


# DVS Enterprise with VMware Horizon View 5.2: Reference Architecture

**Dell Desktop Virtualization Solutions (DVS) Engineering**

**Revision: v4.0**



**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.**

© 2013 Dell Inc. All rights reserved. Dell and its affiliates cannot be responsible for errors or omissions in typography or photography. Dell, the Dell logo, OpenManage, Compellent, Force10, KACE, EqualLogic, PowerVault, PowerConnect, and PowerEdge are trademarks of Dell Inc. Intel and Xeon are registered trademarks of Intel Corporation in the U.S. and other countries. Microsoft, Windows, Hyper-V, and Windows Server are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. VMware, vSphere, ESXi, vMotion, vCloud, and vCenter are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. Linux is the registered trademark of Linus Torvalds in the U. S. and other countries. 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 disclaims proprietary interest in the marks and names of others.

## Document Review and Approval List

Name	Role	Input	Company	Version
Gus Chavira	Solutions Architect (ENG)	Author	Dell	
Peter Fine	Solutions Architect (ENG)	Author	Dell	
Andrew McDaniel	Solutions Development Manager (ENG)	Reviewer	Dell	
Sean Copeland	Solutions Manager	Reviewer	Dell	
David Hulama	Technical Marketing	Reviewer	Dell	
Darpan Patel	Solutions Engineering Lead (ENG)	Contributor	Dell	
Nicholas Busick	Solutions Engineering Lead (ENG)	Contributor	Dell	
Cormac Woods	Solutions Engineering Lead (ENG)	Contributor	Dell	
Mike Leahy	Solutions Engineering Lead (ENG)	Contributor	Dell	
Chhandomay Mandal	Storage Technical Marketing	Storage Section Contributor	Dell	
Sujit Somandepalli	Storage Engineering	Storage Section Contributor	Dell	

## Document Control

Author	Description	Date	Version
Gus Chavira	Initial Draft	7/22/2013	Draft 1.0
Gus Chavira	Final Draft	8/7/2013	Final Draft 2.0
Gus Chavira	Final version	8/12/2013	Final version 3.0
David Hulama, Sean Copeland	Final Review	8/14/2013	Final version 4.0

## Contents

1 Introduction .....	11
1.1 Desktop Virtualization Solutions Overview .....	11
1.2 The DVS Enterprise Solution Today .....	11
1.3 Product Positioning .....	12
1.4 Feature Overview .....	13
1.4.1 Design Principles .....	13
1.4.2 Architecture Scalability .....	13
2 The DVS Enterprise Solution Architecture .....	14
2.1 Introduction .....	14
2.2 DVS Enterprise Solution Layers .....	14
3 VMware .....	16
3.1 VMware vSphere 5.1.....	16
3.2 VMware Horizon View 5.2.....	16
3.3 Summary of Horizon View 5.2 Features .....	16
3.4 VMware Horizon View 5.2 Infrastructure .....	18
4 Storage Infrastructure .....	19
4.1 Storage Networking .....	19
4.2 EqualLogic iSCSI.....	19
4.3 Compellent Fibre Channel.....	20
4.4 Zoning .....	20
4.5 Local Tier 1 Storage.....	21
4.6 Shared Tier 1 Storage.....	21
4.7 Shared Tier 2 Storage.....	21
5 Virtualization Compute Server Infrastructure .....	23
5.1 DVS Enterprise R720 Rack Compute Server.....	23
5.1.1 DVS Enterprise R720 Rack Graphics Compute Server .....	23
5.2 DVS Enterprise M620 Blade Compute Server.....	25
5.3 Management Server Infrastructure.....	26
5.3.1 DVS Enterprise R720 Rack Management Server.....	26
5.3.2 DVS Enterprise M620 Blade Management Server.....	27
6 DVS Enterprise Server Scalability.....	28
6.1 R720 Rack-based Server Scalability.....	28
6.2 M620 Blade-based Server Scalability.....	28
7 Network Infrastructure.....	29
7.1 PowerConnect PC6248 1GbE Switch .....	29

7.2	Force10 S55 1GbE Switch.....	29
7.3	High Performance Ethernet Switching Options.....	30
7.3.1	Force10 S60 1GbE Switch .....	30
7.3.2	Force10 S4810 10GbE Switch .....	30
7.4	Fibre Channel Switching Options.....	31
7.4.1	Brocade 6510 Fibre-Channel Switch .....	31
7.4.2	Brocade M5424 Blade Switch .....	31
7.4.3	QLogic QME2572 Host Bus Adapter.....	32
7.4.4	QLogic QLE2562 Host Bus Adapter.....	32
7.5	Logical Networking .....	33
8	Scaling the Solution .....	34
8.1	Local Tier 1 (RAID enabled local disks on Server with Tier 2 Shared Storage) .....	34
8.1.1	Local Tier 1 – 50 User/ Pilot - (RAID enabled local disks on Server with Tier 2 Shared Storage).....	34
8.1.2	Local Tier 1 (iSCSI) - (RAID enabled local disks on Server with Tier 2 Shared Storage) .....	35
8.1.2.1	Local Tier 1 – Network Architecture (iSCSI).....	35
8.1.2.2	Local Tier 1 Rack Scaling Guidance (iSCSI).....	36
8.2	Shared Tier 1 Rack.....	37
8.2.1	Shared Tier 1 (SAN) – Rack – 500 users (iSCSI – EqualLogic).....	37
8.2.2	Shared Tier 1 (SAN) – Rack – 3000 users (iSCSI – EqualLogic).....	37
8.2.2.1	Shared Tier 1 Rack – Network Architecture (iSCSI).....	38
8.2.2.2	Shared Tier 1 Rack Scaling Guidance (iSCSI).....	38
8.2.3	Shared Tier 1 (SAN) – Rack – 1000 Users (FC – Compellent) .....	39
8.2.4	Shared Tier 1 – Rack (FC – Compellent).....	40
8.2.4.1	Shared Tier 1 Rack – Network Architecture (FC) .....	40
8.2.4.2	Shared Tier 1 Rack Scaling Guidance (FC).....	41
8.3	Shared Tier 1 (SAN) Blade.....	42
8.3.1	Shared Tier 1 – Blade – 500 users (iSCSI – EqualLogic) .....	42
8.3.2	Shared Tier 1 (SAN) – Blade (iSCSI – EqualLogic) .....	43
8.3.2.1	Shared Tier 1 Blade – Network Architecture (iSCSI) .....	43
8.3.2.2	Shared Tier 1 Blade Scaling Guidance (iSCSI).....	44
8.3.3	Shared Tier 1 (SAN) – Blade (FC – Compellent) .....	45
8.3.3.1	Shared Tier 1 Blade – Network Architecture (FC).....	45
8.3.3.2	Shared Tier 1 Blade Scaling Guidance (FC).....	46
8.4	Cabling Diagrams.....	47
8.4.1	Local Tier 1 Cabling (iSCSI and LAN separated) .....	47



8.4.2	Shared Tier 1 Cabling (Rack – EqualLogic)	47
8.4.3	Shared Tier Cabling (Blade - EqualLogic)	48
8.4.4	Shared Tier 1 Cabling (Rack – Compellent)	48
8.4.5	Shared Tier 1 Cabling (Blade – Compellent)	49
8.5	Building a Resilient Infrastructure	50
8.5.1	High Availability Cabling	50
8.5.2	High Availability Networking	51
8.5.2.1	vSphere Virtual Networking	51
9	HA DRS – Load Balancing - DNS	52
9.1	Management Server High Availability	52
9.2	Windows File Services High Availability	52
9.3	SQL Databases	52
9.3.1	SQL Server High Availability	53
9.4	Load Balancing	53
9.5	DNS	53
9.5.1	DNS for SQL	53
9.5.2	DNS for Load Balanced Services	54
10	Customer Provided Stack Components	56
10.1	Customer Provided Storage Requirements	56
10.2	Customer Provided Switching Requirements	56
11	Dell Wyse Cloud Clients	57
11.1	Dell Wyse P25	57
11.2	Dell Wyse Z50D	58
11.3	Dell Wyse D50D	59
12	Dell DVS VMware Horizon View Solution New Feature Sets	59
12.1	High-end VMware vDGA / Pass-thru graphics support	60
12.1.1	Executive Summary	60
12.1.2	Graphics-Specific Performance Analysis Results	61
12.1.2.1	K1 Tests – Fixed Frame-Rate Video Component – Single VM	61
12.1.2.2	K1 Heaven Benchmark Testing – Single VM	62
12.1.2.3	K1 test – Movie Clip + Companion testing	63
12.1.2.4	K1 test – Heaven Benchmark + Companion testing	64
12.1.2.4	K1 Nvidia GPU Results	66
12.1.2.5	K1 Subjective Tests	67
12.1.2.6	K1 Conclusion	67
12.1.2.7	K2 Tests – Fixed Frame-Rate Video Component – Single VM	67
12.1.2.8	K2 Tests – Heaven Benchmark Testing – Single VM	68

12.1.2.9 K1/ K2 Comparison test – Viewperf Single VM.....	69
12.1.2.10 K2 test – Viewperf + Companion testing .....	70
12.1.2.11 K2 Nvidia GPU Results + K1 Comparison .....	74
12.1.2.12 K2 Subjective Tests.....	74
12.1.2.13 K2 Conclusion .....	75
12.1.2.14 K2 P25 Zero Client Single VM Comparison Test .....	75
12.1.2.15 K2 P25 / OptiPlex 7010 Heaven Test with Heaven running on all VMs (4) .....	75
12.1.3 Conclusions.....	80
12.1.3.1 Test Conclusions.....	80
12.1.4 Appendices.....	80
12.1.4.1 Appendix A VMware View 5.1 solution configuration .....	80
12.2 Branch Office Desktop – In Geo location model and Branch office co-located infrastructure model.....	81
12.2.1 Executive Summary .....	81
12.2.2 Introduction .....	82
12.2.2.1 Traditional Approach.....	82
12.2.2.2 VMware Branch Office Desktop Approach .....	82
12.2.3 High Level Solution.....	82
12.2.3.1 VMware Branch Office Desktop Features .....	83
12.2.4 Solution Details .....	83
12.2.4.1 VMware Branch Office Desktop Architecture.....	83
12.2.4.2 How VMware Branch Office Desktop Components Work Together.....	84
12.2.4.3 VMware Branch Office Desktop Topologies .....	84
12.2.4.4 VMware Branch Office Deployment Guidance/Considerations.....	87
12.2.5 Business Benefits.....	88
12.2.6 Summary .....	89
12.3 Business Process Desktop (BOD) Use case.....	89
12.3.1 Executive Summary .....	89
12.3.2 Introduction .....	89
12.3.2.1 Traditional Approach.....	90
12.3.2.2 VMware Business Process Desktop Approach.....	90
12.3.3 High Level Solution.....	90
12.3.4 Solution Details .....	91
12.3.4.1 VMware Business Process Desktop Architecture .....	91
12.3.4.2 How VMware Business Process Desktop Components Work Together .....	92
12.3.5 Business Benefits.....	93
12.3.6 Summary .....	93

12.4	VMware Horizon Suite Bundle V1.0 .....	93
12.4.1	VMware Horizon Workspace .....	93
12.4.2	VMware ThinApp.....	93
12.4.2.1	Executive Summary.....	93
12.4.2.2	Introduction .....	94
12.4.2.3	Solution Overview.....	95
12.4.2.4	Solution Component Details .....	97
12.4.2.5	Business Benefits.....	100
12.4.2.6	Summary.....	100
12.4.3	VMware vShield Endpoint.....	100
12.4.3.1	Executive Summary.....	100
12.4.3.2	Introduction .....	101
12.4.3.3	High Level Solution.....	102
12.4.3.4	Solution Details.....	103
12.4.3.5	Business Benefits.....	106
12.4.3.6	Summary.....	106
12.4.4	Horizon Mirage for Physical desktops .....	107
12.4.4.1	Executive Summary.....	107
12.4.4.2	Introduction .....	107
12.4.4.3	High Level Solution.....	108
12.4.4.4	Solution Details .....	109
12.4.4.5	Business Benefits.....	111
12.4.4.6	Summary .....	112
12.5	VDI within a M1000e Chassis using EqualLogic PS-M4110XS Storage Blades.....	112
12.5.1	Executive Summary .....	112
12.5.2	Objectives.....	112
12.5.3	Audience.....	113
12.5.4	VDI with Dell EqualLogic PS Series Blade Storage .....	113
12.5.5	Infrastructure and test configuration.....	114
12.5.5.1	Host design considerations .....	114
12.5.5.2	Network design considerations.....	115
12.5.5.3	iSCSI SAN configuration.....	116
12.5.5.4	Separation of user data and virtual desktop data.....	117
12.5.5.5	EqualLogic storage array configuration.....	117
12.5.5.6	ESXi host network configuration.....	118
12.5.5.7	Horizon View Configuration.....	119
12.5.5.8	Horizon View pool configuration.....	119

12.5.5.9 Windows 7 Desktop VM configuration .....	119
12.5.6 Horizon View test methodology .....	120
12.5.6.1 Test objectives .....	120
12.5.6.2 Test tools .....	120
12.5.6.3 Test criteria .....	121
12.5.6.4 Test configuration.....	122
12.5.7.1 Test scenarios.....	122
12.5.7.2 One array test for standard users .....	123
12.5.7.3 Two array tests for standard users .....	128
12.5.7.4 Server host performance .....	131
12.5.7.5 User experience monitoring.....	133
12.5.7.6 Results summary .....	134
12.5.8 Sizing guidelines for EqualLogic SANs .....	135
12.5.9 Best Practices .....	135
12.5.9.1 Application layer .....	135
12.5.9.2 Server host layer .....	136
12.5.9.3 Network layer.....	136
12.5.9.4 Storage.....	137
12.5.10 Conclusions .....	137
12.6 PowerEdge VRTX with VMware Horizon View .....	138
12.6.1 Overview.....	138
12.6.1.1 Test objectives.....	138
12.6.2 Test Environment.....	138
12.6.2.1 DVS Stack Infrastructure .....	139
12.6.2.2 Test Infrastructure.....	139
12.6.3 Test Set-up.....	140
12.6.3.1 Shared Storage Raid Configuration for max user density.....	140
12.6.3.2 Network Configuration.....	140
12.6.4 Test Methodology .....	140
12.6.4.1 Test approach .....	140
12.6.4.2 Test Proposal.....	140
12.6.4.4 Test criteria.....	141
12.6.4.5 Test Results .....	142
12.6.5 VRTX Configuration Overview.....	145
12.6.5.1 Storage Volume Overview .....	145
12.6.5.2 Compute Host Configuration Overview.....	145
12.6.5.3 Compute Node Networking Overview .....	146

12.7	Foglight for Virtualization with Horizon View cartridge .....	146
12.7.1	Introducing the Cartridge for VMware View.....	146
12.7.2	Cartridge for VMware View Elements .....	147
12.7.3	Using the VMware View Environment Dashboard .....	147
12.7.4	Exploring the VMware View Environment Details Tab .....	148
12.7.5	Working with the Tiles.....	149
12.7.6	Using the Quick View .....	151
12.7.7	Exploring User Sessions .....	152
12.7.8	Exploring Desktops.....	153
12.7.9	Exploring Horizon Pools .....	154
13	Previous Features Appendix .....	155
13.1	Branch Office Deployments.....	155
13.2	Microsoft Lync 2013 enablement .....	157
13.2.1	Microsoft Lync 2013 Overview and Architecture .....	157
13.2.2	VMware Horizon View 5.2 with Microsoft Lync 2013 Testing Results.....	159
13.2.2.1	Purpose.....	159
13.2.2.2	Test Infrastructure.....	159
13.2.2.3	Lync 2013 Test Methodology .....	159
13.2.2.4	Lync 2013 VDI Plug-in.....	160
13.2.2.5	Lync 2013 VDI Plug-in Testing/Characterization results.....	161
13.2.2.6	VMware View Horizon with VDI Plug-in for Lync 2013 Conclusions .....	164
13.3	View Configuration Tool.....	165
13.4	VMware VCOPS for View (V4V) .....	166
13.4.1	Executive Summary .....	166
13.4.2	Introduction.....	166
13.4.2.1	Traditional Approach .....	166
13.4.2.2	V4V Approach.....	166
13.4.3	High Level Solution.....	167
13.4.3.1	vCenter Operations Manager for Horizon View Features.....	167
13.4.4	Solution Details .....	168
13.4.4.1	vCenter Operations Manager for Horizon View Architecture .....	168
13.4.4.2	How vCenter Operations Manager for Horizon View Components Work Together .....	169
13.4.4.3	Licensing.....	170
13.4.4	Business Benefits .....	171
13.4.5	Summary .....	171
13.5	Graphics Support – NVIDIA K1/K2 GRID cards and VMware vSGA support.....	171

13.5.1	NVIDIA K1/K2 Testing/Characterization results .....	171
13.5.1.1	Environment Summary .....	171
13.5.1.2	Graphic Specific Performance Analysis Results.....	174
13.5.1.3	VMware Horizon View vSGA Conclusions .....	184
13.6	General Feature Updates .....	185
13.6.1	Desktop Storage Reclamation – Storage unmap.....	185
13.6.2	VMware vSphere Virtual Distributed Switches (VDS) .....	203
13.6.1	VMware vSphere Storage and Regular vMotion for Management Tiers.....	220
13.7	AS 50, 200, 800, and 1000 (IOA and MSL).....	241
13.8	Win8 characterization and testing .....	242
13.9	Compellent 6.3 code update.....	254
13.10	Converged Networking.....	255
13.11	Horizon Workspace .....	264
13.12	Dell Mobile Clinical Computing - AlwaysOn Point of Care.....	271
13.12.1	Benefits.....	271
13.12.2	Solution Elements .....	272
13.12.3	Imprivata OneSign .....	272
13.12.4	Dell and VMware Solutions for secure, reliable and continuous access to EMR and patient care applications .....	272
13.13	Mobile Secure Desktop .....	273
13.13.1	Mobility .....	273
13.13.2	Security.....	273
13.13.3	Management.....	274
13.13.4	Solution Elements.....	274
13.13.5	Compliance - VMware vCenter Configuration Manager.....	274
13.13.6	Cortado ThinPrint .....	274
13.13.7	Imprivata OneSign .....	274
14	Reference.....	276

# 1 Introduction

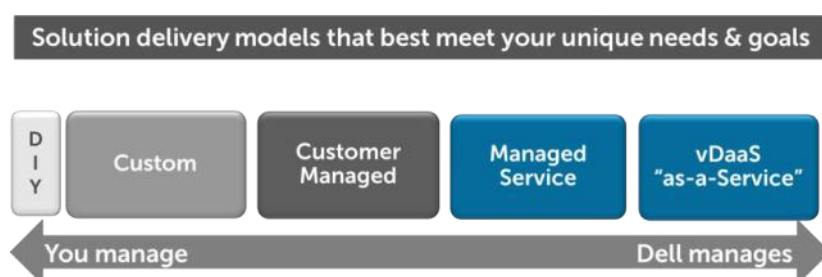
## 1.1 Desktop Virtualization Solutions Overview

Dell Desktop Virtualization Solutions' offers a comprehensive solution portfolio designed to deliver the benefits of virtual end user and Cloud Client Computing. While there are several ways of delivering virtual desktops, this solution is built on the Virtual Desktop Infrastructure (VDI) model. In a VDI environment, user desktops are hosted as virtual machines (VMs) in a centralized infrastructure and delivered over a network to an end user's client device.

Getting the most out of a VDI infrastructure requires a well-developed, reliable architecture. If the VDI architecture is undersized in terms of processing, memory or storage, then performance will suffer and the experience will be less user-friendly than a traditional PC. If the architecture is oversized, then the cost per virtual desktop will be economically unfeasible for VDI adoption.

In order to enable Dell to compete efficiently in the VDI space and to accelerate the sales cycle, a Solution Source Architecture (SSA) approach has been developed. Accelerate your time to benefit and increase your operating efficiency with this tested methodology.

Our extensively pre-tested Dell Desktop Virtualization Solutions leverage purpose built hardware, software and services ingredients to enable a "capable" architecture that maximizes IT control while enhancing the end user experience. Choose a clear path to flexible options, expedient upgrades and support through the solution life cycle with industry standard ingredients. Dell invests in extensive R&D and solution validation to ensure you experience a fine-tuned deployment process that leads to deterministic operational costs. And you can minimize your capital costs with Dell's subscription based Desktop- as-a-Service.



## 1.2 The DVS Enterprise Solution Today

Dell Desktop Virtualization Solutions deliver a range of purpose-built horizontal architectures. These architectures are designed and battle-tested to be modular and scalable for an array of your needs and a defined and tested services methodology. To provide a simplified solution stack we have designed and enhanced the original DVS Enterprise Solution to address the vast majority of your needs and use cases for Desktop Virtualization. Simultaneously, we have worked to make the solution easier to deploy and scale.

Initially there were the ISS Enterprise and ISS Enterprise+ bundles; both with strict guidelines and constraints on sizing and scaling. Subsequently, the DVS Enterprise solution has been refined and enhanced to be custom-tailored and sold as one cohesive stack known as DVS Enterprise. The solution now has the ability to be sold initially as an entry level rack-based solution for as few as 50 users. Alternatively, the solution can either grow into, or be customized and sold as a highly scalable, blade-based solution serving 50,000 users or more.

The DVS Enterprise solution is an architecture incorporating a VMware vSphere hypervisor with a Horizon View 5.2 desktop virtualization solution. On top of this foundation runs Dell's core architecture components for networking, compute and storage designed from a tested and effective selection of ingredients.

### 1.3 Product Positioning

The Dell Desktop Virtualization Solution is a prescriptive, highly scalable, flexible architecture designed to meet the wide array of your VDI needs that exist today. The DVS Enterprise Solution has the ability to scale anywhere from 50 to 50000 users with a high degree of prescription at every user level along the way. This granularity of scale allows you to leverage Dell DVS's accurate pay-as-you-grow model and add VDI capability as their VDI needs increase.

	Minimum Base Configuration	Rack Server, Local Tier 1	Rack Server, Shared Tier 1	Blade Server, Shared Tier 1
Virtual Desktops	Yes	Yes	Yes	Yes
Redundancy	Optional	Yes	Yes	Yes
Recoverability	Optional	Yes	Yes	Yes
Live Migration	No	No	Yes	Yes
High Density	No	No	No	Yes

To provide this level of proven prescription, the DVS Enterprise leverages a core set of hardware and software components that have been tested and proven to provide optimal performance at the lowest cost per user. To provide this level of flexibility, the DVS Enterprise also includes an extended list of optional/upsell components that you can chose from for environments with unique VDI feature, scale or performance needs. Whether you require a Managed Solution from Dell or prefer to manage the solution in-house, the tenants of the DVS Enterprise Solution remain consistent and will be leveraged as the horizontal platform. If the various approved configurations do not meet your requirements then a custom solution can be provided.



## 1.4 Feature Overview

### 1.4.1 Design Principles

The design principles for the flexible computing solution are:

- Secure – Security risks, concerns and policies are addressed or mitigated
- Manageable – The solution includes the tools and software services required to manage the environment
- Standards-based – Makes use of commodity, off-the-shelf components wherever possible
- Distributed – Non-blocking and built with distributed components to maximize the use of available computing resources and eliminate bottlenecks
- Scalable – Capable of scaling up or down to support business needs
- Resilient – Capable of withstanding the failure of a single infrastructure component.

### 1.4.2 Architecture Scalability

The architecture is designed to provide a scalable platform:

- 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)
- The architecture eliminates bandwidth and performance bottlenecks as much as possible
- This scalability enables the reduction of the future cost of infrastructure ownership.

Component	Horizontal scalability	Vertical scalability
Virtual Desktop Compute Servers	Additional servers added as necessary.	Additional RAM or CPU compute power.
View Connection Servers	Additional physical servers added to the Management cluster to deal with additional management VMs.	Additional network and I/O capacity added to the servers.
VMware vCenter	Deploy additional servers and use linked mode to optimize management.	Additional vCenter Management VMs.
Database Services	Migrate databases to a dedicated SQL server and increase the number of management nodes.	Additional RAM and CPU for the management nodes.
File Services	Split user profiles and home directories between multiple file servers in the cluster. File services can also be migrated to the optional NAS device for high availability.	Additional RAM and CPU for the management nodes.

## 2 The DVS Enterprise Solution Architecture

---

### 2.1 Introduction

The DVS Enterprise Solution leverages a core set of hardware and software components in the following categories:


- Networking
- Virtualization Compute Servers
- Management Servers
- Storage Tiers

These components have been tested and proven to provide the optimal balance of high performance and lowest cost per user. Additionally, the DVS Enterprise also includes an approved extended list of optional/upsell components in all the same categories. This extended list of components you can chose from to custom tailor the solution for environments with unique VDI feature, scale or performance needs.


The Enterprise stack is designed to be a cost effective starting point when you want to start your transition to a fully virtualized desktop environment slowly, allowing you to grow the investment and commitment as needed or as you become comfortable with VDI as a solution.

### 2.2 DVS Enterprise Solution Layers


Only a single high performance PowerConnect or 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 user sessions.



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 Fibre Channel to suit your tier 1 and tier 2 scaling and capacity needs. Shared or local Tier 1 will include shared Tier 2 data for user data and management VMs.



## 3 VMware

---

### 3.1 VMware vSphere 5.1

VMware vSphere 5.1 includes the ESXi™ hypervisor as well as vCenter™ Server which is used to configure and manage VMware hosts. Key capabilities for the ESXi Enterprise Plus license level include:

- VMware vMotion™ - VMware vMotion technology provides real-time migration of running virtual machines (VM) from one host to another with no disruption or downtime.
- VMware High Availability (HA) - VMware HA provides high availability at the virtual machine (VM) level. Upon host failure, VMware HA automatically restarts VMs on other physical hosts running ESXi. VMware vSphere 5.1 uses Fault Domain Manager (FDM) for High Availability.
- VMware Distributed Resource Scheduler (DRS) and VMware Distributed Power Management (DPM) - VMware DRS technology enables vMotion to automatically achieve load balancing according to resource requirements. When VMs in a DRS cluster need fewer resources, such as during nights and weekends, DPM consolidates workloads onto fewer hosts and powers off the rest to reduce power consumption.
- vSphere Storage DRS™ and Profile-Driven Storage - New integration with VMware vCloud® Director™ enables further storage efficiencies and automation in a private cloud environment.
- VMware Storage vMotion™ - VMware Storage vMotion enables real-time migration of running VM disks from one storage array to another with no disruption or downtime. It minimizes service disruptions due to planned storage downtime previously incurred for rebalancing or retiring storage arrays.
- Space Efficient Sparse Virtual Disks - SE Sparse Disks introduces an automated mechanism for reclaiming stranded space. SE Sparse disks also have a new configurable block allocation size which can be tuned to the recommendations of the storage arrays vendor, or indeed the applications running inside of the Guest OS. VMware Horizon View 5.2 is the only product that will use the new SE Sparse Disk in vSphere 5.1.
- VMware vCenter Update Manager - VMware vCenter Update Manager automates patch management, enforcing compliance to patch standards for VMware ESXi hosts.
- Host Profiles - Host Profiles standardize and simplify the deployment and management of VMware ESXi host configurations. They capture and store validated configuration information, including host compliance, networking, storage, and security settings.
- vSphere Web Client - The vSphere Web Client is now the core administrative interface for vSphere. This new flexible, robust interface simplifies vSphere control through shortcut navigation, custom tagging, enhanced scalability, and the ability to manage from anywhere with Internet Explorer or Firefox-enabled devices.
- vCenter Single Sign-On - Dramatically simplify vSphere administration by allowing users to log in once to access all instances or layers of vCenter without the need for further authentication.

For more information on VMware vSphere, see [www.vmware.com/products/vsphere](http://www.vmware.com/products/vsphere).

### 3.2 VMware Horizon View 5.2

VMware Horizon View 5.2 is a desktop virtualization solution that delivers virtual desktops as an on-demand service to any user, anywhere. With VMware's desktop delivery technology, Horizon View 5.2 can quickly and securely deliver individual applications or complete desktops to the entire enterprise, whether they are task workers, knowledge workers or mobile workers. Users now have the flexibility to access their desktop on any device, anytime, with a high-definition user experience. With VMware Horizon View 5.2, IT can manage single instances of each OS, application and user profile and dynamically assemble them to increase business agility and greatly simplify desktop management.

### 3.3 Summary of Horizon View 5.2 Features

#### End User Experience

- Support for Windows 8 based desktops
  - View5.2 offers support for Windows 8 desktops.

- Facilitates a smooth transition in rolling out Windows 8.
- Enables IT to leverage the latest Windows capabilities in VDI

#### Hardware Accelerated 3D Graphics

- Horizon View5.2 provides a rich workstation class user experience with high performance graphics
  - Enables shared-access to physical GPU hardware for 3D and high performance graphical workloads.
  - Very cost effective as multiple VMs share the same GPU resource.
  - Offers full compatibility with hosts lacking physical GPUs

#### Improved Video and VOIP communications with Microsoft Lync 2013 support

- View5.2 offers tighter integration with Microsoft Lync and Office applications
  - Full collaboration capabilities with Microsoft Lync on View Desktops.
  - Full support for Unified Communications VoIP and Video using Lync client on View desktops
  - Support PCoIP

#### Streamlined access to View Desktops from Horizon

- View desktops can now be accessed via Horizon gateway.
  - Horizon provides a single point of access for end users to desktops, data and apps.
  - This provides a one-stop shop for all end-user access to their corporate workloads

#### Easily connect to desktops from any device with HTML Access

- View Desktops can now be accessed through a HTML5 capable web browser via Horizon.
  - This provides install-free access to Desktops from ANY modern device.

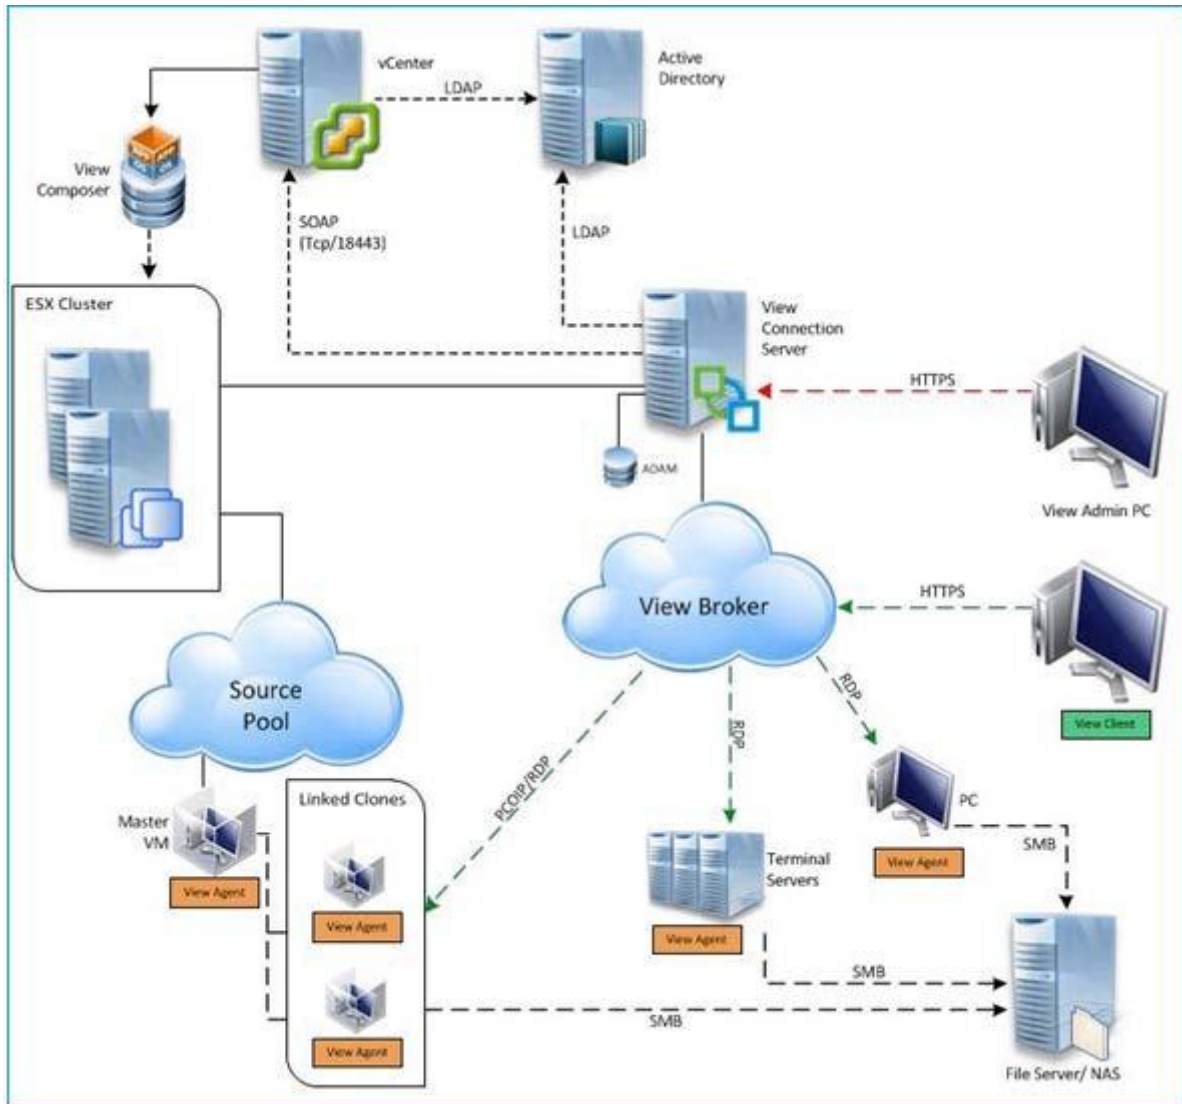
#### Ease of Management

- Large Pool creation with elimination of 8 host limits and multiple vLAN support
  - View5.2 has support for large View pools with more than 8 hosts.
  - This offers OPEX savings with less admin time spent on common operations.
  - Even more OPEX savings with fewer pools to manage in multi-thousand user deployments.
- Tech Preview of a new Integrated Service Console in the VC Web Client
  - View5.2 offers a View plugin into vSphere Web client.
  - Allows for easier desktop support and troubleshooting.
  - Offers a simple interface for novice administrative users which allows for increased efficiency
- Support for VC Virtual Appliance based deployments
  - View5.2 is fully compatible with Virtual Appliance-based VirtualCenter deployments.
  - This eliminates VirtualCenter dependencies on Windows
  - Easier installation and upgrades.

#### Total Cost of Ownership

- Substantial storage capacity savings for persistent desktops with Space Efficient Disks
  - View 5.2 leverages a vSphere capability to offer a new disk format for VMs on VMFS or NFS.
  - Space Efficient disks provide reduced storage capacity requirements (lower CAPEX) for persistent desktops
  - Unused space is reclaimed and View composer desktops stay small
  - IO alignment and grain size - Space Efficient disks guarantee that there will be no misalignment on storage arrays that are 4k aligned

### 3.4 VMware Horizon View 5.2 Infrastructure



The DVS Enterprise architecture with VMware Horizon View 5.2 is designed to provide maximum performance and scalability starting at very low user counts for SMBs and up to tens of thousands of users for large enterprise deployments. VMware Horizon View 5.2 brings with it many new features, as outlined above. This solution architecture follows a distributed model where solution components exist in tiers. The Compute tier is where VDI desktop VMs execute/run, the Management tier being dedicated to the broker management server role VMs. Both tiers, while inextricably linked, scale independently.

## 4 Storage Infrastructure

The leading cause of slow end-user experience in a VDI deployment is inadequate or poorly designed storage architecture. From the moment a user authenticates into a VDI environment, multiple tiers of the VDI storage architecture are being accessed, sometimes heavily. It is this reason that if any element of a storage array; network interface, disk speed, disk type, disk quantity, free space, or similar, is not up to the task of supporting the workload associated with the user population, performance will suffer greatly and failures may occur. The DVS Enterprise solved these potential problems by providing storage options which exceed the performance and \$/user requirements and at the same time provide additional storage options to custom-tailor performance and capacity characteristics to your specific needs.

The DVS Enterprise leverages tiered storage architecture. The tier 2 storage model remains constant throughout the stack no matter what server or networking choices are made. When choosing a rack server-based DVS Enterprise solution stack however, the tier 1 model has two options to choose from:

- Local tier 1 storage (local RAID configured disks on server) for virtual desktop images
- Shared tier 1 storage for virtual desktop images

### 4.1 Storage Networking

The DVS Enterprise solution has greatly expanded tier 1 and tier 2 storage strategy and flexibility over prior releases. You have the choice to leverage best-of-breed iSCSI solutions from EqualLogic or Fibre Channel solutions from Dell Compellent while be assured the storage tiers of the DVS Enterprise solution will consistently meet or outperform user needs and expectations.

### 4.2 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 you 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 that were used as part of the configuration:


- Specify 2 IP Addresses for iSCSI on each host
- Specify NICs (vmnic2, vmnic3)
- Specify Jumbo Frames at 9000 MTU (Broadcom NICs cannot support both iSCSI offload and jumbo frames)
- Initialize iSCSI initiator
- Specify IP for the EqualLogic Storage group.

### 4.3 Compellent Fibre Channel

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 Fibre Channel (FC) is the selected block storage protocol, then the Compellent Storage Center Integrations for the 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, as a best practice recommendation, the environment should be configured with two fabrics to provide multi-pathing and end-to-end redundancy.




When using QLogic HBAs the following BIOS settings should be used:

- The "connection options" field should be set to 1 for point to point only
- The "login retry count" field should be set to 60 attempts
- The "port down retry" count field should be set to 60 attempts
- The "link down timeout" field should be set to 30 seconds.
- The "queue depth" (or "Execution Throttle") field should be set 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.

### 4.4 Zoning

At least 1 port from each server HBA should be zoned to communicate with a single Compellent fault domain. The result of this will be two distinct FC fabrics and four redundant paths per server. Round Robin or Fixed Paths are supported. Compellent Virtual Ports should be leveraged to minimize port consumption as well as simplify deployment. Each controller's front-end virtual ports within a fault domain should be zoned with at least one ESXi initiator per server.





## 4.5 Local Tier 1 Storage

Choosing the local tier 1 storage option means that the virtualization compute 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.

## 4.6 Shared Tier 1 Storage

Choosing the shared tier 1 option means that the virtualization compute servers are deployed in a diskless mode and all desktops 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. Based on the heavy performance requirements of tier 1 VDI execution, it is recommended to use separate arrays for tier 1 and tier 2 above 500 users for EqualLogic and above 1000 users for Compellent. It is recommend using 500GB LUNs for VDI and running 125 VMs per volume to minimize disk contention. Sizing to 1000 basic users, for example, we will require 8 x 500GB volumes per array. A VMware Horizon View replica to support a 1 to 500 desktop VM ratio should be located in a dedicated Replicas volume.

Volume name	Size	Purpose
VDI-BaseImages	100 GB	Storage for Base image for VDI deployment
VDI-Replicas	100 GB	Storage for Replica images created by Horizon View
VDI-Images1	500 GB	Storage for VDI Virtual Machines in Horizon View Cluster
VDI-Images2	500 GB	Storage for VDI Virtual Machines in Horizon View Cluster
VDI-Images3	500 GB	Storage for VDI Virtual Machines in Horizon View Cluster
VDI-Images4	500 GB	Storage for VDI Virtual Machines in Horizon View Cluster
VDI-Images5	500 GB	Storage for VDI Virtual Machines in Horizon View Cluster
VDI-Images6	500 GB	Storage for VDI Virtual Machines in Horizon View Cluster
VDI-Images7	500 GB	Storage for VDI Virtual Machines in Horizon View Cluster
VDI-Images8	500 GB	Storage for VDI Virtual Machines in Horizon View Cluster

For Shared Storage use on Compellent storage it is important to note that you it is assumed that all pre-work for setting up prospering tiering to ensure proper data progression is enabled and ensure optimal performance. General guidance language follows:

Replica (read only data) - SSD

User non-persistent - 15K

User Persistent - Data progression 15K --> 7K

Infrastructure volumes - Data progression "All tiers" (or) 15K --> 7K

## 4.7 Shared Tier 2 Storage

Tier 2 is shared iSCSI storage used to host the Management server VMs and user data. EqualLogic 6100 series arrays can be used for smaller scale deployments (local tier 1 only) or the 6500 series for larger deployments (up to 16 in a group). The Compellent tier 2 array scales simply by adding

disks. The table below outlines the volume requirements for tier 2. Larger disk sizes can be chosen to meet your capacity needs. The user data can be presented either via a file server VM using RDM for small scale deployments or via NAS for large scale deployments. This is the only native NTFS volume on the array. All SQL disks should be presented as VMDKs. RAID 50 can be used in smaller deployments but is not recommended for critical environments. Larger scale and mission critical deployments with higher performance requirements should 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 should 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, View Connection Server, File and SQL	VMFS 5
User Data	2048	Tier 2	File Server/ NAS	NTFS
User Profiles	20	Tier 2	User profiles	VMFS 5
SQL DATA	100	Tier 2	SQL	VMFS 5
SQL LOGS	100	Tier 2	SQL	VMFS 5
TempDB Data	5	Tier 2	SQL	VMFS 5
TempDB Logs	5	Tier 2	SQL	VMFS 5
SQL Witness	1	Tier 2	SQL (optional)	VMFS 5
Templates/ ISO	200	Tier 2	ISO storage (optional)	VMFS 5

## 5 Virtualization Compute Server Infrastructure

A major building block when designing a VDI solution is that of virtualization compute server. This building block is required in all implementations and the choice of server determines what loads can be supported and how agile the environment will be when meeting changing business needs. By offering a broad array of choices, Dell gives you the ability to select the specific platform that is right for your environment.

### 5.1 DVS Enterprise R720 Rack Compute Server

The DVS Enterprise begins with a minimum configuration design based on the R720 rack mounted server. This server is the gold standard and a core server component providing the best performance at the lowest dollar-per-user. However, just as when deciding on the optimal switching architecture, there are certain deployment use cases where an upgrade or change in the server architecture makes sense, either from a management, power consumption, end-user desktop performance, or user density perspective.





PowerEdge R720 Virtualization Host Server For Local Tier 1 Storage	PowerEdge R720 Virtualization Host Server for Shared Tier 1 Storage
2 x Intel Xeon E5-2680 Processor (2.9GHz)	2 x Intel Xeon E5-2680 Processor (2.9GHz)
192GB Memory (12 x 16GB DIMMs @ 1600MHz)	192GB Memory (12 x 16GB DIMMs @ 1600MHz)
10 x 146GB SAS 6Gbps 15k Disks	Diskless
PowerEdge RAID Controller (PERC) H710 Integrated RAID Controller – RAID10	Broadcom 57800 2 x 10Gb SFP+ + 2 x 1Gb NDC
Broadcom 5720 1Gb QP NDC (LAN)	1 x Broadcom 5720 1Gb DP NIC (LAN)
Broadcom 5720 1Gb DP NIC (LAN)	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

In the above configuration, the R720-based DVS Enterprise Solution can also support the following single server user counts with Windows 7 desktops (see below for Windows 8 guidance).

- Basic Workload: 145 Users
- Standard Workload: 116 Users
- Premium Workload: 95 Users

#### 5.1.1 DVS Enterprise R720 Rack Graphics Compute Server

The following is a depiction of DVS Enterprise R720 with NVIDIA Grid card enablement to support vSGA higher-end graphics.



\* K1 vSGA use cases include: Task Worker, Knowledge Worker, Power User

\* K2 vSGA use cases include: 3D/CAD Reviewer

	Task Worker	Knowledge Worker	Power User	3D/CAD Reviewer	High-End Designer / Engineer	Video Editor / Designer
Description	Office worker with little or no dependency on 3D acceleration or Direct X 9 needs	Office worker using productivity apps, video conferencing, rich media Internet apps; using up to two monitors	Multi-Monitor (2+) worker using graphics design and vector graphics, Aero glass effects	Reviews CAD animations, making occasional edits; reviews for accuracy and ordering materials. Low-end rendering.	Dedicated CAD user/content creator. Rotations, heavy rendering. Needs more than 1GB of video memory (VRAM). Open GL 3.x or 4.x CUDA	Dedicated CAD user/content creator. Rotations, heavy rendering. Needs more than 1GB of video memory (VRAM). Open GL 3.x or 4.x. H264 Offload
Typical Apps		Office productivity apps	Google Earth, Adobe Illustrator or Photoshop, MS Visio	Autodesk, AutoCAD, CAD/CAM viewers	Enovia, Siemens NX, Autodesk, CATIA	Adobe Premiere, non-linear editing
Soft 3D	Good	Low-to-mid-range	Stretch	Under-powered	Qualify out	Qualify out
vSGA (K1)	Overkill	Overkill	Good	Stretch	Qualify out	Qualify out
vSGA (K2)	Way Overkill	Overkill	Overkill	Good	<b>Future vDGA</b>	<b>Future vDGA</b>

Data source: "Vmware View Graphics Acceleration Guidelines by Use Case"

## 5.2 DVS Enterprise M620 Blade Compute Server

The Dell M1000e Blade Chassis along with the M620 blade server option can be considered the option of choice for a high-density data center configuration. Although configured similarly to the R720, the M620 server brings with it a host of ancillary benefits.

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, with the introduction of the new Intel Xeon (Sandy Bridge-EP, EP-2600) 2 socket processors and 24 DIMMs (384GB+ RAM) of DDR3 memory in an extremely compact half-height blade form factor.

When comparing a rack versus blade VDI solution there are several considerations; while the initial acquisition costs are higher than comparable rack servers, blades use less power for the same amount of processing. In testing, Dell blades yield 20% more performance per watt or more compared to rack servers. They can be key elements for reducing ever-growing power costs and for implementing environmentally conscious IT initiatives.

A blade-based environment can help reduce total cost of ownership (TCO) with features such as power efficiency, built-in security and advanced systems management capabilities. Designed to handle memory-intensive applications, it is an ideal choice for virtualized workloads. Dell high-density blade solutions let you do more with less — more processing with less space and power, and with fewer resources. In the DVS Enterprise solution, the M620 compute server is configured as follows:

PowerEdge M620 Virtualization Host Server for Shared Tier 1 Storage
2 x Intel Xeon E5-2680 Processor (2.7GHz)
192GB Memory (12 x 16GB DIMMs @ 1600MHz)
VMware vSphere 5 on internal SD
Diskless
Broadcom 57810-k 10Gb DP KR NDC (iSCSI)
1 x Intel i350 1Gb QP mezzanine (LAN)
iDRAC7 Enterprise w/ vFlash, 8GB SD

In the above configuration, the M620-based DVS Enterprise Solution can also support the following single server user counts (with Windows 7 workloads; see below for Windows 8 sizing guidance).

- Basic Workload: 135 Users
- Standard Workload: 110 Users
- Premium Workload: 90 Users

## 5.3 Management Server Infrastructure

### 5.3.1 DVS Enterprise R720 Rack Management Server

In addition to the Virtual Desktop hosts there will be a dedicated management server that will provide compute resources for infrastructure services. These services will be provided by virtual server instances. Initially these servers are also based off a Dell PowerEdge R720 with the following specifications.



PowerEdge R720 Management Server For Local Tier 1 Storage	PowerEdge R720 Management Server for Shared Tier 1 Storage
2 x Intel Xeon E5-2680 Processor (2.9GHz)	2 x Intel Xeon E5-2680 Processor (2.9GHz)
96GB Memory (6 x 16GB DIMMs @ 1600MHz)	96GB Memory (6 x 16GB DIMMs @ 1600MHz)
10 x 146GB SAS 6Gbps 15k Disks	Diskless
PERC H710 Integrated RAID Controller – RAID10	Broadcom 57800 2 x 10Gb SFP+ + 2 x 1Gb NDC
Broadcom 5720 1Gb QP NDC (LAN)	1 x Broadcom 5720 1Gb DP NIC (LAN)
Broadcom 5720 1Gb DP NIC (LAN)	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

The management role requirements for the base solution are summarized below. Data disks should be used for role-specific application files/ data, logs, IIS web files, etc. and should exist in the management volume. Tier 2 volumes with a special purpose (called out above) should be presented in the format specified below:

Role	vCPU	RAM (GB)	NIC	OS + Data vDisk (GB)	Tier 2 Volume (GB)
VMware vCenter	2	8	1	40 + 5	100 (VMDK)
View Connection Server	2	8	1	40 + 5	-
SQL Server	2	8	1	40 + 5	210 (VMDK)
File Server	1	4	1	40 + 5	2048 (RDM)
<b>TOTALS</b>	<b>7</b>	<b>28</b>	<b>4</b>	<b>180</b>	<b>2358</b>

### 5.3.2 DVS Enterprise M620 Blade Management Server

When the Dell M1000e Blade Chassis along with the M620 blade server option is selected, the management servers are also based on the M620 blade. In the DVS Enterprise solution, the M620 management server is configured as follows;

PowerEdge M620 Virtualization Management Server for Shared Tier 1 Storage
2 x Intel Xeon E5-2680 Processor (2.7GHz)
96GB Memory (12 x 16GB DIMMs @ 1600MHz)
VMware vSphere 5 on internal SD
Diskless
Broadcom 57810-k 10Gb DP KR NDC (iSCSI)
1 x Intel i350 1Gb QP mezzanine (LAN)
iDRAC7 Enterprise w/ vFlash, 8GB SD

## 6 DVS Enterprise Server Scalability

As workloads increase, the solution can be scaled to provide additional compute and storage resources independently. Based on the best practices from VMware and DVS Engineering validation work the following scaling guidelines have been established for the virtual management infrastructure and the hosts.

### 6.1 R720 Rack-based Server Scalability

<i>Basic Users</i>	<i>Standard Users</i>	<i>Premium Users</i>	<i>Physical Mgmt Servers</i>	<i>Physical Host Servers - Basic</i>	<i>Physical Host Servers - Standard</i>	<i>Physical Host Servers - Premium</i>
145	116	95	1	1	2	2
1000	815	667	2	7	9	11
2000	1630	1333		14	18	22
3000	2444	2000		21	26	32
4000	3259	2667	3	28	35	43
5000	4074	3333		35	44	53
6000	4889	4000		42	52	64
7000	5704	4667	4	49	61	74
8000	6519	5333		56	69	85
9000	7333	6000		63	78	95
10000	8148	6667	5	69	87	106

### 6.2 M620 Blade-based Server Scalability

<i>Basic Users</i>	<i>Standard Users</i>	<i>Premium Users</i>	<i>Physical Mgmt Servers</i>	<i>Physical Host Servers - Basic</i>	<i>Physical Host Servers - Standard</i>	<i>Physical Host Servers - Premium</i>
135	110	90	1	1	2	2
1000	815	667	2	8	10	12
2000	1630	1333		16	19	24
3000	2444	2000		24	29	36
4000	3259	2667	3	32	38	48
5000	4074	3333		40	48	59
6000	4889	4000		48	57	71
7000	5704	4667	4	56	67	83
8000	6519	5333		64	76	95
9000	7333	6000		72	85	106
10000	8148	6667	5	80	95	118




## 7 Network Infrastructure

The DVS Enterprise initial base configuration begins with either a single Force10 S55 or PowerConnect PC6248 top of rack switch. This is rack-based solution meets the minimum switching requirements, with no switching high-availability option. Routing decisions will be made at the network core.


### 7.1 PowerConnect PC6248 1GbE Switch

The PC6248 is a 48 port, 1GbE, layer 2 edge switch. The PC6248 switch supports up to four 10 Gigabit Fibre (SFP+) & two 10GBase-T copper Ethernet uplinks for connectivity directly to 10GE servers, routers, enterprise backbones and data centers.



### 7.2 Force10 S55 1GbE Switch

The Dell Force10 S-Series S55 1/10 GbE top-of-rack (ToR) 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.




## 7.3 High Performance Ethernet Switching Options

As environments expand and VDI user requirements grow, or a shared tier 1 storage solution chosen as a part of the DVS Enterprise, you can choose to implement alternate top of rack network solutions in lieu of the PowerConnect PC6248 or Force10 S55. Taking DVS top of rack switching to the next level, you have more choices in the newly added switching model options.


### 7.3.1 Force10 S60 1GbE 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 applications and services generating sporadic large bursts of data 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.



### 7.3.2 Force10 S4810 10GbE 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 40Gbps 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.



## 7.4 Fibre Channel Switching Options

### 7.4.1 Brocade 6510 Fibre-Channel Switch

The Brocade® 6510 Switch meets the demands of hyper-scale, private cloud storage environments by delivering market-leading speeds up to 16 Gbps FC technology and capabilities that support highly virtualized environments. Designed to enable maximum flexibility and investment protection, the Brocade 6510 is configurable in 24, 36, or 48 ports and supports 2, 4, 8, or 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.



### 7.4.2 Brocade M5424 Blade Switch



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

### 7.4.3 QLogic QME2572 Host Bus Adapter

The QLogic® QME2572 is a dual-channel 8Gbps 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).



### 7.4.4 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 3200MBps (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.




## 7.5 Logical Networking

The logical network is designed with four VLANs for added security and traffic isolation as outlined below:

We created a number of VLANs to isolate and manage traffic.

- iSCSI VLAN: a switched only VLAN.
- ESXi VLAN: ESXi management traffic – routed
- VDI VLAN: VDI infrastructure traffic – routed.
- Management VLAN: all hardware management traffic, managing via ESXi hosts iDRACs, EqualLogic storage units, network switches etc. routed

Three VLANs have routing interfaces on the core network.



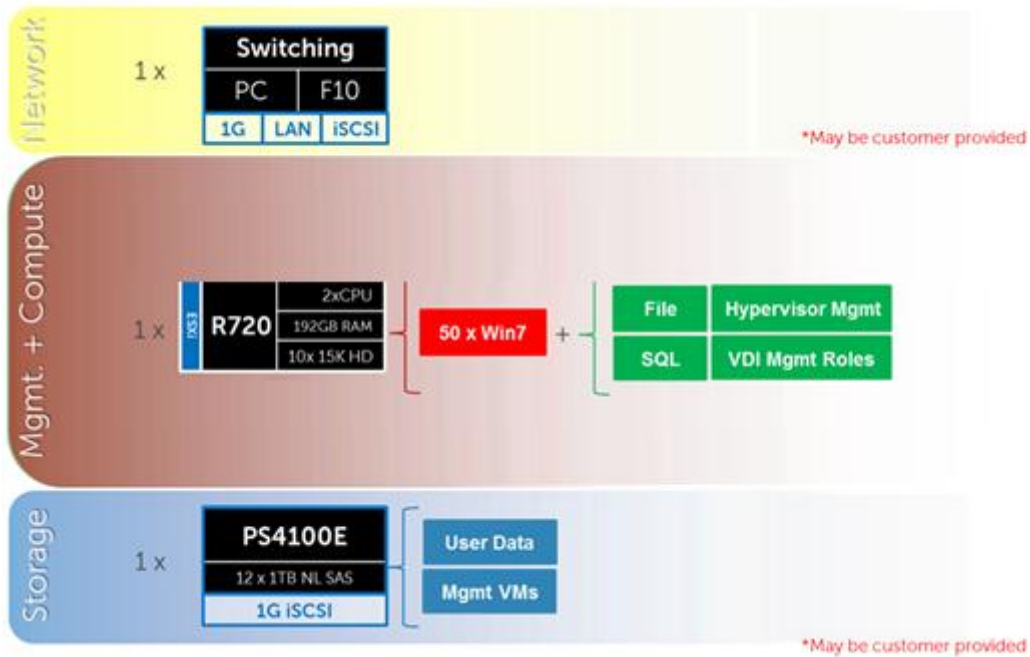
## 8 Scaling the Solution

As your VDI needs grow, so does the depth and breadth of the DVS Enterprise. The various management components exist as virtual server instances running on physical servers. This provides tremendous flexibility when adding resources to the solution while keeping the existing components intact and untouched. This also adds a level of resiliency in being able to backup copies of the server virtual images. The following tables gives an overview of Dell's recommended scaling of management components.

