHz processors. Validation was performed using DVS standard testing



methodology using LoginVSI load generation tool for VDI benchmarking that simulates production user workloads. As a result of preliminary testing of different XenApp server configurations, the optimal supported configuration proved to be:

- Maximum of eight (8) virtual XenApp servers per physical R720 host
- Each virtual XenApp server configured with four (4) vCPUs
- Each virtual XenApp server configured with 10GB RAM

Tier 1 Read/Write Average Ratio	Tier 1 Read/Write Max Ratio	Tier 1 IOPS Average	Tier 1 IOPS Max
7/93	89/11	3.3	8.1

This virtual server-based configuration was able to sustain a user density ranging between 131 and 260 sessions per physical host server depending on the type of user workload chosen as well as the file access and scanning configuration of the anti-virus tools.

User Workload	Anti-Virus Configuration	Max Users Per Physical Server	Max Users Per XenApp Server
Basic User	Read / Write	200	25
Basic User	Write Only	260	33
Standard User	Read / Write	160	21
Standard User	Write Only	205	26
Premium User	Read / Write	131	17
Premium User	Write Only	167	21



8.4.4 XenApp Physical Host Server Scaling (Sandy Bridge)

<i>Basic Users AntiVirus Write</i>	<i>Basic Users AntiVirus Read/Write</i>	<i>Standard Users AntiVirus Write</i>	<i>Standard Users AntiVirus Read/Write</i>	<i>Premium Users AntiVirus Write</i>	<i>Premium Users AntiVirus Read/Write</i>	<i>Physical XenApp Host Servers</i>	<i>Physical XenApp Host Servers w/HA</i>
260	200	205	160	167	131	1	2
500	385	400	312	328	257	2	3
1000	769	800	624	655	514	4	5
1500	1154	1200	937	983	771	6	7
2000	1538	1600	1249	1310	1028	8	9
2500	1923	2000	1561	1638	1285	10	11
3000	2308	2400	1873	1966	1542	12	13
3500	2692	2800	2185	2293	1799	14	15
4000	3077	3200	2498	2621	2056	16	17
4500	3462	3600	2810	2948	2313	18	19
5000	3846	4000	3122	3276	2570	20	21
5500	4231	4400	3434	3603	2827	22	23
6000	4615	4800	3746	3931	3084	24	25
6500	5000	5200	4059	4259	3341	25	26
7000	5385	5600	4371	4586	3598	27	28
7500	5769	6000	4683	4914	3855	29	30
8000	6154	6400	4995	5241	4112	31	32
8500	6538	6800	5307	5569	4368	33	34
9000	6923	7200	5620	5897	4625	35	36
9500	7308	7600	5932	6224	4882	37	38
10000	7692	8000	6244	6552	5139	39	40



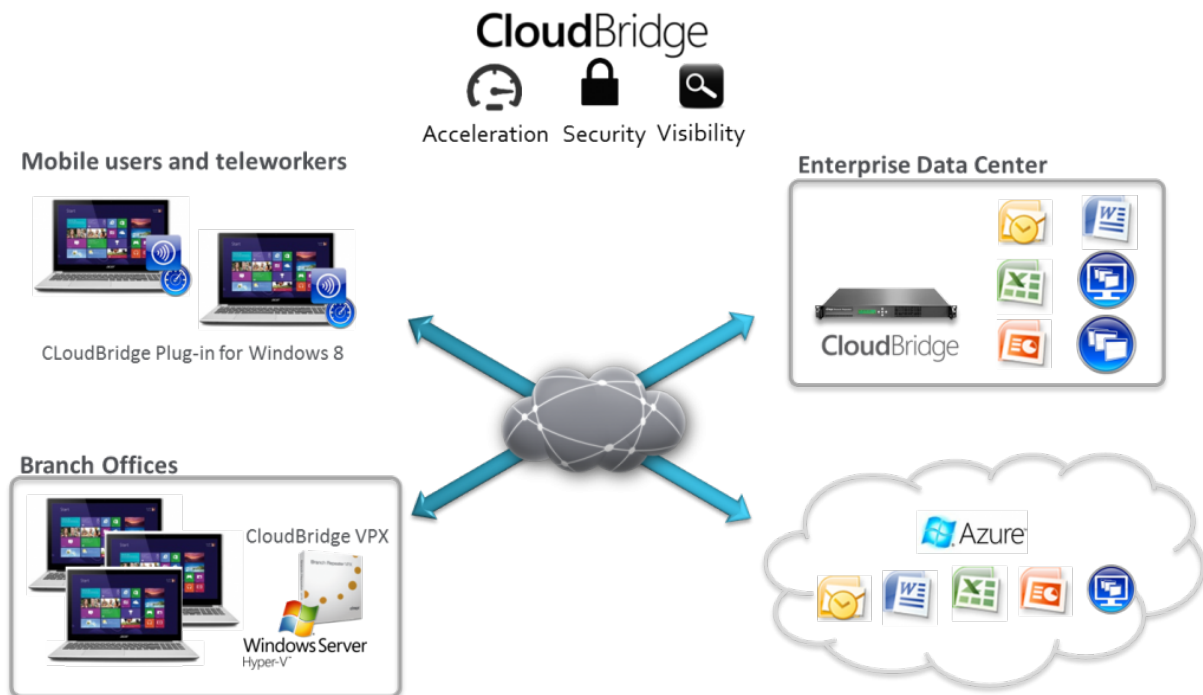
Appendix A – Branch Office WAN Acceleration