### 8.1 Local Tier 1 (RAID enabled local disks on Server with Tier 2 Shared Storage)

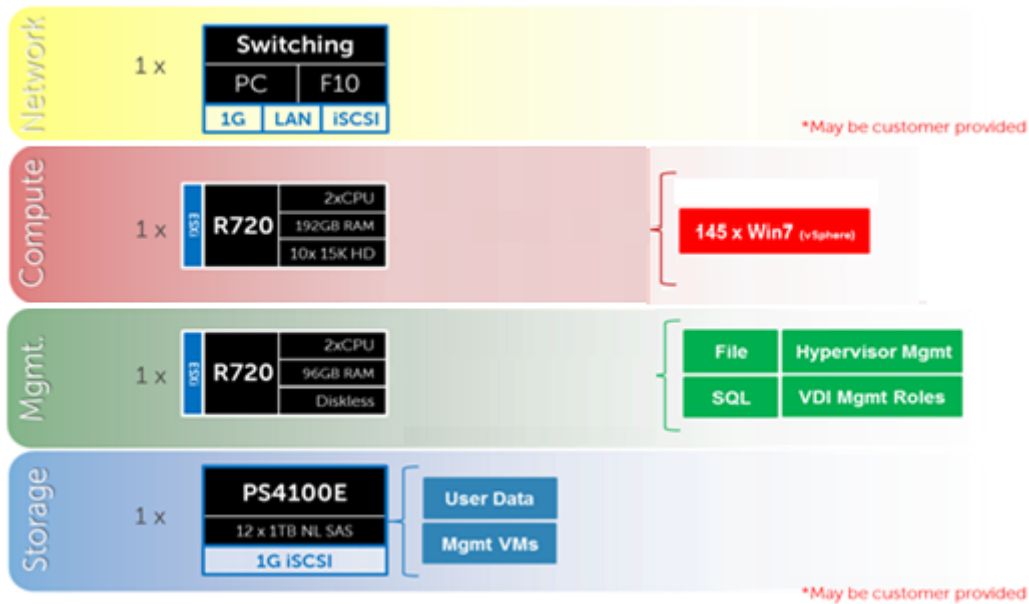
#### 8.1.1 Local Tier 1 – 50 User/ Pilot - (RAID enabled local disks on Server with Tier 2 Shared Storage)

For a very small deployment or pilot effort, we offer a 50 user/pilot solution. This architecture is non-distributed with all VDI and Management functions on a single host. If additional scaling is desired, you can grow into a larger distributed architecture seamlessly with no loss on your initial investment.




### 8.1.2 Local Tier 1 (iSCSI) - (RAID enabled local disks on Server with Tier 2 Shared Storage)

The local tier 1 solution model provides a scalable rack-based configuration that hosts user VDI sessions on local disk in the Compute layer.



#### 8.1.2.1 Local Tier 1 – Network Architecture (iSCSI)

In the local tier 1 architecture, a single PowerConnect or 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 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.



**8.1.2.2 Local Tier 1 Rack Scaling Guidance (iSCSI)**

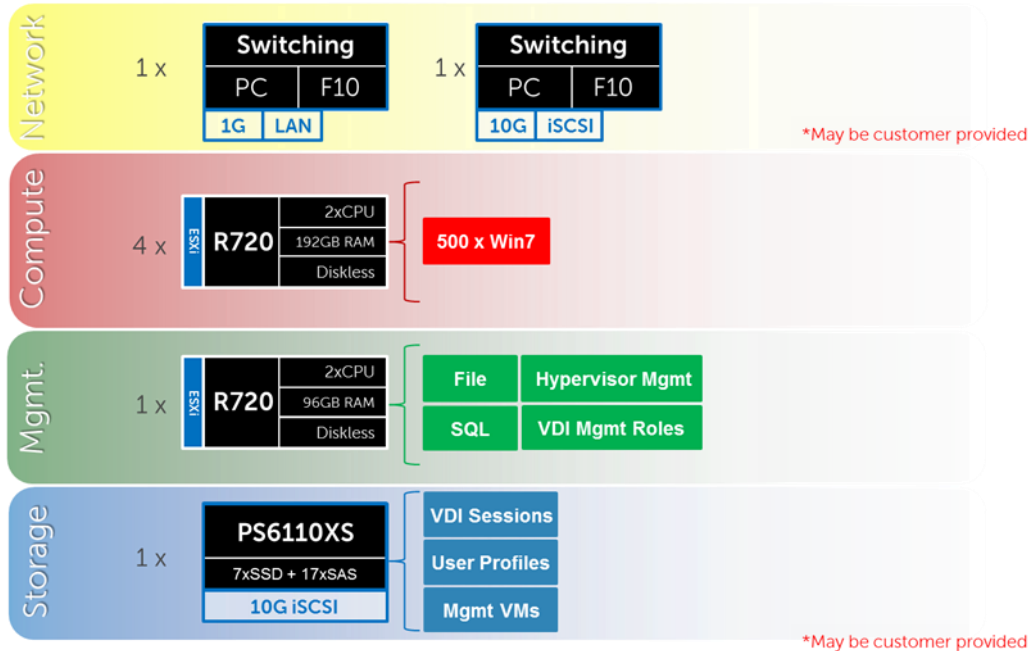
Local Tier 1 HW Scaling (iSCSI)				
User Scale	ToR LAN	ToR 1Gb iSCSI	EQL T2	EQL NAS
0-1000	S55	S55	4100E	
0-1000 (HA)	S55	S55	4100E	FS7600
0-3000	S55	S55	6100E	FS7600
3000-6000	S55	S55	6500E	FS7600
6000+ users	S60	S60	6500E	FS7600



## 8.2 Shared Tier 1 Rack

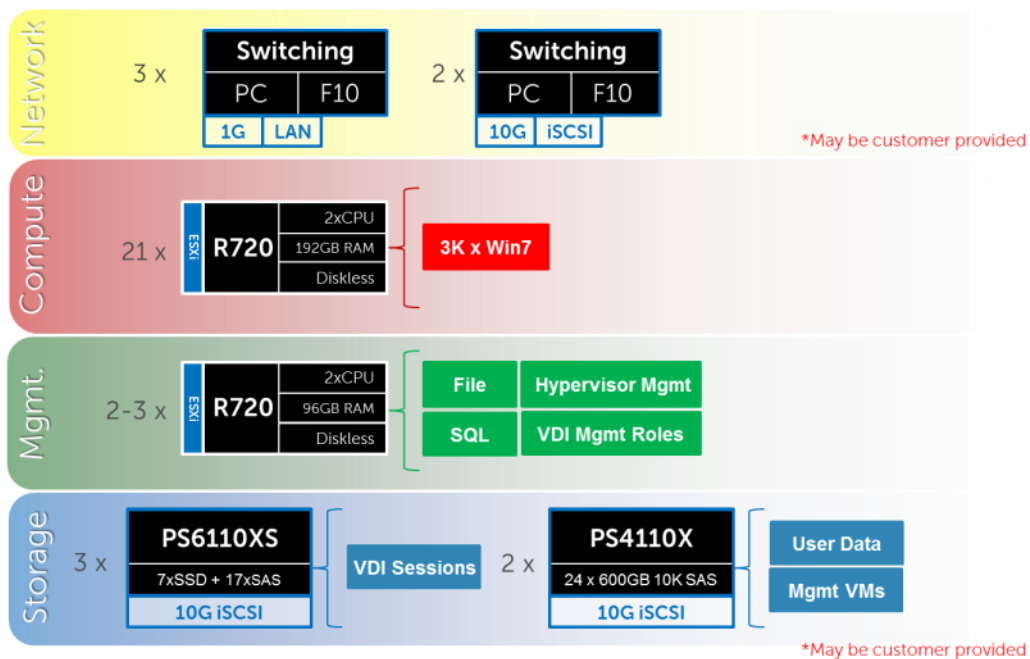
### 8.2.1 Shared Tier 1 (SAN) – Rack – 500 users (iSCSI – EqualLogic)

For POCs or small deployments, tier 1 and tier 2 can be combined on a single 6110XS storage array. Above 500 users, a separate array needs to be used for tier 2.




### 8.2.2 Shared Tier 1 (SAN) – Rack – 3000 users (iSCSI – EqualLogic)

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, 4110X, and 6510E are tier 2 array options. The addition of a NAS head is recommended above 1000 users as well as when optionally when providing high availability to file share services.



### 8.2.2.1 Shared Tier 1 Rack – Network Architecture (iSCSI)

In the Shared tier-1 architecture for rack servers, a single PowerConnect or 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 and making each stack redundant. Both Management and Compute servers connect to all VLANs in this model. All ToR traffic has been designed to be layer 2 / switched locally, with all layer 3 / routable VLANs routed through a core or distribution switch. The following diagrams illustrate the server NIC to ToR switch connections, vSwitch assignments, as well as logical VLAN flow in relation to the core switch.

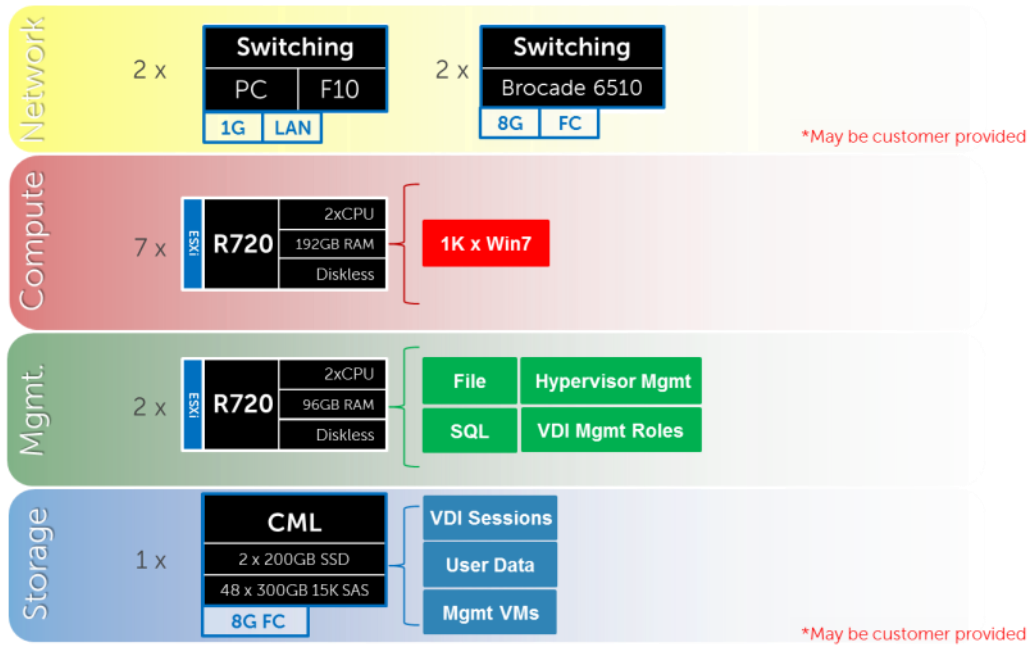


### 8.2.2.2 Shared Tier 1 Rack Scaling Guidance (iSCSI)

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	4110X	NX3300
3000-6000	S55	S4810	6110XS	6510E	NX3300
6000+	S60	S4810	6110XS	6510E	NX3300

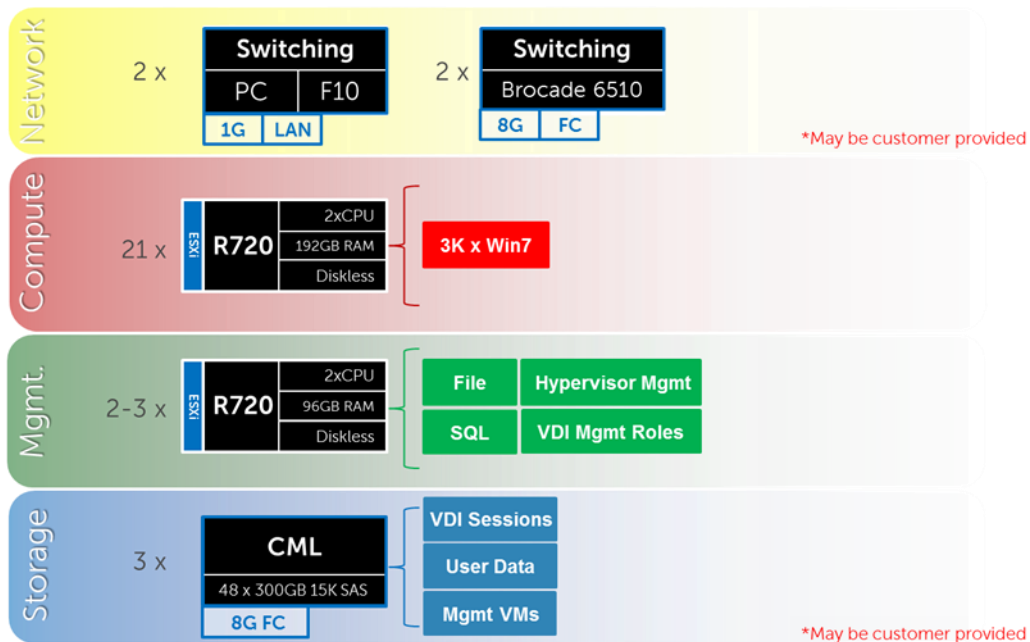
### 8.2.3 Shared Tier 1 (SAN) – Rack – 1000 Users (FC – Compellent)

Utilizing Compellent storage for shared tier 1 provides a FC solution where tier 1 and tier 2 are functionally combined in a single array. Tier 2 functions (user data + Management VMs) can be removed from the array if you have another solution in place. Doing this should net an additional 30% resource capability per Compellent array for user desktop sessions.



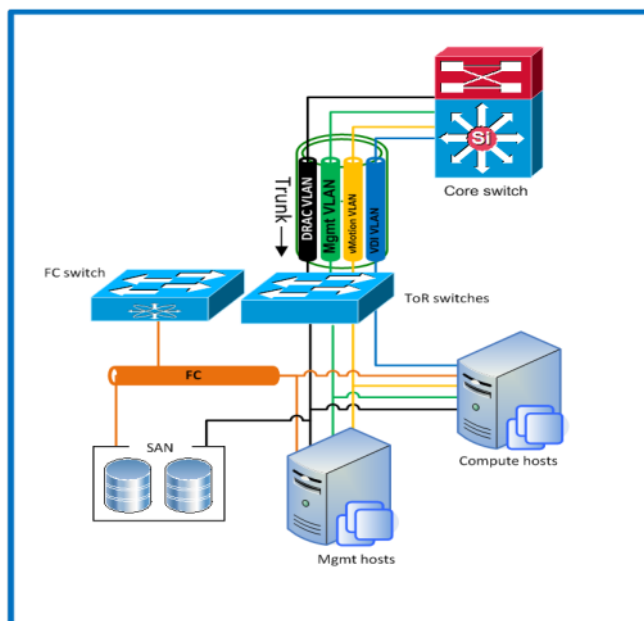
### 8.2.4 Shared Tier 1 – Rack (FC – Compellent)

FC is also supported in this model with discrete Compellent SC8000 arrays in tier 1 and tier 2. The Brocade 6510 is the FC switch of choice using 8Gb along with 8Gb FC IO cards in the Compellent array.



#### 8.2.4.1 Shared Tier 1 Rack – Network Architecture (FC)

In the Shared tier 1 architecture for rack servers using FC, a separate switching infrastructure is required for FC. Management and compute servers 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.



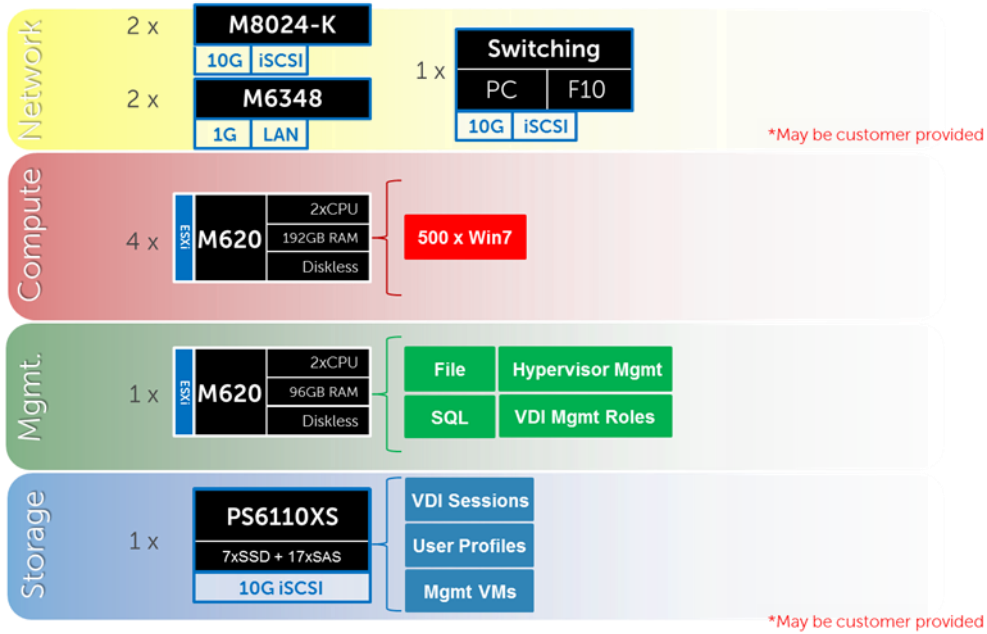
### 8.2.4.2 Shared Tier 1 Rack Scaling Guidance (FC)

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

### 8.3 Shared Tier 1 (SAN) Blade

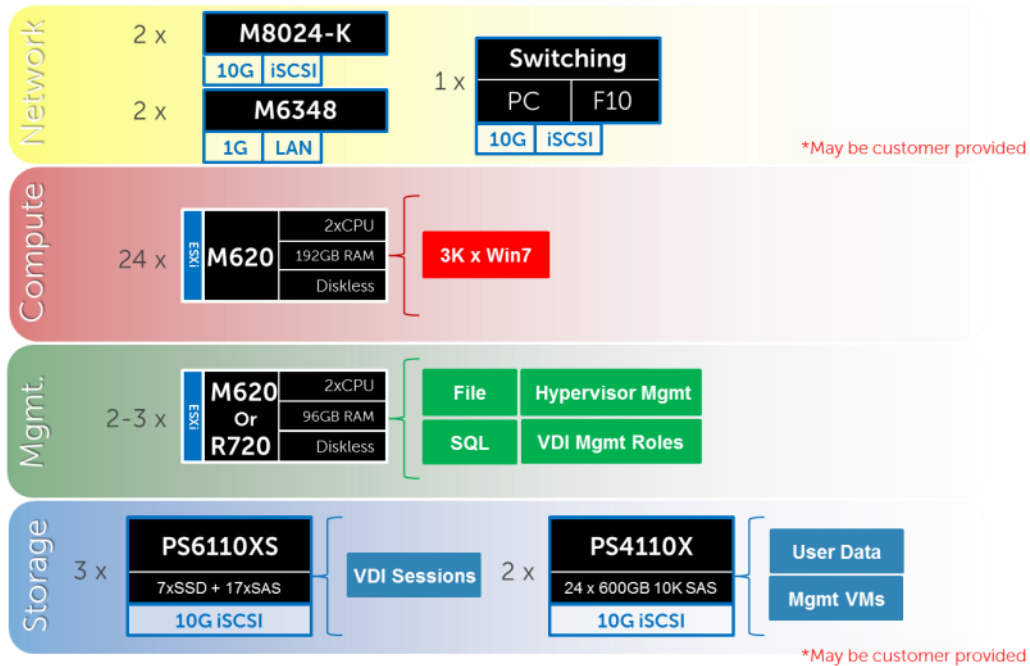
#### 8.3.1 Shared Tier 1 – Blade – 500 users (iSCSI – EqualLogic)

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, tier 1 and tier 2 storage should be separated into discrete arrays.




### 8.3.2 Shared Tier 1 (SAN) – Blade (iSCSI – EqualLogic)

Above 1000 users the Storage tiers need to be separated. At this scale we also separate LAN from iSCSI switching as well as add load balancing and NAS for SMB file shares. The drawing below depicts a 3000 user solution:



#### 8.3.2.1 Shared Tier 1 Blade – Network Architecture (iSCSI)

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



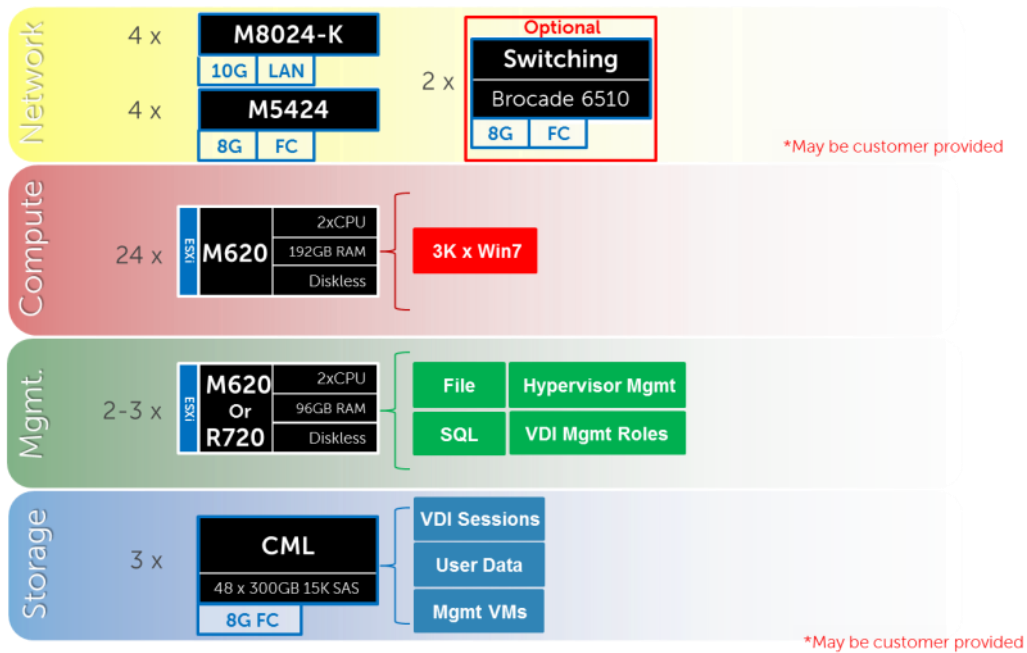
**8.3.2.2 Shared Tier 1 Blade Scaling Guidance (iSCSI)**

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	8024-K	S4810	6110XS	-	-
500-1000	M6348	8024-K	S4810	6110XS	4110E	-
0-1000 (HA)	M6348	8024-K	S4810	6110XS	4110E	NX3300
0-3000	M6348	8024-K	S4810	6110XS	4110X	NX3300
3000-6000	M6348	8024-K	S4810	6110XS	6510E	NX3300
6000+	M6348	8024-K	S4810	6110XS	6510E	NX3300

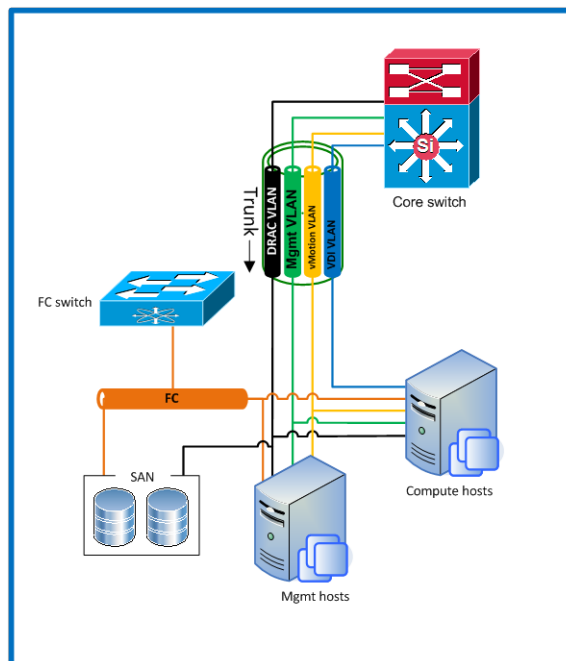


### 8.3.3 Shared Tier 1 (SAN) – Blade (FC – Compellent)

FC 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 there are FC IO cards in the Compellent arrays. ToR FC switching is optional if a suitable FC infrastructure is already in place.



#### 8.3.3.1 Shared Tier 1 Blade – Network Architecture (FC)

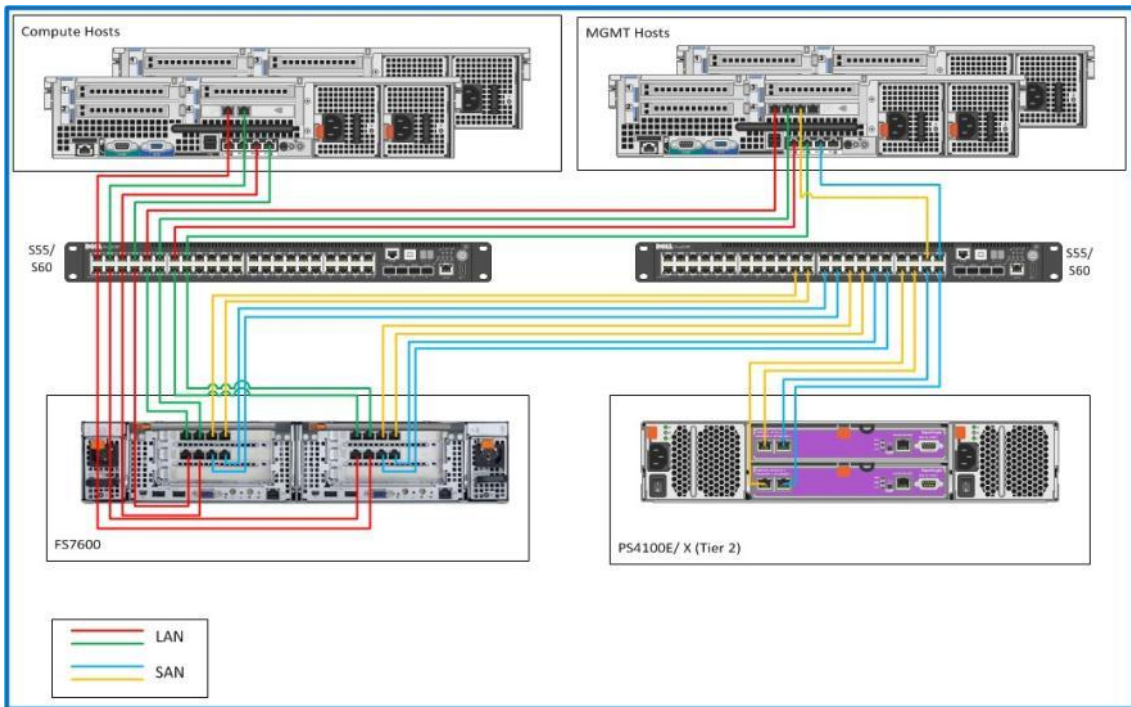


**8.3.3.2 Shared Tier 1 Blade Scaling Guidance (FC)**


Shared Tier 1 HW scaling (Blade - FC)						
User Scale	Blade LAN	Blade FC	ToR FC	CML T1	CML T2	CML NAS
0-500	8024-K	5424	6510	15K SAS	-	-
500-1000	8024-K	5424	6510	15K SAS	-	-
0-1000 (HA)	8024-K	5424	6510	15K SAS	NL SAS	FS8600
1000-6000	8024-K	5424	6510	15K SAS	NL SAS	FS8600
6000+	8024-K	5424	6510	15K SAS	NL SAS	FS8600

## 8.4 Cabling Diagrams


### 8.4.1 Local Tier 1 Cabling (iSCSI and LAN separated)




### 8.4.2 Shared Tier 1 Cabling (Rack – EqualLogic)




### 8.4.3 Shared Tier Cabling (Blade - EqualLogic)



### 8.4.4 Shared Tier 1 Cabling (Rack – Compellent)

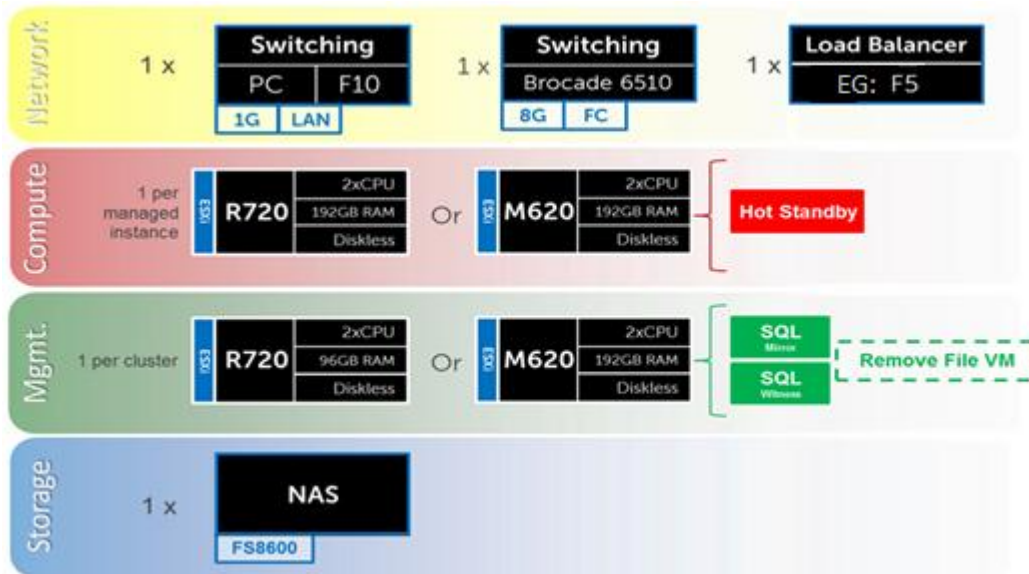


### 8.4.5 Shared Tier 1 Cabling (Blade – Compellent)



## 8.5 Building a Resilient Infrastructure

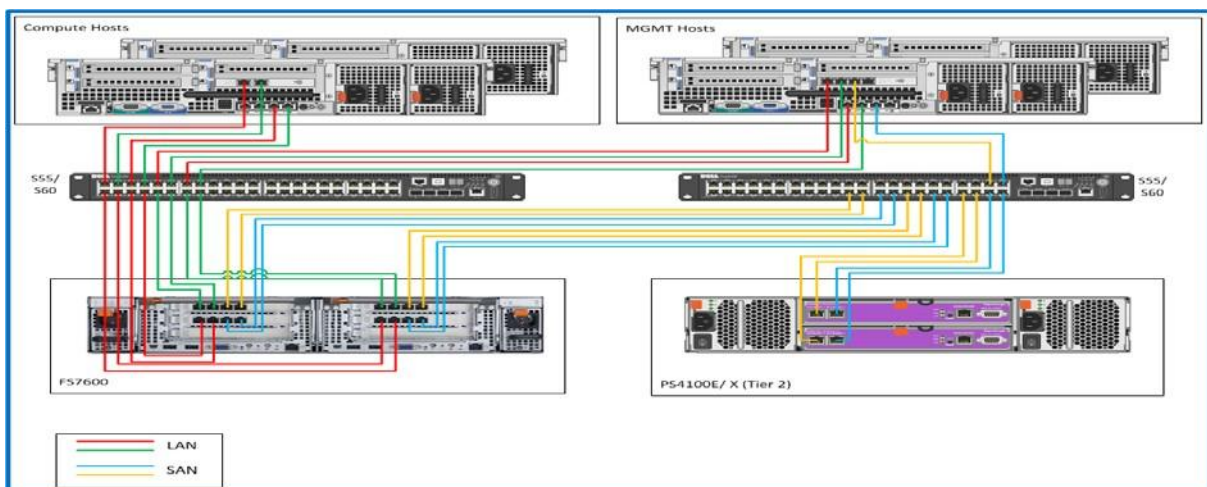
High Availability is achieved with a simple N+1 configuration for each component in the stack, with the exception of the storage arrays which have redundant controllers, network connections and RAID configured disks. The concept is similar for Enterprise blade solutions except top of rack switching would correspond to 10 GbE solutions.



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.
- A number of enhancements occur at the Management tier, the first of which is the addition of another ESXi host. The Management ESXi hosts will then be configured in an HA cluster with vMotion enabled. All applicable VMware server roles should then be duplicated on the new host. SQL will also receive greater protection through the addition and configuration of a SQL mirror with a witness.

### 8.5.1 High Availability Cabling



## 8.5.2 High Availability Networking

### 8.5.2.1 vSphere Virtual Networking

The optional HA bundle adds an additional ESXi host in the Management tier providing the ability to use an HA cluster between the hosts. vMotion will be enabled as well and should be added to the vSwitch housing the VMkernel ports on the Management hosts. Therefore will be no vSphere HA cluster in the compute tier.

Physical network connections should be spread across both switches, for each host, to ensure that each vSwitch has redundancy for each connection. The vSwitch housing the VMkernel ports on the Management hosts will need to be modified to add a VMkernel port. The VLAN used for vMotion should be private (non-routable) and only accessible to each Management host. The physical adapters in the vSwitch should be separated between SC/vMotion and iSCSI ports but configured to provide failover for each other. This can be achieved by specifying an explicit failover order.

On the ESXi Host, the Dell EqualLogic MPIO plugin will be installed to handle load balancing. This module will be added via the command line tool using a Virtual Management Appliance (vMA) in vCenter. This tool will allow for easy configuration of iSCSI on each host. Some key settings that will be used as part of the configuration;

- Sets 2 IP Addresses for iSCSI on each host
- Specifies NIC (vmNIC2, vmNIC3)
- Sets the Jumbo Frame Settings (MTU 1500 – 9000 not supported with iSCSI offload on Broadcom NIC's)
- Initializes software iSCSI
- Sets IP for the EqualLogic Storage group.

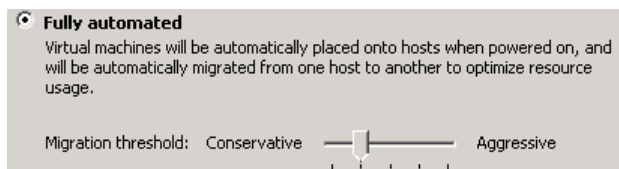
Once the MPIO setup is complete the vSphere host is ready to access the storage and to either create or connect to VMFS data stores. The MPIO plugin will configure the correct multi-pathing for the data stores also.

Note: The iSCSI VLAN's need to be tagged on both iSCSI VMkernel ports, as the iSCSI VLAN is sharing the link with the Management Network. The MPIO script does not complete this step and this has to be performed on each host in vCenter. Each host will have dedicated LUNS assigned that will be visible to all hosts in that specific cluster.



## 9 HA DRS – Load Balancing - DNS

DRS provides the ability for automated VM placement within the cluster when it is powered on. DRS can be optionally used on the Compute layer set to Fully Automated with a Conservative migration threshold set. Affinity rules should be used on the Management servers to ensure that View broker roles are properly distributed.



The vSwitch that carries management traffic will need to be modified to add a VDVS Enterprise kernel port for VMotion on both Management hosts. The VLAN used for VMotion should be private (non-routable) and only accessible to each host in the cluster. The physical adapters in the vSwitch should be separated between SC/VMotion ports but configured to provide failover for each other. This can be achieved by specifying an explicit failover order.

### 9.1 Management Server High Availability

High availability can be enabled on any or all layers in the solution. Following an N+1 methodology protects all layers of the solution architecture. Additional ToR switches can be added to the Network layer, additional Compute and Management hosts are added to each vCenter instance in their respective layers, and a NAS device is added to the Storage layer. Storage arrays are inherently redundant. The applicable core VMware View roles will be load balanced via DNS by default. A load balancing solution, such as F5, is recommended to manage load-balancing efforts for environments requiring HA.

### 9.2 Windows File Services High Availability

High Availability for file services will be provided by the FS7600, FS8600 or PowerVault NX3300 clustered NAS head. See section 6.6.5 for hardware details. To ensure proper redundancy, the NAS should have its cabling split between the switches, please refer to HA cabling diagram in Section 3.1 above as both NAS heads will need to be cabled accordingly.

Unlike the FS8600, the FS7600 and NX3300 do not support for 802.1q (VLAN tagging) so connecting switch ports should be configured 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 original file server VM was configured to use RDMs 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.

### 9.3 SQL Databases

The VMware databases will be hosted by a single dedicated SQL 2008 R2 Server VM in the Management layer, in base form. Care should be taken 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:

- VMware vCenter
- VMware View Composer
- vCenter Update Manager (optional)

Initial placement of all databases into a single SQL instance is fine unless performance becomes an issue, in which case database should be separated into separate named instances. Enable auto-growth for each DB. See evolutionary section below for more information.



Best practices defined by VMware for View should 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, disk partitions should be configured to begin from a sector boundary divisible by 64K.

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

```
DISKPART> select disk x
Disk x is now the selected disk.
DISKPART> create partition primary align=1024
DiskPart succeeded in creating the specified partition.
```

### 9.3.1 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 9.5.1 for more details. The principal VM that will host the primary copy of the data should exist on the first Management host. The mirror and witness VMs should exist on the second or later Management hosts. All critical databases should be mirrored to provide HA protection.



### 9.4 Load Balancing

The applicable core VDI roles will be load balanced by default using DNS. F5 LTM or NetScaler VPX can be added at any time to manage load-balancing efforts. All roles whose configurations are stored in SQL can be protected via an optional SQL mirror. With Storage layer HA option the file server VM is replaced by NAS.

### 9.5 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 VMware 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 should be adhered to.

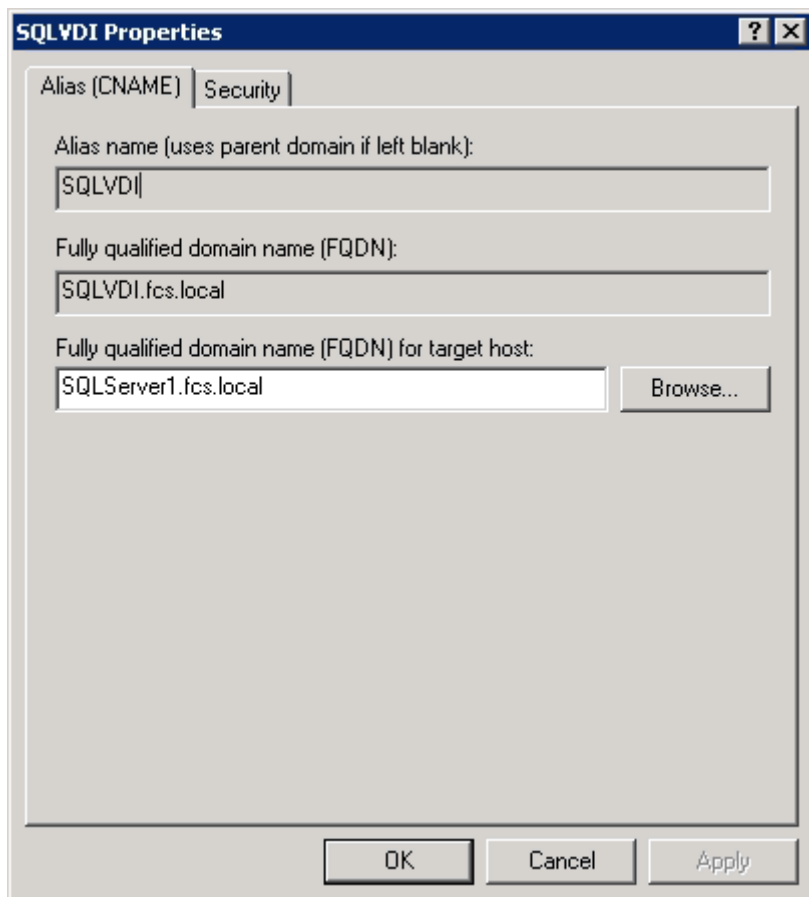
To plan for eventual scaling, access to components that may live on one or more servers (SQL databases, View infrastructure services) should be considered during initial deployment. The use of CNAMEs and the round robin DNS mechanism should be employed to provide a front-end “mask” to the back-end server actually hosting the service or data source.

#### 9.5.1 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, connections should be made to aliases in the form of DNS CNAMEs. So instead of connecting to SQLServer1\

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

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



### 9.5.2 DNS for Load Balanced Services

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

For example, consider a VM called WebInterface that has its own hostname registered in DNS as an A record. You should then also create a new A record to be used should additional servers 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! This case shows the creation of three new A records called "WebInterface", all presumably pointing to three different servers.

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

When a client requests the name WebInterface, DNS will direct them to the 3 hosts in round robin fashion. The following resolutions were performed from 2 different clients:

```
Administrator: Command Prompt
C:\Users\pfine>ping WebInterface
Pinging WebInterface.fcs.local [10.1.1.1] with 32 bytes of data:
Reply from 10.1.1.1: bytes=32 time=1ms TTL=249
Reply from 10.1.1.1: bytes=32 time<1ms TTL=249
Reply from 10.1.1.1: bytes=32 time<1ms TTL=249
Reply from 10.1.1.1: bytes=32 time<1ms TTL=249

Ping statistics for 10.1.1.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 1ms, Average = 0ms

C:\Users\pfine>ping WebInterface
Pinging WebInterface.fcs.local [10.1.1.2] with 32 bytes of data:
Reply from 10.1.1.2: bytes=32 time=1ms TTL=249
Reply from 10.1.1.2: bytes=32 time<1ms TTL=249
Reply from 10.1.1.2: bytes=32 time<1ms TTL=249
Reply from 10.1.1.2: bytes=32 time<1ms TTL=249

Ping statistics for 10.1.1.2:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 1ms, Average = 0ms

C:\Users\pfine>_
```

This method of creating an identical but load-balanced namespace should be repeated for all applicable components of the architecture stack.

## 10 Customer Provided Stack Components

### 10.1 Customer Provided Storage Requirements

In the event that you wish to provide your own storage array solution for an Enterprise 6010-based solution, the following minimum hardware requirements must be met.

Feature	Minimum Requirement	Notes
Total Storage Space	User count and workload dependent	For user workload disk space requirements, see Section 4.1.
Drive Support	7200rpm 3.5" NLSAS	The minimum optimal drive configuration utilizes 6 drives for the best RAID 50 performance.
Tier 1 IOPS Requirement	(Total Users) x 10 IOPS	
Tier 2 IOPS Requirement	(Total Users) x 1 IOPS	
Data Networking	6x 1GbE RJ45	
Drive Controllers	1 with >4GB cache	4GB of cache minimum per controller is recommended for optimal performance and data protection.
RAID Support	50	RAID 50 is used on the external storage array. RAID 10 is leveraged only for local storage on compute servers.

### 10.2 Customer Provided Switching Requirements

In the event that you wish to provide your own rack network switching solution for the Enterprise - based solution, the following minimum hardware requirements must be met.

Feature	Minimum Requirement	Notes
Switching Capacity	180Gbps/switch	
10Gbps Ports	None Required	The base DVS Enterprise rack-based configuration is based on 1Gbps network connectivity.
1Gbps Ports	5x per Management server 5x per Compute Server 6x per Storage Array	
VLAN Support	IEEE 802.1Q tagging and port-based VLAN support.	
Stacking Capability	Yes	The ability to stack switches into a single management view for an HA configuration is highly recommended.

## 11 Dell Wyse Cloud Clients

Dell Wyse Cloud Client devices and software provide superior security, reliability and energy efficiency when compared to a traditional PC. Dell Wyse desktop devices and software help streamline the delivery of VMware Horizon View 5.2 infrastructure to millions of users around the world. Thin Clients create a more secure environment that minimizes or eliminates exposure to data loss, viruses and malware. By utilizing thin clients as the access device for end user, deployments can benefit from centralized management and complete control of all endpoints. Since thin clients eliminate components with high failure rates, deployments can expect reduced costs and improved reliability over the life of a desktop virtualization deployment.

### 11.1 Dell Wyse P25

Experience uncompromised computing, with the benefits of secure, centralized management. The Dell Wyse P25 PCoIP zero client for VMware Horizon View 5.2 is a secure, easily managed zero client that provides outstanding graphics performance for advanced applications such as CAD, 3D solids modeling, video editing and advanced worker-level office productivity applications. Smaller than a typical notebook, this dedicated zero client is designed specifically for VMware Horizon View 5.2. It features the latest processor technology from Teradici to process the PCoIP protocol in silicon, and includes client-side content caching to deliver the highest level of performance available over 2 HD displays in an extremely compact, energy-efficient form factor. The Dell Wyse P25 delivers a rich user experience while resolving the challenges of provisioning, managing, maintaining and securing enterprise desktops.

#### Dell Wyse P25 and Display Recommendations



Click [HERE](#) for more information on the Dell Wyse P25.

Display recommendations for the P25 Zero Client are listed below



The P2412H shown above supports 1920x1080, VGA, DVI and USB. Other options include the E2213 with 1680x1050, VGA and DVI, and the E1913 with 1440x900, VGA and DVI.

## 11.2 Dell Wyse Z50D

Designed for power users, the new Dell Wyse Z50D is the highest performing thin client on the market. Highly secure and ultra-powerful, the Z50D combines Dell Wyse-enhanced SUSE Linux Enterprise with a dual-core AMD 1.6 GHz processor and a revolutionary unified engine for an unprecedented user experience. The Z50D eliminates performance constraints for high-end, processing-intensive applications like computer-aided design, multimedia, HD video and 3D modeling. Scalable enterprise-wide management provides simple deployment, patching and updates. Take a unit from box to productivity in minutes with auto configuration. Delivering unmatched processing speed and power, security and display performance, it's no wonder no other thin client can compare.

### Dell Wyse Z50D and Display Recommendations



Click [HERE](#) for more information on the Dell Wyse Z50D.

Display recommendations for the Z50D are listed below



The P2212H shown above supports 1920x1080, VGA, DVI and USB. Another option includes the E2213 with 1680x1050, VGA and DVI.

### 11.3 Dell Wyse D50D

Designed for power users, the new Dell Wyse D50D is the highest performing thin client on the market. Highly secure and ultra-powerful, the D50D combines Dell Wyse-enhanced SUSE Linux Enterprise with an AMD G-Series T48E Dual Core 1.4GHz processor and a revolutionary unified engine for an unprecedented user experience. The D50D eliminates performance constraints for high-end, processing-intensive applications like computer-aided design, multimedia, HD video and 3D modeling. Scalable enterprise-wide management provides simple deployment, patching and updates. Take a unit from box to productivity in minutes with auto configuration. Delivering unmatched processing speed and power, security and display performance, it's no wonder no other thin client can compare.

#### Dell Wyse D50D and Display Recommendations



Click [HERE](#) for more information on the Dell Wyse D50D.

Display recommendations for the D50D are listed below



The P2212H shown above supports 1920x1080, VGA, DVI and USB. Another option includes the E2213 with 1680x1050, VGA and DVI.

## 12 Dell DVS VMware Horizon View Solution New Feature Sets

## 12.1 High-end VMware vDGA / Pass-thru graphics support

*Note: vDGA support with NVIDIA GRID cards is a tech preview option expected to go RTS/GA very soon. vDGA is supported at the ESXi hypervisor level but not currently with the NVIDIA K1 and K2 GRID cards and with Horizon View 5.1.*

### 12.1.1 Executive Summary

VMware vDGA otherwise known as Pass-thru graphics support is the technology of mapping a VMware Horizon View desktop directly to a GPU on a high-end graphics card such as an NVIDIA K1 or K2 card. These cards along with vDGA modes offer end users using virtual desktop access to run very high end graphics intensive applications such as CAD or other graphics editing/authoring software packages. This solution offers greater density of high end graphics enabled users per server over the traditional one to one model of a graphics workstation. As such cost benefits can be realized by using VMware vDGA, VMware Horizon View, and Dell hardware technologies.

The capability for 3D graphics and video in VMware Horizon View further expands the use cases and target users that IT can deliver with virtual desktops. vDGA allows the IT department to deliver virtual desktops to users who traditionally would have required high performance workstations in order to get the graphics performance that they required for running applications like AutoCAD and eDrawings e.g. engineering and Oil and Gas sectors.

The critical difference between vDGA and vSGA is that when using vDGA the virtual machine has full control and usage of the assigned GPU, it is passed through the hypervisor to the VM and the driver is installed locally on the VM whereas in vSGA the GPU is shared amongst a pool of virtual machines.

vDGA offers DirectX 9, 10 and 11 support and it can also support OpenGL 2.1 / 3.x / 4.1x whereas vSGA can only support DirectX9 and Open GL2.1

vDGA can be costly to implement but can potentially offer a reduction in cost compared to purchasing individual high end workstations. The amount of GPU's that you can install in a server also will have a bearing on the cost per user for vDGA, for all the tests discussed in this document we used two cards at any one time (two K1's or two K2's)

Broadly speaking, users of graphics-intensive applications in a virtual desktop environment can be subdivided into 2 categories, as discussed below:

- "Premium Plus" VDI users are users who may be consuming relatively high-end graphics through relatively high frame-rate applications such as Google Earth, graphics-rich HTML5 pages etc. and also reviewing electrical, mechanical CAD drawings etc. *[The term 'Premium Plus' is used to distinguish this user type from the existing 'Premium' user type that is used in current DVS-Wyse PAAC and Sizing Activities].*
- "Workstation users", as the name implies, are users who would typically have used high-end physical workstations (e.g. Dell Precision); typical activities carried out by these users would include 3D modeling for the oil and gas industry, involving a large amount of resource –intensive activities such as model rotation etc.
- The validation effort described in this document is for Workstation users. In the cases described here the Grid cards are operating in Pass-Through mode, i.e. a VM has direct access to a GPU on the K1 or K2 card. The GPU cannot be shared with other VMs on the Hypervisor. No other VMs are provisioned on the server except those involved in the Workstation user testing. It is assumed that the server will be dedicated to the number of workstation users that can be supported by the appropriate Grid card in pass-through mode.
- This results document describes validation efforts undertaken on NVidia Grid cards K1 and K2 to study their behavior when used with graphics-intensive applications.



- All of the following results were gathered with either two K1 cards or two K2 cards installed in the server. The K1 card allows 4 GPU pass through sessions per card (8 per Server), while the K2 card allows 2 pass-through sessions per card (4 per server).


## 12.1.2 Graphics-Specific Performance Analysis Results

### 12.1.2.1 K1 Tests – Fixed Frame-Rate Video Component – Single VM

When considering EUE in terms of perceived video / graphics smoothness, a useful domain to use for assessment is broadcast video. NTSC (US analog television system) is transmitted at 30 FPS, while PAL (widely used in Europe) is transmitted at 25 FPS (Frame per Second). A movie clip (“The Hobbit” trailer) at 30 FPS has been created and is used in the following tests.


In the graph below, we can see that the OptiPlex 7010 is able to maintain a much higher FPS than the Wyse z90d7 for a ten minute 30 FPS video clip.

Both however were able to maintain an FPS over 25.




In the next graphs, we see the difference in the FPS from both endpoints during the same video trailer when the VM is utilizing a GPU and when **not** utilizing the GPU.

### K1 Single VM - Movie Clip - Optiplex 7010




### K1 Single VM - Movie Clip - Wyse Z90D7



From these two graphs it shows that when the VM is using the GPU then it can maintain a higher FPS (more noticeable in the OptiPlex 7010 rather than the Wyse Z90D7).

#### 12.1.2.2 K1 Heaven Benchmark Testing – Single VM

Heaven Benchmark is a GPU-intensive benchmark that significantly stresses graphics cards. This benchmark tool can be effectively used to determine the stability of a GPU under extremely stressful conditions, as well as validating the characteristics of the card's thermal subsystems. The next graph shows the Heaven Benchmark test on a single VM.



As shown above in the benchmark graph, the Wyse endpoint is not able to maintain a high enough FPS during the Heaven Benchmark so it was not used in any further heaven tests for K1. (The P25/P45 Wyse clients that showed significant performance gains were not used for K1 tests but were characterized with the K2 tests and show to be a perfect match for these K1 workloads as well)

**Note:** Heaven Benchmark produces a score after each benchmark test; however the scores can be unreliable so they are not included in this document.


**Note:** Heaven benchmark was executed in low quality with a resolution of 640 x 480 (lowest possible) – if we ran the program at a higher quality / resolution then the fps would have been under 25 all the time.

### 12.1.2.3 K1 test – Movie Clip + Companion testing

As a companion workload for K1, eDrawings Solid works Viewer was used and set to “Advanced Animation” to rotate indefinitely. The eDrawings activity is not very high in terms of graphics resource utilization. This is appropriate since it is used as a companion workload for the K1 card which is targeted at a lower end graphics scenario.


For the K1 cards, seven VMs running Companion workloads were used.

Companion workloads were executed on the remaining VMs on the host (7) to establish whether or not they would have any impact on the test results of the benchmark VM. In theory they should not have any impact as each virtual machine has access to its own dedicated GPU.






This graph shows almost identical information to the above mentioned graph when just a single VM was used. This confirms that companion workloads on the other VMs do not impact the benchmark VM because the GPU's are assigned to each VM directly and there is little / no contention.


#### 12.1.2.4 K1 test – Heaven Benchmark + Companion testing



Again, this graph is comparable with the earlier graph of the Heaven benchmark test which further shows that the companion workloads did not impact the test results of the benchmark VM.

The following graphs were gathered when the system was running the Heaven benchmark and seven eDrawings Companion tests.







These graphs show there were no host resource issues throughout these tests. CPU spikes a little when the Heaven Benchmark is running but is well within acceptable ranges. Data store latency and network usage are all within the capabilities of the host.

#### 12.1.2.4 K1 NVIDIA GPU Results

Note: The graph below shows the GPU usage during the heaven test – note that total GPU memory Usage is at 100%, this tell us that the K1 card is not capable of running very intensive graphics programs



In contrast – this graph below shows the GPU usage during the Hobbit test – note that Utilization is much lower than the Heaven test.



#### 12.1.2.5 K1 Subjective Tests

We were able to login to the OptiPlex 7010 and Wyse Z90D7 successfully and execute some subjective tests while the companion workloads were running on the host.

Performance was generally good but the OptiPlex was superior to the Wyse terminal.

When using the Wyse, there was occasional jitter and lag which resulted in a poorer user experience than the OptiPlex. There was less lag when using the OptiPlex.


#### 12.1.2.6 K1 Conclusion

The K1 cards work successfully with vDGA. Each VM has direct control over its assigned GPU and workloads on other VMs have little or no impact. However, the K1 cards are not capable of running very intensive graphics programs in high resolution and quality i.e. Heaven benchmark. When running graphics intensive programs, neither the OptiPlex nor Wyse terminal were able to maintain an FPS value of over 25.

Choice of endpoint is also very important with these cards as performance can vary a lot, we found that performance with the OptiPlex 7010 and the K1 card was superior to that of the Wyse Z90D7 and K1.


It should be noted that the Wyse P25 client was not tested with NVIDIA K1 Cards but was tested with K2 cards and that the performance was excellent and on par with the OptiPlex client.

#### 12.1.2.7 K2 Tests – Fixed Frame-Rate Video Component – Single VM




In the above shown graph, the OptiPlex client can maintain a higher FPS when the K2 card is in use as opposed to the K1 whereas the Wyse client is generally the same when using either the K1 or K2.

The graph below shows the average FPS during the Single VM Movie Clip test for both the K1 and K2 GRID cards and the OptiPlex and Wyse terminals. It's clear from this graph that K2 and OptiPlex 7010 offers the best performance.



#### 12.1.2.8 K2 Tests – Heaven Benchmark Testing – Single VM






In this graph above it is shown that when using the K2 card, the OptiPlex is capable of maintaining a higher FPS than the K1 during the Heaven Benchmark.

**Note:** The Wyse terminal used in these tests was incapable of maintaining an FPS over 15-20 so was excluded from the Heaven Benchmark Test.

In the graph below, the K2 GPU was running at a much lower utilization rate while running the Heaven Benchmark than the K1 GPU (running at same resolution / quality as the K1 test) showing greater overall capacity.




**Note:** Using the K2 card, it was possible to run the Heaven Benchmark in high quality mode and a resolution of 1366x 768 while maintaining an FPS of over 25.


#### 12.1.2.9 K1/ K2 Comparison test – Viewperf Single VM

These graphs show the difference in results between the K1 and K2 cards when running the various Viewperf benchmarks. Both endpoints show better results when using the K2 GRID card, further re-enforcing the fact that the K2 is a higher performing card.

## K1 / K2 Comparison Viewperf Single VM OptiPlex 7010



## K1 / K2 Comparison Viewperf Single VM Wyse Z90D7




### 12.1.2.10 K2 test – Viewperf + Companion testing

As a companion workload for K2 tests, we used AutoCAD – we executed a continuous orbit activity on a sample 3D drawing. This is appropriate since it is used as a companion workload for the K2 card which is targeted at higher end graphics scenario.


For the K2 cards, we had three VMs running Companion workloads.

Viewperf reports results in frames per second (FPS). Below the Viewperf score is shown when running the various benchmarks using the Wyse terminal and the OptiPlex 7010, both report similar results.



## K2 Viewperf Score With Companion Workloads Wyse Z90D7





## K2 Viewperf Score With Companion Workloads Optiplex 7010

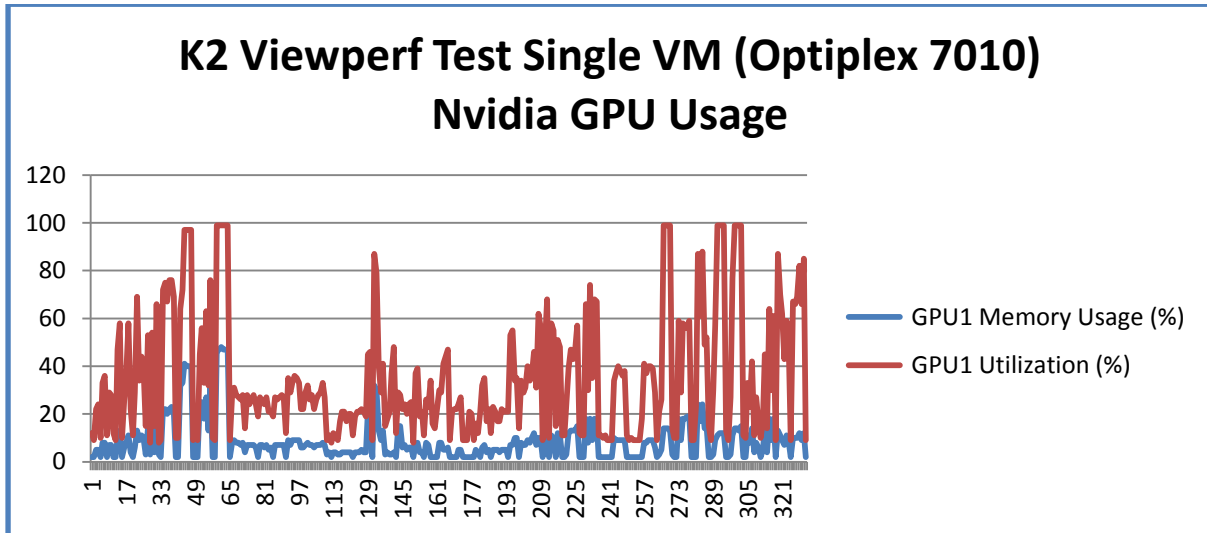


The following graphs were gathered when the system was running the ViewPerf Benchmark in conjunction with three Companion AutoCAD Workloads.






### 12.1.2.11 K2 Nvidia GPU Results + K1 Comparison



These graphs show that the utilization and memory usage of the GPU is much higher when using the K1 as opposed to the K2 when running the Viewperf benchmarks.



### 12.1.2.12 K2 Subjective Tests

Subjective tests were initiated successfully on the OptiPlex 7010 and Wyse Z90D7 while companion workloads were running on the hosts.

Repeated zoom activities were executed in AutoCAD while Viewperf was also running. These activities were designed to generate graphics loads visible to an end-user directly running subjective tests.

Performance using the OptiPlex was better than the Wyse.

There was less jitter and better responsiveness using the OptiPlex client, however both displayed

some "lag" when executing multiple activities simultaneously.

### 12.1.2.13 K2 Conclusion

The K2 cards offer improved performance over the K1 cards. Increased FPS was experienced during the movie clip and in particular the Heaven Benchmark tests.


The K2 cards are capable of running more graphics intensive workloads and are more suited to high end workstation type users than the K1 cards.

As with the K1, choice of endpoints is important as better performance was seen with the OptiPlex 7010 as opposed to the Wyse Z90D7

### 12.1.2.14 K2 P25 Zero Client Single VM Comparison Test

The Dell Wyse P25 PCoIP zero client for VMware View is a secure, easily managed zero client that provides outstanding graphics performance for advanced applications such as CAD, 3D solids modeling, video editing and advanced worker-level office productivity applications. Smaller than a typical notebook, this dedicated zero client is designed specifically for VMware View. It features the latest processor technology from Teradici to process the PCoIP protocol in silicon as opposed to in a soft client, and includes client-side content caching to deliver the highest level of performance available over 2 HD displays in an extremely compact, energy-efficient form factor. The Dell Wyse P25 delivers a rich user experience while resolving the challenges of provisioning, managing, maintaining and securing enterprise desktops.

In the following graph it shows that the FPS of the Wyse P25 is much better than the Wyse Z90D7 and is comparable with the OptiPlex 7010 client.





**Note:** Heaven was run at low quality and 640 x 480 resolutions for the above test.

### 12.1.2.15 K2 P25 / OptiPlex 7010 Heaven Test with Heaven running on all VMs (4)


The following graphs show NVidia GPU usage and host metrics while the Heaven Benchmark was running on all four VMs with a resolution of 1366 x 768 and quality of high. In an effort to increase Host CPU Usage, we used these settings.

Tests were run to compare the Wyse P25 and OptiPlex 7010 to see if there was any difference in the host metrics.


### OptiPlex 7010












## Wyse Z90D7





## Wyse P25





## 12.1.3 Conclusions

### 12.1.3.1 Test Conclusions

vDGA works well and offers very good graphics performance. Setup and installation are straight forward.

The K1 card is definitely more suited to lighter graphics users than the K2 card. The K2 card is capable of running much more graphics demanding applications like Heaven Benchmark and the K2 GPU Usage is much lower than the K1 GPU Usage when running the more demanding applications. The K1 card was only capable of running Heaven in the lowest resolution and lowest quality available in the application, whereas the K2 was able to display a much higher FPS when running Heaven at a higher resolution and quality.

Choice of endpoint is also very important with vDGA because there is a big difference in the frame rate per second between the Wyse Z90D7 and the Wyse P25 for instance. The P25 offers FPS comparable to the OptiPlex 7010 whereas the Z90D7 lags behind. The P25 is much more efficient at rendering the displayed graphics than the Z90D7.

The Apex 2800 PCOIP Offload card has proven in the tests to offer an advantage in terms of CPU Usage. CPU usage declined by approximately 10% when the Apex card was disabled, Note that only two VMs were used in our tests of the Apex 2800. These results (Apex) will be further explored in coming iterations of this RA. The Apex card is not currently a DVS/CCC offering in the VMware Horizon View solution but was tested here to determine future viability.

With vDGA the GPU is passed through to the virtual machine so any workloads running on other virtual machines shouldn't affect the Frame rate on any individual VM. This expected behavior was observed in the tests. Host metrics were also within accepted ranges throughout all the tests.

**Note:** At the time of writing this document, multiple monitor configurations are not supported with VDGA and K1 / K2 GRID cards. However a new firmware version is planned which will support this feature.

## 12.1.4 Appendices

### 12.1.4.1 Appendix A VMware View 5.1 solution configuration

Solution Configuration - Hardware Components:		Description
<b>Virtual Desktops</b>	Windows 7 (64 bit) , 4 x CPU and 32 GB RAM	Configuration of VMs for VDGA
<b>VMware Compute Host</b>	1 x Dell PowerEdge R720 Servers: ESXi 5.1 U1 Intel-(R)-Xeon- (R) CPU E5-2670 @ 2.6 GHZ (per GPU usage spec) 320 GB @ 1600 MHZ 10 x 146GB 15K SAS internal disk drives Broadcom BCM5720 1 GbE NIC PERC H710P RAID Controller 2 x NVidia GRID K1 2 x NVidia GRID K2 Apex 2800	For ESXi Environment 10 x 146 GB drives will be configured on RAID 10
<b>VMware Management Host</b>	1 x Dell PowerEdge R720 Server: ESXi 5.1 U1 Intel-(R)-Xeon- (R) CPU E5-2690 @ 2.7 GHZ 128 GB RAM @ 1600 MHZ 10 x 146GB 15K SAS internal disk drives Broadcom BCM5720 1 GbE NIC PERC H710P RAID Controller	For ESXi Environment 10 x 146 GB drives will be configured on RAID 10  Each VM will host the following workloads on Windows Server 2008 R2:  VMware vCenter  VMware View Connection Server  Microsoft SQL server (View Connection Server and vCenter Database) File Server

## 12.2 Branch Office Desktop – In Geo location model and Branch office co-located infrastructure model

### 12.2.1 Executive Summary

As businesses grow and become more global, IT organizations are faced with a workforce in remote and Branch Office locations that need access to data, applications, and communication methods to effectively collaborate, communicate and be productive. Traditional approaches typically involve solutions that are difficult to manage and troubleshoot remotely and provide a poor end-user experience.

VMware Branch Office Desktop is architected in a way to support the best end user experience at a remote/Branch Office while providing optimized access to centralized resources hosted in a local data center. Several technologies used in this architecture include Dell PowerEdge servers, Wyse thin clients with PCoIP support, Dell SonicWALL WAN acceleration appliances, and VMware Horizon Suite. This architecture will facilitate either local VDI compute resources in the remote Branch Office with Dell SonicWALL WAN acceleration appliances or VDI compute resources accessed remotely in the data center to facilitate optimized access to virtual desktops and applications.

Dell Cloud Client Computing has combined Dell Desktop Virtualization Solutions (DVS), which is composed of best-of-breed Dell data center components with virtualization and management software from VMware, with the world-class portfolio of Dell Wyse virtualization end points to include Dell Wyse thin, zero, and cloud clients. VMware Branch Office Desktop provides a comprehensive approach to addressing multiple requirements within the Branch Office while

ensuring employees have fast, secure access to the applications and data they need to maximize productivity.

## **12.2.2 Introduction**

Today, there are over 11 million branch offices worldwide with 80% of employees accessing their desktops remotely. When it comes to managing IT infrastructure in Branch Offices, it's not uncommon for organizations to look at centralization. The reason being that it simply isn't efficient or cost effective to maintain IT staff and resources locally at each location, especially when the work tasks being supported in each location are almost identical. Not only is the inefficiency of replicated and distributed IT management an unnecessary resource drain, it can expose the organization to greater risks of lost productivity and revenues by creating many more points of vulnerability.

Desktop virtualization is a compelling technology option because it provides a quick and straightforward mechanism for centralizing existing distributed end-user capabilities. The recognized security, high availability and cost streamlining characteristics of desktop virtualization seem ideally suited to the branch office requirement. Unfortunately, there is one major hurdle that has consistently stood in the way: the WAN. Not only is bandwidth expensive, and in some cases a constraint on application performance, but as a single point of failure it presents a risk to user productivity.

To address these different needs, VMware and Dell have partnered to deliver solutions that are tested and validated to deliver the efficiencies and cost savings of centralization. The solution leverages hosted virtual desktops with VMware Horizon View to help enhance security, ensure high availability and streamline management.

### **12.2.2.1 Traditional Approach**

The traditional approach to addressing an increasingly globally distributed workforce has been to implement incremental solutions that address individual components and aspects of the computing experience. Typically, separate management applications had to be installed with separate administrative consoles at each site to manage a variety of devices. Or even worse, IT administrators would have to touch every physical machine. The lack of centralized management and consistency was difficult to manage and required an enormous amount of effort and IT resources to maintain.

There are a multitude of challenges with traditional solutions to addressing remote and Branch Office users. Many of the challenges relate to backups, quality of service, external connectivity, security, management, and maintenance. With traditional approaches, there are no easy answers to these challenges. Because of this, user's workflow and productivity ultimately suffer.

### **12.2.2.2 VMware Branch Office Desktop Approach**

VMware Branch Office Desktop takes a drastically different approach to managing an increasingly global workforce in remote and Branch Office locations. It uniquely addresses many of the challenges of remote and Branch Office computing with comprehensive solutions. IT administrators can leverage the remote management capabilities in VMware vCenter™ Server to monitor and maintain high levels of service across multiple remote and Branch Offices – all from a central point of view.

IT departments can add a higher degree of security and control by using VMware Horizon View to host complete desktop environments for remote office users in virtual machines. This simplifies and streamlines desktop management, reducing costs while providing end users access to their personalized desktops anywhere, anytime, using any device.

## **12.2.3 High Level Solution**

VMware Branch Office Desktop provides a consistent and easy way to manage end-user experience for an increasingly global and distributed workforce. It also streamlines IT management for end-user access by combining multiple technologies to provide a quality desktop experience for the end-user. The user environment is provided in a way that allows employees to be productive, regardless of where they are based - local to the data center or located at a remote facility. For the administrator, the result is fewer management points and less complexity. End users achieve the same user experience as if they were using a physical laptop or desktop and remain productive and connected to the data and information that they need.

### 12.2.3.1 VMware Branch Office Desktop Features

VMware Branch Office Desktop enables enterprises to securely and efficiently provide a virtual desktop infrastructure to users in remote office and Branch Office locations. It allows for a consistent and unified approach for IT administrators to manage an employee's corporate workspace. This solution provides greater security, a better user experience, and a consistent approach. This helps to reduce downtime, configuration issues, and poor application experience that are commonly experienced by end users in remote locations.

Branch Office Desktop is architected with a consistent approach to tackling the difficulties and challenges that are associated with supporting users in remote and Branch Office locations. It allows flexibility for IT architects and administrators to design the solution to meet their requirements and tailor the components to meet their specific business use cases. IT administrators enjoy the benefits of managing the virtual desktop infrastructure from a single pane of glass and in a consistent manner rather than a sprawl of management interfaces and configuration panels. The result is that IT administrators can deliver a more consistent desktop computing environment in less time. The end user receives a consistent, secure, and high quality computing experience and the company benefits from higher efficiency and productivity at all levels.

Typical users of VMware Branch Office Desktop are IT administrators and architects, help desk specialists, and end users. IT administrators can use the management points to provision desktops and support remote user requests. Help desk specialists can quickly manage resources relating to permissions and access control. Remote users get the benefit of a quality computing experience as if they were located locally to the main data center or headquarters.

## 12.2.4 Solution Details

### 12.2.4.1 VMware Branch Office Desktop Architecture

A typical VMware Branch Office Desktop deployment is composed of the following components:

- VMware vSphere and vCenter
- VMware Horizon View
- VMware vCenter Operations Manager (vCOPS)
- Dell SonicWALL WAN Optimization Appliance
- Dell Wyse Thin or Zero Clients with PCoIP support

The purpose of each component is outlined below.

#### **VMware vSphere and vCenter**

The solution is built on top of vSphere, the industry-leading virtualization platform. There are many benefits to using the vSphere platform and more information on this platform can be found at [www.vmware.com/products/vsphere](http://www.vmware.com/products/vsphere).

#### **VMware Horizon View**

The central component of the solution architecture, VMware View is the industry-leading virtual desktop infrastructure (VDI) product. More information on VMware View can be found at [www.vmware.com/products/view](http://www.vmware.com/products/view).

#### **VMware vCenter Operations Manager (vCOPS)**

One of the biggest challenges faced by IT is on-demand management of the entire environment and the need to proactively identify and plan the infrastructure. VMware vCOPs for View provides the management infrastructure required for the environment. More information on VMware vCOPs can be found at [http://www.vmware.com/products/desktop\\_virtualization/vcenteroperations-manager-view/overview.html](http://www.vmware.com/products/desktop_virtualization/vcenteroperations-manager-view/overview.html).

#### **Dell SonicWALL WAN Optimization Appliance**

The Dell™ SonicWALL™ WAN Acceleration Appliance (WXA) Series reduces application latency and conserves bandwidth, significantly enhancing WAN application performance and improving the end user experience for distributed organizations with remote and Branch Offices. After initial data transfer, the WXA Series dramatically reduces all subsequent traffic by transmitting only new or changed data across the network. Deployed in conjunction with a Dell SonicWALL E-Class Network Security Appliance (NSA), NSA or TZ Next-Generation Firewall and Application Intelligence and Control Service, the WXA offers the unique combined benefit of prioritizing application traffic and minimizing it between sites, resulting in optimal network performance.

#### **Dell Wyse Thin Clients with PCoIP support**

Dell Wyse offers a wide selection of secure, reliable, and cost-effective thin and zero clients designed to easily integrate into any virtualized or web-based infrastructure, while meeting the budget and performance requirements for any application.

### **12.2.4.2 How VMware Branch Office Desktop Components Work Together**

VMware vSphere and vCenter provide the basic framework for the virtual desktop infrastructure to reside on. It provides the hypervisor, virtual networking, and management to deliver a mainframe-like resilient environment.

VMware Horizon View provides personalized virtual desktops as a managed service and is built to tightly integrate and take advantage of the benefits and features that VMware vSphere provides.

vCOPs and the vCOPs for View adapter provide a single dashboard for monitoring the entire infrastructure for the branch. By adding Quality of Service monitoring software provided with Dell Wyse thin clients, additional feedback such as application response times can be checked to ensure that the local user experience remains high.

The Dell SonicWALL WXA WAN acceleration appliance optimizes application and network traffic between sites, resulting in optimal network performance and user experience.

The Dell Wyse thin client with PCoIP support is used by the end user to connect to the virtual desktop infrastructure hosted either in the data center or at the Branch Office.

### **12.2.4.3 VMware Branch Office Desktop Topologies**

Topologies that are currently supported by Dell Cloud Client Computing Engineering are discussed in greater detail in the following sections. Currently, a “distributed” and “centralized” compute model are supported.

#### **Distributed Deployment**

While most organizations would prefer to centralize their IT infrastructures, network reliability or functionality requirements may prohibit some from doing so. By deploying VMware vSphere in the remote office, organizations can maintain a local IT infrastructure in the Branch Office and manage it from the central datacenter. By hosting virtual desktops locally in the office, the remote office can continue business operations, even if network connectivity to the datacenter is lost.

Virtualizing workloads at the remote site offers optimal desktop performance and responsiveness. Although technical expertise will remain geographically distant from the remote office IT infrastructure, central IT staff will be able to leverage VMware vCenter Server for automating server maintenance tasks and monitoring resources. These remote management capabilities minimize the need to troubleshoot remote servers and desktops in person.

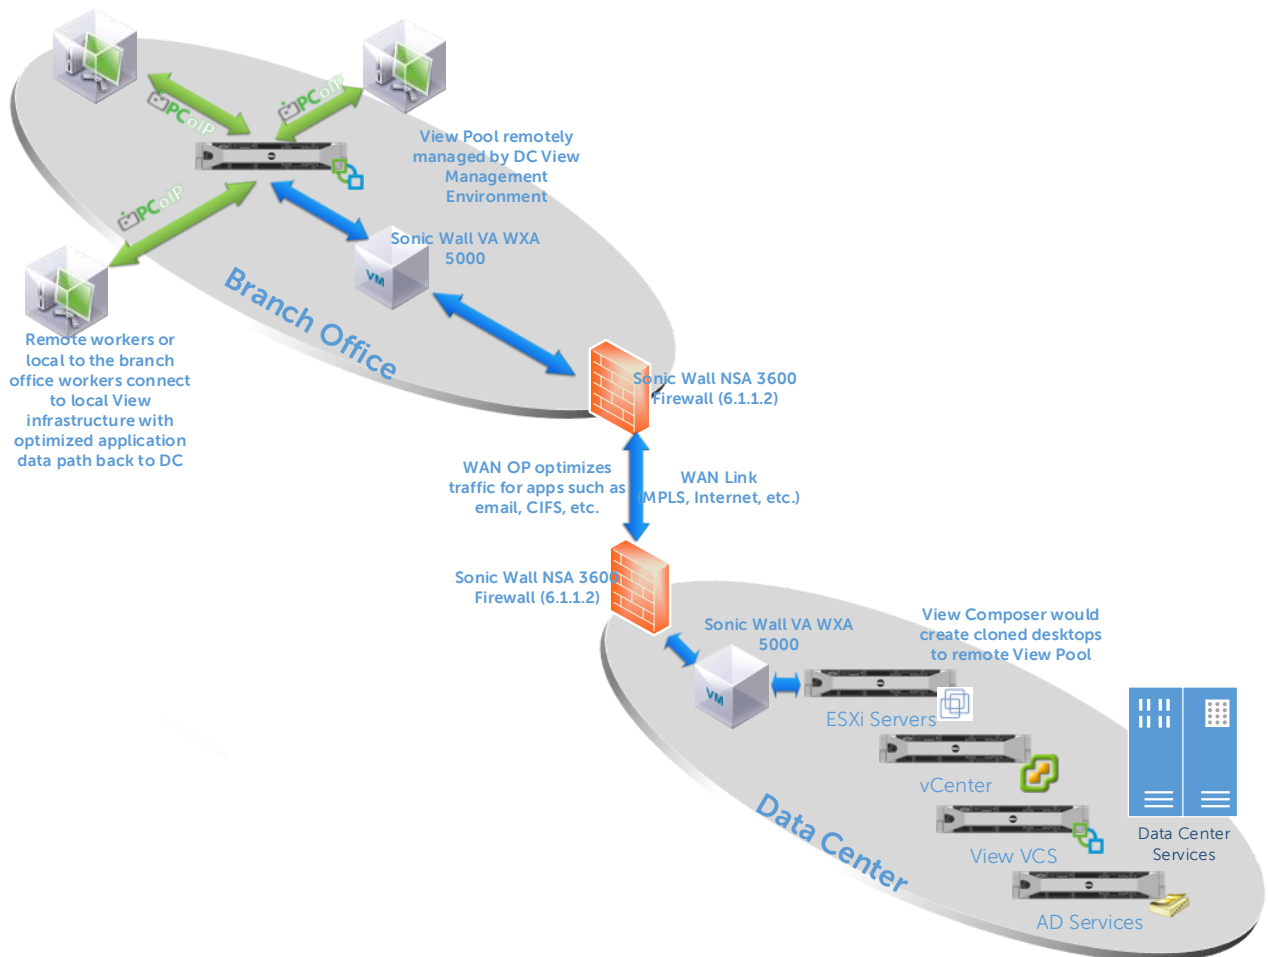
Wyse end point devices that provide access to the virtual desktops are also located on premise at



the Branch Office. The connection to the virtual desktop and its “presentation” traffic are localized. Initial desktop broker connection to VMware View Connection Server redirects Horizon View clients to local desktop connections so that PCoIP traffic is all internal at the Branch Office.

Applications used inside the Horizon View desktop context will take advantage of the Dell SonicWALL WAN acceleration appliances to optimize the connection back to the Data Center resources. In this manner the desktop PCoIP traffic that results in the presentation of the desktop is optimized since the traffic is local to the Branch Office and the applications that run in the desktop also have their data path optimized.

Figure 1 below describes the flow of traffic from the Branch Office to the data center:



**Figure 1**

Figure 2 shows how the initial connection is made to Data Center View Connection Server Broker server. This is a relatively lightweight connection that checks authentication and then directs the PCoIP connection to the appropriate desktop. In this case, the desktop is local to the end user.

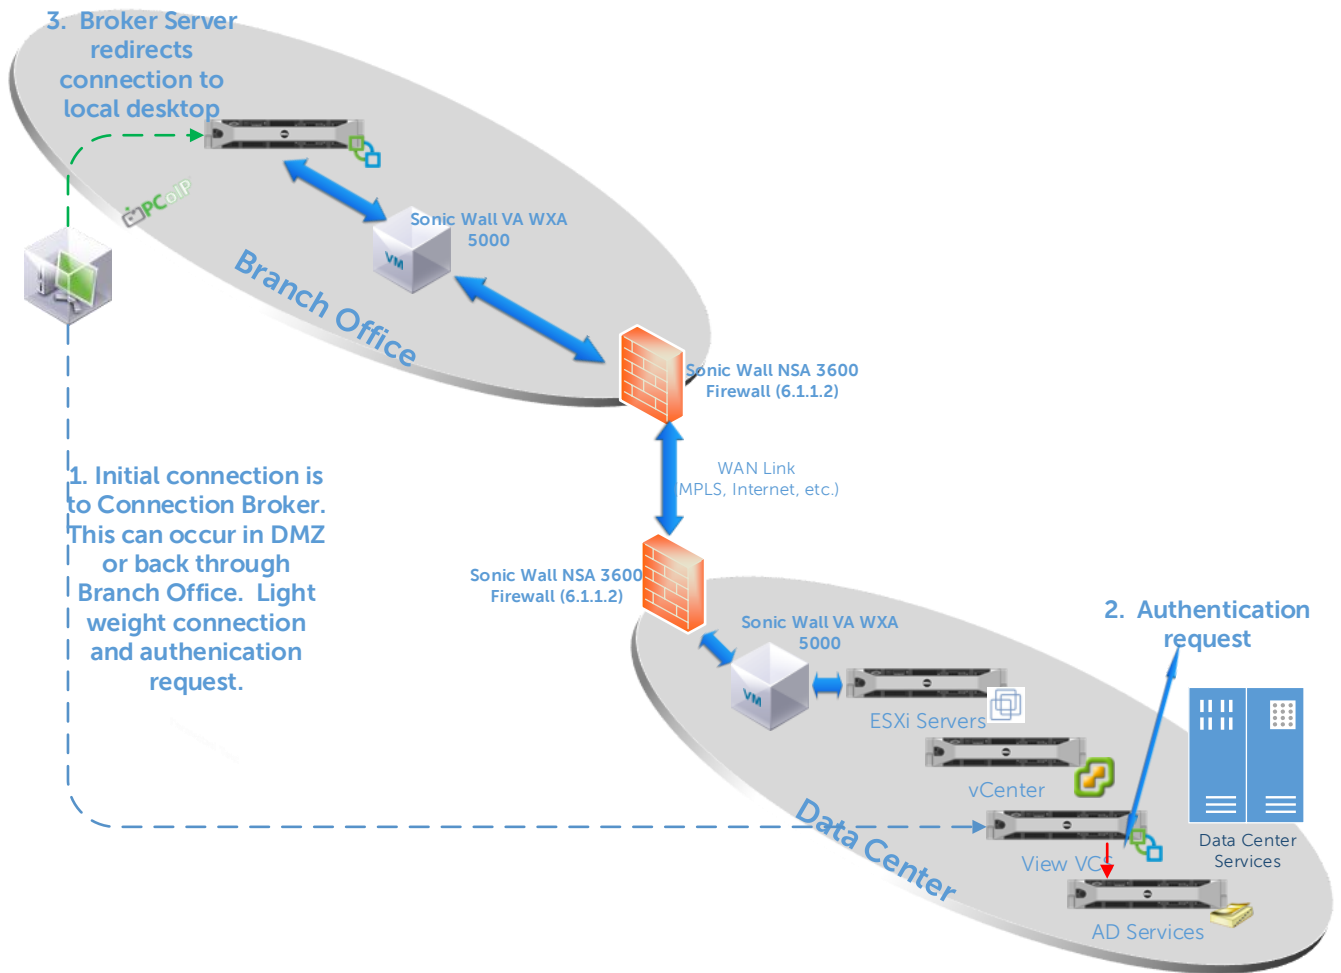


Figure 2

## Centralized Deployment

Organizations with reliable, high-bandwidth, low-latency network links to remote offices, or organizations who have implemented a wide-area data services solution for application acceleration across the WAN, can manage and standardize server and desktop environments in the corporate datacenter, where administrators can perform backups, upgrades and complete maintenance. Administrators can pull servers and desktops out of the remote office, convert them into virtual machines using VMware vCenter Converter and host them on the virtual infrastructure behind a secure firewall in the datacenter. End users in remote offices can then access server and desktop workloads over the network. Administrators can enforce strict control over access to virtual machines by delegating customizable roles and permissions to authorized administrators and end users.

A centralized approach to deployment maximizes consolidation ratios, ensures security and minimizes management complexity. Because the remote office IT infrastructure is located in the datacenter, IT staff with technical expertise can offer more responsive and better support to end-users in remote locations. Additionally, remote office services can leverage datacenter resources, including high-end servers, storage and networking, as well as existing datacenter disaster recovery and backup plans. Centralized deployment not only enhances security and compliance, but local backups can be performed in the datacenter at LAN speeds.

Since end users must access workloads over the WAN, a centralized deployment will increase network traffic between the remote site and the datacenter. Application performance will depend on application type, network bandwidth and distance between the site and the datacenter. Wide-area data services solutions or WAN acceleration products such as SonicWALL WXA can help alleviate performance issues.

In this scenario, Wyse end point devices that provide access to the virtual desktops are located on

premise at the Branch Office. The connection to the virtual desktop and its “presentation” traffic occur over the WAN. The PCoIP protocol used to connect from the Wyse end points to the virtual desktops in the data center offer many enhancements and optimizations that make it ideal for use over the WAN to include caching and bandwidth optimizations.

Applications used inside the Horizon View desktop context traffic are localized in the data center. In this manner, the only traffic that has to traverse the WAN is the desktop PCoIP traffic while the application traffic is local to the data center. This provides for an optimized data path that consumes a minimal amount of WAN bandwidth.

Figure 3 shows how the initial connection is made to Data Center View Connection Server Broker server. This is a relatively lightweight connection that checks authentication and then directs the PCoIP connection to the appropriate desktop.

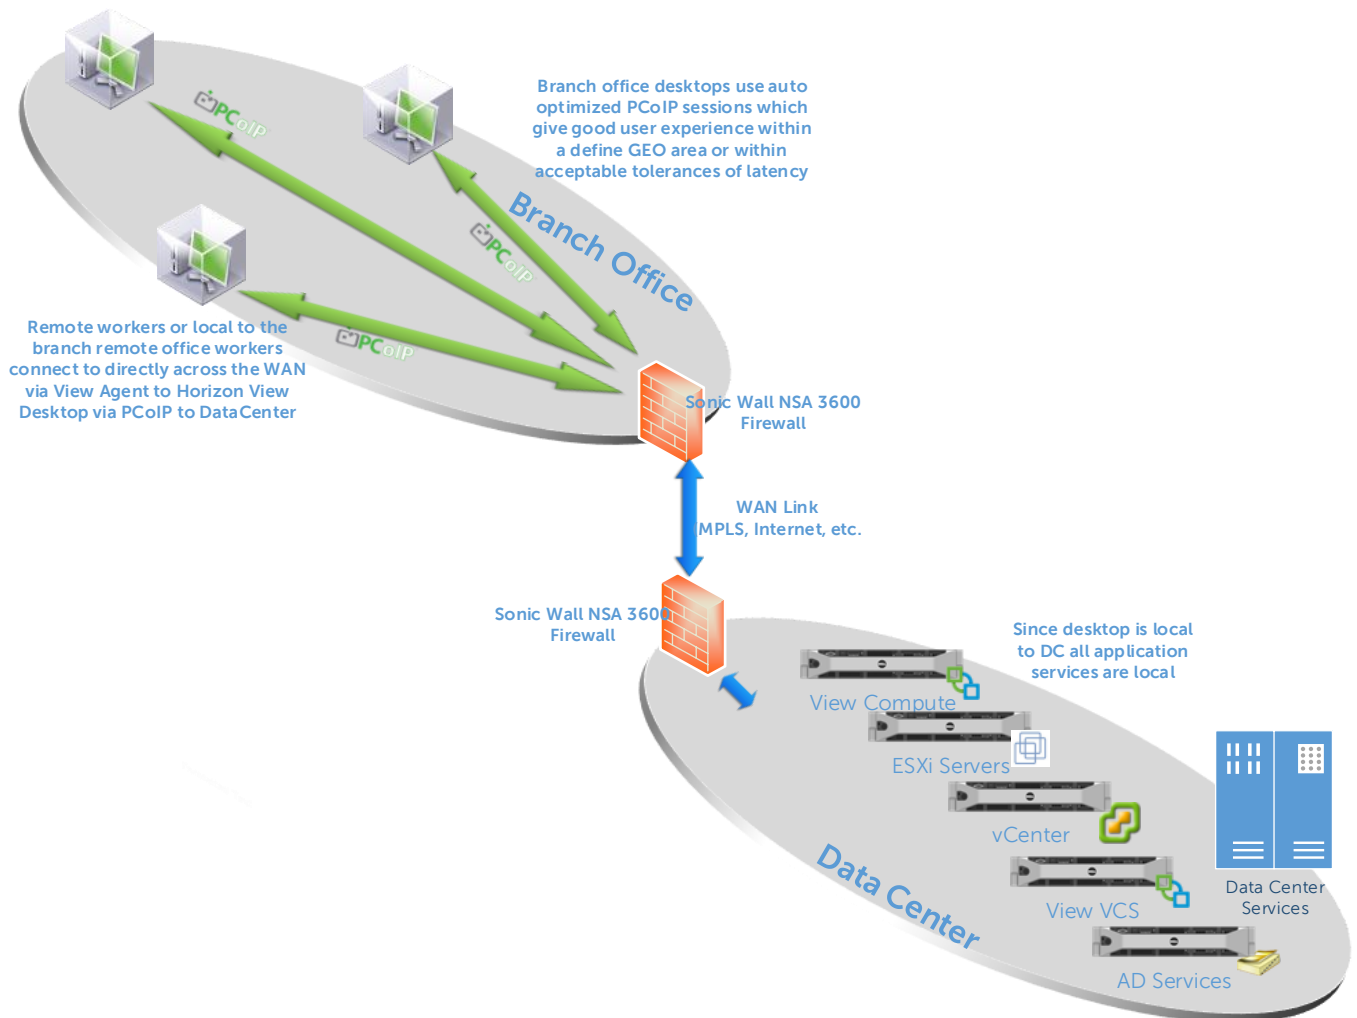


Figure 3

## 12.2.4.4 VMware Branch Office Deployment Guidance/Considerations

### 12.2.4.4.1 Distributed BOD Deployment

The following are general considerations for a Distributed model of Branch Office deployments:

- Datacenter is overseas or has an otherwise very latent connection that would not be conducive to PCoIP “presentation traffic” of a desktop in the datacenter back to a remote Horizon View client.

- Number of remote desktop users is high enough that the bandwidth of outbound WAN traffic would be large and potentially cause contention or the need for a very large WAN pipe to support multiple and concurrent PCoIP desktop connections
- Best possible end-user desktop experience in terms of presentation traffic and “feel” of the desktop is required
- Segmentation is required and remote ESX hosts server remote users to enable certain Business Desktop User cases (i.e. 3<sup>rd</sup> party offshore development work) to more readily ensure security and firewalling.
- Localized services such as print/fax are available in branch office facilitating a more optimal connection to desktop
- Distributes the architecture for Horizon View desktops where there is not a single point of failure across the enterprise
- Still centrally managed from DC Horizon View infrastructure

#### 12.2.4.4.2 Centralized BOD Deployment

The following are general considerations for a Centralized model of Branch Office deployments:

- Datacenter where application servers, data, or services that are consumed by remote desktops is in same GEO area or the latency between the DC and the remote branch office is low. Example would be branch offices in the same country or region to the DC (Datacenter)
- Smaller numbers of remote desktop users where the WAN link size is proportion to the number of users to support a good end user experience of the desktop.
- Where the performance of the application running on the desktops is highly dependent on very low latency access to the data path such that the desktop and the application data are co-located in the same DC.
- Branch office is part of the regular company offices and internal corporate WAN/MPLS access is available and where the link is sufficient to run the remote user’s desktop with a good user experience.
- Support of highly mobile users that do not have a regular “home-base” of operations (i.e. Sales, etc.)

#### 12.2.4.4.3 General BOD Deployment notes

- In scenarios without video or demanding graphics WAN latency of 75 ms or less may deliver acceptable performance for Office-type applications remotely, but is highly dependent on user workload, remote branch office uplink speeds, and the number of concurrent users using the remote branch office pipe to the DC.
- Use of zero clients with embedded PCoIP technology can potentially provide better performance.
- WAN optimization such as that provided by Dell SonicWALL WXA appliances are a very important consideration to supported Distributed models of Branch Office deployments.
- WAN Optimization will not only help the application data path back to the DC (such as CIFS shares, SharePoint, Web applications, etc.) but can facilitate better performance for Horizon View Composer operations and other similar operations or other management functionality.

### 12.2.5 Business Benefits

There are a multitude of business benefits from leveraging VMware Branch Office Desktop in VDI environments. Some of them are measurable by using metrics such as support desk calls, while others are not, such as ease of use or end-user computing experience.

Some of the largest benefits can be summarized by the following points:

- Reduce resources required in the Branch Office to manage and maintain IT equipment

- Reduce IT hardware and operating costs for servers and desktops
- Simplify IT management and accelerate provisioning from a single pane of glass
- Ensure always-on availability and recover quickly from disasters
- Improved service levels, availability, data protection, and platform security
- Centralized management of the entire IT environment through a single interface

## 12.2.6 Summary

VMware Branch Office Desktop is optimized for organizations looking to cost-effectively and reliably support desktops for Branch Office employees accessing applications and data on VMware Horizon View desktops. It is uniquely designed to support a variety of architectures that can be deployed to best address business and user requirements while being comprehensive and cost-effective. It provides an enhanced user experience, ease of management for IT administrators, and the security needed to keep crucial business data safe and secure.

## 12.3 Business Process Desktop (BOD) Use case

### 12.3.1 Executive Summary

The VMware Horizon View Business Process Desktop solution enables customers looking to outsource or offshore business processes to effectively scale their business on demand, streamline and centralize desktop management and provide end users with a standardized and secure desktop experience across the LAN and WAN. Data is replicated and centrally backed up, to further ensure high availability and drive higher service-level agreements (SLAs) across remote locations.

Dell Cloud Client Computing has combined Dell Desktop Virtualization Solutions (DVS), which is composed of best-of-breed Dell data center components with virtualization and management software from VMware, with the world-class portfolio of Dell Wyse virtualization end points to include Dell Wyse thin, zero, and cloud clients. VMware Business Process Desktop provides a comprehensive approach to addressing multiple requirements within remote or third party locations while ensuring security and compliance, simplifying and centralizing management, and improving service-level agreements.

### 12.3.2 Introduction

The use of outsourcing is expanding rapidly and today's business process outsourcing (BPO) buyers and providers are increasingly looking for ways to increase revenues, decrease costs, and bolster worker productivity. For IT organizations, attaining these goals can be challenging. Traditional PC environments are often locally managed which is costly and resource intensive. Data stored locally on endpoints poses a greater security risk for the business and can, in the event of a security breach, jeopardize the reputation of the BPO. Remote access across the WAN is costly and if not sized correctly can often impede productivity and worker performance.

In order to achieve business objectives and remain competitive, BPO IT organizations need to rethink how services are delivered to end users, how data is secured, and how technology platforms are rolled out and integrated.

By virtualizing desktops and hosting them on VMware vSphere, a key component of VMware View, and using this tested architectural design, organizations can now centralize desktop management and provide unparalleled desktop and application access across the LAN and WAN. With the Business Process Desktop, processes are automated and efficient, data is secure, and the total cost of ownership is reduced. And because this solution ties unified communications from the leading vendors to PCs, end users are free to access their data and applications from a VoIP softphone across devices and locations, improving worker access and driving higher levels of productivity.

### 12.3.2.1 Traditional Approach

The traditional approach to addressing an increasingly globally distributed workforce has been to implement incremental solutions that address individual components and aspects of the computing experience. Typically, separate management applications had to be installed with separate administrative consoles at each site to manage a variety of devices. Or even worse, IT administrators would have to touch every physical machine. The lack of centralized management and consistency was difficult to manage and required an enormous amount of effort and IT resources to maintain.

There are a multitude of challenges with traditional solutions to addressing remote and offshore users. Many of the challenges relate to quality of service, external connectivity, management, maintenance, and especially security. With traditional approaches, there are no easy answers to these challenges. Because of this, remote user's workflow and productivity ultimately suffer.

### 12.3.2.2 VMware Business Process Desktop Approach

VMware Branch Process Desktop takes a drastically different approach to managing an increasingly global workforce in remote and third party locations. It uniquely addresses many of the challenges of remote and offshore computing with comprehensive solutions. IT administrators can leverage the remote management capabilities in VMware vCenter Server to monitor and maintain high levels of service across multiple remote locations – all from a central point of view.

IT departments can add a higher degree of security and control by using VMware Horizon View and vShield to host complete desktop environments for remote users in virtual machines. This simplifies and streamlines desktop management, reducing costs while providing end users access to their personalized desktops anywhere, anytime while protecting confidential and private corporate data from compromise.

### 12.3.3 High Level Solution

The VMware Business Process Desktop solution architecture provides a streamlined, cost-effective way for IT organizations to support business process outsourcing by improving user access, centralizing desktop management, enhancing data security, and maximizing employee uptime.

By leveraging state-aware desktops with persona management, Business Process Desktop ensures end users can carry their persona with them across sessions and devices for a more personalized desktop experience. This also allows IT administrators to leverage the same desktop infrastructure for workers sharing desktops and endpoints across shifts.

VMware View Business Process Desktop with PCoIP additionally delivers end users a seamless experience across the LAN and WAN—and supports end users who need unified communications as part of their daily workspace. Integration with unified communications solutions from leading vendors further ensures that end users can easily access their VoIP softphone across devices to drive greater levels of productivity.

End user access via Radius two-factor authentication is secured via the VMware View security server or SSL. vShield products together with VMware Horizon View and leading security vendor solutions, allow IT to offload antivirus operations and provide high levels of isolation between resource pools and networks. This allows IT administrators to apply policies across VMs and pools of users.

With Business Process Desktop, data can be easily replicated across data centers to ensure greater business continuity and to maximize end-user uptime.

Figure 1 below shows a solution that utilizes VMware Business Process Desktop:

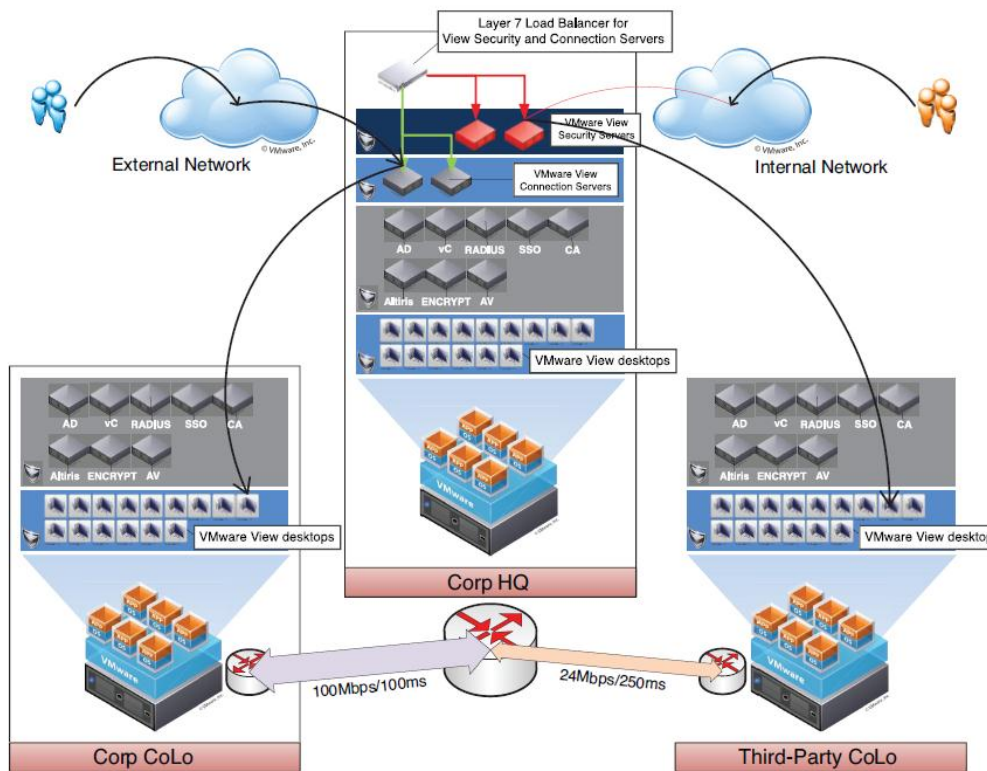


Figure 4

## 12.3.4 Solution Details

### 12.3.4.1 VMware Business Process Desktop Architecture

A typical VMware Business Process Desktop deployment is composed of the following components:

- VMware Horizon View
- Unified Communications
- VMware vShield

The purpose of each component is outlined below.

## VMware Horizon View

The cornerstone of the View Business Process Desktop solution, VMware Horizon View modernizes desktops and applications by moving them to the cloud and delivering them as a managed service. With View, IT has the ability to grant or deny access to desktops, data and applications according to endpoint device configuration, network location and user identity.

## Unified Communications

Unified communications solutions from market leaders are fully integrated with the Business Process Desktop to provide end users with quick, easy access to a VoIP softphone across devices and locations. This provides end users with greater mobility and enhanced access while reducing IT infrastructure costs.

## VMware vShield

The vShield suite of products, including VMware vShield App and VMware vShield Edge, enables IT to effectively firewall virtual machines and partition networks and resource pools. With vShield App, IT can apply rules to virtual machines based on IP addresses as well as business or application requirements. vShield Edge permits segmentation of resource pools and enables IT to provide a common set of services to virtual machines that reside within a defined perimeter. In addition, VMware vShield Endpoint provides antimalware and deep packet inspection. This enables IT to enhance endpoint performance across the desktop environment by offloading antivirus scanning to the hypervisor, eliminating the need to install complex agents inside individual virtual machines.

### 12.3.4.2 How VMware Business Process Desktop Components Work Together

VMware vSphere and vCenter provide the basic framework for the virtual desktop infrastructure to reside on. It provides the hypervisor, virtual networking, and management to deliver a mainframe-like resilient environment.

VMware Horizon View provides personalized virtual desktops as a managed service and is built to tightly integrate and take advantage of the benefits and features that VMware vSphere provides.

VMware vShield integrates seamlessly into the VMware vSphere and Horizon View environments to provide additional layers of security to keep business data safe and secure.

Figure 2 below shows an example BPD configuration and the high level components:

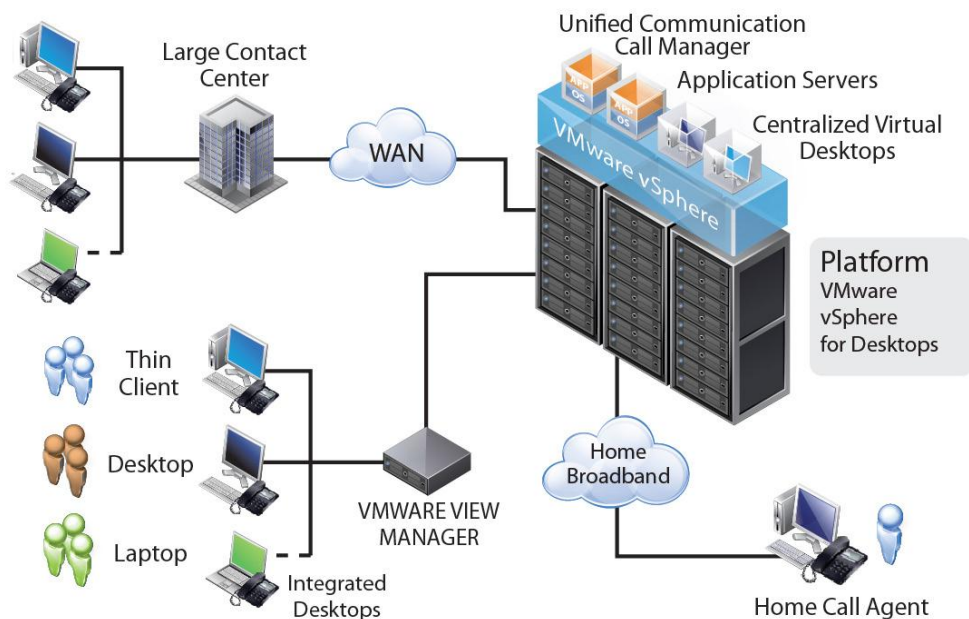


Figure 5



### 12.3.5 Business Benefits

There are a multitude of business benefits from leveraging VMware Branch Process Desktop in virtual desktop environments for outsourced or offshore business processes. Some of them are measurable by using metrics such as support desk calls, while others are not, such as ease of use or end-user computing experience.

Some of the largest benefits can be summarized by the following points:

- Reduce operating costs by centralizing and streamlining desktop management and support
- Provide uninterrupted uptime across remote locations to drive higher SLAs
- Rapidly scale to accommodate changing needs or new client contracts
- Enhance data security and compliance by centralizing data in the datacenter
- Allows companies to focus on their competencies instead of managing IT infrastructure

### 12.3.6 Summary

The Business Process Desktop from VMware is a managed solution optimized for IT organizations looking to more simply and cost-effectively support business process outsourcing. It integrates technology and solutions from VMware, Dell, and other partners. The solution leverages VMware vSphere, Horizon View, vShield, and unified communications components to drive collaboration, manage IT infrastructure, and better protect data across locations.

## 12.4 VMware Horizon Suite Bundle V1.0

### 12.4.1 VMware Horizon Workspace

For information on VMware Horizon Workspace please reference section 13.11 below for information on this feature in the VMware Horizon Suite Bundle.

### 12.4.2 VMware ThinApp

#### 12.4.2.1 Executive Summary

A major issue experienced by IT departments at present is "application sprawl". Various different applications are being installed on all many different kinds of devices e.g. laptops / desktops etc. As a result of this, different users have different versions of software than other users, this leads to version incompatibility and loss of control for IT departments over the software in use on the corporate network.

VMware ThinApp is an application suite designed by VMware that helps separate applications from the operating system. It was designed to eliminate application conflicts and streamline management. It allows for simplified application virtualization and helps reduce the cost and complexity of delivering applications to your networked devices.

Dell Cloud Client Computing has combined Dell Desktop Virtualization Solutions (DVS), which is composed of best of breed data center components with virtualization and management software from VMware, with the world class portfolio of Dell Wyse virtualization end points to include Dell Wyse thin, zero and cloud clients. VMware ThinApp helps add application compatibility to virtual desktop environments and helps reduce the management burden of desktop applications and images.

## 12.4.2.2 Introduction

With the exponential growth experienced in virtualization including new technologies like BYOD, it has become critically important to regain control over applications. IT departments need to know what piece of software is installed on what device and which version it is. The need for this information is a major driver when IT departments start to look at application virtualization. ThinApp allows the applications to be packaged and shared / streamed to multiple users thus allowing much greater control over what applications are installed on what device. This helps reduce IT support and helpdesk costs, helps eliminate application conflicts and increases workforce mobility.

### 12.4.2.2.1 Traditional Approach

In the past, applications had to be installed on each and every individual laptop / desktop. This presented IT departments with many significant problems. If a user had an issue with the application then a member of the IT department would have had to connect to the user device and troubleshoot the application, this was time consuming and frustrating. Another issue this approach presented was that of application compatibility. It was likely that different users had different versions of the software installed on their machine; an example of this would be some users running Microsoft Office 2007 while others were running Microsoft Office 2010.

If a user was receiving a new desktop / laptop then all of the various different applications that he / she uses would have to be re-installed prior to the user receiving the new machine so a lot of time was spent by IT departments on supporting and upgrading various applications on desktops / laptops.

### 12.4.2.2.2 VMware ThinApp Approach

VMware ThinApp helps address many of the problems and challenges surrounding application installation and support. Its agentless application virtualization decouples applications and data from the OS. It allows the delivery of applications' to multiple users from a single network share, this method of application delivery results in no local disk footprint as the application is streamed into memory. ThinApp fits seamlessly into any environment, there is no specific hardware or software needed and plugs into any existing management framework. It also ensures security without compromising user flexibility.

VMware ThinApp enables IT personnel to package the application as an .exe or an .msi file. This can then be deployed to multiple windows operating systems without imposing additional cost and complexity to the server or client. The virtual applications are isolated from each other and executed independently without making changes to the underlying operating system. If there is an issue with the application then it's much easier to troubleshoot because the entire application is stored as a single file.

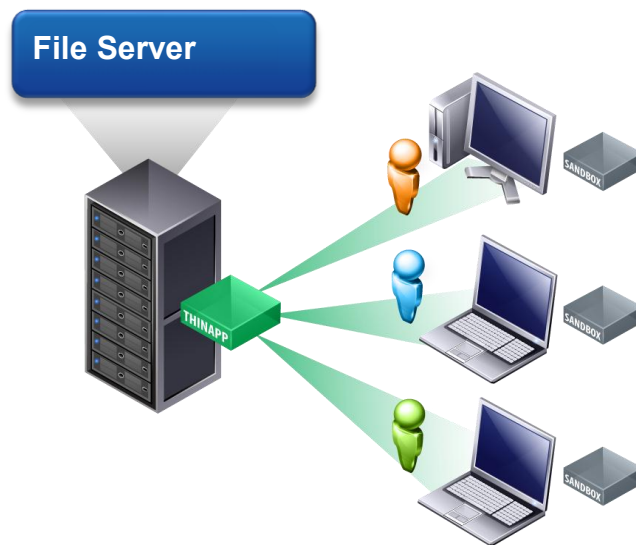
ThinApp virtualizes applications by encapsulating application files and registry settings into a single ThinApp package. The virtualized applications do not make any changes to the underlying OS and behave the same across different desktop configurations, thereby providing a stable, consistent experience.

It's possible to manage and assign ThinApp virtualized applications in the same interface where you deploy and manage virtual desktops, the view administrator console. However it can be difficult to remove applications via view administrator so we recommend using logon scripts for applications registration instead.

With ThinApp 4.7, administrators now have the capability of deploying ThinApp virtualized applications in Horizon Application Manager. VMware Horizon provides a new management platform for entitlement, deploying and monitoring ThinApp packages.

### 12.4.2.3 Solution Overview

In a View environment, the most common kind of application deployment is via a network share on a file server.



Sharing the applications via a network share as in the diagram above offers tremendous advantages to your organization. It allows your users to have applications streamed to their desktops on an on demand basis. The application is installed and packaged on the packaging virtual machine; the resulting project file is copied to the network share and made available to the users via a logon script. The logon script registers the application to the user's desktop using the thinreg.exe command but it's only when the user opens the application that the application gets streamed down from the network share.

When the application is launched, only the necessary blocks of data are streamed into memory for execution. It does not require the local caching of files; instead it only streams into memory what is needed at that time to perform the application function. The amount streamed depends on how much of the application functions are used and which DLL's, registry and files are needed for those functions. As virtualized applications are executed by users, the dynamic application and user settings are redirected to a storage location defined by the administrator which is called the sandbox.

This allows you to have all your applications stored in a central repository and have total control of the applications in use on the user's desktops. Bandwidth usage is also reduced due to only the required files and DLLs being streamed to the desktops.

Here is a screenshot of the package creation workflow



## Updating Applications


When an IT department wants to update an application, they first update the application on the packaging virtual machine. Once the update has taken place and it has been thoroughly tested then it's very easy to update the existing ThinApp package that is currently in use by users.

There are two different method of updating applications. They include the following

- Using ThinApp's built in AppSync feature
- Side by Side file update

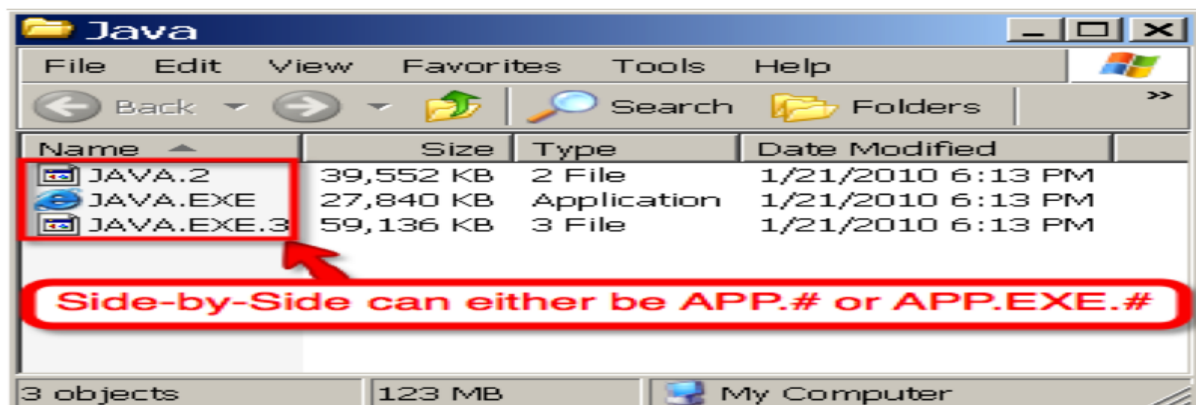
## AppSync

AppSync is an endpoint initiated solution where the ThinApp packaged app on the remote system initiates the update. The new updated package is placed on the AppSync web server and the next time the remote endpoint launches the application then it will detect an update on the web server and download the differences. Once the differences are downloaded, ThinApp recompiles itself and the next time the user of the remote end point launches the ThinApp package they are launching the updated package.



### Side by Side

One of the easiest ways of updating a ThinApp package is by placing the updated package next to the original release on the network share and changing the filename to include a numerical value. Any new execution of the package by a user will result in the package with the highest numerical value in its filename being opened. The advantage of this method is that it can be done during production and is very easily rolled back in the unlikely event an issue occurs.



### 12.4.2.4 Solution Component Details

ThinApp virtualized components include

- ThinApp Compressed Container
- VOS (virtual operating system)
- Sandbox
- Virtual file system
- Virtual Registry

The purpose of each is outlined below

### **ThinApp Compressed Container**

This container is made up of an encapsulated application and its associated dependencies in a single EXE or MSI file. It is comprised of a very small client (~400k) that is built into each EXE package (this allows for clientless execution of ThinApp application packages). The container includes the VOS, VFS and the VREG

### **Virtual Operating System**

The virtual operating system uses the virtual file system and virtual registry to transparently merge the virtual system environment with the physical system environment. It controls the ability of the virtualized application to access the resources of the underlying physical OS; it also allows the ThinApp packages to be OS independent

### **Sandbox**

The sandbox is a private, per-user and per-app directory where all application runtime changes are stored. It contains things like user configurable settings such as web browser home page or favorites, Microsoft Word default template customizations etc. The sandbox allows user settings to persist. Its location is configurable; the default location is %AppData%\Thinstall but you can change this to a network share, a removable USB volume or a VMware view user data disk. Deleting the sandbox returns the application to default behavior; this can be useful in a troubleshooting situation.