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
- Secure, optimized networking between clouds
- 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 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](#)



Appendix B – Load Balancing and Disaster Recovery

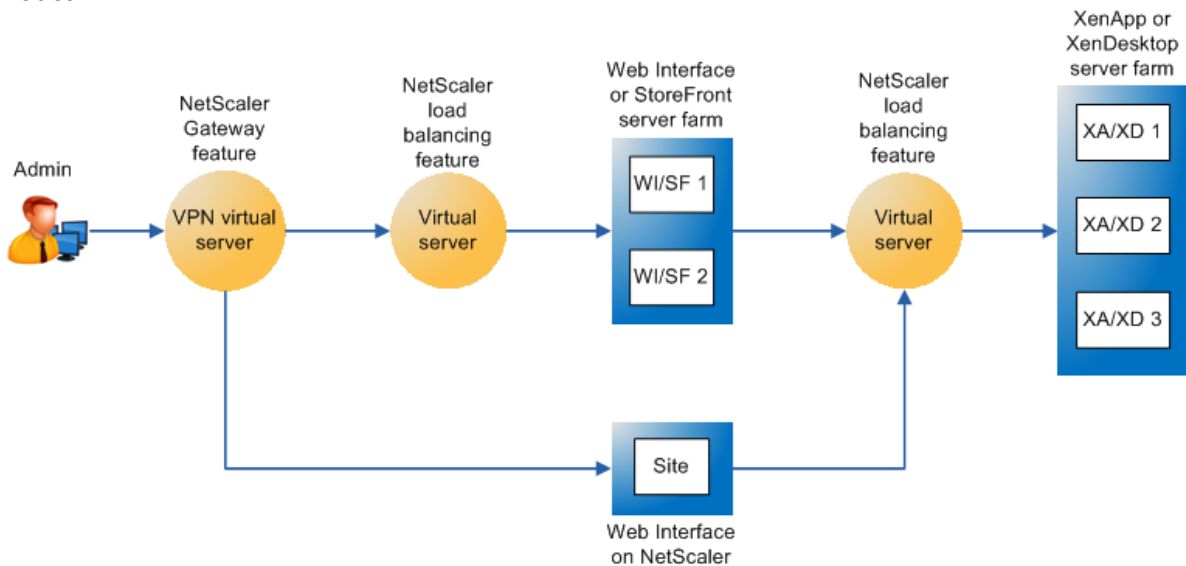
Citrix Netscaler

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 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.



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. This distribution can be 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](#)



Acknowledgements

Thanks to the Darin Schmitz and Damon Zaylskie of the Dell Compellent MSV Solutions team for providing expertise and validation of the DVS 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 for his expertise and continued support validating VDI architectures and Tier 1 shared storage.



About the Authors

Peter Fine is the Sr. Principal Engineering Architect for Citrix-based 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 Desktop Virtualization Solutions Group at Dell building, testing, validating, and optimizing enterprise VDI stacks.

Geoff Dillon is a Sr. Solutions Engineer in the Desktop Virtualization Solutions Group at Dell with deep Citrix experience and validation expertise of Dell's DVS enterprise VDI solutions.

Pranav Parekh is a Sr. solutions engineer at Dell Client Cloud Computing group. 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.