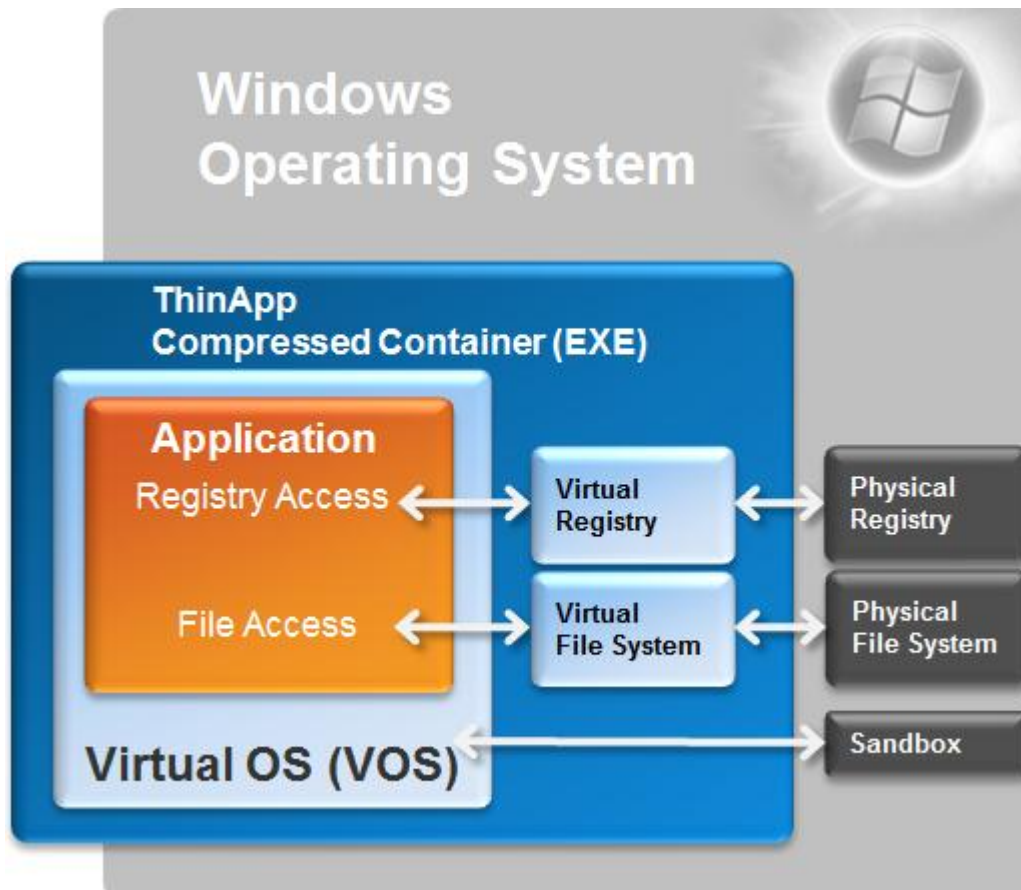
### **Virtual File System**

The virtual file system is embedded in the application package and it is read only. The virtual operating system intercepts requests to open files and redirects them to the virtual file system.

### **Virtual Registry**

Operating systems and applications store a large number of settings in the system registry. The virtual operating system intercepts requests to look up and store values in the registry and redirects them to the virtual registry. By storing settings in the virtual registry, ThinApp ensures that settings required by the application are accessible without actually changing the system registry.

The following diagram shows the interaction between the various components.



VMware ThinApp supports various operating system, applications and systems. 32 bit platforms include Windows 2000, Vista, XP, XPE, 2003, 2008 and Windows 7.

64 bit platforms include XP-64 bit, 2003 64-bit , Vista 64-bit, 2008 64-bit, 2008 R2 64-bit and Windows 7 64-bit.

It also supports 16-bit applications running on 32-bit windows operating systems and 32-bit applications running on 32-bit and 64-bit Windows operating systems.

#### 12.4.2.4.1 VMware ThinApp Licensing

ThinApp is included in Horizon View, Horizon Mirage and Horizon Workspace.

Product	Description	SKU - 10/100 Packs	\$USD
<b>Horizon View Bundle</b>	Bundles include everything a customer needs for a full end-to-end VDI deployment. Includes: - Horizon View Manager - Horizon View Composer - Persona Management - ThinApp - vSphere Desktop - vCenter Desktop	VU5-PR-STR-C VU5-PR-100-C	\$250 /concurrent connection
<b>Horizon View Add-On</b>	For customers who excess vSphere and vCenter from standalone purchases and want Horizon View. Includes: - Horizon View Manager - Horizon View Composer - Persona Management - ThinApp	VU5-PR-A10-C VU5-PR-A100-C	\$190 /concurrent connection
<b>Horizon View Add-on to Bundle Upgrade</b>	For customers who have View Add-ons and want the full bundle. Includes: - vSphere Desktop - vCenter Desktop	VU5-PR-STR10-UG-C VU5-PR-UG-100-C	\$70 /concurrent connection

For further licensing information, please contact your local Dell sales representative

### 12.4.2.5 Business Benefits

- Simplified migration to Windows 7
- Elimination of application conflicts and different versions
- Increased mobility for end users
- Reduced IT support costs
- Integrated application assignment

### 12.4.2.6 Summary

VMware ThinApp offers an efficient method of decoupling applications from the operating system and streaming them to users from a central location. Applications are isolated, security is built in and the architecture is agent less. ThinApp is a feature rich piece of software and offers many advantages over the traditional software installation approach.

### 12.4.3 VMware vShield Endpoint

#### 12.4.3.1 Executive Summary

IT organizations today are faced with an ever increasing number of security threats when deploying a desktop environment to support their business. Virtual desktop infrastructure (VDI) environments are consolidated by nature, which brings both challenges and unique opportunities to increase system security. With an ever increasing landscape of sophisticated and complex security threats by viruses and malware, security, availability, and downtime become larger and larger problems in today's computing environments.

VMware vShield™ Endpoint provides an elegant, integrated, and comprehensive solution that offloads antivirus and anti-malware processing from the operating system in the virtual machine to a secure virtual appliance where antivirus scanning is enforced. This allows for higher efficiency,



greater server density, better performance, and higher quality of service for the end-user.

Dell Cloud Client Computing has combined Dell Desktop Virtualization Solutions (DVS), which is composed of best-of-breed Dell data center components with virtualization and management software from VMware, with the world-class portfolio of Dell Wyse virtualization end points to include Dell Wyse thin, zero, and cloud clients. VMware vShield Endpoint provides a comprehensive framework that allows administrators to protect the entire virtual desktop infrastructure, integrate with their existing antivirus and anti-malware solutions, and provide users with a secure desktop computing experience that they expect.

### **12.4.3.2 Introduction**

As IT computing environments become larger and more complex, security continues to be a large consideration for keeping users, data, and intellectual property safe and secure. Enhanced security solutions are needed that are robust, virtualization aware, and easy to manage. If traditional methods of protecting workstations, or end points, are used, it could potentially affect all of the users that are utilizing those desktops and resources in an adverse manner. VMware has worked extensively with partners to provide robust antivirus and anti-malware solutions, which are discussed in the following sections.

#### **12.4.3.2.1 Traditional Approach**

In a typical enterprise IT organization, antivirus and anti-malware solutions consist of two high level components; a management component and a client agent component. The management component is typically installed on a server or set of servers on which an IT administrator can centrally manage the environment. The client agent is installed on each workstation or server that needs to be protected. Updates are typically downloaded onto the management system and "pushed" out to the client systems via an update mechanism that is built into the management software and client agent. Some examples of updates are software updates for the antivirus or anti-malware software for security or bug fixes and definition updates for newly identified viruses and malware.

In the past, this method of end point protection has worked well. However, it does not address the changes that have been happening in the desktop computing experience. Users are increasingly mobile and expect to be able to access a consistent environment with access to email, files, and data from a multitude of devices and physical locations. Traditional antivirus and anti-malware software is also very resource intensive to each user's workstation. On-demand scanning and scheduled scans of the entire file system on each workstation incurred a lot of overhead and affected overall experience and productivity for the end user. Additionally, if the user's workstation was not accessible due to being turned off or taken remotely, as is common with laptops, definition and software updates were not received leaving the system vulnerable to attacks.

When using a traditional antivirus and anti-malware approach in a virtualization environment, many of the issues that plague system administrators are compounded. Due to the consolidated nature of VDI, antivirus or anti-malware scans of the hard drive and definition updates can cause a storm of disk activity that can overload resources and cause a poor end user experience for all users connected to those shared resources. To compensate, IT administrators would have to decrease the number of users per server; also referred to as server density. This would waste valuable CPU and memory resources and increase the total cost per user.

#### **12.4.3.2.2 VMware vShield Endpoint Approach**

Like the traditional approach described above, an enterprise IT organization that has deployed VMware vShield Endpoint antivirus and anti-malware solutions consists of two high level components; a management component and an agent component. The management component is still typically installed on a server or set of servers on which an IT administrator can centrally manage the environment. In fact, the vShield Endpoint security can integrate directly into the existing antivirus offering. The difference, though, is that there is no special client agent to be installed on each workstation or server that needs to be protected. It is already included with VMware Tools, which is typically already installed on each desktop VM. Instead of managing an agent on each desktop, there is a single agent that resides within a Security Virtual Machine on the ESX host that hosts the desktop virtual machines. Updates are applied to the always-on security Virtual Machine providing up-to-date security for all end point systems.

VMware has worked with antivirus partners and developed a comprehensive solution that not only protects the end user, but also addresses the special requirements introduced by an increasingly mobile workforce. This solution offloads the performance impact of on-demand virus scanning and scheduled virus scans of the entire file system that each desktop VM experiences to a specialized appliance on the VMware ESX host. This provides a much better overall user experience and allows the end user to continue to be productive even while scanning operations are taking place in the background. Additionally, since the specialized appliance on the ESX host is always available, it is always receiving definition and software updates so that the system is never vulnerable to attacks.

When using VMware vShield Endpoint for antivirus and anti-malware in a virtualization environment, many of the issues with using a traditional solution are eliminated or vastly reduced. Virtually all antivirus software activities are offloaded to a specialized appliance that performs the functions that the agent on each system would normally do. This decreases CPU and memory utilization, increases server density, and ultimately drives down the cost per user. The end result is happy users, lower IT operating costs, and greater security and protection against today's viruses and malware.

### **12.4.3.3 High Level Solution**

VMware vShield Endpoint is the solution to the problems inherent in antivirus scanning in a large-scale virtual desktop implementation. In a VMware View environment, vShield Endpoint consolidates and offloads antivirus operations into one centralized virtual appliance.

VMware has partnered with antivirus software vendors to provide this bundled solution to antivirus problems in the VDI environment. VMware partners supply a dedicated, secure virtual appliance. This virtual appliance integrates with vShield Endpoint APIs to protect VMware virtual desktops against viruses and other malware. Instead of installing antivirus agents on each virtual desktop, you connect one virtual appliance to each virtual machine host.

#### **12.4.3.3.1 VMware vShield Endpoint Features**

The management console provided by the VMware partner is used to configure and control the partner's software hosted in the secure virtual appliance. VMware partners can provide a user interface that makes the management experience (including policy management) exactly like managing software hosted on a dedicated physical security appliance.

Virtual infrastructure administrators have a vastly reduced level of effort because virtual machines have no antivirus agents to manage. Instead, the partner's management console is used to manage the secure virtual appliance. This approach also avoids the need to administer frequent updates per virtual machine. For deployment, VMware Tools includes the thin agent, and the ESX module enables hypervisor introspection.

Virtual infrastructure administrators can easily monitor deployments to determine, for example, whether an antivirus solution is operating properly.

Figure 1 below shows the basic protection model of VMware vShield Endpoint:




Figure 6

### 12.4.3.4 Solution Details

#### 12.4.3.4.1 VMware vShield Endpoint Architecture

Instead of installing the antivirus and antimalware software on each virtual machine, you install it only on the single security virtual machine assigned to the vSphere host. Each virtual machine to be protected requires only a small-footprint agent.

VMware vShield Endpoint plugs into vSphere and protects virtual machines against viruses and malware. Administrators can centrally manage VMware vShield Endpoint through the included vShield Manager console, which integrates with VMware vCenter™ Server for unified security management in the virtual datacenter.

Isolating the antivirus scanning engine on the virtual appliance makes it easier to protect the scanning engine than if it were placed on every virtual machine. In addition, detailed logging of activity from the antivirus or antimalware service satisfies auditor compliance requirements.

When viruses or malware are detected, the partner antivirus solution manages the remedial action to the affected virtual machines, based on the administrator's specifications.

VMware vSphere Endpoint is composed of the following components:

- vShield Manager vApp
- Security Virtual Machine (SVM)
- vShield ESX Module
- Virtual Machine Thin Agent

The purpose of each component is outlined below:

##### **vShield Manager vApp**

Used to centrally manage VMware vShield Endpoint through the included vShield Manager console, which integrates with VMware vCenter™ Server for unified security management in the virtual datacenter. The vShield Manager runs as a virtual appliance on an ESX host.

##### **Security Virtual Machine (SVM)**

Antivirus functions are offloaded to a separate Security Virtual Machine (SVM) on each VMware vSphere ESX host. The enterprise antivirus engine and the signature file are located on the virtual appliance, instead of on each virtual machine. This frees up virtual desktop system resources.

##### **vShield ESX Module**

A VMware vSphere ESX hypervisor module that enables communication between the SVM and the Thin Agent on the virtual machines, commonly referred to as "hypervisor introspection".

### Thin Agent

The VMware vShield Agent is now included with VMware Tools, which must be installed on each guest virtual machine to be protected. Virtual machines with VMware Tools installed are automatically protected whenever they are started up on an ESX host that has the security solution installed. That is, protected virtual machines retain the security protection through shut downs and restarts, and even after a vMotion move to another ESX host with the security solution installed. The following versions of Windows client operating systems are currently supported with vShield Endpoint 5.1:

- Windows XP SP3 and above (32 bit)
- Windows Vista (32 bit)
- Windows 7 (32/64 bit)

#### 12.4.3.4.2 How VMware vShield Endpoint Components Work Together

The vShield Manager vApp is a virtual appliance that runs on an ESX host and is managed through VMware Virtual Center. The vShield Manager is used to manage communicate with the Security Virtual Machine and the antivirus or anti-malware solution software.

The Security Virtual Machine (SVM) is a virtual machine that must be installed on each ESX host that is to be protected. All of the VDI VMs that are resident on that ESX host are protected when this VM is configured and installed. All antivirus and anti-malware functions are offloaded to the SVM. The SVM communicates with both the vShield ESX Module and the vShield Manager.

The vShield ESX Module is a hypervisor module that enables communication between the SVM and the Thin Agent on the virtual machine. It is included as part of the install in VMware vSphere 5.1. It provides a secure mechanism for offloading the antivirus and anti-malware functions to the SVM.

The Thin Agent is an agent that is included in VMware Tools that is installed on each VDI VM that needs to be protected. Typically, VMware Tools is already installed on every VDI VM as it provides a high level of integration between the guest operating system and VMware vSphere. The Thin Agent utilizes the vShield ESX Module to offload all antivirus and anti-malware functions to the SVM.

Figure 2 outlines the basic relationships between the components of VMware vShield Endpoint:

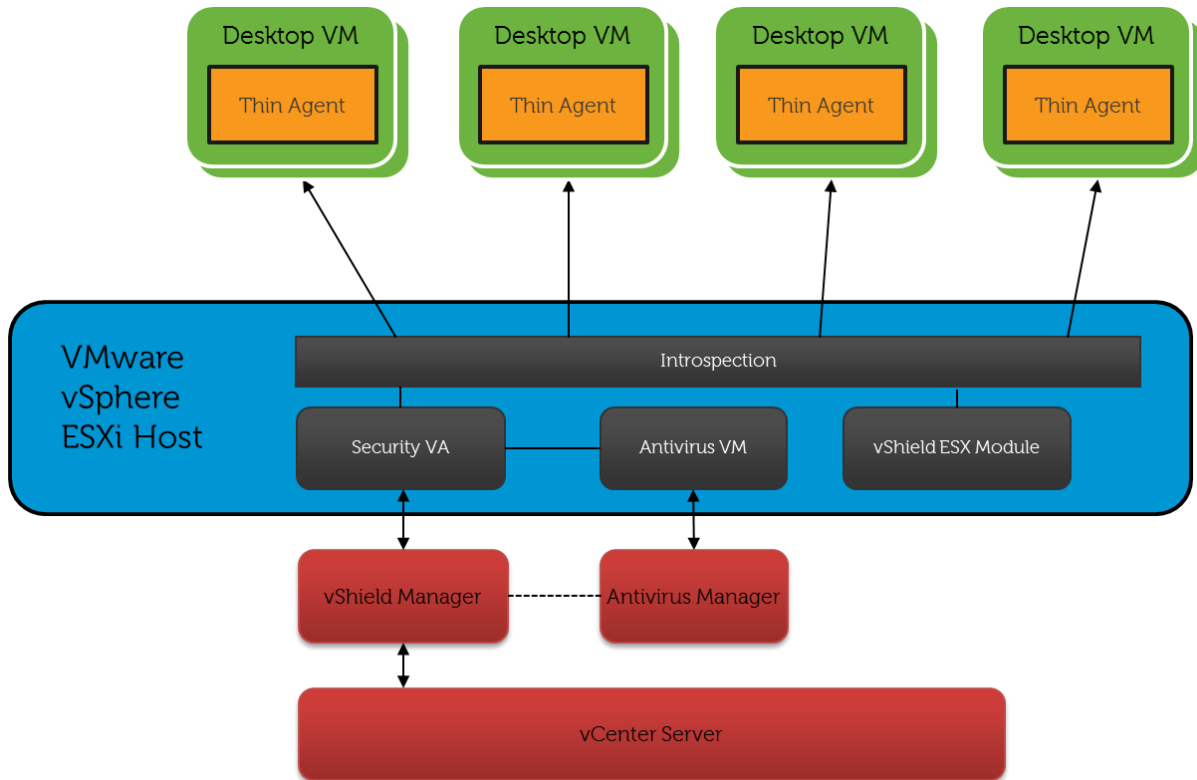


Figure 7

#### 12.4.3.4.3 VMware vShield Endpoint Licensing

Licensing of VMware vShield Endpoint 5.1 is as follows:

- With the launch of VMware vSphere 5.1, customers with valid Support and Subscription (SnS) contracts for vSphere Essentials Plus or higher editions are entitled to vShield Endpoint functionality at no extra cost.
- VMware vSphere 5.1.x license enables the vShield Endpoint 5.1.x functionality and no additional Endpoint license is required.
- Additional licenses may be required for the antivirus and anti-malware solutions that are used with the partner secure virtual appliance. Please contact your antivirus or anti-malware partner for additional licensing details.

Note: The vShield Endpoint asset may show up as unlicensed or with an evaluation license in the vCenter Server console, but functionality is fully enabled.

#### 12.4.3.4.4 Integrated Antivirus Partner Solutions

Customers are free to use the antivirus and anti-malware solutions (that vShield endpoint supports) that are best suited for their environment. The following are some support products that support the VMware vShield Endpoint environment.

##### Trend Micro Deep Security

Trend Micro Deep Security software provides compliance by providing a comprehensive server security platform designed to protect your data center and cloud workloads from data breaches and business disruptions, and achieve cost-effective compliance across these environments. Tightly integrated modules including anti-malware, web reputation, firewall, intrusion prevention, integrity monitoring, and log inspection easily expand the platform to ensure server, application, and data security across physical, virtual, and cloud environments. Deep Security simplifies security operations while accelerating the ROI of virtualization and cloud projects.

Additional information about Trend Micro Deep Security can be found here:

<http://www.trendmicro.com/us/enterprise/cloud-solutions/deep-security/index.html>

### **McAfee MOVE Anti-Virus**

McAfee Management for Optimized Virtual Environments (MOVE) Anti-Virus for virtual desktops and servers is uniquely designed to relieve the overhead of traditional endpoint security, yet provide the protection and performance essential for success.

Additional information about McAfee MOVE Anti-Virus can be found here:

<http://www.mcafee.com/us/products/move-anti-virus.aspx>

### **Symantec Endpoint Protection**

Symantec Endpoint Protection offers comprehensive defense against complex attacks for both physical and virtual environments. It integrates nine essential security technologies in a single, high performance agent with a single management console. Endpoint Protection integrates with VMware® vShield Endpoint and provides leading protection without slowing you down.

Additional information about Symantec Endpoint Protection can be found here:

<http://www.symantec.com/endpoint-protection/>

## **12.4.3.5 Business Benefits**

VMware has partnered with antivirus software vendors to provide a bundled solution to antivirus problems in the VDI environment. VMware partners supply a dedicated, secure virtual appliance. This virtual appliance integrates with vShield Endpoint APIs to protect VMware virtual desktops against viruses and other malware. Instead of installing antivirus agents on each virtual desktop, you connect one virtual appliance to each virtual machine host.

The vShield Endpoint product offers many benefits for small and large-scale antivirus protection. Some of the largest benefits can be summarized by the following points:

- Enables VMware partners to eliminate antivirus storms that affect performance of virtual machines in a virtual desktop environment. Signature file updates and antivirus scanning are isolated and offloaded to a virtual appliance.
- Improves consolidation ratios of virtual desktops by offloading antivirus functions to a separate security virtual machine. The enterprise antivirus engine and the signature file are located on the virtual appliance, instead of on each virtual machine. This frees up virtual desktop system resources.
- Instead of an antivirus agent installed on each desktop to be protected, vShield Endpoint utilizes a thin agent that is already a part of the VMware Tools installation. The antivirus scanner and virus signatures are installed only in the virtual appliance.
- Ease of maintenance of the desktops to be protected: Any changes to the antivirus software are configured only in the virtual appliance, not in each desktop. You can change the configurations for the antivirus solution in the virtual appliance without reconfiguring the desktop driver. You do not have the responsibility of maintaining, patching, and updating antivirus agents on all of the desktops; you direct all changes to the virtual appliance instead.
- Simple addition or subtraction of antivirus vendors: Administrators can add or change partner solutions by adding or removing the virtual appliances. There is no need to reconfigure any antivirus agents on the desktop VMs.
- Satisfies audit requirements by providing detailed logging of antivirus tasks.

## **12.4.3.6 Summary**

VMware vShield Endpoint provides an elegant and comprehensive end point security solution that is virtualization aware and can integrate into existing antivirus and anti-malware implementations to allow for higher efficiency, greater availability, better performance, higher server density, and ultimately a better end-user experience.

## 12.4.4 Horizon Mirage for Physical desktops

### 12.4.4.1 Executive Summary

IT organizations today are faced with a number of challenges when deploying a desktop environment to support their business needs. Two of the greatest challenges when addressing end user computing environments are user migration to new or replacement hardware and data protection for those end point devices.

VMware Horizon Mirage provides a desktop and user data management solution that complements VMware Horizon View virtual desktop environments while addressing these challenges. It allows for the migration of a user's entire desktop environment from one device to another; whether it is a physical desktop or virtual machine. VMware Mirage also provides granular backup and restore capabilities that protect a user's data and entire desktop environment to provide excellent data protection and disaster recovery (DR) capabilities.

Dell Cloud Client Computing has combined Dell Desktop Virtualization Solutions (DVS), which is composed of best-of-breed Dell data center components with virtualization and management software from VMware, with the world-class portfolio of Dell Wyse virtualization end points to include Dell Wyse thin, zero, and cloud clients. VMware Horizon Mirage provides a comprehensive desktop management framework that allows administrators to protect the each user's desktop, files, and documents while also providing ways to migrate that desktop experience from physical to virtual devices and back again.

### 12.4.4.2 Introduction

As IT computing environments become larger and more complex, data protection and disaster recovery continue to be a large consideration for keeping users, data, and intellectual property protected. End users have become increasingly reliant on computing resources being available and data at their disposal. To address these challenges, enhanced migration and backup solutions are needed that are robust, work seamlessly, and easy to manage. If traditional methods of protecting workstations, or end points, are used, there can be adverse consequences ranging from reduced productivity to data loss. VMware has developed and integrated Mirage into Horizon Suite to address these challenges, which are discussed further in the following sections.

#### 12.4.4.2.1 Traditional Approach

In a typical enterprise IT organization, there are many challenges in supporting a desktop computing environment. The end user computing environment may consist of IT provided laptops, desktops, virtual desktops, or any combination thereof. Each computing device is typically installed with an IT developed image using a specific operating system version and pre-loaded software. The process of imaging the computing devices with an operating system is typically streamlined, but not a lot of thought is given to what happens after that.

Over time, a multitude of changes occur in the user computing environment relating to both hardware and software. Computing devices become obsolete or need to be repaired or replaced. Operating systems need to be reinstalled or upgraded. Hard drives, which could contain significant amounts of important data, will fail and need to be replaced. The backup of important files and data is left as a responsibility of each user to address. Different methods and techniques are used for protecting data; some which work and some that don't work quite as well. Viruses and malicious software can also delete or destroy important business data and information whether it is on a physical or virtual desktop.

All of these scenarios present significant disruptions in workflow for an end user and ultimately the business or organization. When a computer is lost, stolen, or damaged the data on it is potentially lost forever. Replacing the computer and restoring the data requires a significant effort from IT administrators. Using traditional tools, the process is also very time consuming and a very manual process. Once the IT administrator is finished restoring or imaging a new computer system, the user will also have to spend a significant amount of time restoring their computing environment to include applications, files and data, favorites, email settings, and preferences.

#### **12.4.4.2 VMware Horizon Mirage Approach**

In an enterprise IT organization that has deployed VMware Horizon Mirage, many of the primary challenges surrounding desktop computing environments are addressed with a comprehensive and integrated solution. The entire lifecycle of the user's desktop and data is protected and maintained from the time that it is deployed to an endpoint device to when it needs to be migrated or restored. A common desktop image can be used across both virtual and physical infrastructure to maintain a consistent computing environment while also giving each user the flexibility of customization.

This approach greatly reduces that amount of disruptions and interruptions in a typical user's workflow. Operating system upgrades post minimal interruptions. End user data is protected on the endpoint, whether it has been deleted, lost, or corrupted. Each user's desktop customizations are preserved across upgrades, migrations, and disruptive activities such as hardware refreshes.

#### **12.4.4.3 High Level Solution**

VMware Horizon Mirage provides a simplified approach to endpoint migration and protection with unique capabilities that integrate and complement virtual desktop environments and especially VMware Horizon View. When Mirage is installed on a Windows PC, either physical or virtual, it centralizes a complete virtual copy of that end point to the data center and keeps it synchronized. The synchronization includes changes from a user's Windows PC being uploaded to the data center, and changes from IT being downloaded and applied directly to the user's Windows PC.

#### **12.4.4.3.1 VMware Horizon Mirage Features**

VMware Horizon Mirage provides many features and benefits to IT organizations that are supporting physical or virtual desktop environments on a variety of endpoint devices. Two of the largest issues that are addressed by VMware Horizon Mirage in a Horizon View virtual desktop environment are migration of users to different hardware and endpoint data protection.

##### **Hardware Migration**

Horizon Mirage provides extensive capabilities and flexibility when migrating users from one endpoint device to another while maintaining their user data and personalization. Some examples of migration possibilities with Horizon Mirage are:

- Operating system upgrade (e.g. Windows XP to Windows 7 migrations)
- Windows 7 - 32 bit to 64 bit
- Physical to Virtual device
- Virtual to Physical device
- Virtual to Virtual device

These options give IT administrators the flexibility to dynamically address an ever changing computing environment. Some examples of projects that Mirage are a huge benefit in are:

- Migrations from physical endpoints to VMware Horizon View.
- Migration from Windows XP or Vista to Windows 7 operating systems.



- Migration of an office worker to a laptop to support a remote requirement.

## Total Endpoint Protection

When using Horizon Mirage for endpoint data protection, many of the issues with data loss are reduced or eliminated. When Mirage is used to protect an endpoint device, a snapshot of the entire system is taken and replicated back to the datacenter on an hourly, daily, weekly, monthly, or yearly basis. The IT administrator has full control over the frequency and retention of the snapshots for each endpoint. Mirage has many intelligent mechanisms built in to optimize this process. Some of the optimizations are:

- Intelligent file system filter driver that tracks changes to all files and directories.
- Only snapshot deltas or changes are replicated across the LAN or WAN.
- Recently transmitted blocks are cached to reduce network utilization.
- User activity monitoring to ensure operations do not impact user experience.

Files, directories, or even the entire operating system can be restored from any snapshot. Additionally, files and directories can be restored by an end user to allow for "self-service" restores. The end result is fast recovery from accidents or incidents lower IT operating costs, and ultimately happy users.

Figure 1 below shows how a centralized PC image can be deployed to a variety of endpoint devices:

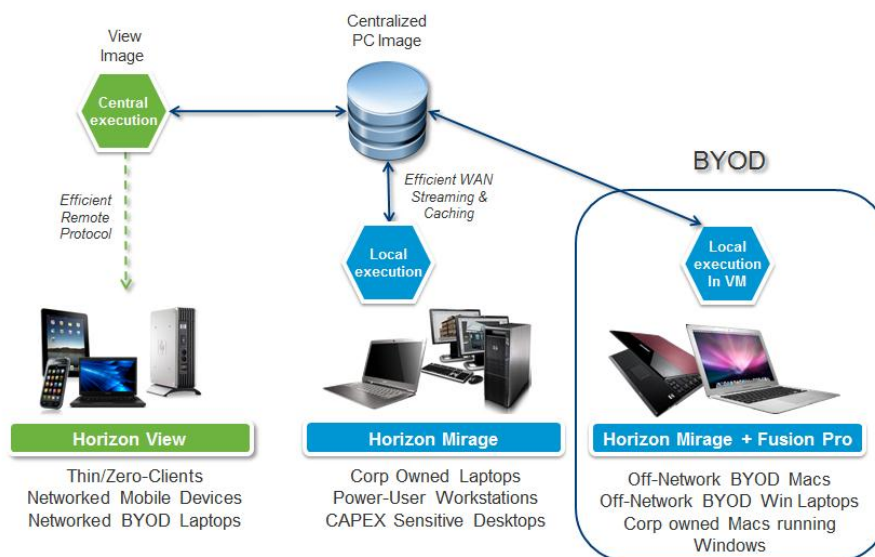


Figure 8

## 12.4.4.4 Solution Details

### 12.4.4.4.1 VMware Horizon Mirage Architecture

VMware Horizon Mirage is composed of the following components:

- Mirage Server
- Mirage Management Server
- Mirage Management Console
- Mirage Database
- Mirage Web Access
- Mirage Client

The purpose of each component is outlined below.

### **Mirage Server**

The Mirage Server is used to control all operations and objects within Mirage. It efficiently manages the storage and delivery of Centralized Virtual Desktops (CVDs), Base Layers, and Application Layers to clients, and consolidates monitoring and management communications. A Base Layer is used as a template for desktop content, cleared of specific identity information, and made suitable for central deployment to a large group of endpoints. Multiple Mirage Servers can be deployed as a server cluster to manage endpoint devices for large enterprise organizations.

### **Mirage Management Server**

The Mirage Management Server controls and manages the Mirage Server cluster. It also creates and interfaces with the Mirage database.

### **Mirage Management Console**

The Mirage Management Console is the graphical user interface used to perform scalable maintenance, management, and monitoring of deployed endpoints.

Through the Mirage Management Console, the Administrator configures and manages Clients, Base and App Layers, and reference machines, performs operations such as update and restore, and monitors the system operation through the dashboard and event logs.

### **Mirage Database**

The Mirage database contains pointers to the base and application layers and desktop images in storage, and also an inventory of what is on the endpoints. The database catalogs the information while the actual information is placed on storage.

### **Mirage Web Access**

Mirage Web Access provides access to the Administrative and File Portals via a web browser. The Horizon Mirage File Portal enables end users to view their files within historical snapshots of their datacenter desktop image. Users can access their files through a web browser from any device. Because these files are stored in the datacenter, users can view their files even if the normal Mirage-managed endpoint is damaged or lost.

### **Mirage Client**

The Mirage Client is installed on the endpoint device. This software executes in the base operating system, making sure the image at the endpoint and the server are fully synchronized.

The Mirage Client has a role in managing uploads and downloads between the datacenter desktop image and the endpoint. The Mirage Client helps to:

- Track changes on each endpoint's file system
- Synchronize and securely replicate changes on the local system to the image in the datacenter
- Monitor user activity to ensure good user experience
- Restore endpoint devices to a previous operating state when corrupted or destroyed

The following versions of Windows client operating systems are currently supported with VMware Horizon Mirage:

- Windows XP SP2 and above (32 bit)
- Windows Vista Business or Enterprise (32/64 bit)

- Windows 7 Professional or Enterprise (32/64 bit)
- Guest Virtual Machine support on Macs and Linux-based systems

#### 12.4.4.4.2 How VMware Horizon Mirage Components Work Together

Mirage Server (or a cluster of Mirage Servers) manages desktop images in the datacenter and orchestrates uploads and downloads between datacenter desktop images and Mirage-managed endpoints. The Mirage Server components require a connection to a Microsoft SQL database.

Storage disks in the datacenter contain the desktop images, and base and application layers. The Mirage database contains pointers to the base and application layers and desktop images in storage, and also an inventory of what is on the endpoints. The database catalogs the information; the storage contains the actual information.

Each Mirage-managed physical or virtual computer has the Mirage Client installed, which communicates closely with the Mirage Server to synchronize the datacenter desktop image with changes to the endpoint.

Figure 2 outlines the basic relationships between the components of VMware Horizon Mirage:




Figure 9

#### 12.4.4.4.3 VMware Horizon Mirage Licensing

Licensing of VMware Horizon Mirage is as follows:

- Horizon Mirage is priced and licensed on a per-named-user model.
- Horizon Mirage is available for purchase—a la carte or bundled in the Horizon Suite—directly from VMware or any VMware authorized reseller partner.

#### 12.4.4.5 Business Benefits

The VMware Horizon Mirage product offers many benefits for organizations with small and large-scale virtual desktop environments. Some of the largest benefits can be summarized by the following points:

- Easily migrate branch and remote office employees without added infrastructure costs.
- Complete desktop protection for all users on LAN or WAN.
- Enable desktop backup and recovery.

- Reduce user downtime.

#### 12.4.4.6 Summary

VMware Horizon Mirage provides an elegant and comprehensive end point migration and disaster recovery solution that can integrate into new or existing physical and virtual desktop environments to allow for greater availability, seamless migrations, faster recoveries, and ultimately a better end-user experience.

### 12.5 VDI within a M1000e Chassis using EqualLogic PS-M4110XS Storage Blades

The information contained within this section was collaborative effort between Dell Enterprise Storage Solutions Group and Dell Cloud Client Computing teams. For full documentation and information on this topic please go to document titled "VMware Horizon View 5.2 on Dell EqualLogic PS-M4110XS Hybrid Blade Storage Arrays" [here](#). Questions about these reference architectures can be directed to both the teams.

#### 12.5.1 Executive Summary

In this section, we present a unified compute, storage, and switching all-in-blade-form-factor platform for hosting and running VDI workloads. This unified and integrated Dell solution delivers the performance, scalability, management simplicity, and desktop density that are essential for delivering the VDI potential.

A key component of this unified VDI solution is Dell EqualLogic PS-M4110XS hybrid arrays, which provide a blade form factor suitable for a complete, self-contained VDI solution within a modular and compact blade enclosure. Together with Dell PowerEdge™ blade servers and Dell Force10™ blade switches, these hybrid blade arrays create a "data center in a box" for VDI deployments. This approach helps organizations reduce virtual desktop deployment and operational costs through efficient use of switching resources, minimized cabling, and consolidated management.

This section demonstrates how a modular 1000 standard user virtual desktop environment – all self-contained within a Dell PowerEdge M1000e blade chassis – can be deployed in a VMware Horizon View 5.2 (Horizon View) VDI infrastructure leveraging 12 PowerEdge M620 blade servers, four Force10 MXL blade switches, and two EqualLogic PS-M4110XS hybrid blade arrays. Details are provided for the storage I/O characteristics under various VDI workload scenarios like boot and login storms along with performance characteristics throughout the VDI stack (e.g. ESXi server performance as well as user experience as determined by Liquidware Stratusphere UX).

#### 12.5.2 Objectives

The primary objectives of the tests conducted for this paper were:

- Develop best practices and sizing guidelines for a Horizon View based VDI solution deployed within a single Dell PowerEdge M1000e blade chassis
- Determine how many virtual desktops can be deployed in this environment using a single Dell EqualLogic PS-M4110XS blade storage array with acceptable user experience indicators for a standard user workload profile
- Analyze the performance impact with an additional Dell EqualLogic PS-M4110XS blade storage array
- Determine the performance impact on the storage array of peak I/O activity such as boot and login storms
- Determine the optimal compute, storage and switching infrastructure for a VDI deployment that is modular and completely self-contained within a blade chassis.

The test infrastructure used includes:

- VMware Horizon View 5.2
- VMware vSphere 5.1 hypervisor
- Dell PowerEdge M620 blade servers
- Dell Force10 MXL switches
- Dell EqualLogic PS-M4110XS blade storage arrays

### 12.5.3 Audience

This paper is intended for solution architects, storage network engineers, system administrators, and IT managers who need to understand how to design, properly size, and deploy Horizon View based VDI Solutions using Dell EqualLogic blade storage. It is expected that the reader has a working knowledge of the Horizon View architecture, VMware vSphere system administration, iSCSI SAN network design, and Dell EqualLogic iSCSI SAN operation.

### 12.5.4 VDI with Dell EqualLogic PS Series Blade Storage

The Dell EqualLogic PS-M4110 blade arrays offer virtualized, enterprise-class storage in a consolidated blade form factor. These blade storage arrays offer intelligent self-optimization, automation, ease-of-use, and data protection. Integrating with Dell PowerEdge M-series blade servers and Dell Force10 MXL blade switches within a Dell PowerEdge M1000e blade chassis, these arrays enable modular, self-contained VDI solutions within the blade chassis form factor that help organizations to simplify management, enhance efficiency, and deploy and scale VDI solutions quickly.

The double-wide, half-height blade form factor of the EqualLogic PS-M4110 blade array plugs into the PowerEdge M1000e enclosure. It features dual, hot-pluggable 10GbE controllers. The array is available in a variety of disk configurations, with the EqualLogic PS-M4110XS hybrid blade array being optimal for VDI deployments. This hybrid blade configuration consists of 14 2.5" drives – five 400 GB SSDs and nine 600 GB 10,000 rpm SAS drives – for 7.4 TB of raw storage capacity. Like the other EqualLogic hybrid arrays, EqualLogic PS-M4110XS offers automated load balancing and data tiering within the SSDs and HDDs, and is highly adaptive to the utilization spikes of the VDI workload.

For more information on Dell EqualLogic hybrid array load balancer, see:

<http://en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/20349614/download.aspx>

The modular VDI architecture with compute, storage, and switching all contained within a single PowerEdge M1000e chassis creates a comprehensive, end-to-end desktop virtualization solution that enables IT organizations to do more while consuming less space, power, and cabling than the traditional rack-mounted storage approaches. Specifically, shared resources such as chassis-based power and cooling and backplane connectivity help reduce cabling and minimize space requirements cost-effectively. Additionally, unified management through the Dell Chassis Management Controller (CMC) enables rapid deployment and infrastructure provisioning without requiring specialized expertise.

The figure below shows the modular, all-within-the-blade-chassis VDI building block architecture for 1000 standard user desktops for Horizon View-based environments with 12 PowerEdge M620 servers, four Force10 MXL switches, and two EqualLogic PS-M4110XS blade arrays. Two of the twelve blade servers were used to for hosting infrastructure VMs with the remaining ten blade servers being used to host virtual desktop VMs.





Horizon View Composer server, Microsoft Windows Server 2008 R2 based file server, and SQL Server 2008 R2.

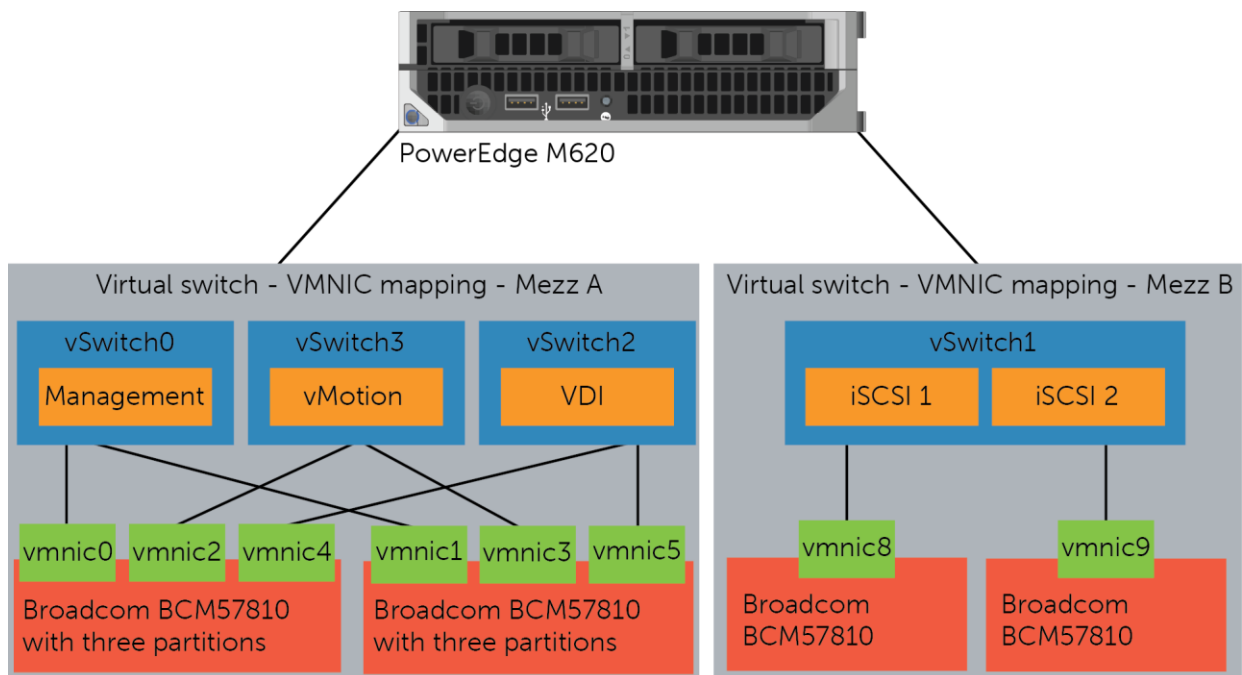
- **Horizon View Client Cluster:** Ten M620 blade servers hosting virtual desktops.

**Note:** VMware has removed the limitation of having a maximum of eight hosts per vCenter cluster in a non-NFS shared datastore with release 5.2 of Horizon View. More information can be found here: <http://www.vmware.com/support/view52/doc/horizon-view-52-release-notes.html>

In addition to the above servers, three Dell PowerEdge R810 rack servers were used for VDI load generation purposes but are not part of the design or architecture herein described and are solely used for testing.

### 12.5.5.2 Network design considerations

The Figure below shows the network layout at one of the 12 PowerEdge M620 blade servers that has ESXi 5.1 installed on it.



The M1000e blade chassis consisted of the following switches:

- Two Force10 MXL blade switches in fabric A for connectivity to the Management LAN, VDI client LAN, and a vMotion LAN.
- Two Force10 MXL blade switches in fabrics B for connectivity to the dedicated iSCSI SAN.

Network Interface Card Partitioning (NPAR) was used to divide the physical NICs in Fabric A into multiple logical NICs. This allows for dynamic allocation of bandwidth for the different partitions and this helps reduce the total cost of ownership for the solution. For the purposes of our testing, all partitions were given access to 100% of the available 10 Gb bandwidth. This allows the partitions to use all the bandwidth when required.

More information on NIC partitioning with Broadcom NICs is available in the white paper titled "Enhancing Scalability through Network Interface Card Partitioning", available here: <http://www.dell.com/downloads/global/products/pedge/en/Dell-Broadcom-NPAR-White-Paper.pdf>

The following partitions were created on the Broadcom NICs on Fabric A.

Partition Name	Purpose
Management	Management access to the ESXi hosts and infrastructure VMs.
vMotion	VMware vMotion access to provide live migrations with zero downtime as well as load balancing on ESXi hosts.
VDI	VDI LAN over which the clients access their desktops.

The networks are segregated at the Force10 MXL switch using VLANs to separate different types of traffic, namely:

- **Management LAN:** This network provides a separate management network for all the physical ESXi hosts. It also allows communication between various infrastructure components such as Microsoft Active Directory Server, Microsoft SQL Server, and VMware vCenter server.
- **VDI Client LAN:** This is the network over which the clients access the virtual desktops in the Horizon View desktop pools. The connectivity to the existing network of the client is provided by uplink modules on the Force10 MXL switches.
- **VMware vMotion LAN:** This is the network over which the VMs are migrated to provide high availability and load balancing between the hosts. Since the 12 M620s are divided into two clusters, only hosts in the same cluster can provide HA and load balancing.

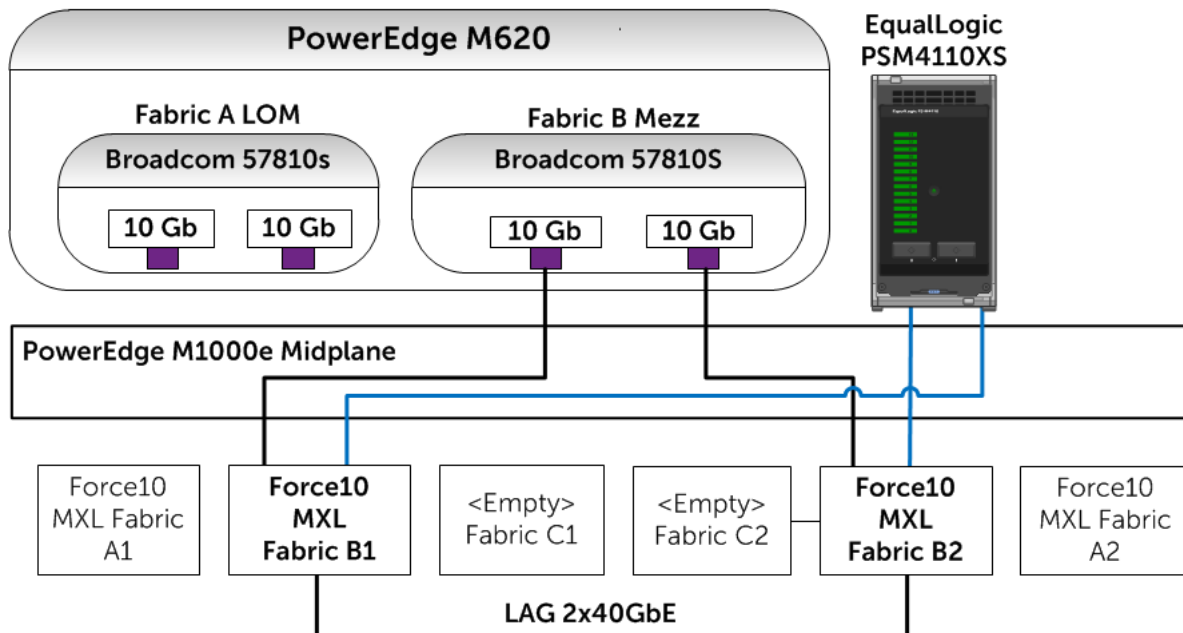
A pair of dedicated switches provide access to a **dedicated iSCSI SAN** through which all the virtual desktops and other infrastructure components access the EqualLogic storage arrays.

### 12.5.5.3 iSCSI SAN configuration

The Figure below shows the network connectivity between a single M620 blade server and the blade storage array through the blade server chassis. The figure shows only one M620, though all 12 of the blades were used in the testing. The topology is identical for all the remaining blades in the chassis.

- Each PowerEdge M620 blade server is configured with two Broadcom 57810S-k Dual Port 10 Gb NIC cards. One card was assigned as the fabric A LOM and the other as fabric B on the blade chassis.
- Fabric A is entirely used to provide Management LAN, VDI Client LAN, and vMotion LAN functionality while fabric B is entirely used to provide dedicated iSCSI SAN connectivity.
- Dual Force10 MXL switches were installed in fabrics A and B on the blade server chassis. The NIC cards on the blade servers are internally connected to each of these switches through the mid-plane on the blade server chassis.
- The Force10 MXL switches in fabric B are interconnected using two 40 GbE links to provide high availability and redundancy of the iSCSI fabric.
- The EqualLogic blade storage arrays are setup to communicate on fabric B. These are connected to the servers through the MXL blade switch internally.
- Fabric C is unused.





#### 12.5.5.4 Separation of user data and virtual desktop data

Typically, user data can be maintained on a separate file share through roaming profiles and folder redirections. In this approach, a desktop – comprised of the shared read-only base image and the individual differential data from temporary changes – is always stateless. When a user logs in, the personal data is mapped to the desktop and all the changes to the data are maintained on the file share over the network. When the user logs off, the personal data is no longer on the desktop and the desktop can be put back in the original state with the temporary data discarded.

This approach has two benefits. First, the advantages of the non-persistent desktop deployment model can be leveraged while using the persistent desktop model. Second, the performance needs of VM data and user data are distinctly different, with the former needing very high performance to handle the I/O storms. The VM data can be placed in a high-performance storage array, while the user data can be served from the file shares running on a capacity-oriented storage array.

Additionally, if a single storage array can cost-effectively serve both the performance and capacity needs of VM data and user data, then separate volumes (maintaining the VM data and user data separation) from the same storage array can be used for this deployment model. EqualLogic hybrid arrays are ideal for this approach as it automatically tiers data based on the I/O workload thereby lowering the storage costs for VDI deployments. For these reasons, the user data and the VM data were stored on the same EqualLogic hybrid blade arrays (but on different volumes) in the test environment. Each desktop was assigned about 2 GB of user data space on a 2 TB volume that can support up to 1000 user desktops.

#### 12.5.5.5 EqualLogic storage array configuration

Dell EqualLogic PS-M4110XS blade storage arrays hosted all the virtual desktops as well as the infrastructure virtual machines used in this solution. Initially, tests were conducted with a single PS-M4110XS hybrid array. Once all the tests were performed with one array, an additional PS-M4110XS hybrid array was added to the same pool and the tests were repeated. The volume layout used for the infrastructure functions including user data is shown below.

EqualLogic volumes layout for hosting infrastructure components and user data

Volume name	Size	Purpose
Infra-1	500 GB	Storage for Active Directory, SQL Server, vCenter Server, Horizon View Server-1
Infra-2	500 GB	Storage for File server, Horizon View Server-2, Horizon View Composer

UserSpace	2 TB	Storage for User profiles and folder redirection space (Average, 2 GB per user)
-----------	------	---

In addition to the infrastructure volumes, the storage arrays also provided shared storage for hosting the virtual desktops. The volume layout used for configuring the base image, replicas, and VDI volumes on the two arrays is shown below.

#### EqualLogic layout for volumes hosting virtual desktops

Volume name	Size	Purpose
VDI-Baselmages	100 GB	Storage for Base image for VDI deployment
VDI-Replicas	100 GB	Storage for Replica images created by Horizon View
VDI-Images1	500 GB	Storage for VDI Virtual Machines in Horizon View Cluster
VDI-Images2	500 GB	Storage for VDI Virtual Machines in Horizon View Cluster
VDI-Images3	500 GB	Storage for VDI Virtual Machines in Horizon View Cluster
VDI-Images4	500 GB	Storage for VDI Virtual Machines in Horizon View Cluster
VDI-Images5	500 GB	Storage for VDI Virtual Machines in Horizon View Cluster
VDI-Images6	500 GB	Storage for VDI Virtual Machines in Horizon View Cluster
VDI-Images7	500 GB	Storage for VDI Virtual Machines in Horizon View Cluster
VDI-Images8	500 GB	Storage for VDI Virtual Machines in Horizon View Cluster

### 12.5.5.6 ESXi host network configuration

VMware ESXi 5.1 hypervisor was installed on all 12 blades. The network configuration on each of those hosts is described below. Each ESXi host was configured with four virtual switches, vSwitch0, vSwitch1, vSwitch2, and vSwitch3 to separate the different types of traffic on the system.

#### vSwitch configuration in ESXi hosts

vSwitch	Description	Virtual NICs used
vSwitch0	Management Network	vmnic0, vmnic1
vSwitch1	iSCSI SAN	vmnic8, vmnic9

vSwitch2	VDI LAN	vmnic4, vmnic5
vSwitch3	vMotion LAN	vmnic2, vmnic3

### 12.5.5.7 Horizon View Configuration

Horizon View 5.2 was installed by following the documentation provided by VMware.

Horizon View 5.2 Documentation: <http://pubs.vmware.com/view-52/index.jsp>

Specific configuration decisions used in the configuration:

- Two Horizon View servers were configured to provide load balancing and high availability.
- The Horizon View servers were installed as VMs on two separate hosts with four virtual CPUs, 12 GB of RAM and one 40 GB virtual hard drive.
- The first Horizon View Server was configured as a "View Standard Server" during the installation, while the second Horizon View Server was installed as a "View Replica Server".
- Horizon View Composer was installed in a separate VM with the same properties as the Horizon View servers.
- Self-signed SSL certificates were applied to the VMware vCenter server VM, Horizon View servers, and the Horizon View Composer server.

### 12.5.5.8 Horizon View pool configuration

The Add pool wizard in Horizon View was used to choose the following specific configuration options to create the virtual desktop pool.

Option	Selected setting
Virtual desktop pool type	Automated Pool
User Assignment	Floating
vCenter Server	View Composer linked clones
View Composer Disks	Redirect disposable files to a non-persistent disk of size 4096 MB
Storage Optimization	Select separate datastores for replica and OS disk
Advanced Storage Options	Use host caching
Guest Customization	Sysprep

More information about Horizon View Pool configuration can be found in the VMware Horizon View 5.2 documentation:  
<http://pubs.vmware.com/view-52/topic/com.vmware.view.administration.doc/GUID-0A9CA985-3A78-428A-BCFB-B3E2DCCA90AD.html>

### 12.5.5.9 Windows 7 Desktop VM configuration

Following the guidelines from VMware and Login VSI, the Windows 7 base image was generated based on a generic base VM with the following properties:

- VMware Virtual Hardware version 8
- One virtual CPU
- 1.5 GB RAM
- 25 GB virtual hard drive
- One virtual NIC connected to the VDI Network
- Windows 7 64 bit OS

Additionally, the base image was customized using the VMware Optimization guide for Windows 7, available here: <http://www.vmware.com/files/pdf/VMware-View-OptimizationGuideWindows7-EN.pdf>

## 12.5.6 Horizon View test methodology

This section outlines the test objectives along with the test tools and criteria used to determine the sizing guidelines and best practices for deploying Horizon View on EqualLogic storage.

### 12.5.6.1 Test objectives

- Develop best practices and sizing guidelines for a Horizon View based VDI solution deployed within a single Dell PowerEdge M1000e blade chassis
- Determine how many virtual desktops can be deployed in this environment using a single Dell EqualLogic PS-M4110XS blade storage array with acceptable user experience indicators for a standard user workload profile
- Determine the performance impact on the storage array of peak I/O activity such as boot and login storms
- Determine the performance impact on the user experience indicators of scaling with an additional Dell EqualLogic PS-M4110XS blade storage array

### 12.5.6.2 Test tools

All tests were conducted using Login VSI 3.7 as the workload generator and user experience analyzer tool. Login VSI, developed by Login Consultants, is a VDI benchmarking and load generation tool.

**Note:** Login VSI is a benchmarking tool to measure the performance and scalability of centralized desktop environments such as Server Based Computing (SBC) and VDI. More information can be found at: <http://www.loginvsi.com>

#### 12.5.6.2.1 Load generation

The "Medium" workload from Login VSI was used to simulate the standard user workload. The characteristics of the Medium workload are:

- Up to five applications are open simultaneously.
- Applications include Microsoft Internet Explorer, Microsoft Word, Microsoft Excel, Microsoft PowerPoint, PDF reader, 7-Zip compression software, Movie player.
- Once a session is started, the medium workload repeats approximately every 12 minutes.
- During a loop, the response time is measured every two minutes.
- Idle time is about two minutes in each 12 minute loop.
- Type rate is approximately 130 ms per character.

Although Login VSI provides other workloads, the Medium workload was used in the testing because it closely resembles the workload of a "Standard" user.

#### 12.5.6.2.2 Monitoring tools

The following monitoring tools were used:

- Dell EqualLogic SAN Headquarters (SAN HQ) for monitoring storage array performance.
- VMware vCenter statistics for ESXi performance.
- Login VSI Analyzer and Liquidware Labs Stratusphere UX for end user performance statistics.  
Detailed performance metrics were captured from the storage arrays, hypervisors, virtual desktops, and the load generators during the tests.

### 12.5.6.3 Test criteria

The primary focus of the tests was to determine the maximum number of desktops which can be deployed on a single PS-M4110XS blade storage arrays in this environment while using VMware Horizon View Composer to provide Linked Clone virtual desktops in an automated pool. The tests were then rerun to find the maximum number of desktops that can be hosted on two PS-M4110XS blade storage arrays.

VDI configurations involve many components at different layers – application, hypervisor, server, network, and storage. As a result, multiple metrics need to be captured at different layers to ensure that the environment is healthy and performing optimally and appropriately for all users.

The specific test criteria are described in the following sections.

#### 12.5.6.3.1 Storage capacity and I/O latency

The typical industry standard latency limit for storage disk I/O is 20 ms. Maintaining this limit ensures good user application response times when there are no other bottlenecks at the infrastructure layer. In addition to this, it is also recommended to maintain a 10% spare capacity on the storage array for optimal performance.

#### 12.5.6.3.2 System utilization at the hypervisor

Even though the primary focus of these tests was storage characterization, additional metrics at the hypervisor infrastructure layer were defined to ensure solution consistency. These are:

- CPU utilization on any ESXi server should not exceed 85%.
- Minimal memory ballooning on the VMs.
- Total network bandwidth utilization should not exceed 90% on any one link.
- TCP/IP storage network retransmissions should be less than 0.5%.

#### 12.5.6.3.3 Virtual desktop user experience

Stratusphere UX was used to gather data for user experience and desktop performance. Data gathered from the hosts in VMware vCenter and the virtual desktops (via software installed on the VM) is reported back to the Stratusphere Hub. The Stratusphere Hub was used to generate the comprehensive report “Desktop Virtualization Performance Analysis.” This report includes information about the host performance, virtual machine performance and user performance. The report also provides a scatter plot that shows the performance of all the users in the system.

Liquidware Labs Stratusphere UX can generate a variety of reports that compare and validate user experience metrics. More information about these reports, including sample reports can be found here:

<http://www.liquidwarelabs.com/products/stratusphere-ux-validation-reports.asp>

Liquidware Lab Stratusphere UX user experience metric was used to ensure that all desktops had acceptable levels of application performance.

Additionally, Login VSI Analyzer was also used to gather metrics on the user experience at the virtual desktop layer to ensure that all the desktops had acceptable levels of application performance. See Appendix C for details about Login VSI’s VSIMax (Dynamic) parameter results.

## 12.5.6.4 Test configuration

A single virtual desktop pool was configured using the Horizon View Administrator interface. Each pool was built from a Windows 7 base image.

Desktop pool properties:

- One Horizon View desktop pool.
- Pool with 500 desktops, spread across four 500 GB volumes (VDI-Images1 through 4).
- 500 desktops were deployed across five hosts (100 desktops per host).
- Base Images are stored on a separate 100 GB volume (VDI-BaseImage).
- Replica images are stored on a separate 100 GB volume (VDI-Replicas).
- Storage over-commit for all volumes was set to aggressive.
- Horizon View Composer disposable disk size was 4096 MB.
- Disposable disk drive letter was set to auto.
- Host caching (Horizon View Storage Accelerator) was enabled for all hosts.

When an additional PS-M4110XS blade storage array was added, the desktop pool was modified so that the additional 500 desktops were hosted on four additional VDI-Images volumes.

## 12.5.7 Test results and analysis

This section presents the results from the different Horizon View VDI characterization tests and the key findings from each test. The standard user workload represents a majority of the VDI users in the industry today, and the testing was focused on this workload profile.

### 12.5.7.1 Test scenarios

The following tests were conducted to gather results and analysis on the solution stack.

#### 1. One array tests for standard users

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

#### 2. Two array tests for standard users

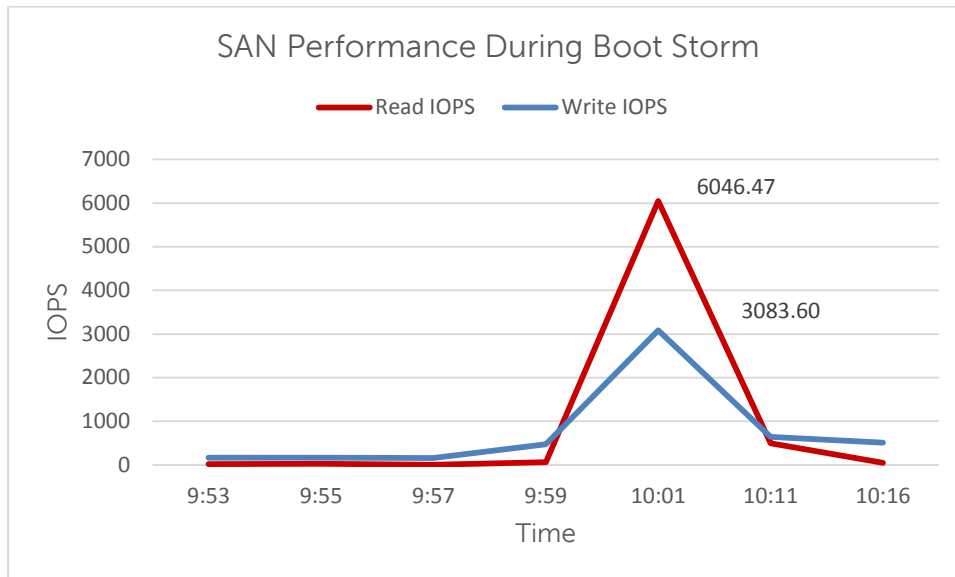
In these tests, an additional EqualLogic PS-M4110XS blade storage array was added to the chassis and the number of desktops was doubled. The boot storm, login storm, and steady state tests were repeated for this configuration. The intent of this test was to show the linear scalability of adding an EqualLogic array.

## 12.5.7.2 One array test for standard users

The following sections provide results from the boot storm, login storm, and steady state testing for a standard user on one PS-M4110XS array.

### 12.5.7.2.1 Boot storm I/O


To simulate a boot storm, the virtual desktops were reset simultaneously from the VMware vSphere client. The figure below shows the storage characteristics during the boot storm – the PS-M4110XS array delivered 9130 IOPS (14-18 IOPS per VM) under the peak load during this test and all 500 desktops were available in about 10 minutes.



SAN HQ data showing PS-M4110XS array IOPS during boot storm

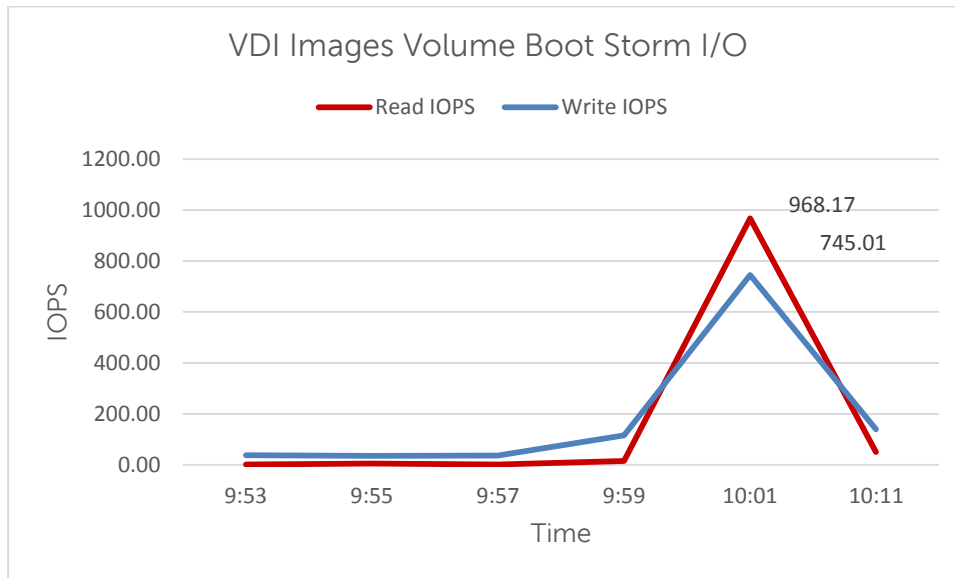
The spike seen in the figure above was primarily due to read operations that occur on the replica volume. This is because the boot operation causes many simultaneous reads to the replica image. The read and write ratios during the boot storm were about 66% reads to 34% writes.

On the replica image volume, nearly 100% reads and negligible writes were observed, as seen in graph below. This is because the replica volume only hosts a copy of the base image that all the VMs are linked to and due to this almost all of the I/O operations on this volume are reads.



Read and write IOPS on the replica volume

The read and write ratios on the remaining volumes that hosted the virtual machines were about 56% reads and 44% writes. The mix of reads and writes here is because changes made to the OS are written to these volumes. The IOPS on one such volume are shown in the graph below. The other volumes that hosted the VMs showed similar IOPS as the graph below.



Read and write IOPS on one of the VDI-Images volumes

The ESXi hosts did not show any bottlenecks with respect to meeting CPU and memory resource demands during the boot operation. Other infrastructure servers, such as the Horizon View Servers and Active Directory servers, did not show any exceptional load during the boot storm.

The table below shows that most of the operations during the boot storm were handled by SSDs in the hybrid array. This ensured that the boot storm lasted for a short time and all the desktops were available quickly. The SSD drives also acted as an accelerated write cache on the storage array providing additional caching for the write I/O load.

Disk usage on PS-M4110XS during boot storm

Member	Disk	Description	Average IOPS	Read I/O rate	Write I/O rate	Status
PS-M4110XS	0	SSD 400GB SAS	2266.16	32.60 MB/sec	16.82 MB/sec	online
PS-M4110XS	1	SSD 400GB SAS	2261.68	32.58 MB/sec	16.79 MB/sec	online
PS-M4110XS	2	SSD 400GB SAS	2269.99	32.49 MB/sec	16.80 MB/sec	online
PS-M4110XS	3	SSD 400GB SAS	2267.08	32.61 MB/sec	16.75 MB/sec	online
PS-M4110XS	4	SSD 400GB SAS	2296.26	32.56 MB/sec	16.77 MB/sec	online
PS-M4110XS	5	10K 600GB SAS	2.53	3.52 KB/sec	0 KB/sec	online
PS-M4110XS	6	10K 600GB SAS	0	0 KB/sec	0 KB/sec	online
PS-M4110XS	7	10K 600GB SAS	0	0 KB/sec	0 KB/sec	online
PS-M4110XS	8	10K 600GB SAS	<1.0	29.96 KB/sec	0 KB/sec	online
PS-M4110XS	9	10K 600GB SAS	0	0 KB/sec	0 KB/sec	online
PS-M4110XS	10	10K 600GB SAS	0	0 KB/sec	0 KB/sec	online
PS-M4110XS	11	10K 600GB SAS	0	0 KB/sec	0 KB/sec	online
PS-M4110XS	12	10K 600GB SAS	0	0 KB/sec	0 KB/sec	online
PS-M4110XS	13	10K 600GB SAS	0	0 KB/sec	0 KB/sec	spare



Storage network utilization was well within the available bandwidth. The peak network utilization during the boot storm reach approximately 7.7% of the total network bandwidth and then gradually declined once all the VMs were booted up. There were also no retransmissions on the iSCSI SAN.

These results show that the EqualLogic PS-M4110XS hybrid array can handle a heavy random I/O load like a boot storm with no issues.


#### **12.5.7.2.2 Login storm I/O**

Login VSI was programmed to launch 500 virtual desktops over a period of about 15 minutes after pre-booting the virtual desktops. The peak IOPS observed during the login storm was about 4900 IOPS (7-10 IOPS per VM).

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

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

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



SAN HQ data showing login storm I/O

The peak latency seen on the storage array is less than 20 ms, and the storage array is able to handle the 500 users logging in over such a short duration with no performance issues. The table below shows the overall usage of the disks in the array during a login storm as collected by SAN HQ.

Disk usage on PS-M4110XS during login storm

Member	Disk	Description	Average IOPS	Read I/O rate	Write I/O rate	Status
PS-M4110XS	0	SSD 400GB SAS	1838.53	19.41 MB/sec	26.28 MB/sec	online
PS-M4110XS	1	SSD 400GB SAS	1838.99	19.47 MB/sec	26.22 MB/sec	online
PS-M4110XS	2	SSD 400GB SAS	1852.14	19.55 MB/sec	26.34 MB/sec	online
PS-M4110XS	3	SSD 400GB SAS	1840.26	19.53 MB/sec	26.24 MB/sec	online
PS-M4110XS	4	SSD 400GB SAS	1836.91	19.49 MB/sec	26.29 MB/sec	online
PS-M4110XS	5	10K 600GB SAS	0	0 KB/sec	0 KB/sec	online
PS-M4110XS	6	10K 600GB SAS	0	0 KB/sec	0 KB/sec	online
PS-M4110XS	7	10K 600GB SAS	0	0 KB/sec	0 KB/sec	online
PS-M4110XS	8	10K 600GB SAS	0	0 KB/sec	0 KB/sec	online
PS-M4110XS	9	10K 600GB SAS	<1.0	4.06 KB/sec	0 KB/sec	online
PS-M4110XS	10	10K 600GB SAS	0	0 KB/sec	0 KB/sec	online
PS-M4110XS	11	10K 600GB SAS	0	0 KB/sec	0 KB/sec	online
PS-M4110XS	12	10K 600GB SAS	0	0 KB/sec	0 KB/sec	online
PS-M4110XS	13	10K 600GB SAS	0	0 KB/sec	0 KB/sec	spare

The table above clearly shows that most of the login storm I/O is handled by SSDs and therefore the array is able to provide the best possible performance.

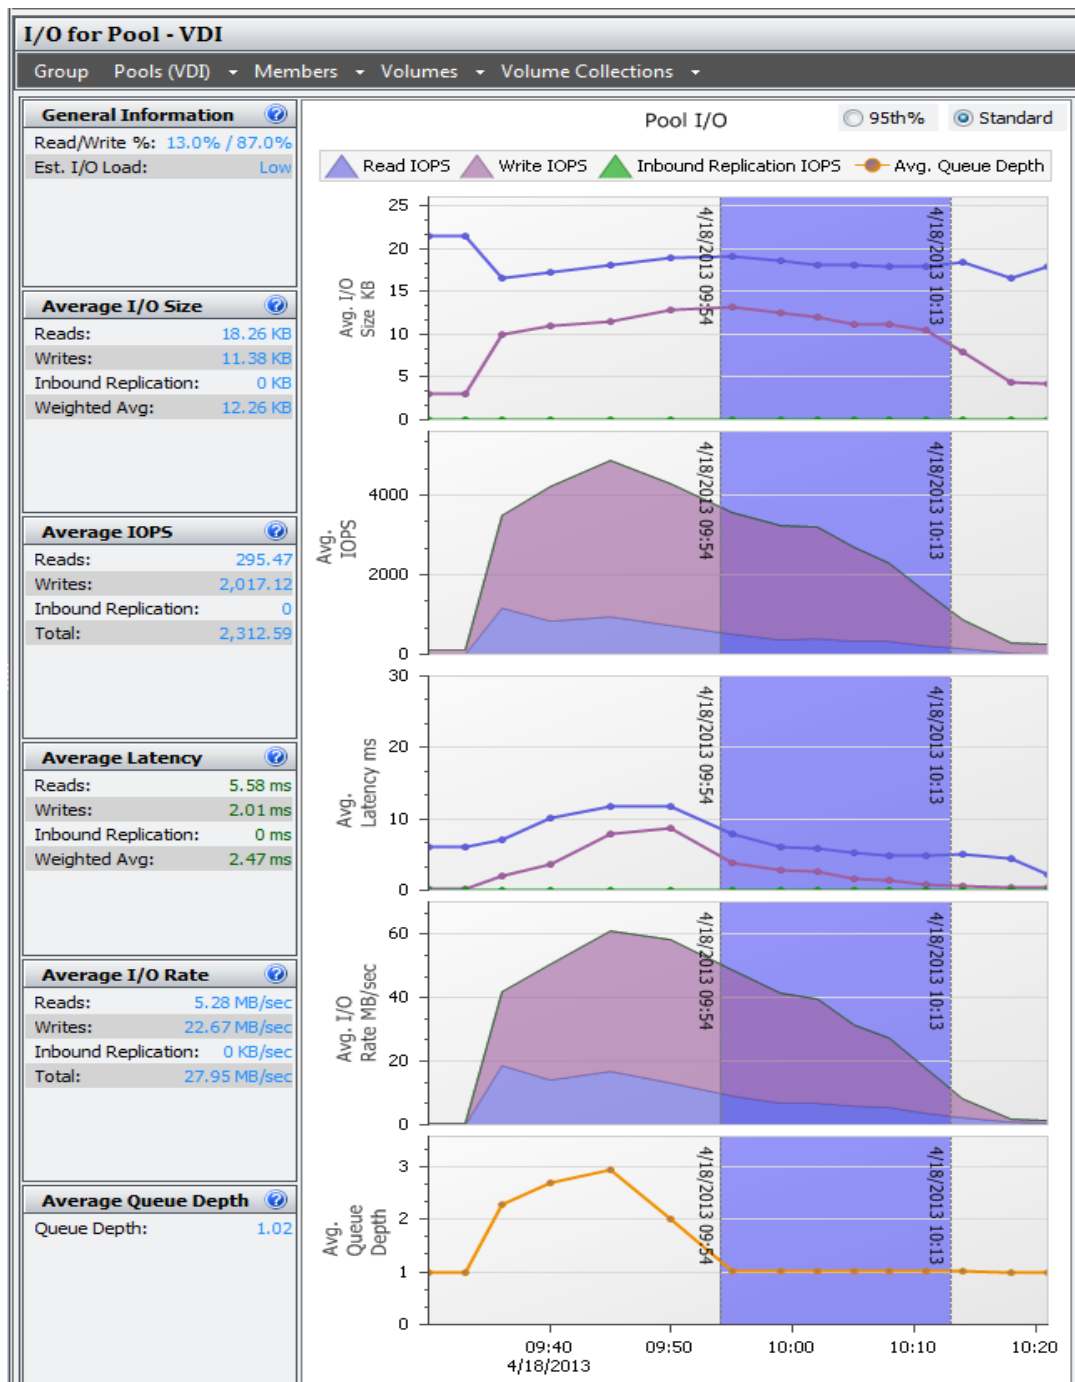
### 12.5.7.2.3 Steady state I/O

The total IOPS on the EqualLogic PS-M4110XS array averaged during the steady state with all the users logged in was around 2,300 (6-8 IOPS per VM). Of this, the read IOPS accounted for about 295 (approximately 13% of the total I/O load) and the remaining 2000 IOPS or 87% were write IOPS. Read and write latencies were also well below 20 ms throughout the test.

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

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

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



SAN HQ data showing steady state I/O


### 12.5.7.3 Two array tests for standard users

The following sections provide results from the boot storm, login storm, and steady state testing for a standard user on two PS-M4110XS arrays.

#### 12.5.7.3.1 Boot storm I/O

The boot storm test from Section 0 was repeated after adding another PS-M4110XS array and an additional 500 desktops.

The figure below shows the I/O pattern on the pool comprised of the two PS-M4110XS arrays.



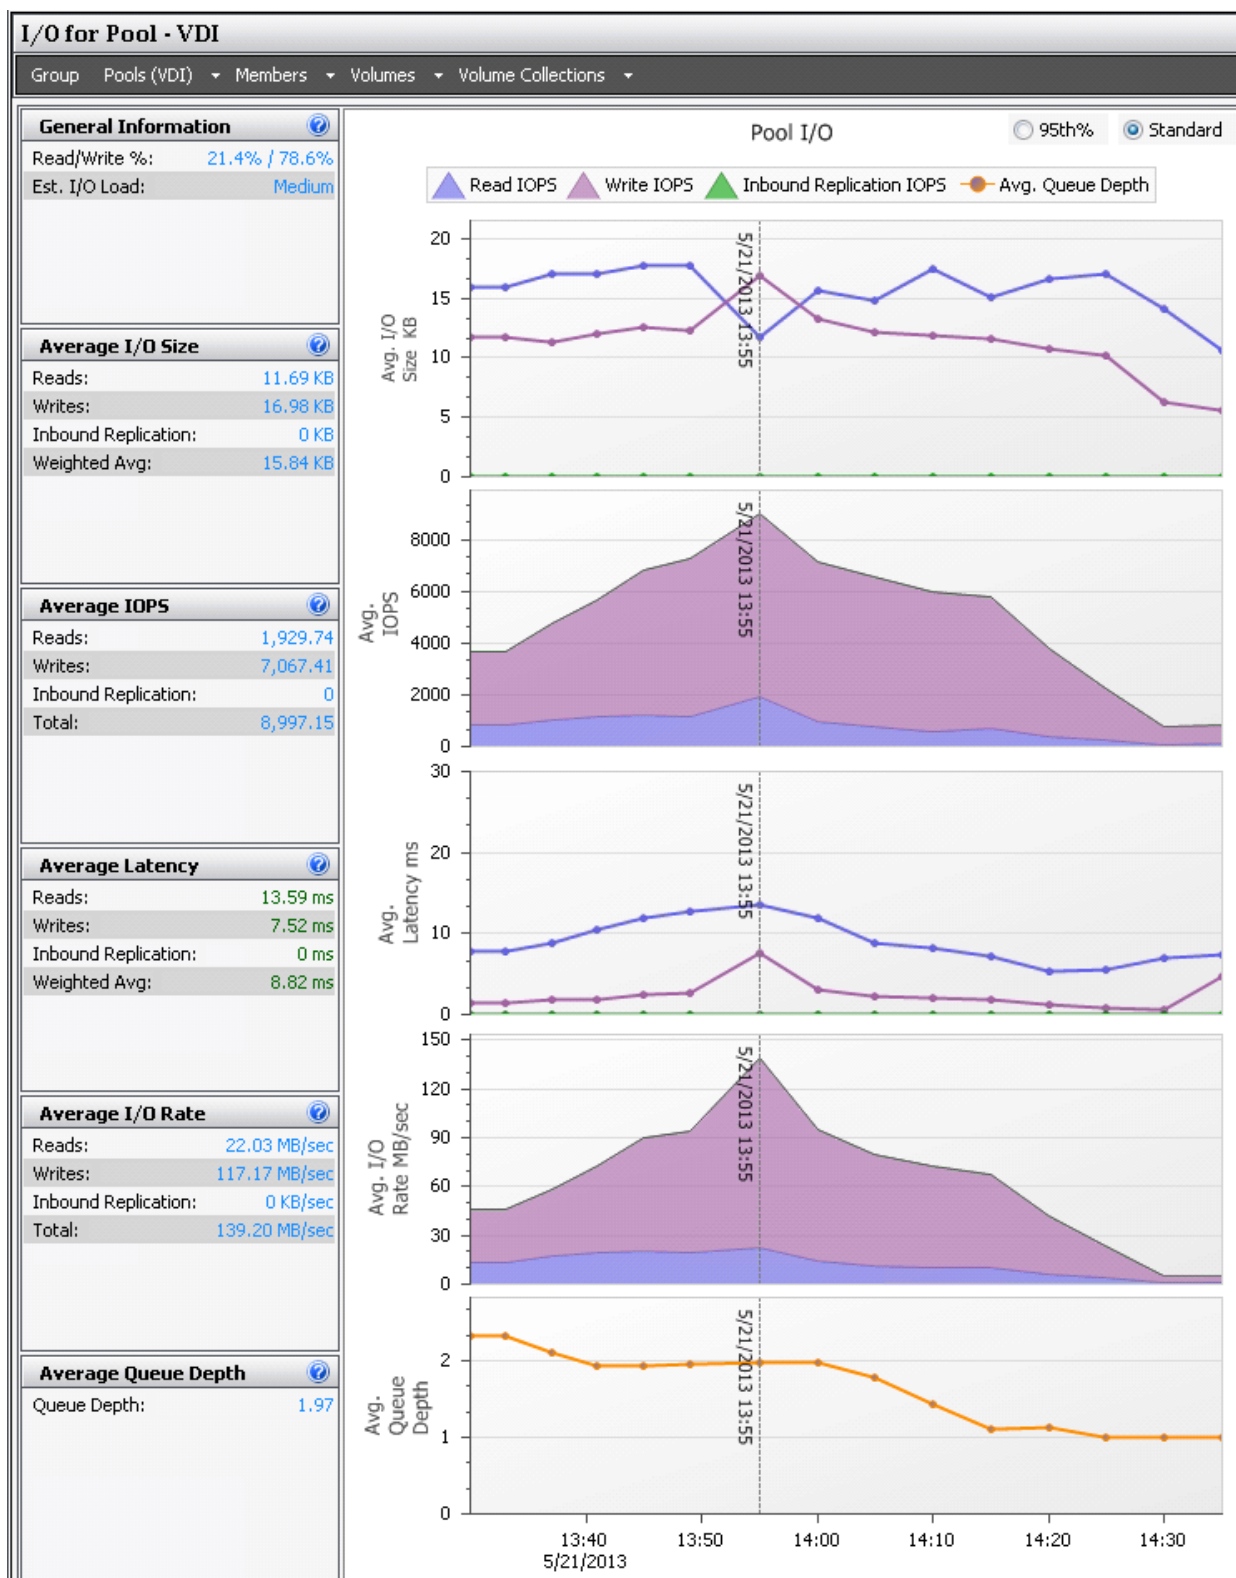
SAN HQ data showing boot storm performance with two PS-M4110XS arrays

With 1000 desktops, the boot storm generated over 19,000 total IOPS with a majority of them being read operations. All desktops were available in less than 15 minutes. This follows the pattern as described in Section 0. The read-write ratio for the pool was 75% reads to 25% writes.

The replica volumes and the volumes hosting the virtual desktops showed similar behavior as in the single array test. The network utilization on the SAN with two PS-M4110XS hybrid arrays reached a maximum of about 9.1% when the 1000 desktops were reset.

### 12.5.7.3.2 Login storm I/O

The tests were repeated after adding a second PS-M4110XS hybrid array to the pool. The results from the login storm with 1000 desktops can be seen in the figure below.



SAN HQ data showing login storm I/O with two PS-M4110XS arrays

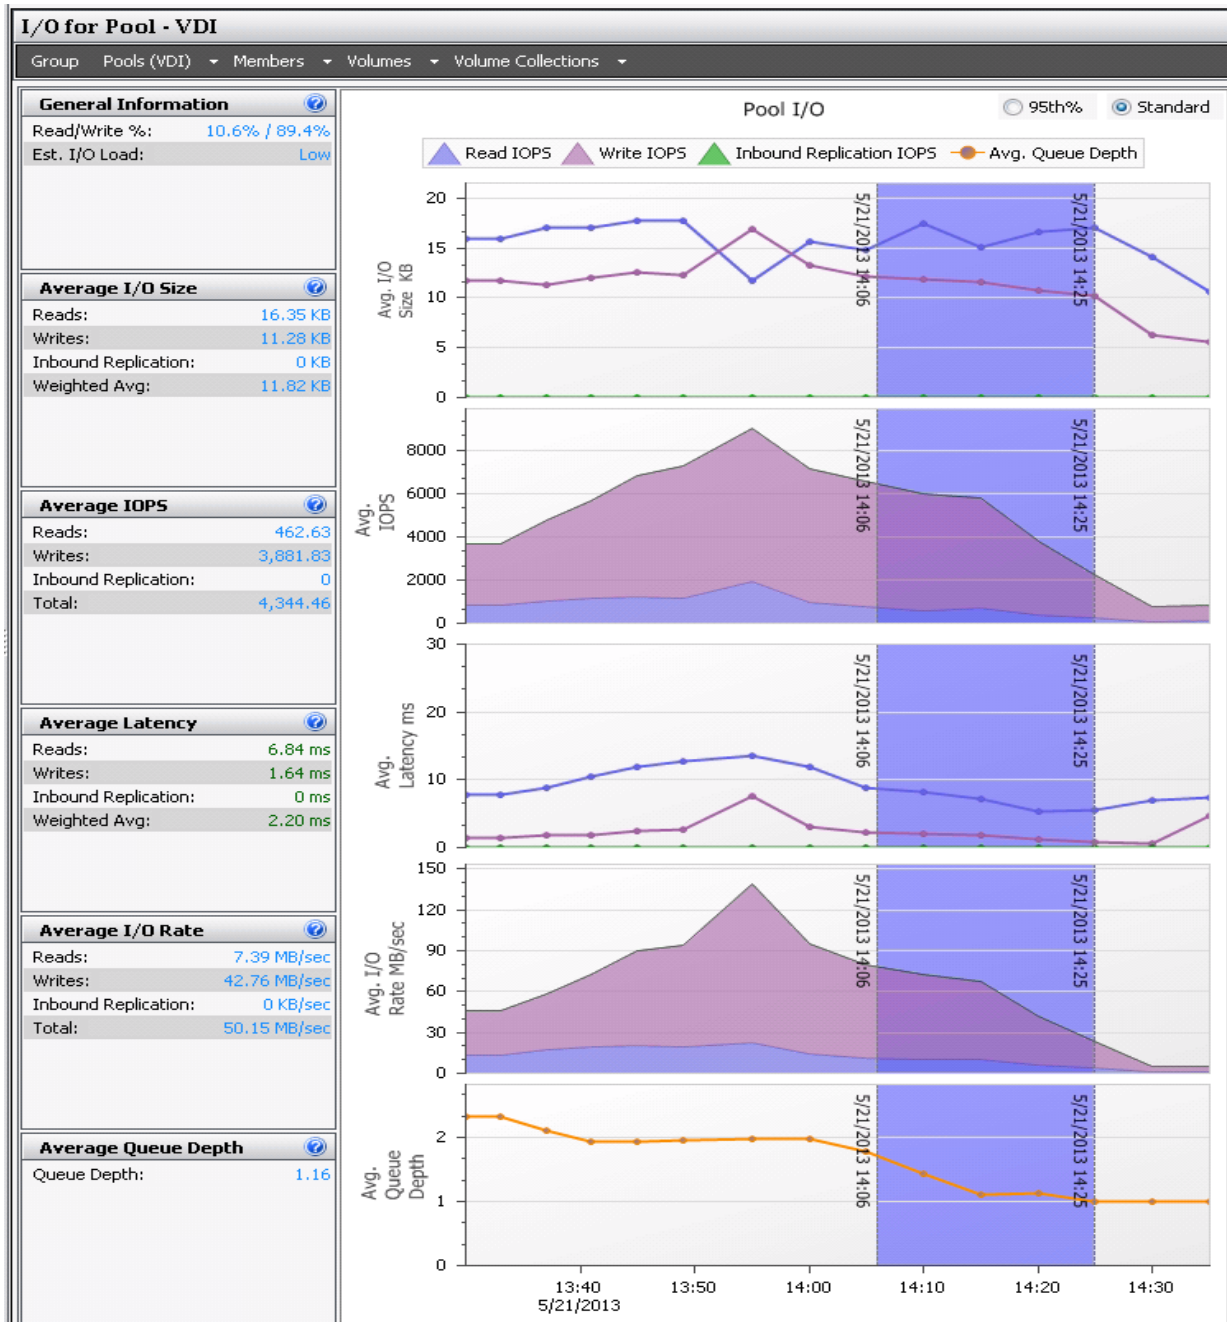
As seen in the 500 desktops case, the maximum latency seen in this test was also well below the 20 ms criteria and the two arrays were able to sustain about 9000 total IOPS for the login storm. This shows that the arrays are able to support the 1000 desktops with no issues.

### 12.5.7.3.3 Steady state I/O

Just like the boot storm and the login storm, the test was repeated after adding another PS-M4110XS hybrid array to the pool. During the steady state for 1000 desktops, the array was able to sustain about 4,400 IOPS (6-8 IOPS per VM). The read - write ratio changed marginally to 11% reads (460 IOPS) and 89% writes (3,900 IOPS). The latency was well below 20 ms indicating that

the two arrays were able to sustain the 1000 desktops without any issues.

The figure below shows the performance of the two array pool during a steady state of 1000 desktops.




SAN HQ data showing steady state I/O for two PS-M4110XS arrays


### 12.5.7.4 Server host performance

During the login storm and steady state of the test, the ESXi host CPU, memory, network, and storage performance was measured on all the servers hosting the virtual desktops. The performance of one such ESXi server is presented here. The other ESXi servers had similar performance characteristics.

Statistics for the ESXi hosts were captured using VMware vCenter server. The figures below show the CPU, memory, and network utilization for boot storm, login storm and steady state of one of the ESXi servers hosting the virtual desktops. The results shown below are for a test run with 1000 desktops in the desktop pool.




Average CPU Performance per core on one ESXi host during login storm and steady state




Memory usage during login storm and steady state





Overall Network performance during login storm and steady state



Storage adapter performance during login storm and steady state

The key observations from the statistics were:

- CPU utilization was well below the 85% threshold throughout the test.
- Active memory usage was about 80% during the boost storm and about 60% during login storm and steady state. There was minimal or no memory ballooning observed.
- Network utilization was about 30% which included all the networks including iSCSI SAN, VDI LAN, Management LAN and vMotion LAN.
- Average read and write latencies at the storage adapter level were very close to the observed latencies in SAN HQ.

### 12.5.7.5 User experience monitoring



The Liquidware Stratusphere UX scatter plot for user experience shows virtually all the users are in the Good category. This shows that the EqualLogic PS-M4110XS arrays are capable of easily

providing adequate user experience for all the users in the VDI environment.

Note that Stratusphere UX registered 961 of the total 1000 desktops for the scatter plot in figure below, and the performance of the remaining desktops was also verified in other reports to be in the Good category.

## Average VDI UX - Users

VDI UX, All Users	Count
Good	961
Fair	1
Poor	0



Average VDI user experience for all users

### 12.5.7.6 Results summary

The key observations from the test results are listed below.

- A single EqualLogic PS-M4110XS was able to host 500 virtual desktops and support a standard user type of I/O activity.
- The VDI I/O was mostly write-intensive I/O with more than 85% writes and less than 15% reads.
- None of the system resources on the ESXi servers hosting the virtual desktops reached maximum utilization levels at any time.
- During the boot storm simulation, nearly 9,100 IOPS were observed and all the 500 desktops were available in Horizon View within 20 minutes of the storm.
- With 1000 desktops, the arrays observed nearly 19,000 IOPS and all desktops were available within 20 minutes of the storm.
- A single PS-M4110XS array was able to sustain a login storm of 500 desktops that lasted around 15 minutes and had most of the I/O served by the SSDs on the array.
- Two PS-M4110XS storage arrays were able to handle an aggressive login storm for 1000 desktop users.
- The user experience for 500 and 1000 desktops was well within acceptable limits and virtually all the desktops were in the good category on the Liquidware Stratusphere UX scatter plot.
- Adding a second EqualLogic PS-M4110XS array to the solution linearly scaled the number of desktops supported by the solution to 1000 virtual desktops.
- A PowerEdge M1000e blade chassis, fully populated with 12 PowerEdge M620 blade servers and two EqualLogic PS-M4110XS hybrid blade storage arrays was able to provide a self-contained 1000 desktop VDI solution.

## 12.5.8 Sizing guidelines for EqualLogic SANs

The storage array selected for the solution should be able to handle various I/O patterns that occur throughout the day for a VDI solution. These include the login storm at the beginning of a shift or a work day when employees login to their virtual desktops in a relatively short period of time. Once they are logged in, the virtual desktops reach a steady state where they generate predictable IOPS as the employees go about their work day. The same storage array needs to handle recovery situations due to unexpected events such as power outages which might cause boot and login storms.

A good way to deploy the storage infrastructure for VDI is to understand the VDI workload characteristics including the performance characteristics of the applications being deployed. These include:

- Capacity requirements
- Performance requirements
  - IOPS
  - Average disk latency
  - Read and write ratios
  - Type of I/O (Sequential or Random)

For more information on sizing guidelines, see the paper "Sizing and Best Practices for Deploying VMware View 5.1 on VMware vSphere 5.0 U1 with Dell EqualLogic Storage" at:

<http://en.community.dell.com/dell-groups/dtcmedia/m/mediagallery/20219029/download.aspx>.

## 12.5.9 Best Practices

This section provides the best practices derived from the testing and analysis.

### 12.5.9.1 Application layer

This sub-section provides best practices for the implementation of VDI using VMware Horizon View.

#### 12.5.9.1.1 Implement roaming profiles and folder redirection

It is highly recommended that all users in the VDI environment be configured with roaming profiles and folder redirection. This preserves user profiles and user data across boots while using non-persistent virtual desktops.

It is also recommended to use a high performance file service to provide the profile and folder redirection. A separate array may be used to host these volumes for best performance.

#### 12.5.9.1.2 Boot and login storm considerations

To avoid I/O bursts due to boot storms, it is recommended that all desktops be pre-booted, preferably with the boots staggered over time before users begin login at the start of a shift or a workday.

It is important to size storage based on the IOPS needs of boot and login storms. The storage subsystem should be designed to handle these storms in addition to the steady state IOPS. Sizing a storage system only on the steady state IOPS is not recommended because this can cause degraded user experience and performance during a boot or login storm.

#### 12.5.9.1.3 Windows 7 master image for desktop VMs

It is recommended that the operating system be customized in order to provide the best

performance in a VDI environment. This includes disabling some services which may not be required. This can improve performance for the end user. VMware has a specific set of recommendations and settings for Windows 7 that allow for faster logins, quicker screen refreshes, and generally better performance.

The VMware recommendations for Windows 7 image customization can be found here: <http://www.vmware.com/resources/techresources/10157>

#### 12.5.9.1.4 SSL certificate requirements

It is recommended that all servers use either signed SSL certificates or use self-signed SSL certificates to provide the best security for your infrastructure. Horizon View servers and Horizon View Composer servers require valid SSL certificates to communicate with each other and to operate correctly.

#### 12.5.9.1.5 Horizon View recommendations

Depending on the actual applications and the actual usage of the virtual desktop, it is recommended to adjust the Adobe Flash settings for the remote sessions.

More information on Adobe Flash quality and throttling settings can be found here: <http://pubs.vmware.com/view-52/topic/com.vmware.view.administration.doc/GUID-8CE3908A-95B2-4D4D-9E00-B924E0D6D400.html>

Enabling Horizon View Storage Accelerator is recommended to get the best performance from the storage layer and the best user experience on the virtual desktops. Additionally, use blackout times to ensure that the cache is not regenerated when the users are active on the system because this process puts additional load on the ESXi host and the storage array causing a performance drop that could affect the user experience.

### 12.5.9.2 Server host layer

The ESXi servers hosting the infrastructure service providers and the virtual desktops are recommended to be configured as follows:

- Follow VMware and Dell best practices for installing and configuring ESXi
- Install and configure EqualLogic Multipathing Extension Module (MEM) for vSphere 5.1 to get the best performance from the storage array
- Separate virtual switches to segregate iSCSI SAN traffic, VDI traffic, vMotion traffic, and Management network traffic
- Each network path should be assigned to a minimum of two physical NICs for high availability

VMware KB article on best practices for installing ESXi 5.1: <http://kb.vmware.com/kb/2032756>

Installing and configuring the Dell EqualLogic MEM for VMware vSphere 5 (login required): <http://en.community.dell.com/dell-groups/dtcmmedia/m/mediagallery/19991633.aspx>

### 12.5.9.3 Network layer

It is recommended that at least two physical NICs on each server be dedicated to each of the following logical networks:

- Management network for Infrastructure service, vMotion services, and VDI LAN
- iSCSI SAN

This allows the solution to work even in the event of the failure of a single blade switch, rack switch, or individual NIC port on the storage array or on the ESXi host server.

Use VLANs to segregate different types of network traffic on the same physical network. In this case, it is recommended to separate the infrastructure, vMotion, and VDI LAN traffic into separate VLANs.

Do not use VLANs to segregate iSCSI SAN traffic. It is required that SAN traffic be on a separate physical network to provide the best performance. The recommended way to provide a converged fabric for iSCSI SAN traffic is through the use of Data Center Bridging (DCB) technologies which allows the iSCSI SAN to have lossless end to end connections while guaranteeing a minimum bandwidth.

Virtual switches in ESXi have a default limit of 120 ports. If the number of virtual desktops on each host exceeds the available ports, vSwitch properties should be changed to support the required number of virtual desktops. This change requires a reboot of the host ESXi server.

On iSCSI SAN switches, Spanning tree should be disabled on switch ports connected to end devices for server and storage ports. The Portfast setting should be enabled in the switch configuration for these ports. Jumbo frames and Flow Control (if the NICs support it) should be enabled for all components of the iSCSI network.

More information on configuring Dell Networking switches for use with EqualLogic iSCSI SANs is available here:

<http://en.community.dell.com/techcenter/storage/w/wiki/4250.switch-configuration-guides-by-sis.aspx>

#### 12.5.9.4 Storage

It is recommended to use the EqualLogic PS Series hybrid arrays, which consist of SSD drives and 10,000 rpm SAS drives within a single chassis, for VDI environments. These hybrid arrays automatically move hot data to the SSD tier, which improves performance in VDI environments in a cost-effective way.

It is recommended to have separate volumes for base images, replica images, and virtual desktops. This aids in better manageability of the volumes and easier performance monitoring as well as allows for easy future growth.

It is recommended to use a separate high performance file service to provide file shares for roaming profiles and user shares.

#### 12.5.10 Conclusions

The section demonstrates how a modular, 1000 standard user virtual desktop environment – all self-contained within the blade chassis – can be deployed using Horizon View VDI platform leveraging 12 PowerEdge M620 blade servers, four Force10 MXL blade switches and two EqualLogic PS-M4110XS hybrid blade arrays. The storage I/O characteristics under various VDI workload scenarios (boot storm, login storm and steady state) along with performance characteristics throughout the VDI stack (ESXi server performance as well as user experience as determined by Liquidware Stratusphere UX) demonstrate the optimal configuration of all the elements across this VDI building block.

Additionally, this testing of the EqualLogic storage platform showed that the EqualLogic PS-M4110XS arrays can easily support very fast access to high-demand data in a VDI environment. It can easily handle high IOPS spikes that can occur during boot and login storms.

The ability to support 500 desktops on one array and 1000 standard user desktops with two arrays means that the EqualLogic PS-M4110XS can help organizations that are limited by the cost-per-desktop hurdle start their VDI deployment plans much sooner. Moreover, the peer storage architecture of the EqualLogic arrays achieves linear scalability of both controller throughput and storage capacity as new arrays are added with the VDI environment growth. This linear scaling in

both performance and capacity keeps the storage-per-virtual-desktop cost low even in very large VDI deployments.

In addition to improved performance and flexible scalability, the PS-M4110XS hybrid arrays help reduce VDI deployment costs by handling workloads on highly cost-effective storage media, enhanced VDI performance by placing read-sensitive replica images on low-latency SSDs, and easier VDI deployment through automated and intelligent on-board tiering software. Combined with the ease of administration of EqualLogic arrays, these benefits reduce the TCO for VDI in many organizations.

For full documentation and information on this topic please go to document titled "VMware Horizon View 5.2 on Dell EqualLogic PS-M4110XS Hybrid Blade Storage Arrays" [here](#).

## 12.6 PowerEdge VRTX with VMware Horizon View



### 12.6.1 Overview

This section contains performance and characterization that was performed on a PowerEdge VRTX platform. As seen below, the PowerEdge VRTX platform shows exceptional performance with VMware Horizon View and M620 blades to power a large branch office deployment or other similar sizing use cases. This platform is very capable of running a very dense quantity of Horizon View desktops in a smaller footprint and with an all in one chassis approach, shared storage is also provided.

Note: Drivers noted below are available at <http://www.dell.com/support/troubleshooting/us/en/19/Product/poweredge-vrtx>

#### 12.6.1.1 Test objectives

The primary objectives of the testing are:

- Determine the maximum user density for a fully loaded VRTX chassis of M620 Servers.
- Optimize shared storage drive configuration to support max users above at a lowered cost per seat basis
  - Verify Hybrid drive performance and value prop characteristics.
- Determine the user density for a minimum VRTX configuration

The following tests will be delivered:

- 1 x VMware View Density Pipeclean (designed to get testing at a normal state and negate any caching type effects)
- 3 x VMware View 5.2 user recordable sets.

Results obtained from the testing including IOPs, CPU utilization, latency, Memory, Network, Stratusphere UX reports and subjective user testing will be used to determine the capacity of the system.

#### 12.6.2 Test Environment

The VRTX chassis was racked and connected to the DVS Core test network in our testing datacenters. Network connectivity between the PowerEdge VRTX and the network was delivered

via a PowerConnect 6248 ToR configured in line with best practices for VDI deployments.

The standard DVS PAAC test process was applied to the validation which uses Login VSI as a test tool.

### 12.6.2.1 DVS Stack Infrastructure

#### 12.6.2.1.1 Compute VRTX Nodes for max user density

The compute host system configurations are outlined below:

Compute Host Server Configuration	
PowerEdge VRTX, 4x M620	
2x Intel Xeon E5-2680 (2.7GHz)	
192GB Memory (8 x 16GB DIMMS @ 1600MHz plus 8x 8GB DIMMs @ 1600MHz)	
SD card (ESXi configurations only); 2x HDD for Hyper-V	
Broadcom 57810-k 10Gb DP KR NDC (iSCSI)	
Table 1 BRCM 5720 1Gb DP NIC (LAN)	
iDRAC7 Enterprise w/ vFlash, 8GB SD	

#### 12.6.2.1.2 Management Host

The above mentioned Compute Hosts will also supply compute resources to the Management VMs that support the VMware Horizon View infrastructure such as VMware vCenter, VMware View Connection Server, SQL Server, and a file server. These Management VMs will not be prescribed to a dedicated server within the VRTX chassis but instead will be setup and allowed to be "floating" amongst the cluster and utilizing DRS, HA, and a resource pool for management VMs to establish priorities.

#### 12.6.2.1.3 Additional Core Hardware

Hardware Components	
Component	Description
PowerConnect 6248	1GB 48 Port Switches

#### 12.6.2.1.4 Core Software

Software Components (VMware View PAAC Configuration)	
Component	Description
VMware ESXi 5.1	Hypervisor
VMware Horizon View 5.2	VDI Software
Microsoft SQL 2008 R2	Database

### 12.6.2.2 Test Infrastructure

In addition to the core infrastructure for the DVS stack, a corresponding test infrastructure must exist to host the core network services as well as the Login VSI and Stratusphere UX environments. These are used as part of the testing functions and are isolated so as to not impart load on the infrastructure/systems under test.

## 12.6.3 Test Set-up

### 12.6.3.1 Shared Storage Raid Configuration for max user density

#### “Tier 1” Linked Clone and replica Disk

- RAID 10
- 20x 300GB SAS 15K Disks (*fully populated it would use a max of 25 x 2.5in drives but did not have that in place for this test*)

#### Local “Tier 2”

- Tier 2 will be initially placed on one blade configured as a floating assignment between all 4 hosts.
- RAID 5
- 5 x 146GB SAS 15K Disks

### 12.6.3.2 Network Configuration

- The VRTX Chassis was connected to the Test network via F10 ToR switches.
- 1GB Connectivity was used for testing
- DVS Networking best practices were applied to ToR configuration
- Please note that only GB Pass-through and Gb Switch IOMs are available at RTS. 10Gb will be available post-RTS

## 12.6.4 Test Methodology

### 12.6.4.1 Test approach

The testing process was broken into two pieces using the VSI Basic User configuration that includes a “pipe-clean” operation to clear caching behaviors in testing. 3 x VMware View Basic User tests were carried out to determine the max user density of a full chassis and ensure a consistent metrics gathering exercise.

Workload type is listed below.

Workload Type	Desktop Quantity VMware
Basic	500

### 12.6.4.2 Test Proposal

#### 12.6.4.2.1 Load Generation

Login VSI is a VDI benchmarking tool which can be used to determine the maximum number desktops that can be run on a physical or virtual server. Login VSI simulates a realistic VDI workload using the Auto IT script within each desktop session to automate the execution of generic applications like Microsoft Office 2007, Internet Explorer, Acrobat Reader, Notepad, and other software.

For example, the characteristics of “Standard” workload are as follows:

- Applications used include Microsoft Internet Explorer, Adobe Acrobat Reader, Adobe flash player, Adobe shockwave player, Freemind, Kid-Keylock and Bullzip PDF printer
- Only two applications are open simultaneously.
- Total time is about 1:45 minutes

#### 12.6.4.2.2 Liquidware Labs Stratusphere UX

The complex setup and different configuration layers involved in VDI makes it very challenging to



measure the user experience. VDI performance and functionality depends on various factors such as resource requirements like CPU, memory, network, and storage. The other key consideration is the user experience with respect to time spent while trying to login and response times of applications launched within the desktop.

Stratusphere UX by Liquidware Labs is designed for desktop administrators and support personnel to measure and analyze the end user experience and application response times.

#### **12.6.4.2.3 Monitoring tools**

The system resource utilization and performance metrics at the VMware infrastructure layer and ESXi hypervisor layer were monitored using VMware vCenter.

Detailed performance metrics were captured from the hypervisor, virtual desktop, storage array, and Stratusphere UX and load generator.

#### **12.6.4.4 Test criteria**

The criteria listed below outline what constitutes a passing test run. This is outlined for reference during the full PAAC test runs and is not applicable to the cursory runs.

##### **12.6.4.4.1 Storage capacity and I/O latency**

The following performance criteria are used for the storage layer:

- Maintain less than 20 ms disk latency
- Maintain at least 15% spare capacity on the storage array
- Ensure CPU does not exceed 90%

The typical industry standard latency limit for storage disk I/O is approximately 20 ms. Maintaining this limit will ensure good user application response times when there are no other bottlenecks at the infrastructure layer.

##### **12.6.4.4.2 System resource utilization on the hypervisor infrastructure**

The following metrics at the hypervisor infrastructure layer are used to ensure solution consistency:

- CPU utilization on any ESXi servers not reaching 90% at any point in time
- No memory ballooning on ESXi servers
- Total network bandwidth utilization not to exceed 90%

##### **12.6.4.4.3 Virtual desktop user experience**


VDI configurations involve many components at different layers – application, hypervisor, server, storage and network. As a result, multiple metrics need to be monitored at different layers to ensure that the environment is healthy and performing appropriately for users. Liquidware Labs Stratusphere UX is to be used to assess the performance of each desktop (application level).

Please use the following user experience criteria:

- At least 90% of the virtual desktops are in the 'Good' performance category.
- No virtual desktop should be in the 'Poor' performance category.


## 12.6.4.5 Test Results

### 12.6.4.5.1 Test: VMware 500 Users 100% Pre-booted




Pictured above is the CPU performance graph for a 500 User configuration. All hosts ran 125 compute VMs and management VMs moved around between hosts with their locations determined by DRS. Typically CPU is the mitigating factor in density and this graph shows that we are below the CPU thresholds with a bit more capacity per server available which gives indication that more than 500 users in the Basic workload is most likely attainable.


Memory performance was excellent. Active memory remained very low throughout the test process and shows the servers have plenty of capacity to server memory requirements of the Infrastructure and Horizon View desktop VM needs.



The following graph shows network utilization across all hosts. Two physical NICs were used and the graph shows there was no network congestion during the 500 users test.





The next graph shows the internal shared storage performance and the note that the host registering the highest utilization contained the File Server infrastructure VM.



A total of 5805 IOPS were recorded at peak of steady state with latency not exceeding 20ms.

Finally the Stratusphere UX graph showed excellent machine and user performance for the test.



490 of the VMs called back to the hub with all 500 sessions logging in.

### 12.6.4.5.2 Summary

User experience for the entire test was excellent. This section showed results for 500 users. Latency was within thresholds for 500 desktops. (Note that the test lacked 5 more drives in the chassis that would increase performance). The exercise shows a 4 node 500 user solution is viable on the PowerEdge VRTX platform.

Network utilization using 1Gb networking is more than sufficient as is the use of a pass through module which performed optimally during the testing period.

A single PERC was shared across all 4 nodes. The user experience for each host proves that a shared PERC with 1GB cache can comfortably drive 4 M620 servers.

CPU and regular stats including network and memory were good during the tests with CPU not exceeding 80% on any node. Each node had management VMs running on them in addition to their VDI desktop workload with "blade 10" retaining the file server infrastructure VM at all times.

The VRTX platform has a similar feel to the M1000E blade chassis in terms of management and as such that skillset can be leveraged to support this platform.

## 12.6.5 VRTX Configuration Overview

Note: All firmware files can be downloaded from

<http://www.dell.com/support/troubleshooting/us/en/19/Product/poweredge-vrtx>

### 12.6.5.1 Storage Volume Overview

*Note: Before the storage can be presented you must install the supported Mega Raid controller ESXi VIB.*

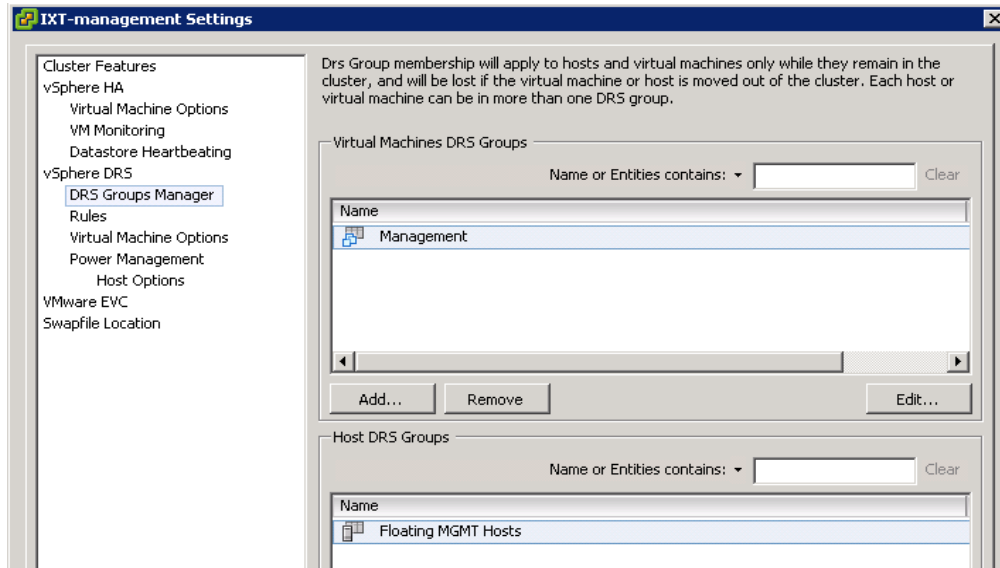
- Management VMs were provisioned in a RAID5 5 Disk volume.
- Compute VMs were provisioned in a single RAID10 volume using 14 Disks for this test (note: another 5 disks would be available in a full disk configuration). The total compute volume size was 2.18TB.

### 12.6.5.2 Compute Host Configuration Overview

Management VMs were deployed as a "floating assignment" in a DRS Pool to allow them to move amongst the cluster. See information below:

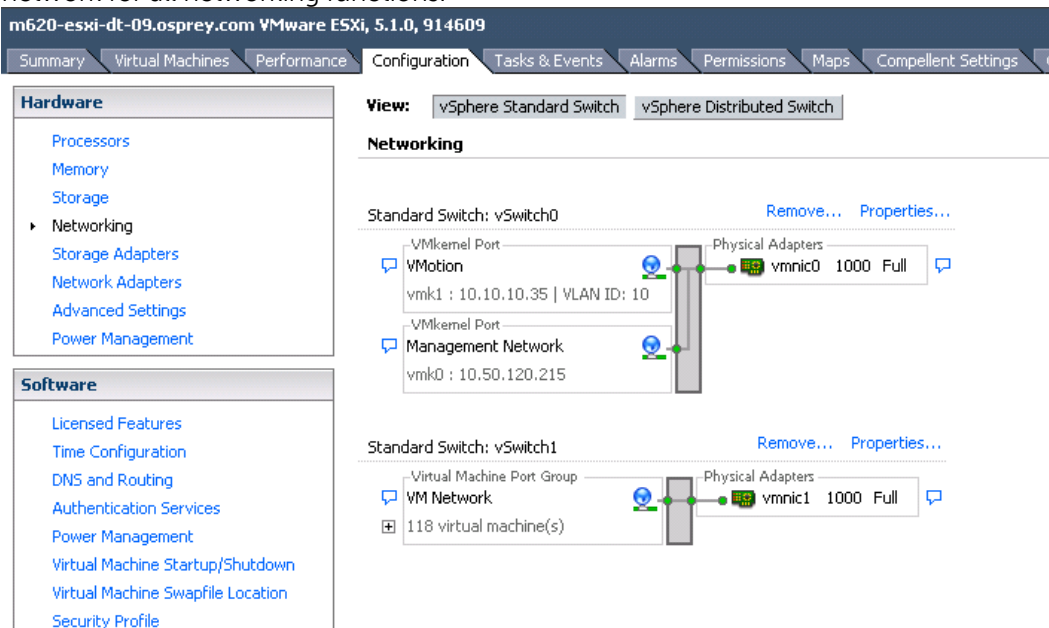
The screenshot displays the vSphere Management console for host IXT-DC-01. The left-hand navigation pane shows a tree structure with 'IXT-DC-01' expanded to 'IXT-management', where four ESXi instances (m620-esxi-dt-09 through -12) and a 'Management' folder are listed. The 'Management' folder contains 'ClientMaster', 'IXT-FILE', 'IXT-SQL', 'IXT-WC', 'IXT-VIEW', and 'VDI'. The main console area shows the 'Management' summary for the host. The 'General' tab is active, displaying statistics for the host's DRS pool: 'Virtual Machines and Templates: 5/5', 'Powered on Virtual Machines: 4/4', 'Child Resource Pools: 0/0', and 'Child vApps: 0/0'. Below this, the 'CPU' section shows 'Host CPU' utilization at 156832 MHz.

Management VMs are able to move between all hosts as per DRS recommendations



### 12.6.5.3 Compute Node Networking Overview

In this test each blade had 2 x Physical network interfaces connected to a PowerConnect 6248 ToR switch via the pass through module installed in the VRTX Chassis. In a production deployment full scale deployment of network interfaces and redundant ToR switches would ensure HA at the network for all networking functions.



## 12.7 Foglight for Virtualization with Horizon View cartridge

### 12.7.1 Introducing the Cartridge for VMware View

The Cartridge for VMware View monitors a virtual desktop environment, providing performance metrics, alarms, and alerts to help you better manage your environment, from user desktops to related infrastructure.

VMware View gathers extensive performance metrics and presents that data in a graphical interface, utilizing architectural diagrams, graphs, and drilldown screens to quickly identify virtual desktop problems. This allows users to browse through their virtual infrastructure to identify desktop performance, connection, and infrastructure issues.

It fully supports VMware, vCenter, and vSphere environments.

### 12.7.2 Cartridge for VMware View Elements

The Cartridge for VMware View provides monitoring capabilities so that all elements of a virtual desktop environment are considered. A typical VMware View environment contains one or more user sessions and related infrastructures

- View Infrastructure—the VMware View infrastructure is composed of one or more connection servers, transfer servers, and security servers.
- User Session—shows the connected users, their session details, and the session type (for example, VDI user session, Terminal Server, or physical desktop). Each user will have one or more sessions.
- vSphere—this underlying infrastructure is broken down by cluster, datastore, ESX hosts, and resource pools.
- Desktops—shows the resource utilization values for a selected VM, Terminal Server Host, or Physical Host.
- Horizon Pools—shows the health of pools in the View Instance.

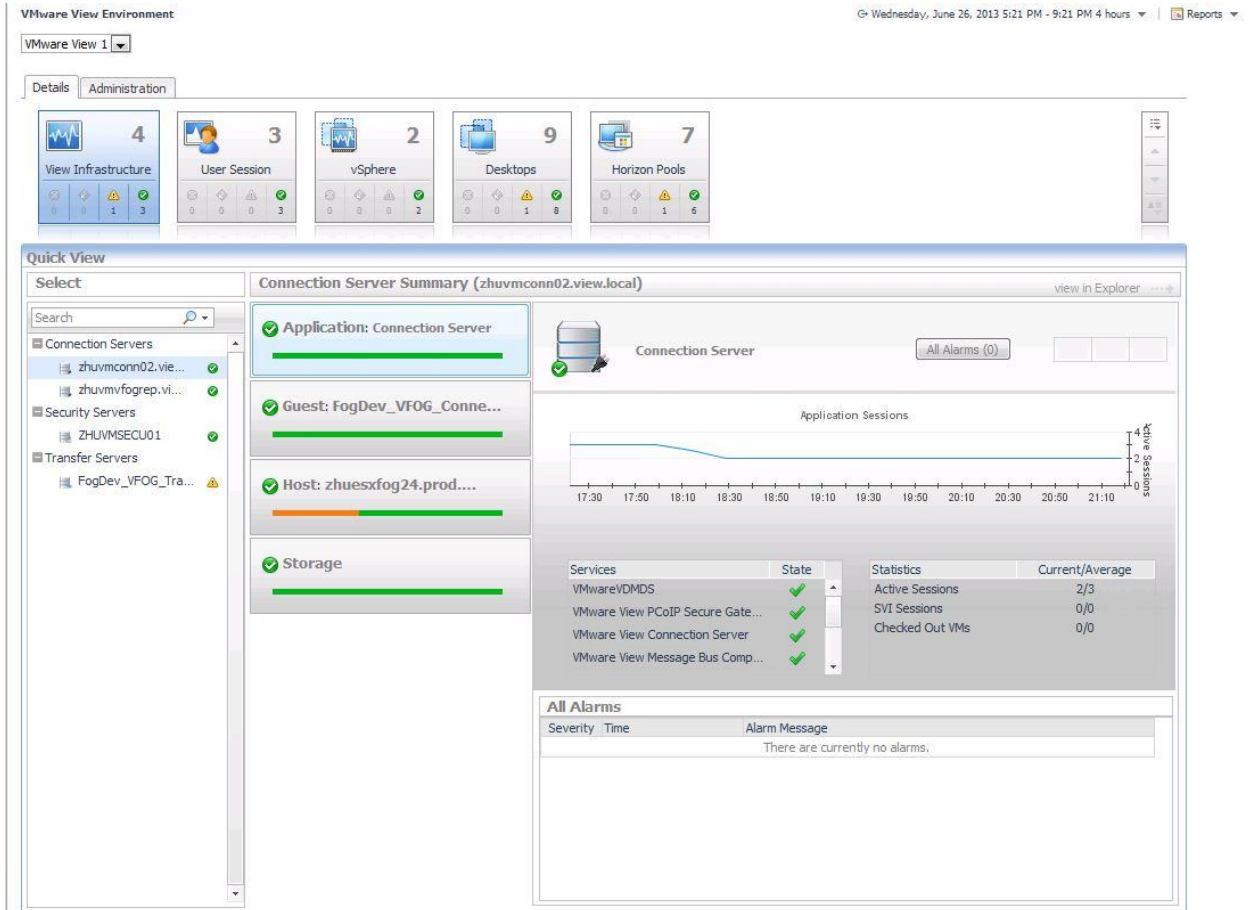
You can view the overall state of all these components on the VMware View Environment dashboard.

### 12.7.3 Using the VMware View Environment Dashboard

A typical VMware View environment contains a combination of physical and virtual components. A physical component can be a Connection Server, Terminal server, or a user desktop. You can view the overall state of all VMware View components on the VMware View Environment dashboard. To access the VMware View Environment dashboard:

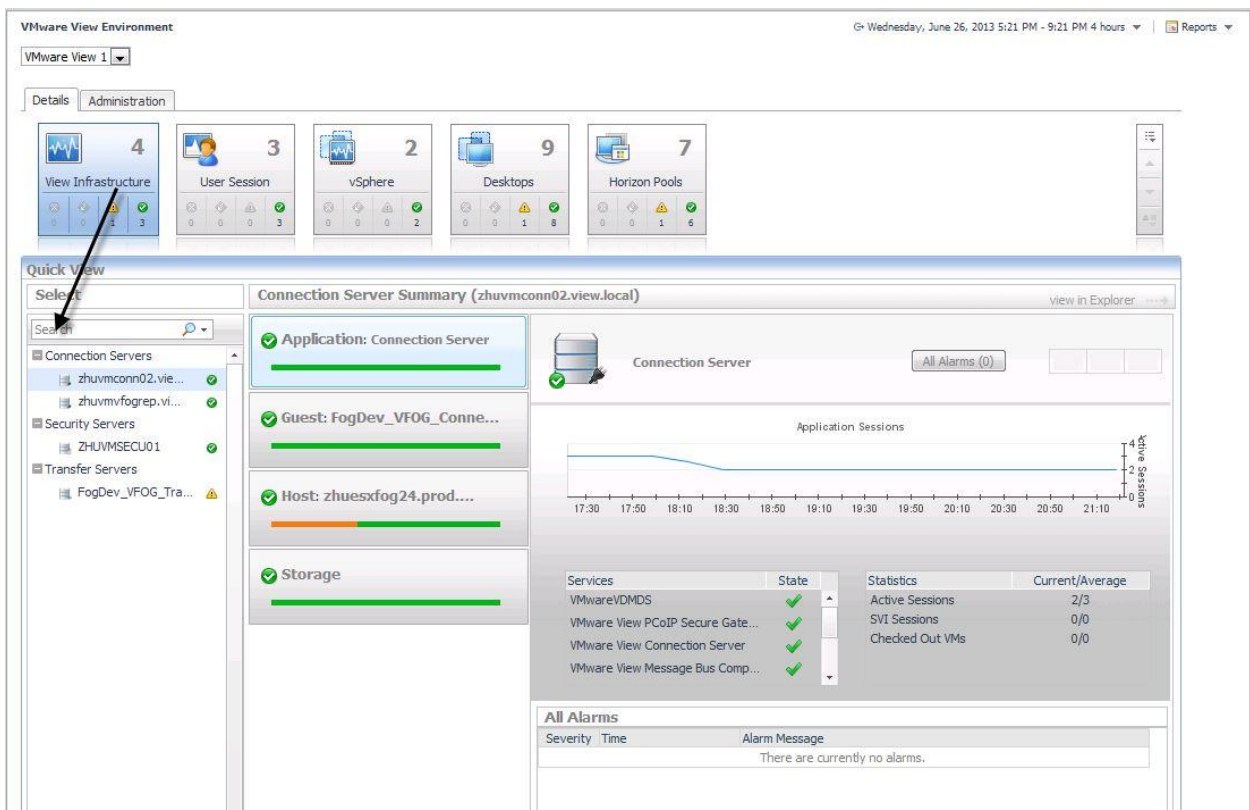
- Log in to the Foglight browser interface
- Ensure that the navigation panel is open. To open the navigation panel, click the right-facing arrow on the left.
- On the navigation panel, under Dashboards, choose VMware View

The Details and Administration tabs are available for selection.



## 12.7.4 Exploring the VMware View Environment Details Tab

The VMware View Environment Details tab contains tiles that summarize status. Selecting a tile changes the content displayed in the Quick View area. This content varies depending on the tile that you select. For example, selecting the View Infrastructure tile displays the connection server, terminal server, and security server details associated with a View Server instance





## 12.7.5 Working with the Tiles

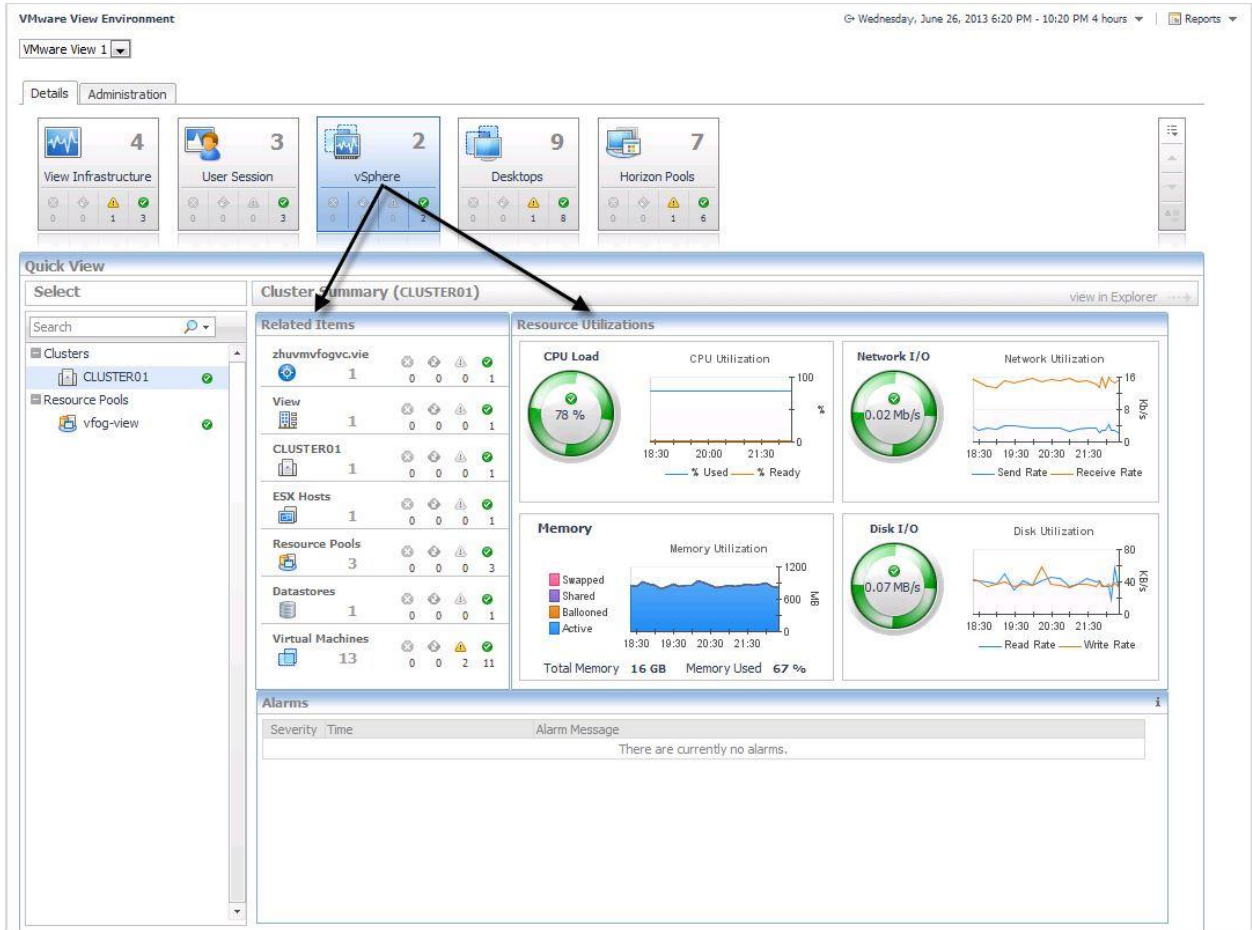
The upper part of the tile displays the desktop management component and a total count of these entries in the environment



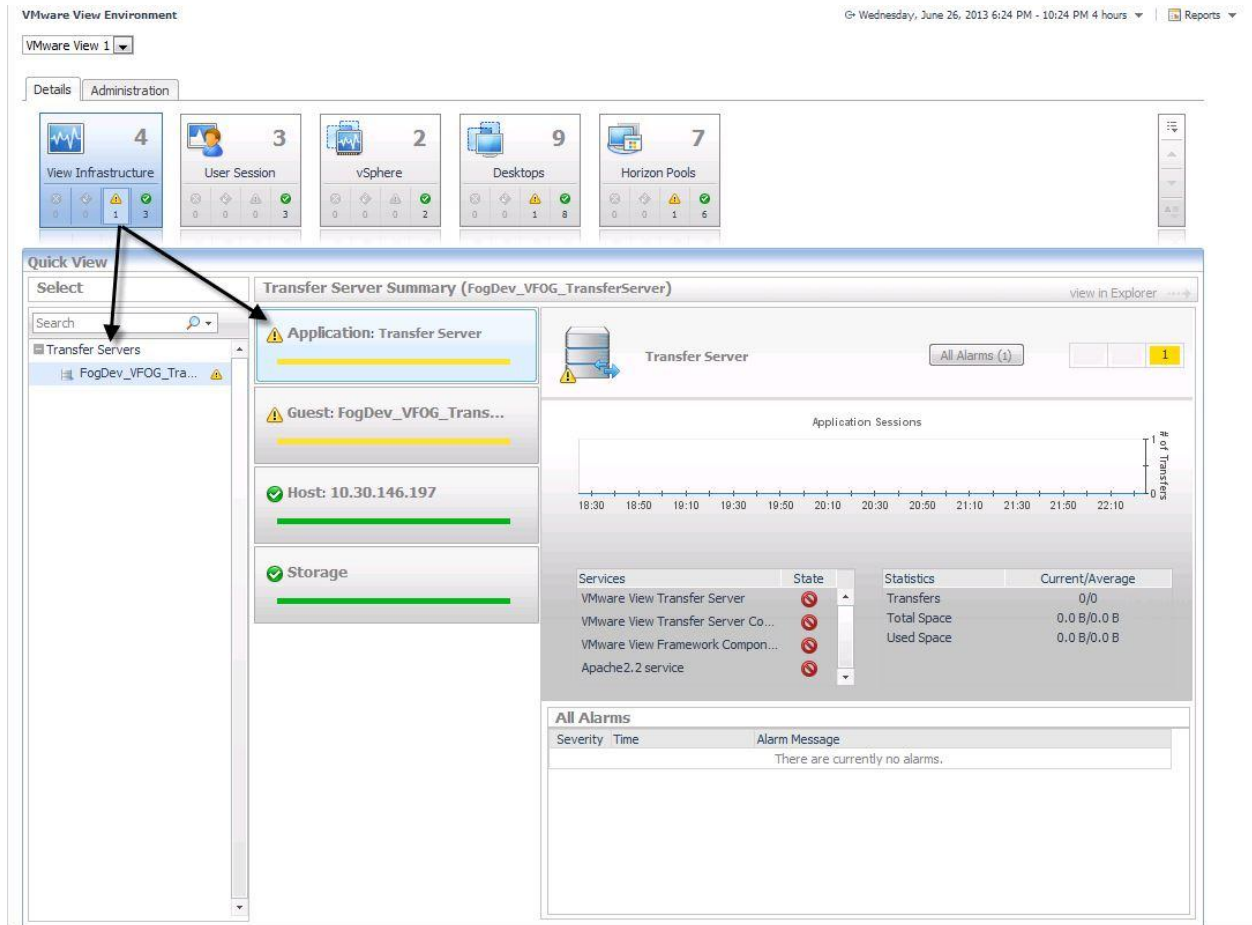
The lower part of the tile displays the count of entities at each severity level, based on the alarms currently active for those entries.



Clicking the label in the tile, for example, vSphere, displays summary and alarm information for all components of that type in the Quick View area.



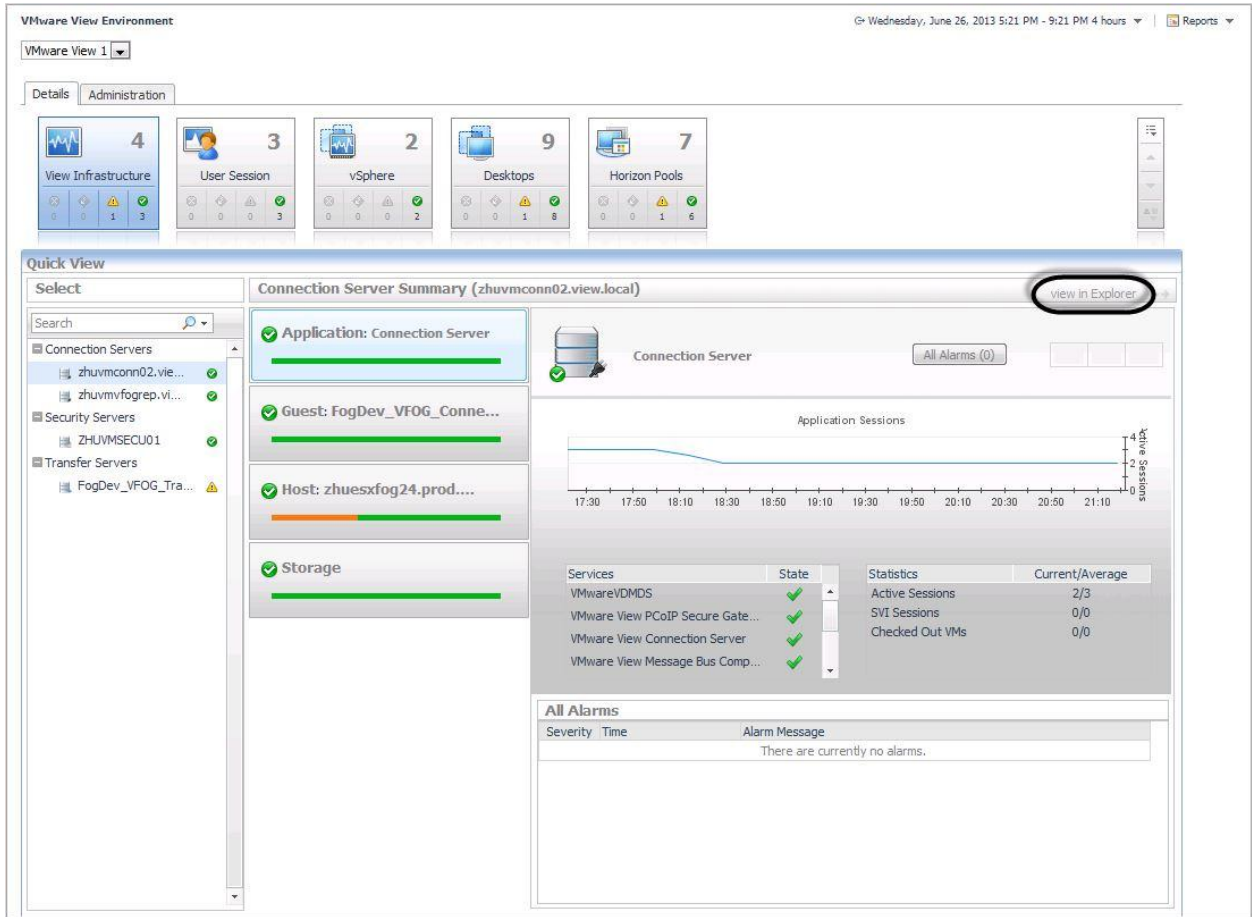
Clicking the alarm count in the lower part of the tile displays summary and alarm information for only the components with that status.



### 12.7.6 Using the Quick View

Selecting a specific object in the navigation panel displays key summary information about child component health, resource utilization, and alarms.

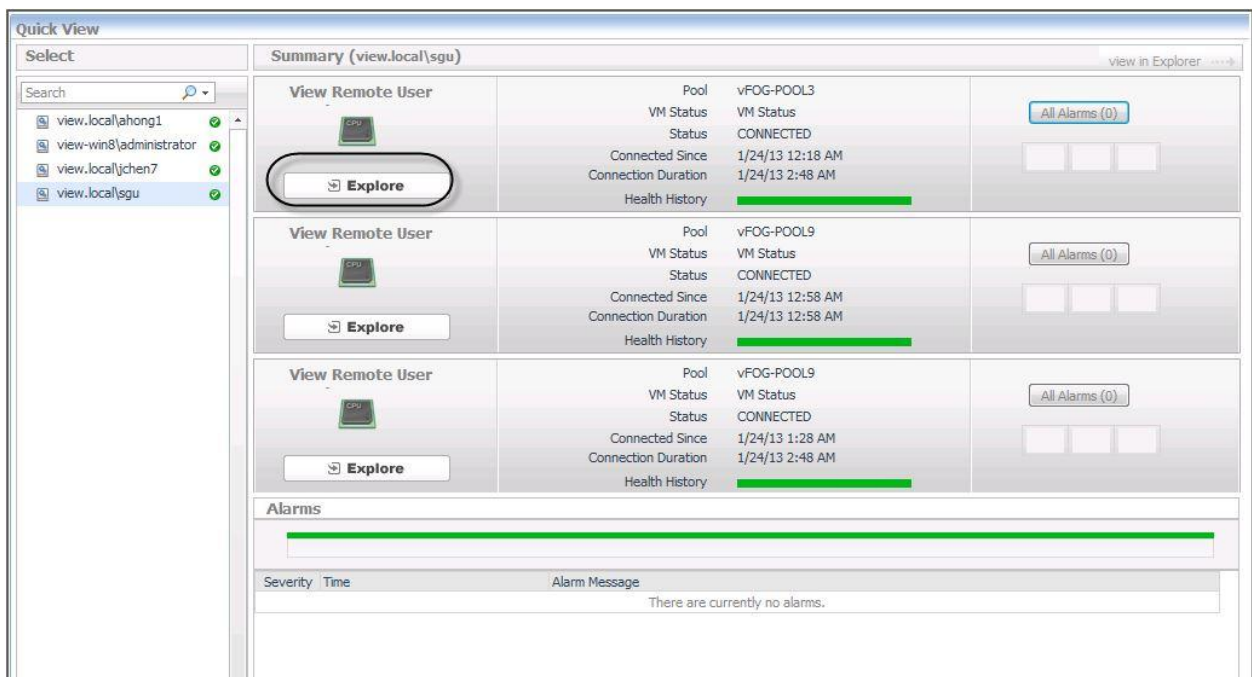
Click the view in Explorer link in the upper-right corner to drill down to the VMware Explorer view of the object, to see detailed information about the selected object and its components. If the selected object is a monitoring VM, the VMware Explorer page appears. If the selected object is a host, the Host Monitor page appears.



## 12.7.7 Exploring User Sessions

When you click the User Session tile, a list of the connected users is shown in the Quick View. Selecting a user from the Quick View displays the session types for the user, and a summary of the session details depending on the session type. Each user will have one or more sessions in a View instance.

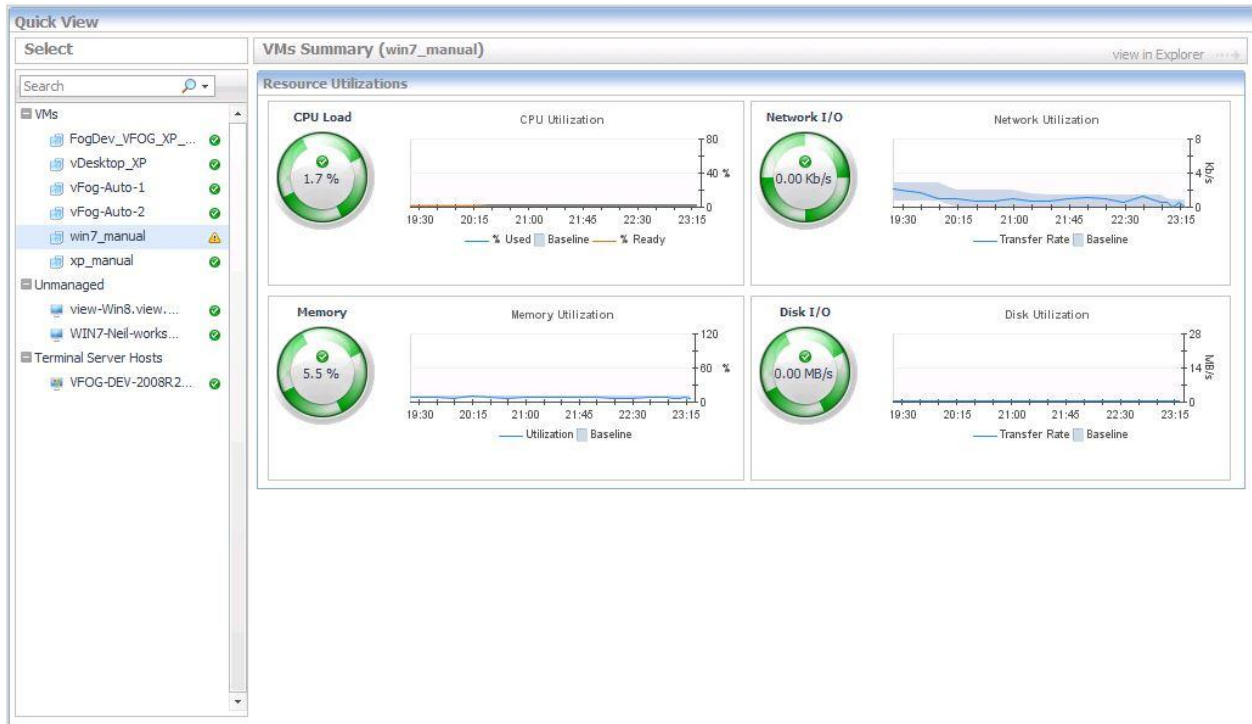
An Explore link is located below the icon of the session type. Click the link to view detailed information about the session type.



## 12.7.8 Exploring Desktops

When you click the Desktops tile, a list of Virtual Machines, Terminal Server Hosts, and Physical Hosts is shown in the Quick View.

Selecting a Virtual Machine from the Quick View list displays the resource utilization values for that VM in the Summary pane. Click the view in Explorer link in the upper-right corner to drill down to the VMware Explorer view of the object, to see detailed information about the selected object and its components. If the selected object is a monitoring VM, the VMware Explorer page appears.



Selecting a Terminal Server or Physical host from the Quick View list displays the CPU, Memory, Network, and Storage utilization charts for the selected host in the Summary pane. Clicking any data series on the charts allows you to drill down into the Metric Analyzer dashboard for the associated metric and view the metrics collected for that topology object. Click the view in Explorer link in the upper-right corner to drill down to the Host Monitor page for the selected host

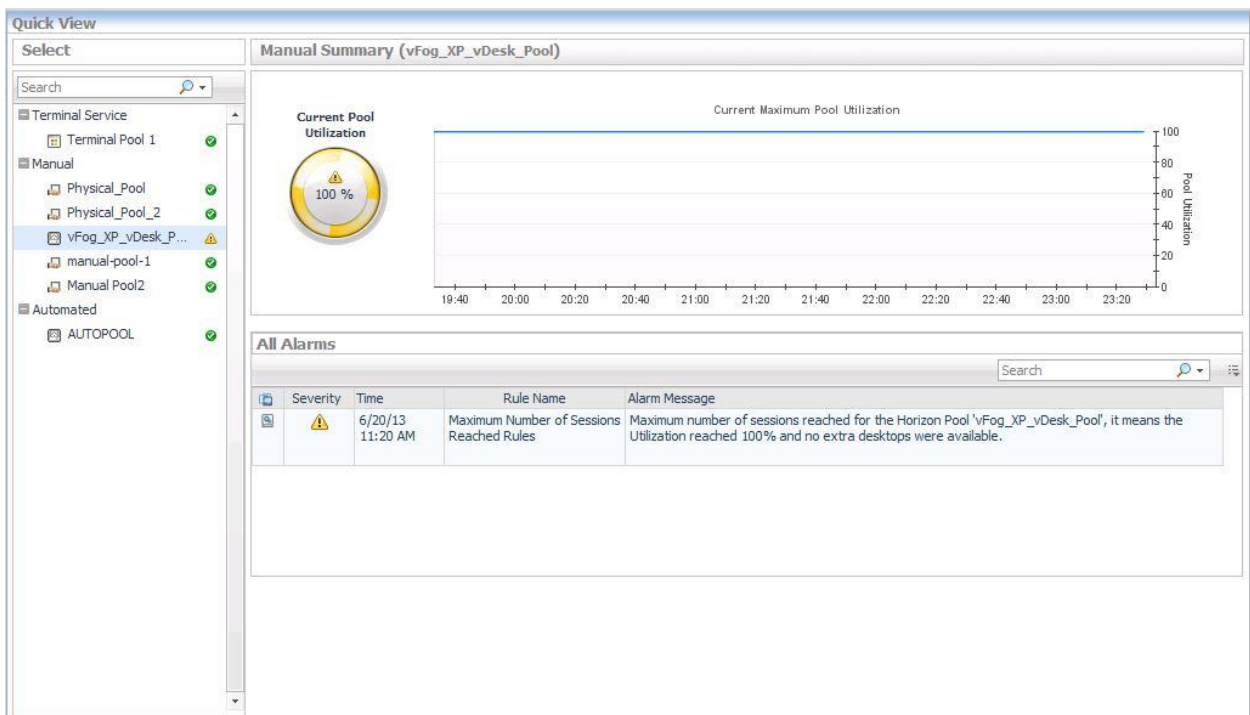


## 12.7.9 Exploring Horizon Pools

When you click the Horizon Pools tile, the Quick View shows a list of pools available in the View Instance. The pools are grouped by type: Automated, Manual, and Terminal Services.

When selecting a pool on the Quick View list, the Summary pane displays the following information for the selected pool:

- Pool name and type
- For Automated and Manual pools: the current maximum number of sessions possible, and a chart of session utilization percentage
- For Terminal Services pools: a chart of session count.
- All Alarms table showing all the alarms triggered for the pool.



## 13 Previous Features Appendix

### 13.1 Branch Office Deployments

#### Executive Summary

Remote branch office deployments are architected in a fashion to support the best end user experience at a remote/branch office while providing optimized access back to Data Center resources such as internal intranet websites, mail exchanges services, CIFS shares, SharePoint, etc. These technologies include WAN OP appliances from Dell SonicWALL WXA appliances. This architecture will facilitate local VDI compute resources to the remote branch office with WANOP appliances to facilitate optimized access to applications accessing remote data or services.

#### Introduction

Supporting users at branch offices can create a large challenge for IT to properly support from both infrastructure management and the fact that there may not be IT staff onsite. In addition, providing the end-users with the highest quality desktop experience and application performance is a must. Centralizing the management of the infrastructure will ensure easier management and enforce compliance.

Branch office deployments have had many difficulties that are mitigated here:

- Presentation traffic has not been optimal for deployments of remote desktop directly back to the Data Center.
- Application performance has been slow or degraded due WAN latency.
- Difficult to support for IT due to remoteness of environment or lack of local IT staff.
- Large WAN bandwidth requirements for some deployments.

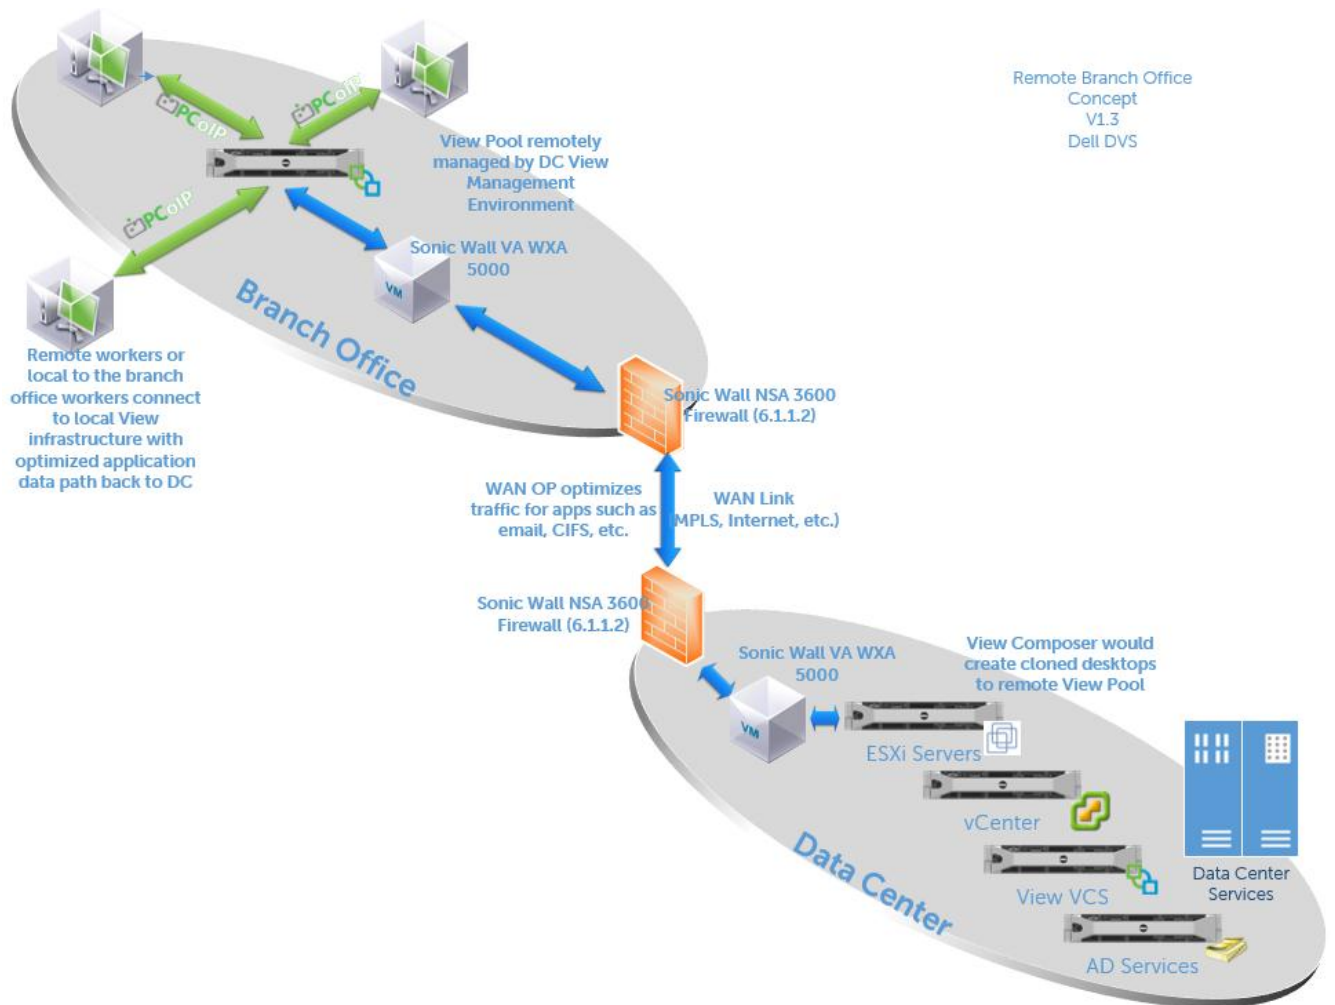
#### Architecture

Actual desktops are located on premise at the Branch office. So connection to the desktop and its "presentation" traffic are localized. Initial desktop broker connection to VCS redirects View client to local desktop connection so PCoIP traffic is all internal at branch office. Once connected even if WAN connection is compromised, then desktop and its connection still persist.

The applications used inside the desktop context will take advantage of the Dell SonicWALL Firewall/WXA WANOP devices to optimize the connection back to the Data Center resources. In this manner the desktop PCoIP traffic that results in the presentation of the desktop is optimized since the traffic is local to the branch office and the applications that run in the desktop also have their data path optimized.

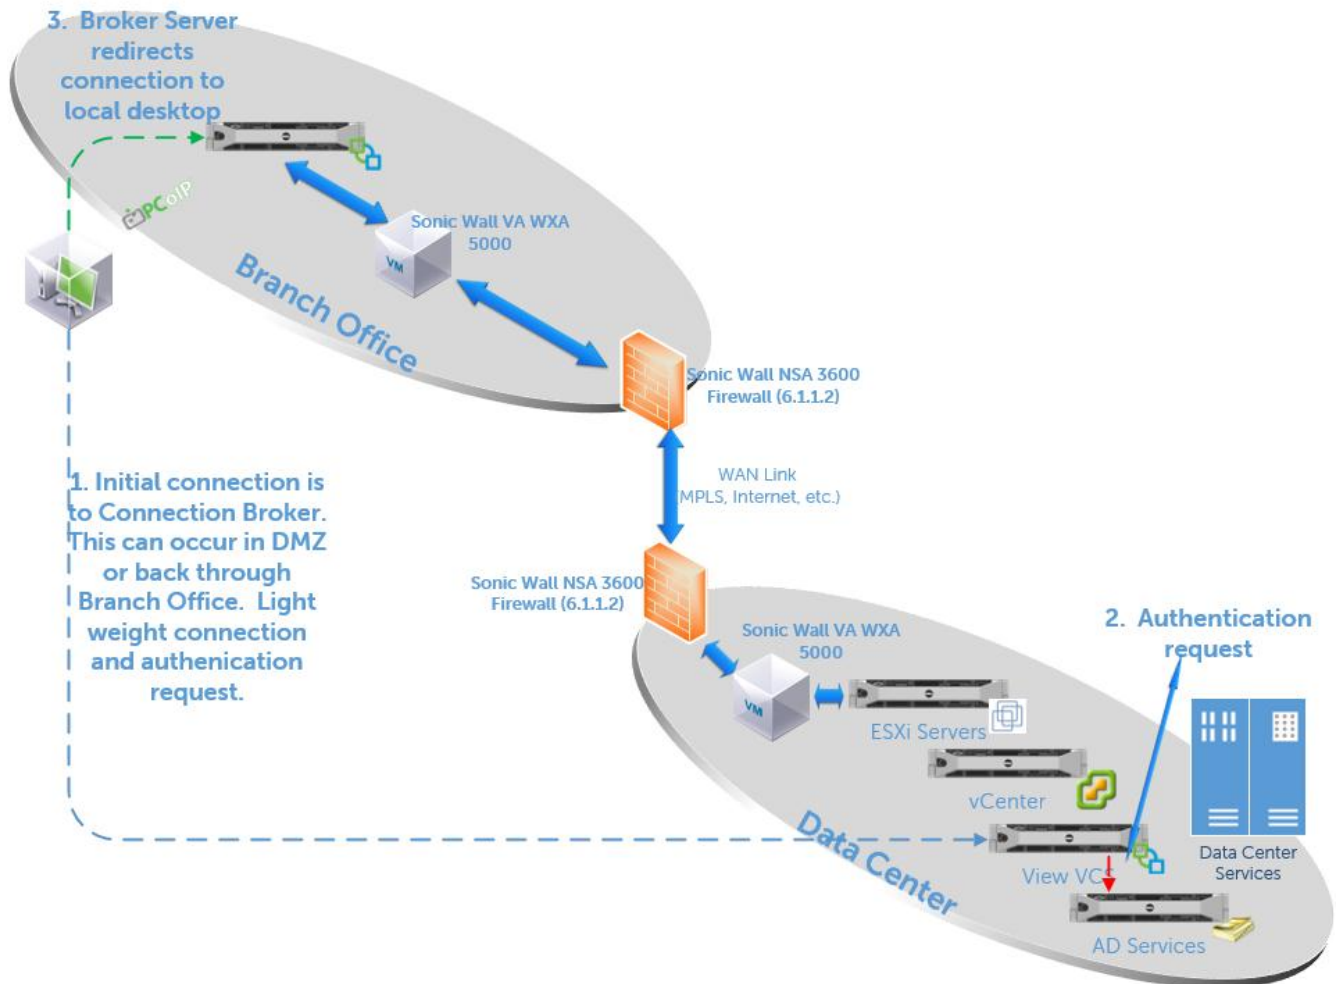
The architecture diagram below provides visualization:





The next diagram shows how the initial connection is made to Data Center VCS Broker server. This is a relatively lightweight connection that checks authentication and then directs the PCoIP connection to the appropriate desktop. In this case the desktop is local to the end user. Once this desktop connection is made if the WAN link to the Data Center is lost the connection persists from the View client to the local branch office desktop.



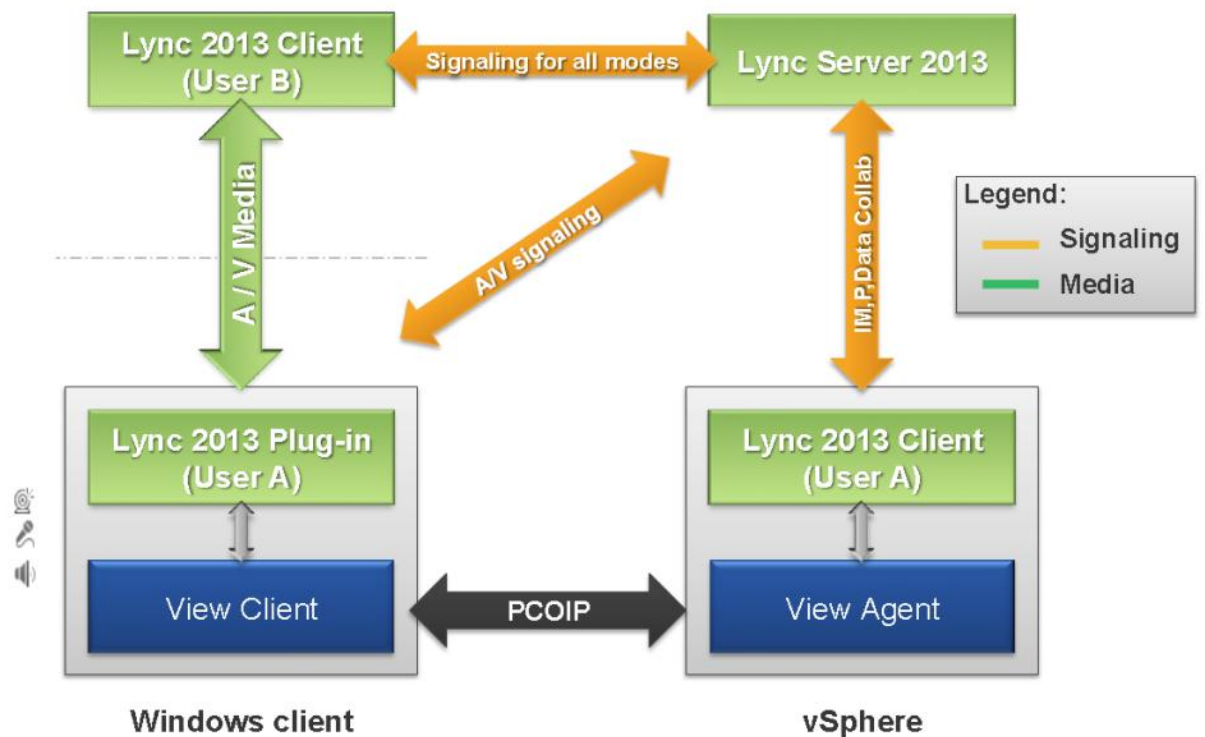


## 13.2 Microsoft Lync 2013 enablement

### 13.2.1 Microsoft Lync 2013 Overview and Architecture

Specifically Microsoft Lync 2013 has been optimized by VMware for PCoIP. A DVC channel is facilitated via PCoIP and a point to point redirect occurs on Lync 2013 plug-in enabled clients (Windows only) that keeps valuable VDI compute host resources from being consumed.

## Lync 2013 Architecture



## Microsoft Lync 2013 Supported Features with VDI

Features	Supported
Presence	✓
Instant Messaging	✓
Desktop sharing	✓
Application Sharing/Powerpoint sharing	✓
Whiteboards	✓
File Transfers	✓
Online Meetings	✓
Office Integration	✓
VoIP	✓
Video Chat	Yes, but multi-party video chat is not supported

## 13.2.2 VMware Horizon View 5.2 with Microsoft Lync 2013 Testing Results

### 13.2.2.1 Purpose

The purpose of this validation is to provide guidance when running workloads that are voice and video intensive. The application validated as part of this use case is Microsoft Lync 2013. Microsoft Lync 2013 is a suite of applications that provides access to unified communication services such as instant messaging, audio/video conferencing, desktop sharing and VoIP telephony. Communications applications such as Microsoft Lync 2013 require fast computing environments which are also optimized for low latency and able to quickly render video and audio data. The main purpose of this testing is to determine the impact of this workload type on a Dell optimized virtual desktop infrastructure. The focus of this testing was desktop sharing, audio and video conferencing parts of the unified communications suite. These components were selected specifically as these workloads put an intensive load on the infrastructure.

### 13.2.2.2 Test Infrastructure

The test infrastructure included:

- VMware vSphere 5.1 Hypervisor, VMware Horizon View 5.2
- 1 x Dell R720 rack servers (Compute Host to support two desktop sessions for Lync 2013 test)
- 1 x Dell R710 rack servers (Management Host)
- 1 x Dell R710 Windows 2013 Lync Server and AD host

<b>VMware Compute Host</b>	<ul style="list-style-type: none"><li>• 1 x Dell PowerEdge R720 Server:<ul style="list-style-type: none"><li>○ ESXi 5.1</li><li>○ Intel-(R)-Xeon- (R) CPU E5-2690 @ 2.9GHz</li><li>○ 196 GB @ 1600MHz</li><li>○ 10 x 146GB 15K SAS internal disk drives</li><li>○ Broadcom BCM5720 1 GbE NIC</li><li>○ PERC H710P RAID Controller</li></ul></li></ul>	For ESXi Environment 10 x 146 GB drives will be configured on RAID 10 This host will be hosting the Windows 7 VDI desktops.
----------------------------	---	--

### 13.2.2.3 Lync 2013 Test Methodology

#### 13.2.2.3.1 Test Objectives

The primary objective of the testing are:

- Determine the CPU, Memory, Disk IOPS and Network impact on the desktop of using the Lync 2013 client with and without the Lync 2013 Plugin for VDI for both desktop sharing, audio and video sessions.
- Determine the CPU, Memory, Disk IOPS and Network impact on the View Horizon Compute hosts of using the Lync 2013 client with and without the Lync 2013 Plugin for VDI for both desktop sharing, audio and video sessions.
- Audio tests would be conducted with USB headsets on the client endpoints that will initiate a Lync audio call to each desktop session. A radio newscast streamed via a cell phone was the input to the headsets to create a sustained audio conversation for ~30 minutes.
- Video tests would be conducted via integrated client endpoint video device where the same above mentioned radio newscast created the audio portion and two subjects make continued head and hand motions for the video portion for ~30 minutes.

- Desktop sharing was initiated from one client to another and multipage PowerPoint document was auto moved from page to page and set to an infinite loop for ~30 minutes. The Power Point contained a variety of technical diagrams and text.

The testing will focus on the all CPU, memory, Disk IOPs, and networking aspects of VDI compute hosts and detailed statistics and metrics captured.

### 13.2.2.3.2 Test Approach

The key VDI use cases which are to be validated are listed below:

- Two 2 vCPU and 2.5GB RAM VMs to test sessions to each other via manual Lync initialization.
- Measure above mentioned compute items during messaging, voice, video and desktop sharing.
- per VM overhead associated with running Lync 2013 clients in VDI sessions as well as measuring overhead on the host as discussed above
- Analyzing "point to point" connectivity and what that does to the VDI infrastructure.
- Provide guidance on a per VDI host scale measurement
- Provide guidance on a per VDI host overhead associated with running Lync 2013 clients in VDI sessions
- Provide guidance on assessing the type of clients that best suites the solution (zero client, thin client or fat clients).

### 13.2.2.3.3 Load Generation

Two desktops in the View Horizon infrastructure will manually connect to each other and test a desktop sharing, audio and video call. In each of these types of Lync activities all data described above (4 areas) will be captured with and without the Lync 2013 VDI plugin and guidance will be provided.

### 13.2.2.3.4 Monitoring Tools

The system resource utilization and performance metrics at the VMware ESXi hypervisor layer will be monitored using VMware vCenter at the ESXi host level as well as local Windows Performance Monitor metric captures in the guest VDI session.

### 13.2.2.3.5 Test criteria

Determination of a per VM session density change with Lync 2013 on VMware Horizon View.

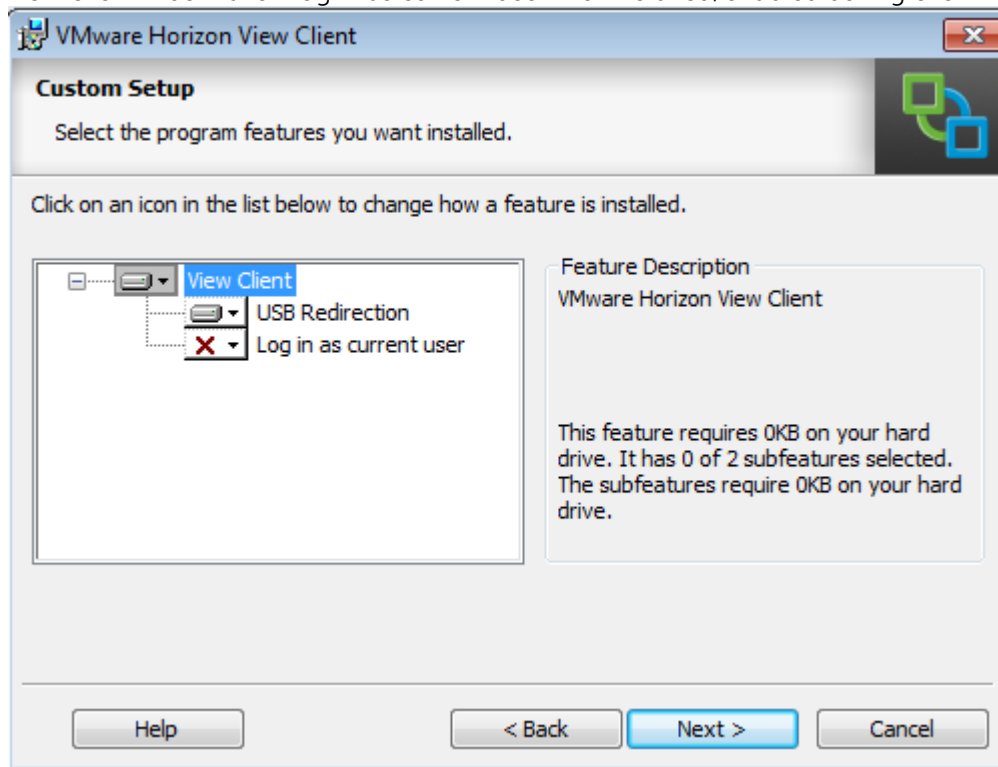
### 13.2.2.4 Lync 2013 VDI Plug-in

The Lync 2013 VDI Plugin is an optimization component built by Microsoft for Lync 2013 in VDI. It installs only on a Windows client endpoint and provides pairing capabilities with the Horizon View desktop that optimizes the locally installed Lync 2013 client. The VDI plug-in provide "point to point" connectivity for functions such as desktop sharing, audio and video calls. With this functionality it offloads the connection through the VDI compute host thus reducing CPU, memory, IOPS, and network consumption from the VDI compute host. Later in this document we will show the impact on the VDI compute architecture showing the deltas in resource consumption with and without the Lync 2013 VDI plugin.

#### 13.2.2.4.1 Prerequisites for VDI plugin with VMware Horizon View 5.2

1. View Agent 5.2
2. Endpoint client must be Windows.

3. View client for Windows 5.3 (available at: [https://my.vmware.com/web/vmware/info/slug/desktop\\_end\\_user\\_computing/vmware\\_horizon\\_view\\_clients/2\\_0](https://my.vmware.com/web/vmware/info/slug/desktop_end_user_computing/vmware_horizon_view_clients/2_0))
4. View client must have "Log in as current user" not installed/enabled during client install.



5. Lync Server media redirection enabled (described in detail below)
6. Client certificate from AD given to Lync Server installed on client endpoint where Lync 2013 VDI plugin will be installed. (described in detail below)
7. Lync 2013 plugin on Windows Client endpoint (must be 32 bit). Currently located at: <http://www.microsoft.com/en-us/download/details.aspx?id=35457>
8. Matched Lync 2013 bit version with OS version of VDI desktop (example 64 bit Lync client on 64 bit Windows 7 SP1).
9. VDI desktop OS must be Windows 7 SP1 as of June 7, 2013

### 13.2.2.5 Lync 2013 VDI Plug-in Testing/Characterization results


Note: Full results are located below in the Appendix section of this document. The entire testing results are located there. Below is a summary of the 3 different Lync tests showing CPU with and without the Lync 2013 VDI Plug-in. CPU is summarized here because it is the mitigating factor in user density.

In order to test out the functionality and provide VDI compute server overhead, two VDI sessions were manually initiated that were enabled with Lync 2013 on each desktop. The sessions used independent Windows clients that were enabled with USB headsets and integrated video devices. The test results were run with the Lync 2013 VDI Plug-in and without to verify the delta of changes and all applicable information captured on the VDI compute server AND on the desktop session. The data on the VDI compute host should be interpreted as a divisible of 2 since there were two sessions running concurrently. In addition all tests were conducted using the PCoIP protocol which controls to a certain extent the amount of data traffic due to its ability to optimize the connection. The observed idle state of the ESXi hypervisor was less than 1% as such is negligible in data inferences below.


### 13.2.2.5.1 Overall VDI Compute Host Performance information

Audio performance with and without the Lync 2013 VDI Plug-in:

#### *Audio, CPU Performance without the Lync 2013 VDI Plug-in*




#### *Audio, CPU Performance with the Lync 2013 VDI Plug-in*




In the initial test, which was conducted with no VDI Plug-in enabled in the desktop, Processor utilization (CPU) for audio showed an increase. In the no VDI Plug-in enabled test, the peak processor utilization was around 3.8% (for two sessions). In the VDI Plug-in enabled test, the CPU utilization peaked around 1.1% (for two sessions). Based on these results, it can be concluded that CPU utilization, while using the Lync 2013 VDI Plug-in with Horizon View, was reduced on the Compute Host by approximately 74% for the two sessions and that for **each** session with active Audio Lync calling would account for an ~.55% CPU increase.

Video performance with and without the Lync 2013 VDI Plug-in:

#### *Video, CPU Performance without the Lync 2013 VDI Plug-in*




**Video, CPU Performance with the Lync 2013 VDI Plug-in**




In the initial test, which was conducted with no VDI Plug-in enabled in the desktop, Processor utilization (CPU) for video showed an increase. In the no VDI Plug-in enabled test, the peak processor utilization was around 10.29% (for two sessions). In the VDI Plug-in enabled test, the CPU utilization peaked around 1.22% (for two sessions). Based on these results, it can be concluded that CPU utilization, while using the Lync 2013 VDI Plug-in with Horizon View with a video call, was reduced on the Compute Host by approximately 88% for the two sessions and that for **each** session with active Video Lync calling would account for an ~ .61% CPU increase.

**Lync Desktop Sharing performance with and without the Lync 2013 VDI Plug-in:**

***Lync Desktop Sharing, CPU Performance without the Lync 2013 VDI Plug-in***



***Lync Desktop Sharing, CPU Performance with the Lync 2013 VDI Plug-in***



In the initial test, which was conducted with no VDI Plug-in enabled in the desktop, Processor utilization (CPU) for Lync Desktop Sharing showed an increase. In the no VDI Plug-in enabled test, the peak processor utilization was around 2.25% (for two sessions). In the VDI Plug-in enabled test, the CPU utilization peaked around 1.62% (for two sessions). Based on these results, it can be concluded that CPU utilization, while using the Lync 2013 VDI Plug-in with Horizon View with a Lync Desktop Sharing session, was reduced on the Compute Host by approximately 28% for the two sessions and that for **each** session with an active Lync Desktop Sharing session would account for an ~.81% CPU increase.

**13.2.2.6 VMware View Horizon with VDI Plug-in for Lync 2013 Conclusions**

- The overall results showed that the execution of Microsoft Lync 2013 in a VMware Horizon View session using the Microsoft Lync 2013 VDI Plug-in has significant reductions in all 4 major categories (CPU, memory, network, and disk) on the VDI compute host and as such should be used when a Windows endpoint client is being used. Of the 4 measured compute categories the mitigating resource was CPU. In all of the three major workloads (basic, standard, and premium), use of Lync 2013 for all sessions on top of these workloads would result in guidance of the following for these workloads. (The reduction noted is a reduction from regular Login VSI workloads without a Lync workload running. Without the Lync Plug-in installed and configured, the reductions would be much larger to the scaled numbers)



<b>Workload</b>	<b>Percentage Reduction</b>	<b># Desktops Reduction</b>
Basic – 145 W7 users per PE R720	28%	41
Standard – 116 W7 users per PE R720	23%	27
Premium – 95 W7 users per PE R720	19%	18

In the guidance information above note that in the calculations the following were used:

- An average of CPU consumed per session by Microsoft Lync 2013 (.55% for Audio, .61% for Video with Audio, and lastly .81% for Desktop Sharing) of .65
- That .65 average CPU was then reduced by 50% to provide a level of concurrence at any given time of 50% of the users on a VDI compute server. 25% is probably closer to real world concurrence but a conservative number was chosen.
- Use of the VDI plugin dramatically reduces all compute hosts resources as opposed to not having it installed. The reduction noted above is a reduction from regular Login VSI workloads without a Lync workload.

### 13.3 View Configuration Tool

View Configuration Tool, also known as "VCT", is a tool created by VMware to facilitate building out a VMware Horizon View infrastructure in an easy and quick to deploy method. It has been tested and will deploy on Dell PowerEdge R720 hardware including the other prescribed DVS Enterprise Solution hardware bundles called out in this Reference Architecture. It is made up of two easy to deploy Linux based virtual appliances (VMware Studio and VMware View Configuration tool). Once networking settings are configured on the two appliances then access to the View Configuration Tool via a web browser to configure it is enabled. Once the relevant information is input into the fields and hit submit, a pre-created set of scripts automate the full deployment of VMware View. There are two different workflows: one which integrates VMware Horizon View into an existing Active Directory environment and another that includes the automated setup of a new Active Directory infrastructure including DNS / DHCP, etc.

If the option to integrate into your existing AD infrastructure is selected, then the following VMs will be deployed:

**Virtual Center Appliance**  
**View Connection Broker**  
**View Composer**

If the option to create a new AD is selected, then the followings VMs will be created:

**Windows 2008 R2 server running AD / DNS / DHCP**  
**Virtual Center Appliance**  
**View Connection Broker**  
**View Composer**

VCT is a useful piece of software that will be relevant in several situations:

- An IT department that already has AD setup will be able to deploy VMware View rapidly and easily without having to download each individual piece of software.
- The workflow that includes the new AD setup will be very useful in small offices or branch office setup.
- It would even be possible for non-trained IT personnel to deploy the tool in a branch office and a few hours later they would have a fully functional AD / DNS / DNCP / VMware View environment.

- Test environments were quick setup with a resulting teardown for testing or POC purposes

*Note1: At present you need Firefox to perform the configuration but IE will be supported soon.*

*Note2: Relevant Windows 2008 R2 licenses are needed to perform the deployment*

*Note3: The Virtual Center that is deployed is the Virtual Center appliance (Linux based) (not the windows version)*

## **13.4 VMware vCenter Operations Manager for View**

### **13.4.1 Executive Summary**

IT organizations today are faced with increasing complexity when deploying virtual desktop infrastructure (VDI) environments to support their business requirements. With additional complexity in the environment, inefficiency, reactive troubleshooting, and downtime become larger and larger problems.

VMware vCenter Operations Manager for View or "vCOPS", provides an elegant, integrated, and comprehensive end-to-end solution that enables great visibility into all aspects of the underlying virtual infrastructure to allow for higher efficiency, greater availability, better performance, proactive troubleshooting, and higher quality of service.

### **13.4.2 Introduction**

As IT computing environments become larger and more complex, additional resources are required to configure, manage, and maintain the environment to keep it in optimal operating condition. The resources that are required are, but not limited to, employees, servers, storage, switching infrastructure, and management software. All of these are very important pieces to a computing infrastructure, especially in a VDI environment. If any piece of the underlying infrastructure is not working correctly, it will most likely affect all of the users that are utilizing those desktops. These problems can be tackled in a number of ways, which are discussed in the following sections.

#### **13.4.2.1 Traditional Approach**

In a typical enterprise IT organization, the computing environment will be split up and assigned to different groups who are responsible for maintaining their respective infrastructure. Management tools that are created or purchased are typically designed to have a very narrow scope or view of the environment to report issues with those specific infrastructure components. The teams will have a silo view of the environment and aren't able to adapt to a changing and dynamic environment. When an issue is discovered, it can also be very difficult to track down. Locating which part of the infrastructure that is having an issue can often take longer than the actual fix.

Overall monitoring solutions may be used to get a high level view of the environment, but are commonly lacking in a lot of areas. The sheer amount of monitoring data and quantity of alerts may overwhelm administrators. Additionally, the errors that are received may be cryptic or uninformative. They may even be completely incorrect, which is referred to as a false positive. Because of the overwhelming amount of data that has to be sifted through, many critical issues can be masked by notifications or errors that have varying importance.

The lack of insight and visibility into the desktop computing environment has far reaching affects. Administrators tend to be reactive instead of proactive when addressing issues. End users tend to report issues first by calling and overwhelming the helpdesk. Productivity suffers because end users are unable to reach their desktop and files or the quality of service is bad, causing frustration and lack of confidence in the capability to complete work and projects in a timely manner.

#### **13.4.2.2 V4V Approach**

In an enterprise IT organization that has deployed VMware vCenter Operations Manager for View, the individual operating components in the environment can still be managed by the respective teams or as a single team of systems administrators and engineers. V4V provides the flexibility to manage the desktop computing environment in your environment based on business needs while also leveraging the advanced visibility, monitoring, and alerting capabilities built specifically for VDI.

V4V provides visibility, analytics, and alerting to provide robust and comprehensive monitoring, troubleshooting, and performance information. The monitoring data is presented in a very clear and concise manner in the form of a dashboard. This dashboard provides comprehensive visibility into the VDI environment. The dashboard provides a simple, easy to use interface that helps an administrator quickly assess the health and performance of the overall desktop computing environment, identify bottlenecks, and improve infrastructure efficiency of the entire View environment.

The capabilities provided by introducing V4V into the desktop computing environment will have a direct impact on the productivity of both administrators in the organization and the end users who are actually using the desktops. V4V has patented self-learning analytics that adapt to each computing environment and continuously analyzes thousands of metrics for server, storage, networking, and end-user performance. Dynamic threshold and "smart alerts" will notify administrators when there are potential issues with the system and allow them to be proactive instead of reactive. This means that end-users will get a very high and consistent quality of service and confidence in a system that they have confidence in and can rely on.

### **13.4.3 High Level Solution**

vCenter Operations manager for Horizon View extends the functionality of vCenter Operations Manager to monitor and manage VMware Horizon View environments.

vCenter Operations Manager for Horizon View collects data from managed virtual desktops and presents that data in preconfigured dashboard for real-time and predictive analysis.

#### **13.4.3.1 vCenter Operations Manager for Horizon View Features**

vCenter Operations Manager for Horizon View extends the functionality of vCenter Operations Manager Enterprise, and enables IT administrators and help desk specialists to monitor and manage Horizon View Virtual Desktop Infrastructure (VDI) environments.

vCenter Operations Manager for Horizon View is based on vCenter Operations Manager Enterprise. It collects performance data from monitored software and hardware resources in your enterprise and provides predictive analysis and real-time information about problems in your VDI infrastructure. It presents data through alerts, in configurable dashboards, and on predefined pages in the Custom user interface.

The View Adapter obtains the topology from the Horizon View environment, collects metrics and other types of information from the desktops, and passes the information to vCenter Operations Manager.

Typical users of vCenter Operations Manager for Horizon View are IT administrators and help desk specialists. IT administrators can use vCenter Operations Manager for Horizon View to get a quick overview of how the Horizon View environment is behaving, and to view important metrics associated with their environment. Help desk specialists need to quickly see resources related to end users sessions, and perform basic troubleshooting to view, analyze, and resolve problems.

Figure 1 below shows how the computing environment can be managed by distinctly different groups while using VMware vCenter Operations Manager for View to monitor the complete environment.

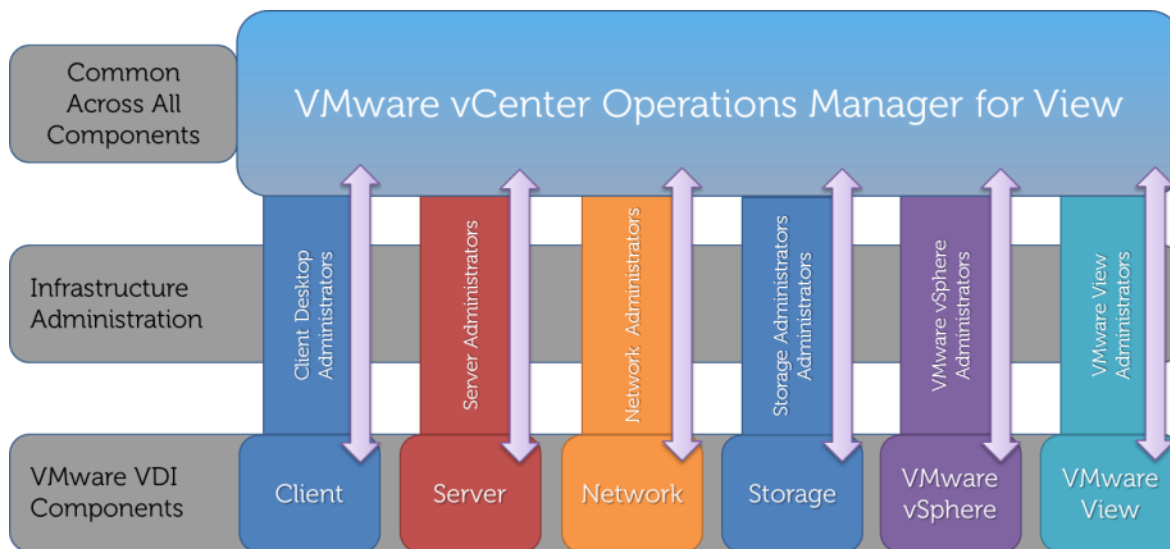


Figure 10

## 13.4.4 Solution Details

### 13.4.4.1 vCenter Operations Manager for Horizon View Architecture

vCenter Operations Manager for Horizon View is based on the vCenter Operations Manager Custom user interface, which you can use to monitor your Horizon View 5.0, 5.1, or 5.2 Virtual Desktop Infrastructure (VDI) environments.

vCenter Operations Manager for Horizon View is composed of the following components:

- vCenter Operations Manager vApp
  - vCenter Operations Manager for Horizon View Adapter
  - vCenter Operations Manager for Horizon View Broker Agent
  - vCenter Operations Manager for Horizon View Desktop Agent
- The purpose of each component is outlined below.

#### vCenter Operations Manager vApp

The vCenter Operations Manager vApp consists of two VMs: the Analytics VM and the UI VM. The vCenter Operations Manager vApp is deployed in the View Management Block.

The Analytics VM is responsible for collecting data from vCenter Server, vCenter Configuration Manager, and third party data sources such as metrics, topology and change events. The components within the Analytics VM are Capacity and Performance Analytics, Capacity Collector, File System Database and Postgres DB.

The UI VM allows you to access the results of the analytics in the form of badges and scores using the web-based console. The applications in the UI VM are vSphere Web Application, Enterprise Web Application and Administration Web Application.

#### vCenter Operations Manager for Horizon View Adapter

The vCenter Operations Manager for Horizon View Adapter is now integrated in the Analytics VM for the most part and all you need to do is deploy the package to install the custom View dashboards in the environment. The View Adapter pulls in the metrics from the View Desktops, View Connection Server and the View Events Database and feeds that information back to the vCenter Operations Server.

#### vCenter Operations Manager for Horizon View Broker Agent

The vCenter Operations Manager for Horizon View broker agent is a Windows service that carries out the connection between the vCenter Operations Manager for Horizon View adapter and the agent

components on the Horizon View desktops. The broker agent collects the Horizon View inventory for the adapter, collects events from the database, and configures the desktop agents when used in Horizon View 5.2 environments.

### **vCenter Operations Manager for Horizon View Desktop Agent**

The vCenter Operations Manager for Horizon View desktop agent is an agent that is installed on each virtual desktop with the View 5.2 agent installation. It allows for data collection and aggregation from the Horizon View desktops. For View 5.0 or 5.1 desktops, the desktop agent must be deployed manually.

#### **13.4.4.2 How vCenter Operations Manager for Horizon View Components Work Together**

The vCenter Operations Manager for Horizon View adapter runs on the vCenter Operations Manager vApp, or on a remote vCenter Operations Manager Collector server. The adapter collects performance data from the broker agent and the desktop agents and sends it to vCenter Operations Manager for analysis and visualization.

The vCenter Operations Manager for Horizon View broker agent runs on a Horizon View Connection server. The broker agent collects the Horizon View inventory, monitors the Horizon View infrastructure, and sends the collected data to the adapter. The broker agent pushes configuration data to the desktop agents so that they can communicate with the adapter. If you configure an events database in your Horizon View environment, the broker agent can connect to that database and send Horizon View events data to the adapter.

The vCenter Operations Manager for Horizon View desktop agent runs on every virtual desktop that you want to monitor on vCenter Operations Manager dashboards. The desktop agent uses the configuration data that the broker agent provides to connect to the adapter. The desktop agent sends performance information for the desktop it is running on directly to the vCenter Operations Manager for Horizon View adapter.

For Horizon View versions 5.0 and 5.1, you must install the desktop agent manually. vCenter Operations Manager for Horizon View desktop agent is installed as part of the Horizon View Agent in Horizon View 5.2 and later.

Figure 2 outlines the basic relationships between the different components of the vCenter Operations Manager for Horizon View.

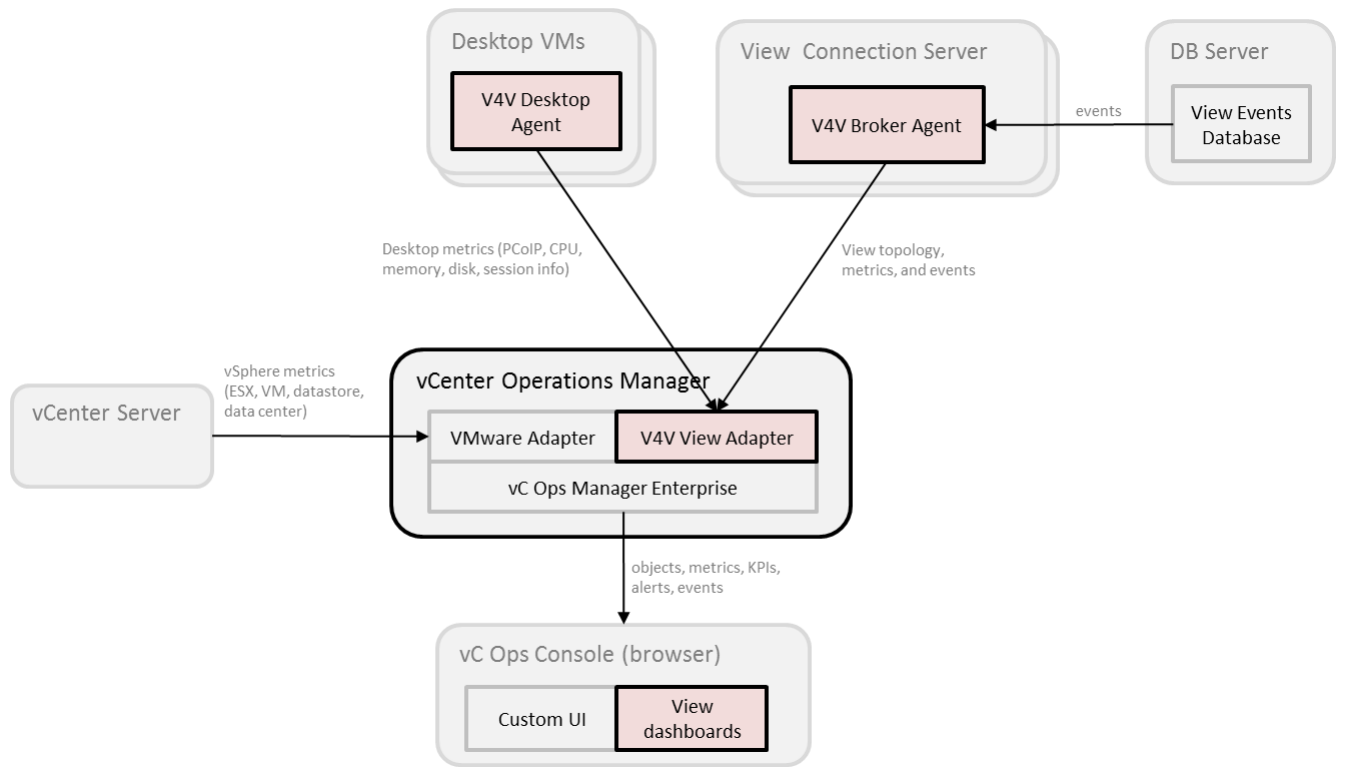


Figure 11

### 13.4.4.3 Licensing

Licensing of vCenter Operations Manager for Horizon View 1.5 is as follows:

- vCenter Operations Manager for Horizon View 1.5 is licensed by "concurrent Horizon View user"
- Licenses are available in 10 or 100-user packs
- The vCenter Operations Manager component license that is included in per-user licensing is for View-supporting infrastructure
- A separate vCenter Operations Manager license is required for non-View infrastructure monitoring

#### 13.4.4 Business Benefits

There are a multitude of business benefits from installing and utilizing VMware vCenter Operations Manager for View in VDI environments. Some of them are measurable by using metrics such as support desk calls, while others are not, such as ease of use.

Some of the largest benefits can be summarized by the following points:

- Comprehensive visibility into the performance and health of VMware View infrastructure which eliminates bottlenecks and improves infrastructure efficiency.
- Intelligent automation of root-cause analysis and autocorrelation of monitoring data across the entire stack reduce troubleshooting times and improve team productivity.
- Self-learning analytics that notify desktop administrators of impending issues before they impact end users enable proactive management and process improvements.
- Lower operational costs by improving staff productivity and allowing smaller teams to manage the infrastructure.
- Providing better and more consistent end-user experience from any device by allowing desktop and infrastructure administrators comprehensive visibility into virtual desktop operations and infrastructure, enabling them to eliminate performance problems before they impact end users.

#### 13.4.5 Summary

VMware vCenter Operations Manager for View provides an elegant and comprehensive end-to-end solution that enables great visibility into all aspects of the underlying virtual infrastructure to allow for higher efficiency, greater availability, better performance, proactive troubleshooting, and higher quality of service.

### 13.5 Graphics Support – NVIDIA K1/K2 GRID cards and VMware vSGA support

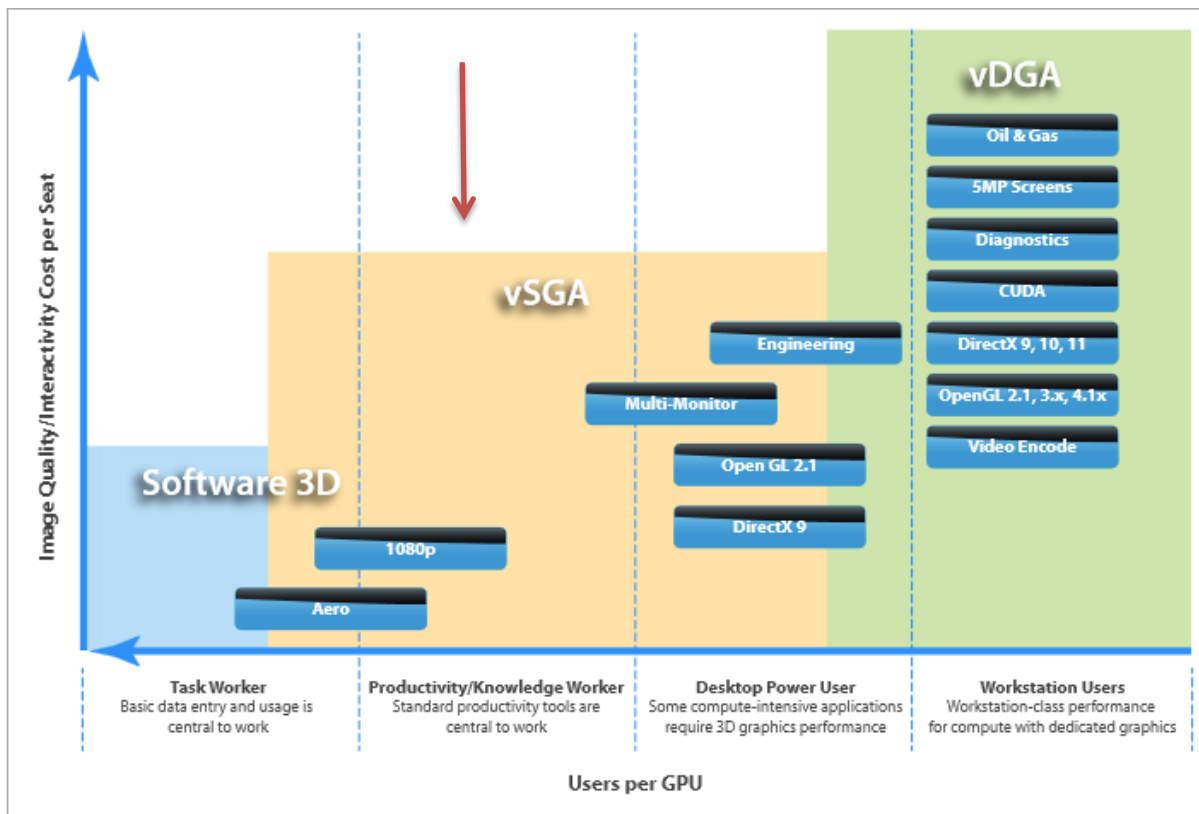
#### 13.5.1 NVIDIA K1/K2 Testing/Characterization results

##### 13.5.1.1 Environment Summary

###### Solution Overview

These results describe validation efforts undertaken on NVIDIA Grid cards K1 and K2 to determine density on an R720 Compute host using an Intel Xeon CPU E5-2670 @ 2.6GHz (highest currently qualified CPU with NVIDIA GRID cards) running ESXi 5.1. The aim of the tests were to determine the maximum number of VMware View desktops that can run on the compute host while running a Multimedia Login VSI workload with a graphics load for the K1 tests including using eDrawings and The Hobbit movie trailer for the K2 cards.

The capability for 3D graphics and video in VMware® Horizon View™ further expands the use cases and target users that IT can deliver with virtual desktops to. In addition to expanding the target use cases, 3D augments the virtual desktop user interface by enabling a more graphically rich experience



In the testing and characterization described in this section the NVIDIA GRID cards are configured in shared mode or vSGA which results all the VMs sharing access to the collective of GPUs in a similar fashion that CPUs are shared. In this scenario a GPU is not directly assigned to a VM as it is in vDGA.

The target for NVIDIA K1 tests were lighter 3D graphics. For NVIDIA K2, with its higher end GPUs we targeted more intensive graphics tests. Again both were tests against vSGA functionality.

For the purposes of this testing and characterization, users of graphics-intensive applications in a Virtual Desktop environment are divided into 2 categories: Premium Plus users and Workstation users. These are defined below:

- “Premium Plus” VDI users are typically users who may be consuming relatively high-end graphics through activities such as Google Earth and HTML5 graphics and also reviewing/consuming electrical, mechanical CAD drawings etc. *[Note: The term ‘Premium Plus’ is used to distinguish this user type from the existing ‘Premium’ user type that is used in current Cloud-Client Computing (CCC) PAAC and Sizing Activities].*
- Workstation users, as the name implies, are users who would typically have used high-end physical workstations (e.g. Dell Precision). Typical activities carried out by these users would include 3D modeling for the oil and gas industry, involving a large amount of resource –intensive activities such as model rotation etc.
- The Premium-Plus use case is targeted at sVGA (shared GPU model)

The NVIDIA K1 card represents a graphics card which hosts GPUs that fit into this mid-range GPU category. These mid-range GPUs are targeted at users who execute activities such as Google Earth, HTML5 based video/graphics, etc. There is a larger quantity of GPUs per K1 Card (4) than K2 Card (2). The NVIDIA K2 card represents a graphics card which hosts GPUs that fit into this high-end GPU category. These high-end GPUs are targeted at “Workstation Light” users (in vSGA mode) who execute activities such as reviewing 2-D and 3-D CAD designs. There is a smaller quantity of GPUs per K1 Card (4) than K2 Card (2).

## Hardware and Software Environment



Solution Configuration - Hardware Components:		Description
<b>Dell Host (K1)</b>	1 x Dell PowerEdge R720 Servers: <ul style="list-style-type: none"> <li>• VMware ESXi 5.1.0 (Build 914609)</li> <li>• Intel Xeon CPU E5-2670 @ 2.6 GHz</li> <li>• 192GB @ 1600 MHz</li> <li>• 10 x 146GB 15K SAS internal disk drives</li> <li>• Broadcom BCM5720 1 GbE NIC</li> <li>• PERC H710P RAID Controller</li> <li>• 2 * x16 PCIe Risers</li> <li>• 2 * 1100 W Power supplies</li> <li>• 2 * NVIDIA K1 / K2 Grid Card</li> </ul>	For ESXi environment the Hypervisor is installed on an internal SD card.
<b>Virtual Desktops</b>	<ul style="list-style-type: none"> <li>• Windows 7 SP1 x32</li> <li>• 1GB memory (2.5GB for K2 test)</li> <li>• 1 vCPUs (2 for K2 test)</li> <li>• SCSI – LSI Logic SAS</li> <li>• Thin provisioned 25GB hard drive</li> <li>• 3d support enabled</li> </ul>	
<b>End Point Devices</b>	<ul style="list-style-type: none"> <li>• Dell OptiPlex 7010</li> <li>• Dell Wyse Z90D7</li> </ul>	
<b>Performance Monitoring</b>	<ul style="list-style-type: none"> <li>• Subjective User Experience</li> <li>• Login VSI</li> <li>• Perfmon Frames per second counter</li> <li>• ESX Performance Data</li> </ul>	

The setup of the ESXi hosts and configuration of the desktop pool is covered in the run book section below.

## Workload Generation

For the K1 cards, the primary workload was Login VSI (multimedia workload) and a 60 second HTML5 Fishbowl graphics load generator test.

For the K2 cards the workload was eDrawings (to represent higher end graphics loads) and The Hobbit movie trailer video which was played at 30 FPS (frames per second) on one VM.

**Note: More information on fishbowl can be found here**

<http://ie.microsoft.com/testdrive/performance/fishbowl/>

**Note: More information on eDrawings can be found here**

<http://www.eDrawingsviewer.com/ed/download.htm>

For the K1 cards we reached the desktop density number by continually adding desktops until one of the host metrics moved through its threshold i.e. CPU over 85%

For the K2 cards we reached the desktop density number by continually adding desktops until the Video FPS value dropped below 25 FPS (in the case of K2) or the host metrics went over the thresholds below.

The success criterion of the video trailer is the FPS value consistently measuring over 25 FPS while the companion workloads were in use on other VMs.

The following tests were run:

1. K1 cards - Login VSI with multimedia workload and a 60 second HTML5 Fishbowl test

2. Both cards - Subjective User Experience tests using various workloads
3. K2 cards - eDrawings + Hobbit trailer

### 13.5.1.2 Graphic Specific Performance Analysis Results

#### Density tests

#### K1 card test results (Automated Login VSI multimedia and Fishbowl tests)


We loaded the host with virtual desktops while running these tests until the host metrics reached maximum accepted performance levels – these metrics are:

Parameter	Pass / Fail Threshold
Compute Host CPU Utilization	85%
Compute Host Memory Utilization	85%
Network Throughput	85%
Tier 1 Storage Latency	20ms
Tier 1 Storage IOPS	Monitor in conjunction with Tier 1 storage latency.
GRID Card GPU Utilization	85%
GRID Card Memory Utilization	85%

The results showed that up to 42 VDI desktops was the optimum number of desktops in order to maintain acceptable VDI host performance. With this many desktops all of them properly logged off successfully after the Login VSI tests were complete.

The graph on the next page shows the FPS on one representative VM during the tests.

#### FPS With K1 GPUS's enabled




#### FPS With K1 GPU's disabled






The average FPS with the GPUs enabled is 9.2 and without the GPUs it is 7.7. So it appears that using the K1 GPUs offer a little better performance in FPS at a rate of approximately 16.3% better with K1 GPUs.

The next few graphs show the host utilization with 42 desktops running the multimedia Login VSI workload and a 60 second fishbowl test

### Host metrics with GPU's enabled











These graphs show the CPU was the host resource that was being impacted most from these tests.


### Host metrics without GPU's enabled









**Stratosphere Results with GPU's enabled.**



As you can see on this graph, the user experience is very good when the GPU's are enabled.

**Stratosphere Results with GPU's disabled.**






As you can see on this graph, the user experience is also good when the GPU's are disabled.


### **K2 Card results (eDrawings and Hobbit)**



K2 testing was designed to test the capacity of a slightly higher end 3D graphics experience. As such testing determined that up to 36 VDI desktops was the optimum number of desktops in order to maintain acceptable host performance. With this many desktops, FPS on the desktop VM running the Hobbit video stayed consistently above 25 fps and the host metrics were acceptable. These results are approximately 14% less desktops than the K1 workload testing results but are a higher end graphics load.




The graph above shows the FPS on the VM running the Hobbit trailer while 35 other desktop VMs run eDrawings. Once a 37<sup>th</sup> desktop was powered up and eDrawings was started then the FPS on the benchmark VM dropped under 25 consistently.

**Host metrics during K2 density tests**







All host metrics tested within tolerances during this test period.

### Subjective End-User Experience

The type of endpoint can have a significant impact on the user experience regardless of the GPU being used.

There were successful connections to desktops and execution of graphics applications with thin clients such as the Wyse Z90D7 and the OptiPlex 3010 during these tests.

With the K2 card tests, success was achieved in opening applications (Office and Google earth) and interaction with the desktop. Performance on the Wyse terminals was satisfactory as was the OptiPlex 3010. During heavier FPS activity on the OptiPlex during the video trailer the FPS value was maintained at a higher rate than the Wyse client. This is changing and evolving rapidly as the Wyse client portfolio has many changes and client options.

With the K1 card tests, successful connections to the desktops and execution of applications were achieved with the thin clients and the applications were responsive with good overall experience. Again the FPS value on the OptiPlex, during the video trailer was maintained at a higher rate than the Wyse client.

### 13.5.1.3 VMware Horizon View vSGA Conclusions

The K1 cards show improvements in host resource consumption in light graphics usage as described in the tests afore mentioned.

The K2 cards also show improvements in host resource consumption in heavier graphics usage as well.

When attempts were made during testing for a "lighter" heavy graphics load the vSGA mode proved to not be able to provide a suitable end-user experience. For example in the least intensive mode the Heaven Benchmark application was barely able to sustain a running mode in the lowest resolution and lowest FPS modes. As such the recommendation is that a pass-thru mode (vDGA would be more suitable to a higher end spectrum of graphics 3D applications).

The current guidance is that on DVS specified VDI compute host (PE 720 specified above):

- 42 desktops can be deployed with tolerance specs in a lighter 3D graphics loads such as the Login VSI Multi-media workload with an HTML 5 Graphics load generated payload with two NVIDIA K1 cards.
- 36 desktops can be deployed with tolerance specs in a heavier (but still lighter) 3D graphics loads such as an eDrawings sample with rotation actions and a 30FPS movie trailer (The Hobbit movie

trailer) generated payload with two NVIDIA K2 cards.

It is important to note that the choice of endpoint also makes a difference when considering the FPS value which translates to the perception of smoothly flowing video.

## **13.6 General Feature Updates**

### **13.6.1 Desktop Storage Reclamation – Storage UNMAP**

#### **Executive Summary**

With the release of VMware vSphere 5.1 and View 5.2, a new VAAI SCSI primitive has been introduced called Dead Space Reclamation as part of the overall thin provisioning primitive. Another name for Dead Space Reclamation is "UNMAP". A critical component of the new feature is the Space Efficient Virtual Disks, which is a new virtual disk (vmdk) format that adds the benefit of growing and shrinking dynamically. The new SE sparse disk implements a space reclaim feature to reclaim blocks that were previously used but now are unused on the guest OS. These are blocks that were previously written but currently are unaddressed in a file system/database due to file deletions, temporary files, and so on. Before the introduction of this feature, LUNs would forever grow and consume space. Without manual administrative intervention, virtual machines would eventually consume all space on the LUN effectively making it "thick". VMware states that with their internal testing, virtual machine unused capacity would grow about 1GB per VM per week.

This testing involved qualifying the VMware space reclamation feature with the Dell Compellent SC8000 and Dell EqualLogic PS6110XS storage arrays. The testing was to ensure that the feature worked as expected, but does not extensively study performance impact or interoperability with other products or features.

## Technology Deep Dive

The new SE sparse disk implements a space reclaim feature to reclaim blocks that were previously used but now are unused on the guest OS. These are blocks that were previously written but currently are unaddressed in a file system/database due to file deletions, temporary files, and so on.

There are two steps involved in the space reclamation feature: The first step is the wipe operation that frees up a contiguous area of free space in the virtual machine disk (VMDK); the second step is the shrink, which unmaps or truncates that area of free space to enable the physical storage to be returned to the free pool.

### Wipe

- Initiate a call to VMware Tools to scan the guest OS file system.
- Mark the unused blocks as free.
- Run the SCSI UNMAP command in the guest to instruct the virtual SCSI layer in the VMkernel to mark the blocks as free in the SE sparse disk.

### Shrink

- SCSI device – VMware® ESXi™ issues a SCSI UNMAP command to the array.

The wipe operation is initiated via an API call to VMware Tools. VMware Tools initiates a scan of the guest OS to find stranded space and mark the file system blocks as free. The first SCSI UNMAP operation is then run from within the guest OS, instructing the VMkernel as to which blocks can be reclaimed. The VMkernel captures these SCSI UNMAP commands and does not pass them through to the array. When the VMkernel detects which blocks are free, it uses its virtual SCSI layer to reorganize the SE sparse disk by moving blocks from the end of the disk to unallocated blocks at its beginning. This creates a contiguous area of free space within the VMDK. The shrink operation then sends a SCSI UNMAP command to the array to free the space.

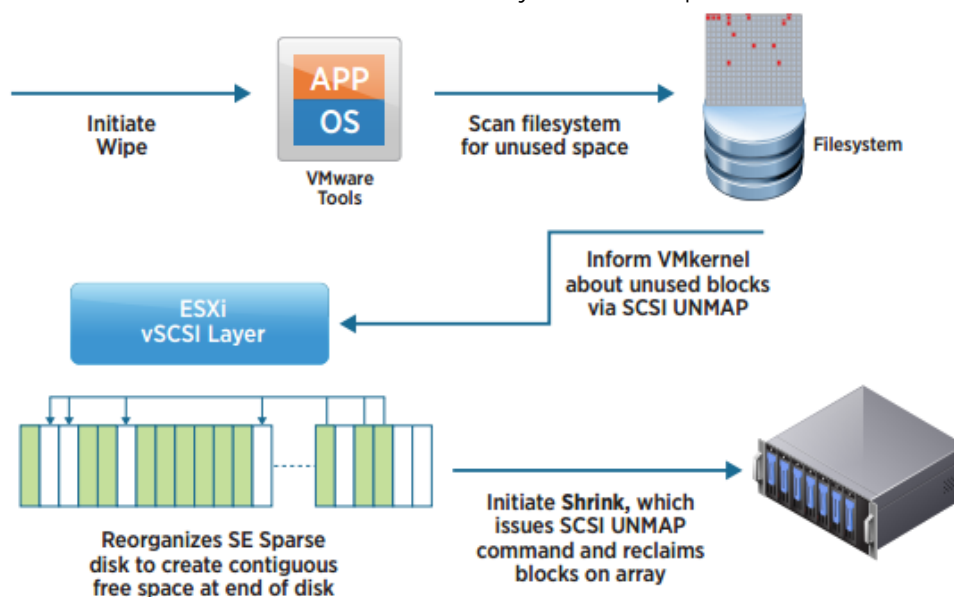


Figure 12

## Prerequisites

This document assumes that you already have a working VMware® Virtual Center Cluster with VMware Horizon View 5.2 installed and configured and you have a working knowledge of both.

The UNMAP feature is a primitive that is not available on all Storage Arrays / Firmware. Review the VMware HCL for compatibility of your storage array:

<http://www.vmware.com/resources/compatibility/search.php>

Configure the Compellent SC8000 as per the best practice guide available here:

<http://www.dellstorage.com/WorkArea/DownloadAsset.aspx?id=2847>

Configure the EqualLogic storage array as per the best practice guide available here:

<http://en.community.dell.com/techcenter/storage/w/wiki/2639.equallogic-configuration-guide.aspx>

- This new View feature requires vSphere 5.1 as well as hardware version 9 for the VMs.
- The Space Efficient format is only for linked clone pools; and only for the OS disk.
- Horizon View automatically chooses the space efficient format when possible and there is no official way to change this behavior.
- The Space Efficient Format is not yet supported for Windows 8.
- The space reclamation process is kicked off automatically by Horizon View when the reclaimable space for the VM exceeds the threshold set by the administrator.
- The space reclamation statistics (reclaimed amount, reclaimed time) are displayed in the Pool Summary, and Desktop Summary.
- SESparse disks are compatible with the Horizon View Storage Accelerator technology (Understanding CBRC) from View 5.1 / vSphere 5.
- If you are doing a fresh install of Horizon then the reclaim feature is enabled automatically. If you are upgrading then you will have to enable the feature for VC in View administrator. Click view configuration – servers – right click vc and edit – storage and tick “reclaim VM disk space”
- The LUNs on the storage array should be thin provisioned and support VAAI Delete

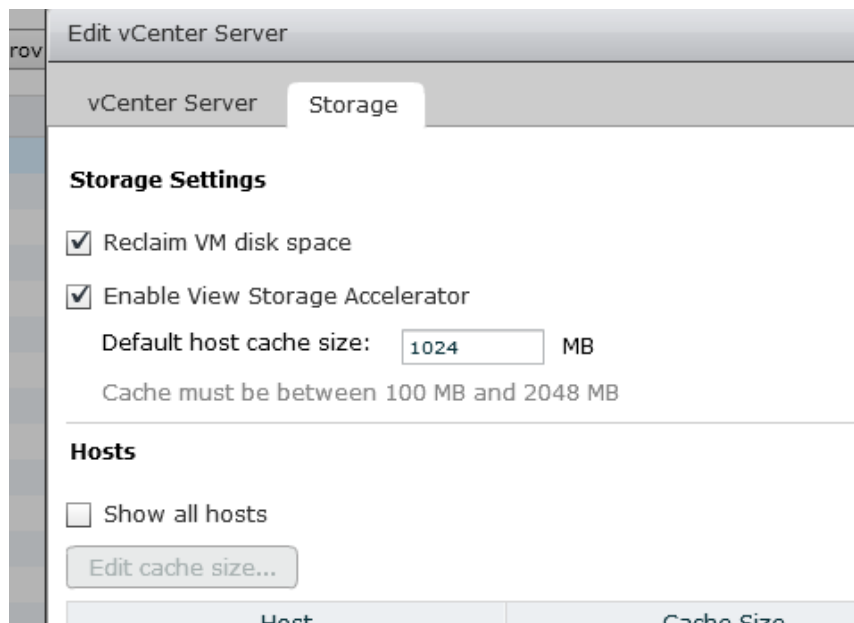


Figure 13

## VMware View 5.2 UNMAP with Compellent SC8000

### Components Overview

#### Hardware

- Dell PowerEdge R720 servers
- 1x Compellent SC8000 (dual controller)
  - Firmware 6.3
- Power Connect Switches

#### Software

- VMware® ESXi 5.1 Update 1

- VMware® Virtual Center 5.1 Update 1
- VMware® Horizon View 5.2
- Microsoft SQL Server 2008 R2
- Microsoft Windows Server 2008 R2

### Procedure

Firstly, make sure that your data store supports the UNMAP primitive using the following command:  
**esxcli storage core device vaai status get -d naa.6000d31000eced000000000000000024**

```


~ # esxcli storage core device vaai status get -d naa.6000d31000eced000000000000000024
naa.6000d31000eced0000000000000000024
  VAAI Plugin Name:
  ATS Status: supported
  Clone Status: supported
  Zero Status: supported
  Delete Status: supported
  
```

**Figure 14**

As you can see above, the delete status is supported.

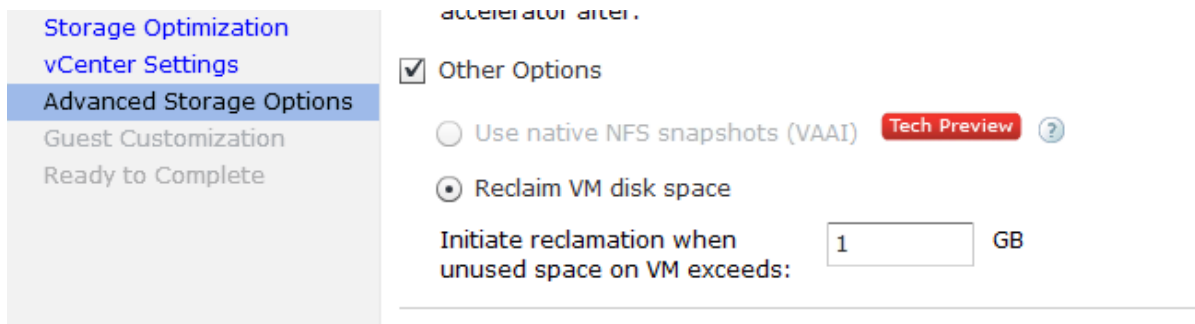
Next, create a pool using the following criteria.

1. Create a Windows XP or Windows 7 VM in vSphere using the latest virtual hardware version, install the Tools & View Agent, and snapshot it to serve as a base image for a View Composer based View Desktop.



**Figure 15**

2. Create a new Automated, Dedicated, and Horizon View Composer linked clone based pool and provision & entitle 1 or more desktops in the pool, this test was performed with 5.
3. Ensure that space reclamation is enabled for the desktop pool.



**Figure 16**

You can select Blackout times on this screen also where you can prevent the VM Space reclamation procedure running at certain times.



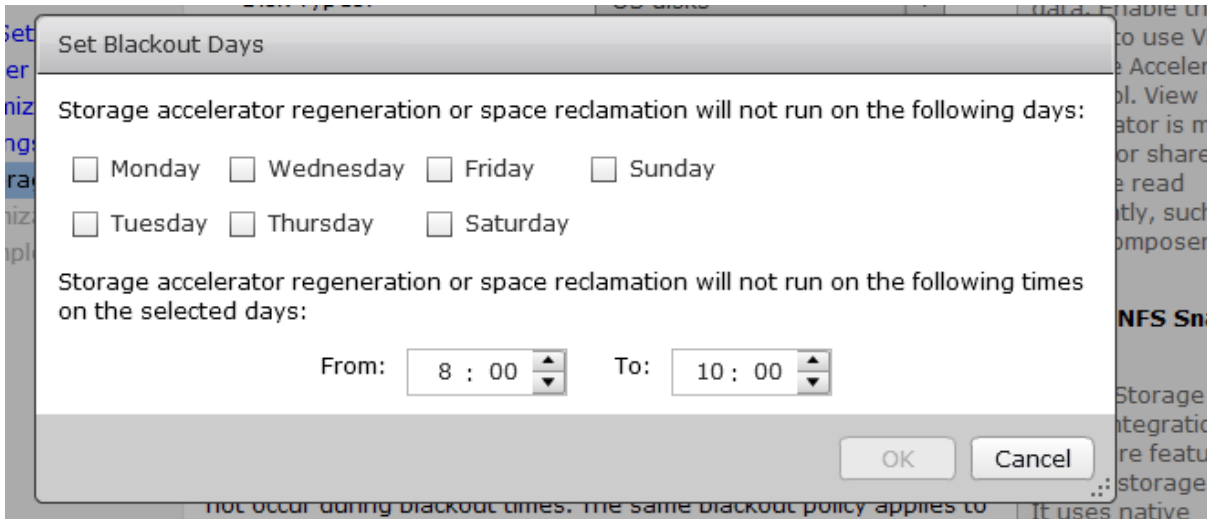


Figure 17

- Once the pool is deployed then monitor the VMs used storage in VC. You can see this in VM in VC under Summary -> Resources -> Used Storage. Make a note of how much it is using. For our example we will show one VM for brevity.

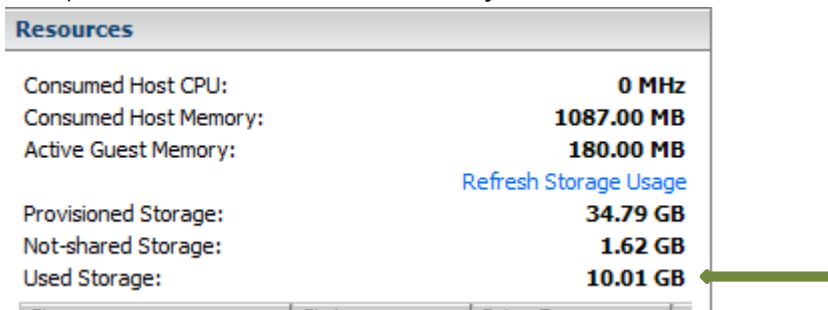


Figure 18

We can also check the Total data store usage for the VDI desktops:




Figure 19

Now, check the data store usage on the Storage Center 60653:

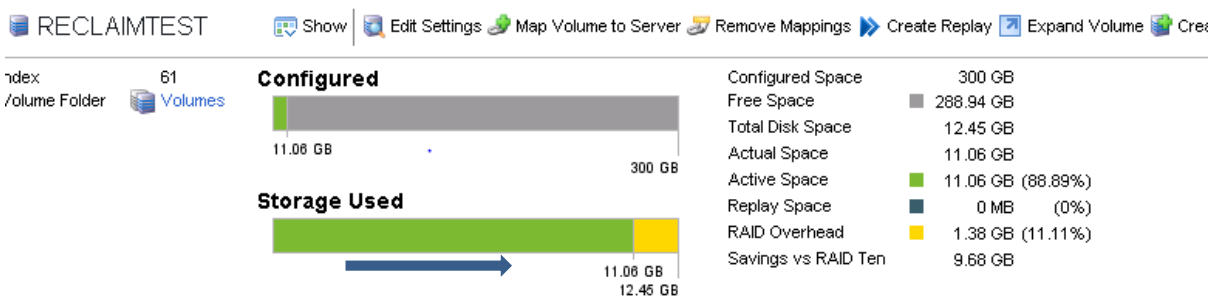


Figure 20

- At this point, large files are copied to each of the virtual desktops.

After copying in the large file(s), monitor the VM usage again, as you can see vm -1 has now increased its used storage:

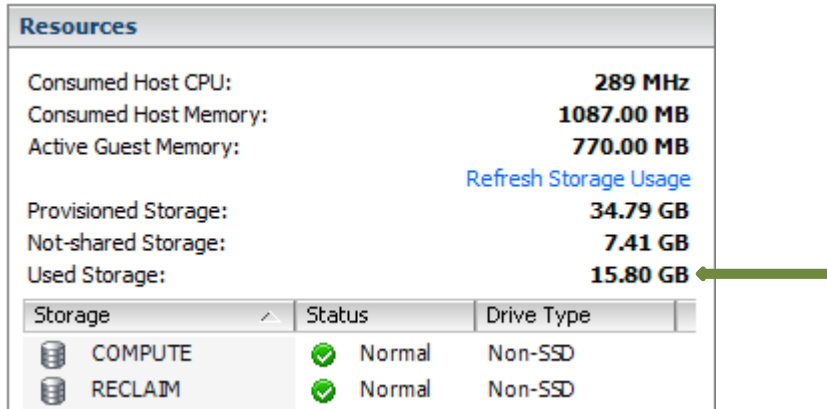


Figure 21

We can also see the increased usage at the data store view level, free space has decreased from 310 to 290 GB (due to 5 desktop file copies):



Figure 22

And the used space has also increased on the Compellent SC8000:

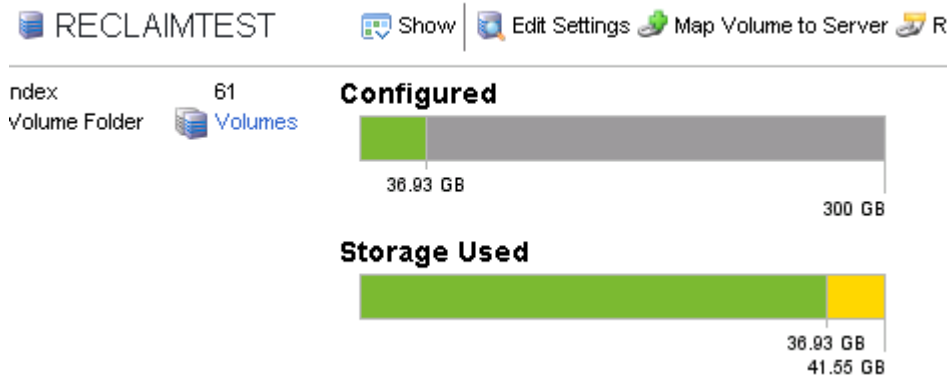


Figure 23

6. Now, delete the large files from the VDI desktops.
7. Re-observe the VMs disk consumption shrink following reclamation. You can see this event occur by either:
  - a) Waiting up to 1 hour for the automated process to kick in
  - b) Kicking the reclamation process of manually using the vmdadmin command

```
C:\Program Files\VMware\VMware View\Server\bin>vmdadmin.exe -M -d View -m vm-1 -markForSpaceReclamation
```

Figure 24

Once the process starts we can see it happening in vSphere events - first a wipe and then a shrink.

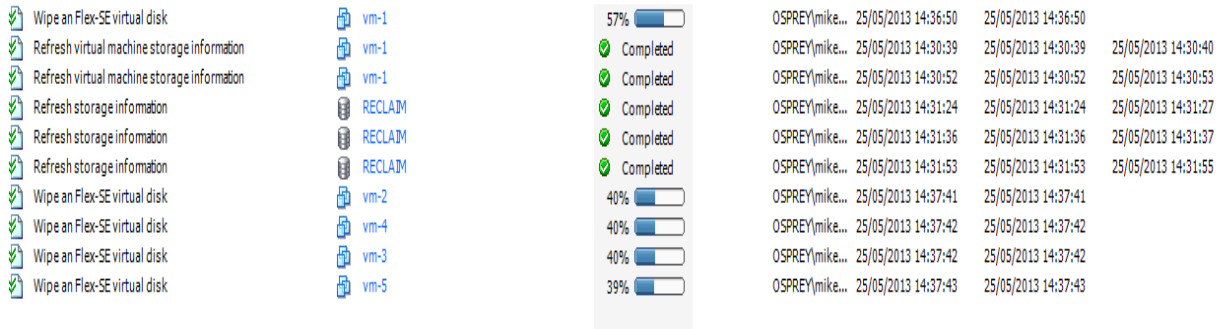


Figure 25

We can see the delete (UNMAP) commands being sent to the array:

DEVICE	CLONE RD	CLONE WR	CLONE F	MBC RD/s	MBC WR/s	ATS	ATSF	ZERO	ZERO F	MBZERO/s	DELETE DF
mpx.vmhba32:C0:T0:L0	0	0	0	0.00	0.00	0	0	0	0	0.00	0
mpx.vmhba36:C0:T0:L0	0	0	0	0.00	0.00	0	0	0	0	0.00	0
naa.6000d31000eced00000000000000003e	6842	6842	0	0.00	0.00	10233	0	7296	0	0.00	0
naa.6000d31000eced00000000000000003f	0	0	0	0.00	0.00	15070	0	17678	0	0.00	50
naa.6d4ae52086a11400191f9e940c76dfe2	0	0	0	0.00	0.00	0	0	0	0	0.00	0

Figure 26

We can also see the event in the view connection broker log:

User	Severity	Time	Module	Message	Objects
	Info	25/05/2013 13:42:44	Agent	The agent on machine VM-3 sent a resume message	2
	Info	25/05/2013 12:46:26	Connection Server	Successfully reclaimed 2.52 GB of space from machine vm-4 in Pool view on datastore RECLAIM	1

Figure 27

In VMware vCenter we can see the increase in the average write requests as the space is reclaimed:




Figure 28

On the Compellent SC8000, we can also see the increase in Total I/O per second and an increase in write latency:




Figure 29

This process does result in higher CPU usage as can be seen here:




Figure 30

After the process has completed, we check the storage used by VM-1. We can see the space has returned to pre-test level:

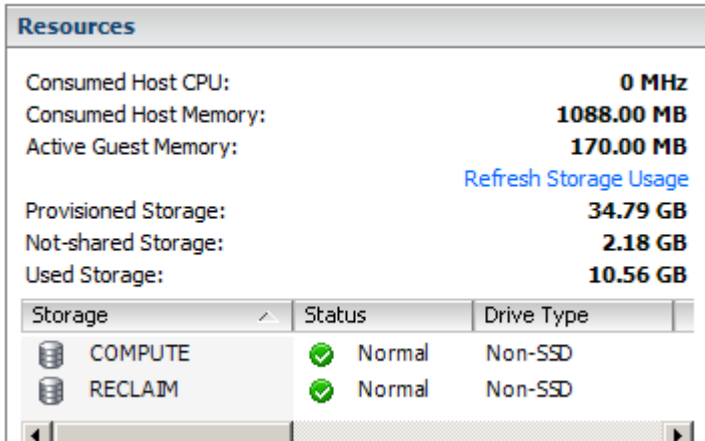


Figure 31

Data store usage has also returned to pre-test levels:




Figure 32

And lastly the Compellent also has its used storage reduced:




Figure 33

You can also view statistics on space reclamation in View administrator by selecting the pool:

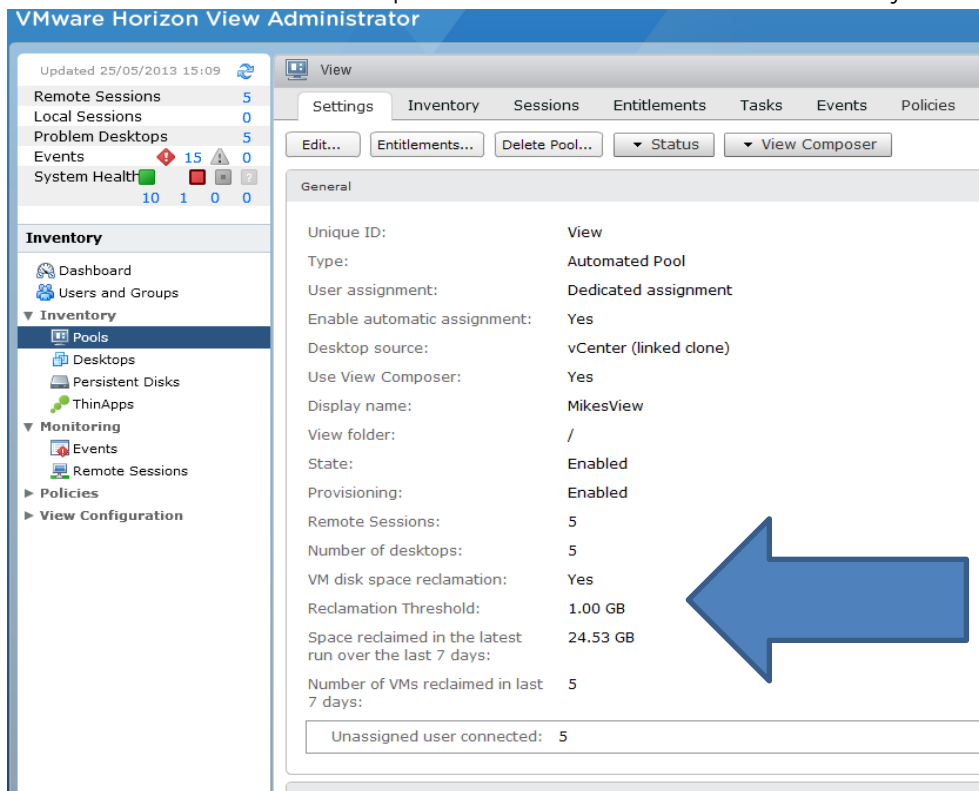


Figure 34

The test has completed as expected and all components involved see reclaimed space.

# VMware View 5.2 UNMAP with EqualLogic PS6110XS

## Components Overview

### Hardware

- Dell PowerEdge R710 servers
- 1x EqualLogic PS6110XS
  - Firmware 6.0.4
- PowerConnect 6248 (Client switch)
- PowerConnect 8024 (iSCSI switch)

### Software

- VMware® ESXi 5.1 Update 1
- VMware® Virtual Center 5.1 Update 1
- VMware® Horizon View 5.2
- Microsoft SQL Server 2008 R2
- Microsoft Windows Server 2008 R2
- EqualLogic SANHQ 2.5

## Procedure

Firstly, make sure that your data store supports the UNMAP primitive using the following command:  
**esxcli storage core device vaa1 status get -d naa.6000d31000eced000000000000000024**

```
~ # esxcli storage core device vaa1 status get -d naa.64ed2ab578036f37d719f50000000a053
naa.64ed2ab578036f37d719f50000000a053
  VAAI Plugin Name: VMW_VAAIP_EQL
  ATS Status: supported
  Clone Status: supported
  Zero Status: supported
  Delete Status: supported
```

Figure 35

As you can see above, the delete status is supported.

Next, create a pool using the following criteria.

1. Create a Windows XP or Windows 7 VM in vSphere using the latest virtual hardware version, install the Tools & View Agent, and snapshot it to serve as a base image for a View Composer based View Desktop.




Figure 36

2. Create a new Automated, Dedicated, and Horizon View Composer linked clone based pool and provision & entitle 1 or more desktops in the pool, this test was performed with 5.
3. Ensure that space reclamation is enabled for the desktop pool.



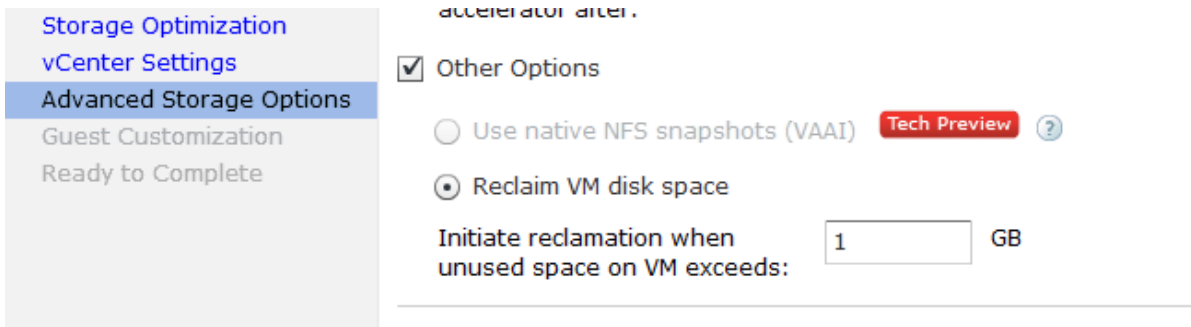


Figure 37

You can select Blackout times on this screen also where you can prevent the VM Space reclamation procedure running at certain times.

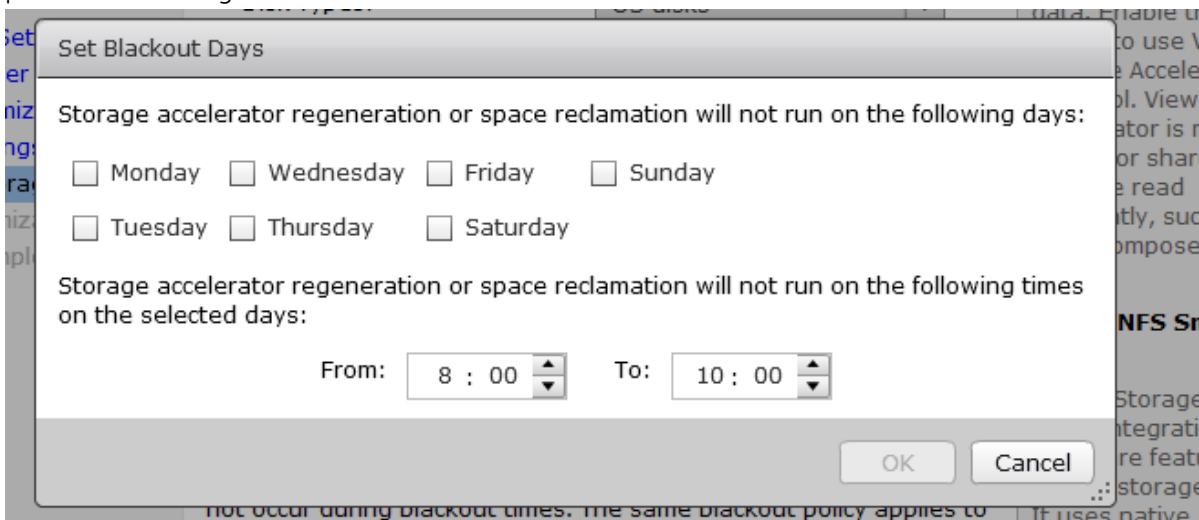


Figure 38

- Once the pool is deployed then monitor the VMs used storage in Virtual Center. You can see this in VM in Virtual Center under Summary -> Resources -> Used Storage. Make a note of how much it is using. For our example we will show one VM for brevity.

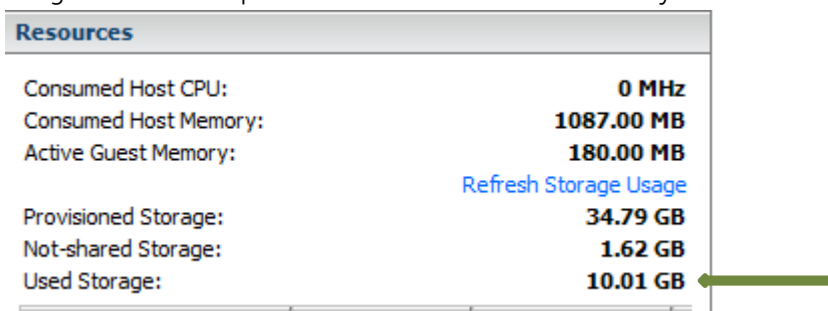


Figure 39

- The EqualLogic LUN is checked for space consumed with the 5 VMs deployed. Approximately 24.2 GB is reported as being used from the array:




Figure 40


6. The VMware vSphere datastore is checked for space usage as well:



**Figure 41**

The file system reports approximately 33.08 GB of used space.

- 7. At this point, large files are copied to each of the virtual desktops.
- 8. Space is checked on the EqualLogic LUN again after the files are copied (Unfortunately, this screenshot is very poor resolution for some reason. Space used was ~45 GB):



**Figure 42**

9. Space is also checked on the VMware vSphere datastore:



**Figure 43**

The file system reports approximately 55.95 GB of used space.

10. The 5 files are deleted from the VMs:
- 11.
12. Re-observe the VMs disk consumption shrink following reclamation. You can see this event occur by either:
  - a) Waiting up to 1 hour for the automated process to kick in
  - b) Kicking the reclamation process of manually using the vdmadmin command

```
C:\Program Files\VMware\VMware View\Server\bin>vdmadmin.exe -M -d View -m vm-1 -markForSpaceReclamation
```

**Figure 44**

13. Once the process starts we can see it happening in vSphere events - first a wipe and then a shrink:

Shrink an Flex-SE virtual disk	Desktop-02	Completed	DVSTEST\svc...	vcenter.dvstest...	5/23/2013 4:53:48 PM
Refresh virtual machine storage information	Desktop-02	Completed	DVSTEST\svc...	vcenter.dvstest...	5/23/2013 4:53:47 PM
Shrink an Flex-SE virtual disk	Desktop-01	Completed	DVSTEST\svc...	vcenter.dvstest...	5/23/2013 4:53:46 PM
Shrink an Flex-SE virtual disk	Desktop-03	Completed	DVSTEST\svc...	vcenter.dvstest...	5/23/2013 4:53:45 PM
Refresh virtual machine storage information	Desktop-01	Completed	DVSTEST\svc...	vcenter.dvstest...	5/23/2013 4:53:45 PM
Shrink an Flex-SE virtual disk	Desktop-05	Completed	DVSTEST\svc...	vcenter.dvstest...	5/23/2013 4:53:44 PM
Refresh virtual machine storage information	Desktop-03	Completed	DVSTEST\svc...	vcenter.dvstest...	5/23/2013 4:53:44 PM
Shrink an Flex-SE virtual disk	Desktop-04	Completed	DVSTEST\svc...	vcenter.dvstest...	5/23/2013 4:53:44 PM
Refresh virtual machine storage information	Desktop-05	Completed	DVSTEST\svc...	vcenter.dvstest...	5/23/2013 4:53:44 PM
Refresh virtual machine storage information	Desktop-04	Completed	DVSTEST\svc...	vcenter.dvstest...	5/23/2013 4:53:43 PM
Wipe an Flex-SE virtual disk	Desktop-05	Completed	DVSTEST\svc...	vcenter.dvstest...	5/23/2013 4:51:53 PM
Wipe an Flex-SE virtual disk	Desktop-04	Completed	DVSTEST\svc...	vcenter.dvstest...	5/23/2013 4:51:53 PM
Wipe an Flex-SE virtual disk	Desktop-03	Completed	DVSTEST\svc...	vcenter.dvstest...	5/23/2013 4:51:52 PM
Wipe an Flex-SE virtual disk	Desktop-02	Completed	DVSTEST\svc...	vcenter.dvstest...	5/23/2013 4:51:52 PM
Wipe an Flex-SE virtual disk	Desktop-01	Completed	DVSTEST\svc...	vcenter.dvstest...	5/23/2013 4:51:51 PM

**Figure 45**

14. After the process has completed, we can check the logs on the View Connection Server:

Successfully reclaimed 4.08 GB of space from machine Desktop-05 in Pool test on datastore vdi-1

Successfully reclaimed 4.04 GB of space from machine Desktop-01 in Pool test on datastore vdi-1


Successfully reclaimed 4.23 GB of space from machine Desktop-03 in Pool test on datastore vdi-1

Successfully reclaimed 4.15 GB of space from machine Desktop-04 in Pool test on datastore vdi-1

Successfully reclaimed 4.08 GB of space from machine Desktop-02 in Pool test on datastore vdi-1


**Figure 46**

15. After the process has completed, we then check the storage used in the datastore. We can see the space has returned to pre-test levels:




**Figure 47**

16. And lastly the EqualLogic also has its used storage reduced:




**Figure 48**

17. Performance during the file copy process can be reviewed through VMware Virtual Center (SCSI cmds/s). The files were copied starting at ~3:35 PM and deleted shortly after 4 PM in the screenshot below:



**Figure 49**

18. And during the UNMAP operation (SCSI cmds/s):



Performance Chart Legend

Key	Object	Measurement	Rollup	Units	Latest	Maximum	Minimum	Average
■	EQLOGIC iSCSI Disk (naa.64ed2ab578036f37d719f500000a053)	Average commands issued per second	Average	Number	7	10204	0	501.478
■	EQLOGIC iSCSI Disk (naa.64ed2ab578036f37d719f500000a053)	Average read requests per second	Average	Number	0	4953	0	187.2
■	EQLOGIC iSCSI Disk (naa.64ed2ab578036f37d719f500000a053)	Commands issued	Summation	Number	146	204097	8	10038.839
■	EQLOGIC iSCSI Disk (naa.64ed2ab578036f37d719f500000a053)	Maximum queue depth	Average	Number	128	128	128	128
■	EQLOGIC iSCSI Disk (naa.64ed2ab578036f37d719f500000a053)	Average write requests per second	Average	Number	7	6261	0	313.322

Figure 50

19. VMware Virtual Center LUN Latency during UNMAP operation:




Figure 51

20. Additionally, performance impact can be noted via SANHQ:




Figure 52

The test has completed as expected and all components involved see reclaimed space.

## Conclusion

The shrink / wipe and UNMAP feature in Horizon View 5.2 works as expected and is a very useful feature. However, the procedure does increase CPU usage and also results in increased I/O to the storage array.

We recommend using this feature during the blackout period so it will not result in performance degradation during peak usage in your environment.

## References

<http://www.vmware.com/files/pdf/techpaper/Whats-New-VMware-vSphere-51-Storage-Technical-Whitepaper.pdf>

<http://blogs.vmware.com/euc/2013/03/space-efficient-virtual-disks-in-practice.html#more-2817>

<http://www.vmware.com/files/pdf/view/VMware-View-Evaluators-Guide.pdf>

### 13.6.2 VMware vSphere Virtual Distributed Switches (VDS)

#### Executive Summary


VMware vSphere Distributed Switches (VDS) were introduced with VMware vSphere 4.0 as a way to centrally manage and maintain virtual switching infrastructure in a largely distributed environment. In addition to simplifying management of the networking configuration, VMware VDS provide additional advanced capabilities that VMware vSphere Standard vSwitches do not. These features include, shaping inbound (RX) traffic, Private VLANs, (PVLANs), and customization of Data and Control Planes. vSphere 5.x also provides the additional improvements of increased visibility of inter-virtual machine traffic through Netflow, improved monitoring through port mirroring (dvMirror), and support for LLDP (Link Layer Discovery Protocol), which is supported on Dell PowerConnect and Force10 switches.

The testing involved qualifying the VMware vSphere Distributed Switch functionality with regards to VMware Horizon View 5.2 VMs. The testing was to ensure that the feature worked as expected, but does not extensively study performance impact or interoperability with other products or features.

#### Technology Deep Dive

VMware vSphere Distributed Switches (VDS) extends the features and capabilities of virtual networks while simplifying provisioning and the ongoing process of configuration, monitoring, and management.

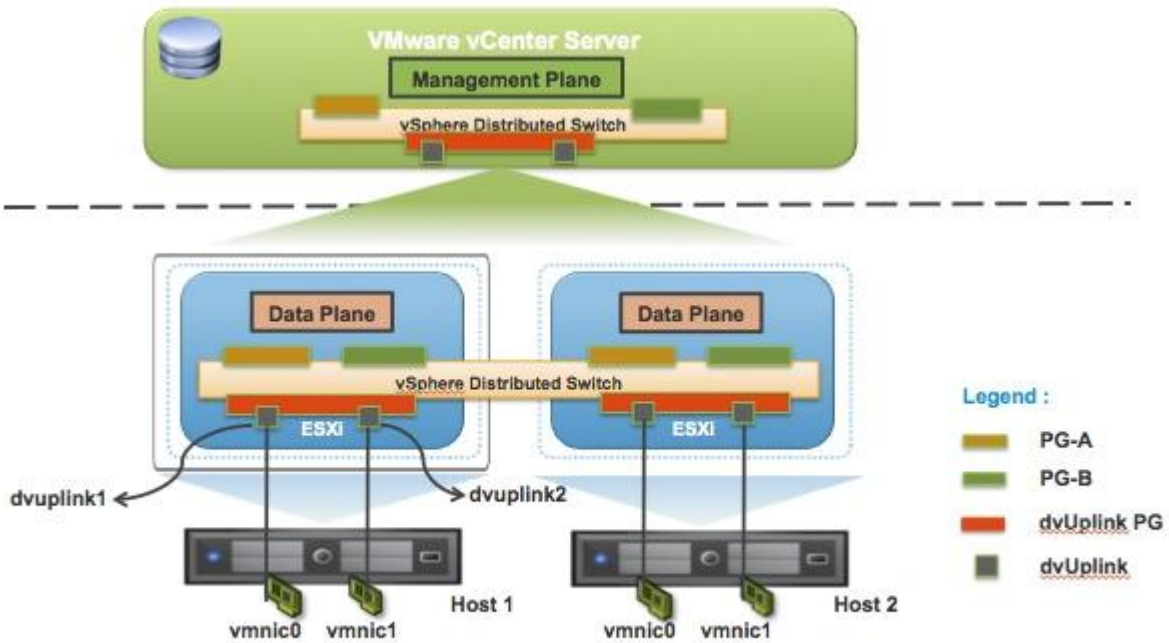
vSphere network switches can be broken into two logical sections. These are the data plane and the management plane. The data plane implements the actual packet switching, filtering, tagging, etc. The management plane is the control structure used to allow the operator to configure the data plane functionality. With the vSphere Standard Switch (VSS), the data plane and management plane are each present on each standard switch. In this design, the administrator configures and maintains each VSS on an individual basis.



**Management Plane :** Allows to configure various parameters of the virtual switch  
**Data Plane :** Handles the packet switching function

Figure 53

With the release of vSphere 4.0, VMware introduced the vSphere Distributed Switch. VDS eases the management burden of per host virtual switch configuration by treating the network as an aggregated resource. Individual host-level virtual switches are abstracted into a single large VDS that spans multiple hosts at the Datacenter level. In this design, the data plane remains local to each VDS, but the management plane is centralized with vCenter Server acting as the control point for all configured VDS instances.



**One Management Plane :** Allows to configure various parameters of the distributed switch  
**Data Plane :** Handles the packet switching function

Figure 54



Each vCenter Server instance can support up to 128 VDSs and each VDS can manage up to 500 hosts. Many of the concepts involved in configuring and managing a Standard Switch are carried forward with the VDS, with changes made to enable managing multiple switches.

**Distributed Virtual Port Groups (DV Port Groups)** are port groups associated with a VDS and specify port configuration options for each member port. DV Port Groups define how a connection is made through the VDS to the Network. Configuration parameters are similar to those available with Port Groups on Standard Switches. The VLAN ID, traffic shaping parameters, port security, teaming and load balancing configuration, and other settings are configured here. Each VDS supports up to 10000 static port groups.

**Distributed Virtual Uplinks (dvUplinks)** are a new concept introduced with VDS. dvUplinks provide a level of abstraction for the physical NICs (vmnics) on each host. NIC teaming, load balancing, and failover policies on the VDS and DV Port Groups are applied to the dvUplinks and not the vmnics on individual hosts. Each vmnic on each host is mapped to a dvUplinks, permitting teaming and failover consistency irrespective of vmnic assignments.

### **Private VLANs**

Private VLAN (PVLAN) support enables broader compatibility with existing networking environments using Private VLAN technology. Private VLANs enable users to restrict communication between virtual machines on the same VLAN or network segment, significantly reducing the number of subnets needed for certain network configurations.

### **Network vMotion**

Network vMotion is the tracking of virtual machine networking state (e.g. counters, port statistics) as the virtual machine moves from host to host on a VDS. This provides a consistent view of a virtual network interface regardless of the VM location or vMotion migration history. This greatly simplifies network monitoring and troubleshooting activities where vMotion is used to migrate VMs between hosts.

### **Bi-directional Traffic Shaping**

VDS expands upon the egress only traffic shaping feature of Standard Switches with bi-directional traffic shaping capabilities. Egress (from virtual machine to network) and now ingress (from network into virtual machine) traffic shaping policies can now be applied on DV Port Group Definitions. Traffic shaping is useful in cases where you may wish to limit the traffic to or from a VM or group of VMs to either protect a VM or other traffic in an oversubscribed network. Policies are defined by three characteristics: average bandwidth, peak bandwidth, and burst size.

Figure 3 provides a comparison of the features supported with a standard and distributed vSwitch:

	VMware vSphere vSwitch	VMware vSphere DVS
<b>Switching Features</b>		
Associate VMs to Network	✓	✓
CDP v1/v2	✓	✓
IEEE 802.1Q VLAN Trunking	✓	✓
L2 Forwarding	✓	✓
<b>Management Features</b>		
Multicast Support	✓	✓
Netflow v5	✓	✓
Network Monitoring	✓	✓
Network Provisioning	✓	✓
Network VMotion (Port state follows VM)		✓
NIC Teaming/Port Channels	✓	✓
Port Profile	✓	✓
Private VLAN		✓
Management Location	vSphere ESXi Host	vSphere vCenter
<b>System Features</b>		
VLAN Segmentation	✓	✓

VM Rx Rate Limiting		✓
VM Tx Rate Limiting	✓	✓
VMsafe Compatibility	✓	✓
VMware Port Mirroring (Promiscuous)	✓	✓
VMware Port Security	✓	✓
VMware Remote CLI	✓	✓

**Figure 55**

## Prerequisites

This document assumes that you already have a working VMware® Virtual Center Cluster with VMware Horizon View 5.2 installed and configured and you have working knowledge of both.

In addition to the above requirements, VMware vSphere Enterprise Plus licensing is required:

	VSPHERE WITH OPERATIONS MANAGEMENT		
	Standard	Enterprise	Enterprise Plus
<b>Entitlements per CPU License</b>			
vCPU/VM	8-way	32-way	64-way
<b>Features</b>			
Health Monitoring and Performance Analytics	•	•	•
Capacity Management and Optimization	•	•	•
Operations Dashboard and Root Cause Analysis	•	•	•
Hypervisor	•	•	•
vMotion®	•	•	•
Storage vMotion	•	•	•
High Availability and Fault Tolerance (1 vCPU)	•	•	•
Data Protection™ and Replication	•	•	•
vShield Endpoint™	•	•	•
Distributed Resource Scheduler™ and Distributed Power Management™		•	•
Storage APIs for Array Integration, Multipathing		•	•
Distributed Switch™			•
Storage DRS™ and Profile-Driven Storage			•
I/O Controls (Network and Storage) and SR-IOV			•
Host Profiles and Auto Deploy			•

Figure 56

# VMware vSphere Distributed Switch Test Results

## Components Overview

### Hardware

- Dell PowerEdge R710 servers
- 1x EqualLogic PS6110XS
  - Firmware 6.0.4
- PowerConnect 6248 (Client switch)
- PowerConnect 8024 (iSCSI switch)

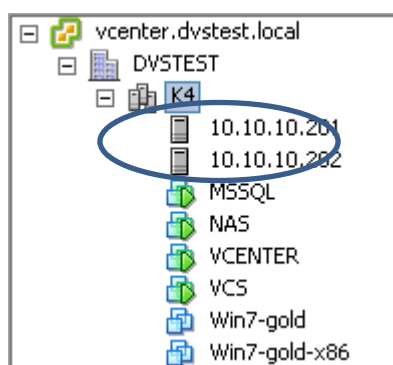
### Software

- VMware® ESXi 5.1 Update 1
- VMware® Virtual Center 5.1 Update 1
- VMware® Horizon View 5.2
- Microsoft SQL Server 2008 R2
- Microsoft Windows Server 2008 R2
- EqualLogic SANHQ 2.5

### Procedure

In the following procedure, the VMware vSphere 5.1 and Horizon View 5.2 environment have already been set up and configured using standard vSwitches. The VDI vSwitch is then migrated to a vSphere Distributed Switch and the uplink ports are migrated.

- 1) In VMware vSphere, two ESXi servers are placed into a cluster for ease of management:




- 2) The existing network was configured with 3 Standard Switches. vSwitch1 and all of the associated virtual machines will be migrated to a VMware vSphere Distributed Switch in this procedure.


**View:** vSphere Standard Switch vSphere Distributed Switch

### Networking

Standard Switch: vSwitch0 [Remove...](#) [Properties...](#)



Standard Switch: vSwitch1 [Remove...](#) [Properties...](#)



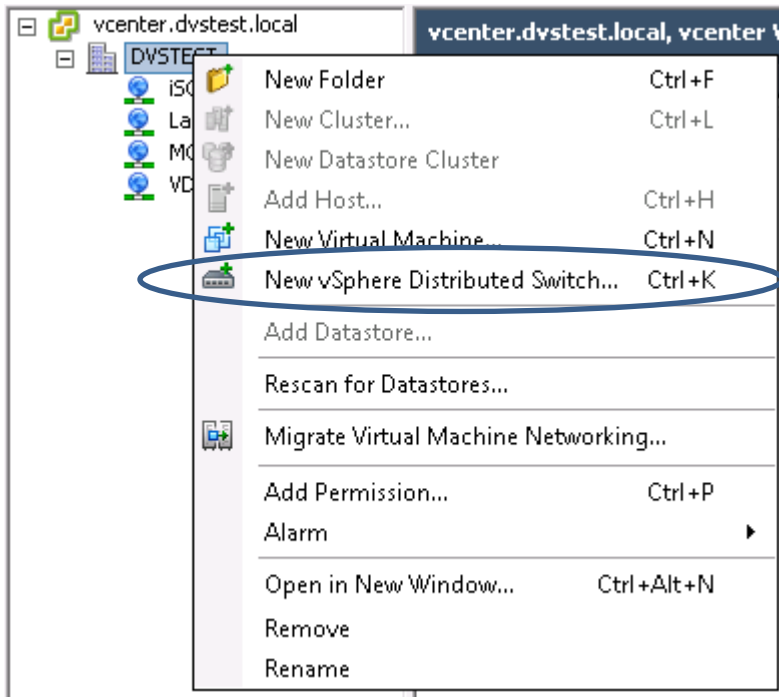
Standard Switch: vSwitch2 [Remove...](#) [Properties...](#)



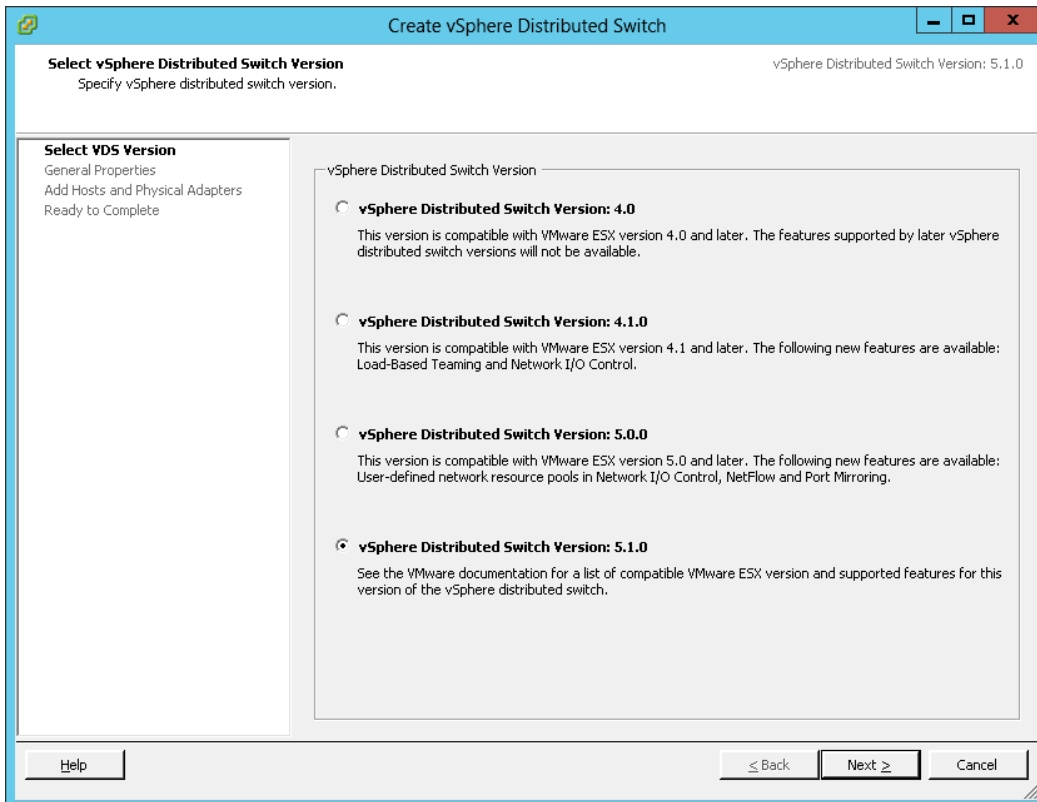
- 3) In VMware vCenter, go to Home -> Networking

Network	Type	Status	Alarm Actions	Number of hosts	Number of VMs
Lab Network	Standard port group	Normal	Enabled	1	2
MGMT Network	Standard port group	Normal	Enabled	3	4
VDI Network	Standard port group	Normal	Enabled	3	10
iSCSI Network	Standard port group	Normal	Enabled	1	1

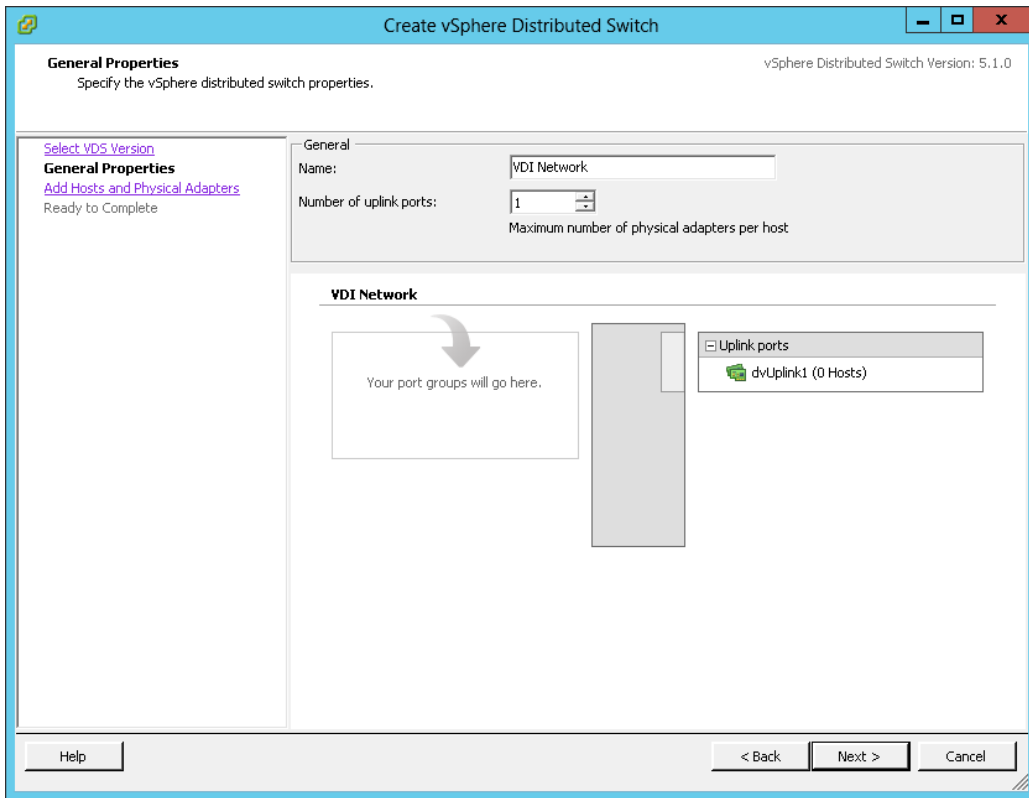
4) Right click on the Datacenter (DVSTEST here) and click "New vNetwork Distributed Switch"



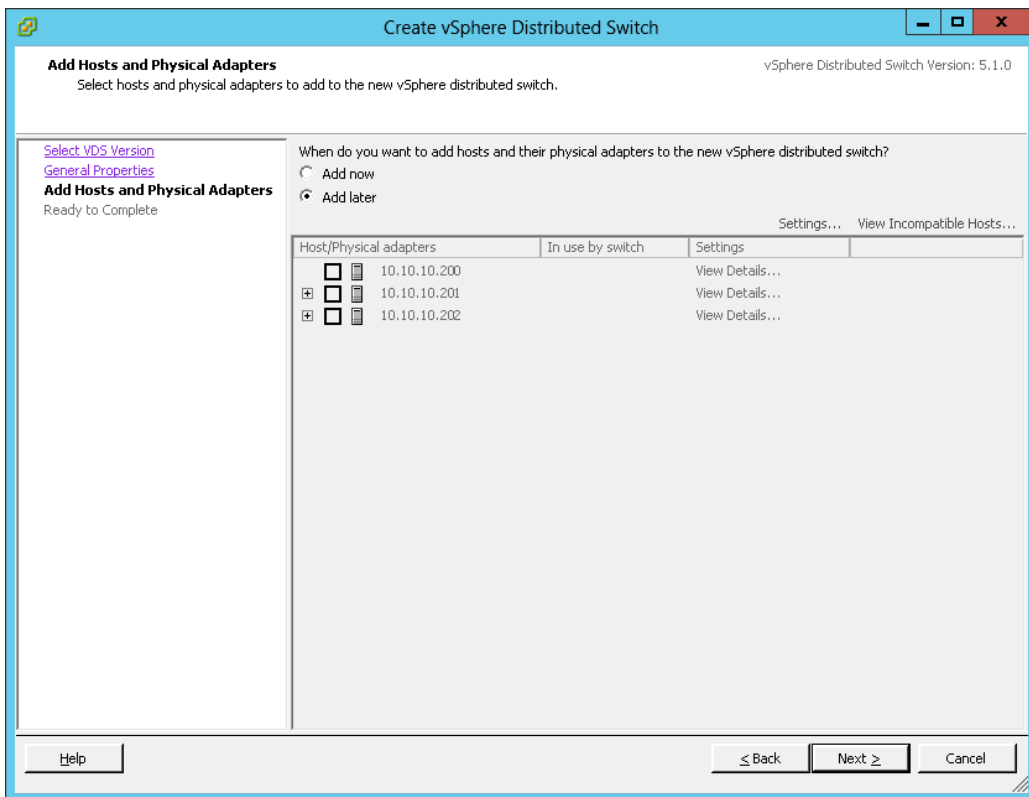
5) Choose the appropriate vSphere DVS version and click Next



6) Provide a name for the DVS, set the number of uplink ports, and click Next.

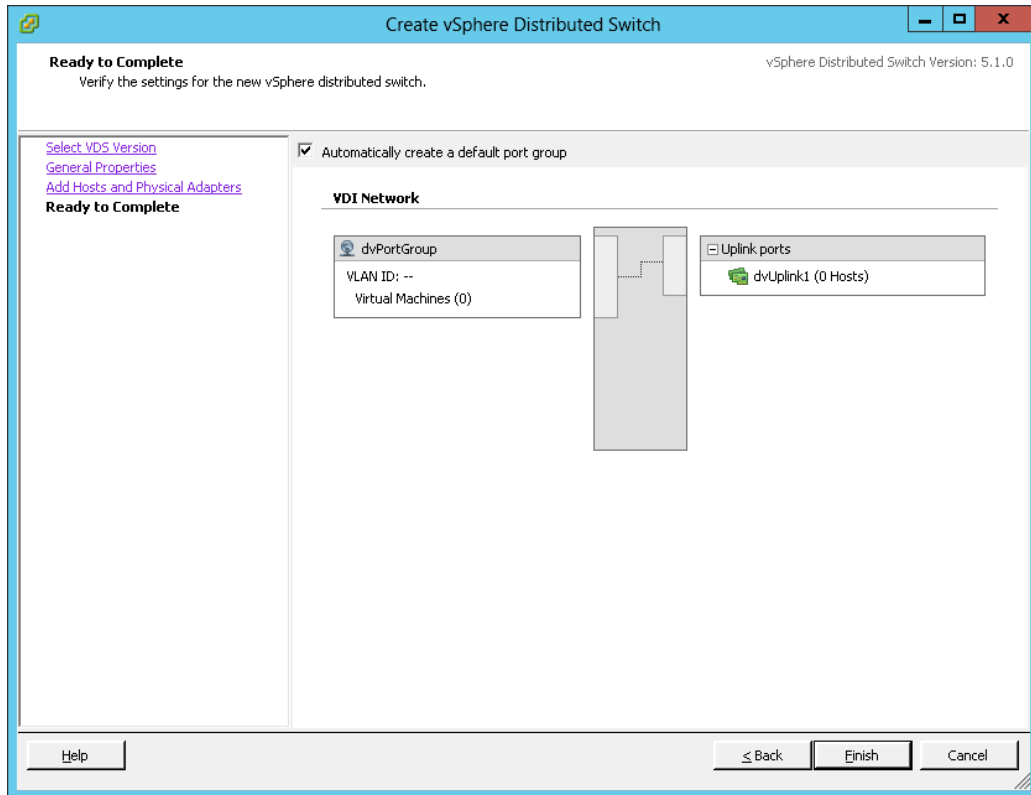


7) Select "Add later" and click Next.




8) Check settings and click "Finish".

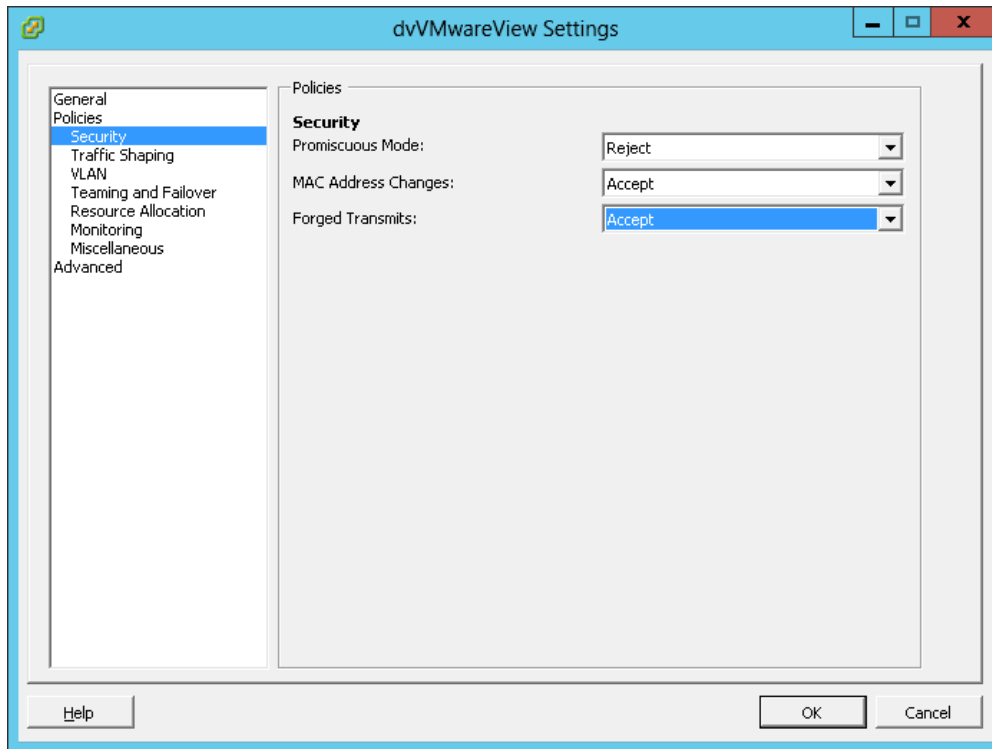




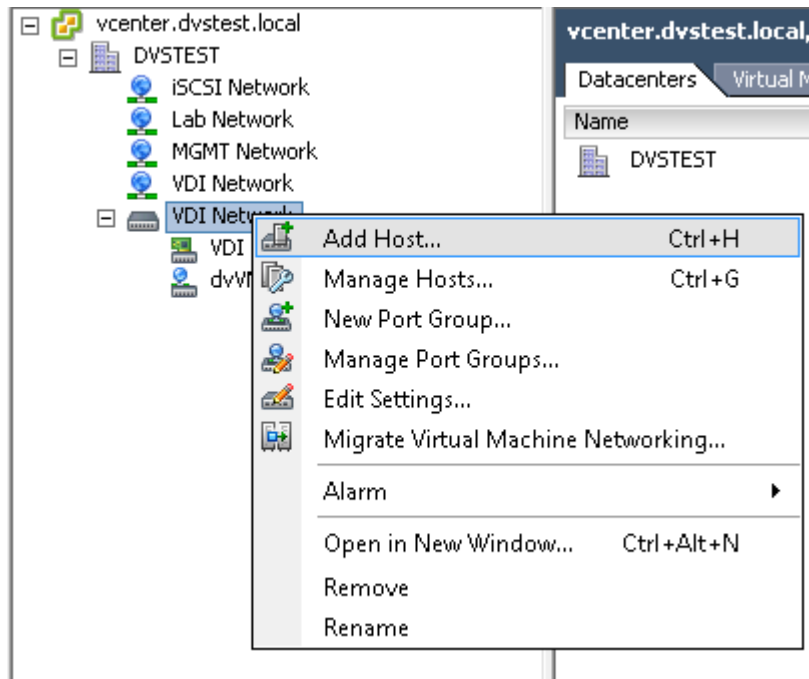
9) There should now be a Distributed vSwitch and port group.



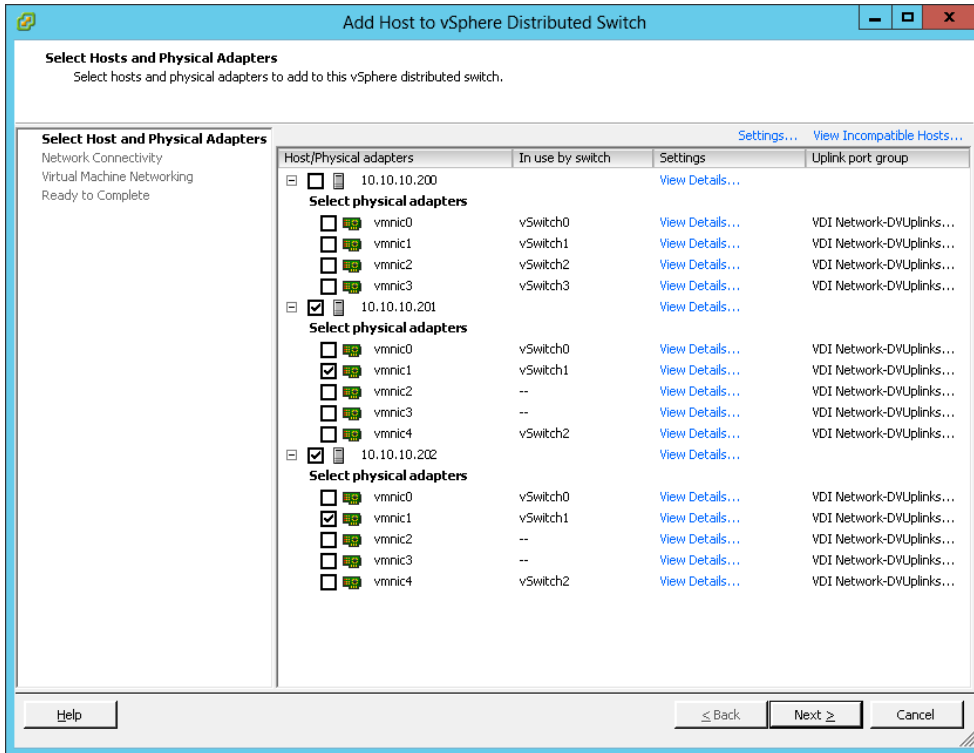
10) Right click the Distributed port group and edit settings.



11) Right click the DVS and select "Add Host...".

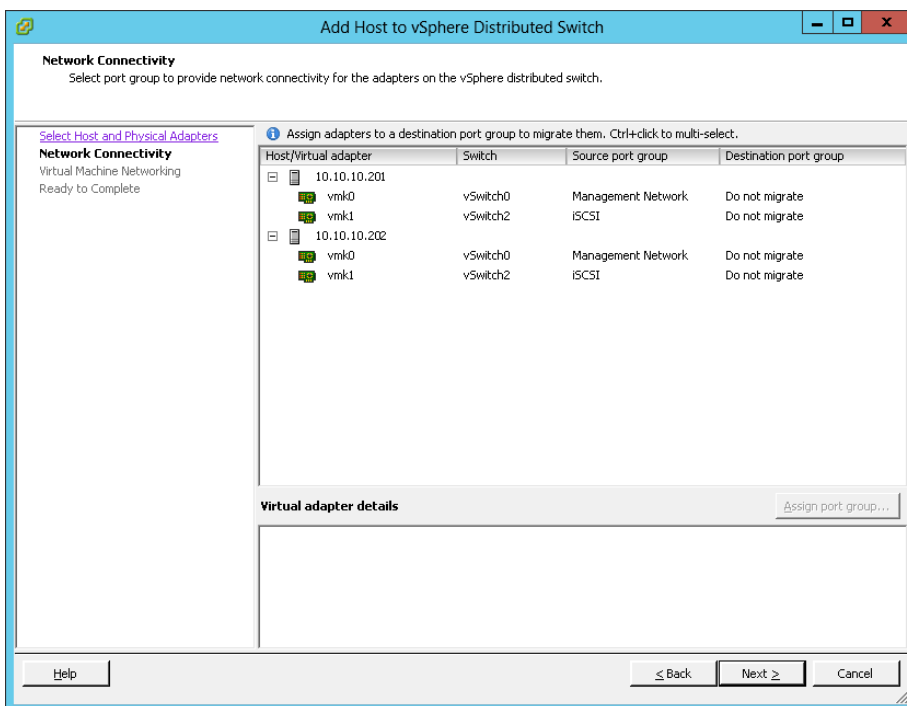


12) Select the ESXi Servers and the physical adapters that will be migrated to the VDS.

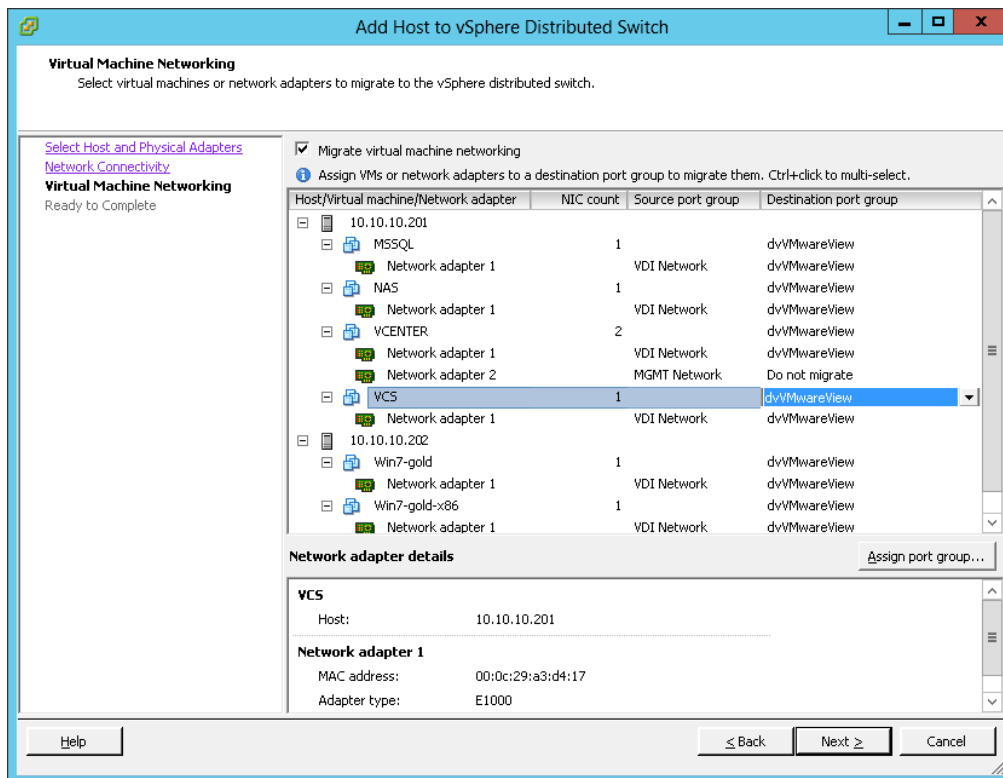


\*\* Note: Make sure that the vCenter, SQL, and View servers are communicating on another management network, otherwise the management infrastructure will go down. There are special procedures with redundant networking and uplinks to migrate management interfaces that are available in VMware best practices guides and KB articles.

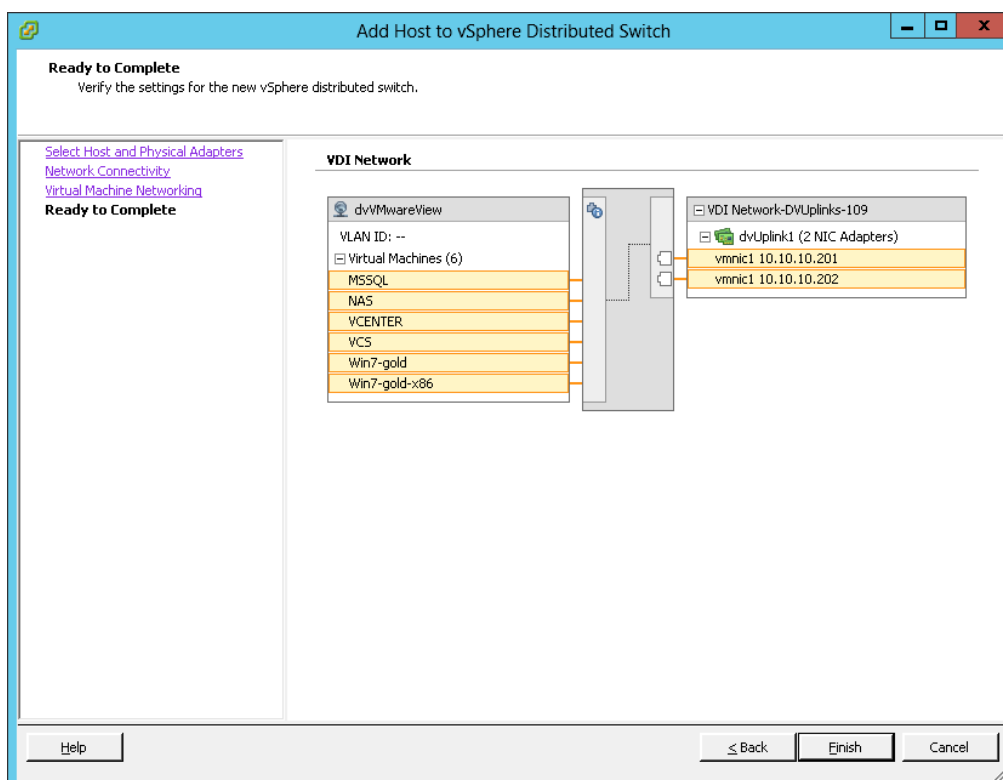
13) Click Next to leave the Management and iSCSI networks alone



14) Select the "Migrate virtual machine networking" and ensure that everything on the old vSwitch is migrated to the DVS. Be careful not to select NICs that were on other networks.



15) Click Finish



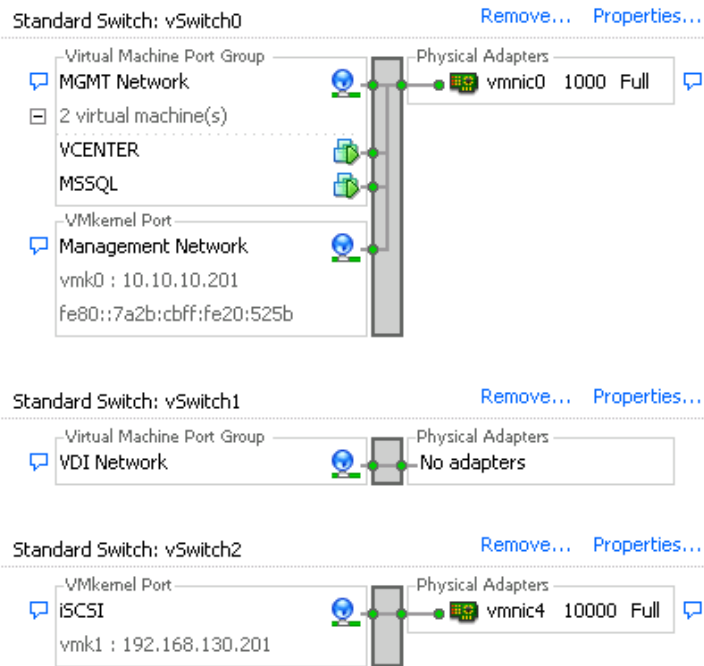
16) The process can be monitored in the "Recent Tasks" pane at the bottom of the window.

Name	Target	Status	Details	Initiated by	vCenter Server	Requested Start Time	Start Time	Completed Time
Reconfigure vSphere Distributed Switch	VDI Network	Completed		DVSTESTYnc...	vcenter.dvstest...	5/29/2013 2:18:10 PM	5/29/2013 2:18:10 PM	5/29/2013 2:18:15 PM
Reconfigure Distributed Port Group	dvVMwareView	Completed		DVSTESTYnc...	vcenter.dvstest...	5/29/2013 2:11:44 PM	5/29/2013 2:11:44 PM	5/29/2013 2:11:44 PM
Reconfigure Distributed Port Group	dvVDI	Completed		DVSTESTYnc...	vcenter.dvstest...	5/29/2013 2:09:59 PM	5/29/2013 2:09:59 PM	5/29/2013 2:09:59 PM
Reconfigure Distributed Port Group	dvPortGroup	Completed		DVSTESTYnc...	vcenter.dvstest...	5/29/2013 2:09:45 PM	5/29/2013 2:09:45 PM	5/29/2013 2:09:45 PM
Add Distributed Port Groups	VDI Network	Completed		DVSTESTYnc...	vcenter.dvstest...	5/29/2013 2:08:25 PM	5/29/2013 2:08:25 PM	5/29/2013 2:08:25 PM
Create a vSphere Distributed Switch	DVSTEST	Completed		DVSTESTYnc...	vcenter.dvstest...	5/29/2013 2:08:24 PM	5/29/2013 2:08:24 PM	5/29/2013 2:08:24 PM

17) The ESXi server NIC configuration will now look like this.

**View:** vSphere Standard Switch vSphere Distributed Switch

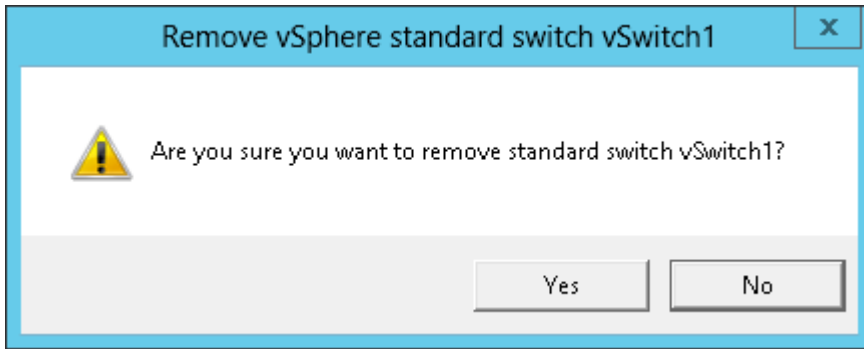
**Networking**



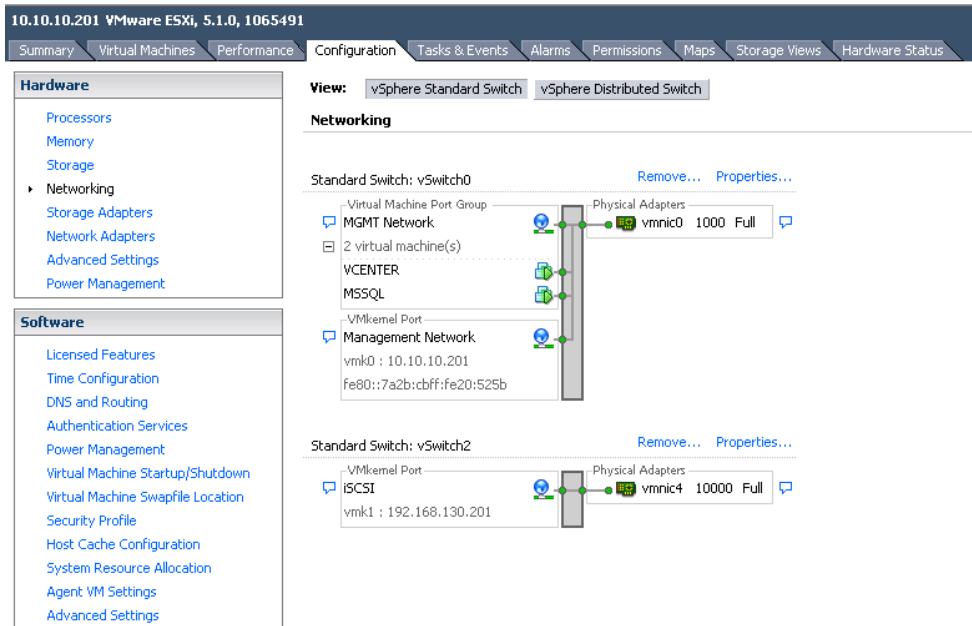
The screenshot shows the ESXi configuration page for vSwitch1 (VDI Network) under the Networking section. The left sidebar lists Hardware and Software settings. The main content area shows:

- View:** vSphere Standard Switch vSphere Distributed Switch
- Networking:** Distributed Switch: VDI Network (Manage Virtual Adapters... Manage Physical Adapters... Properties...)
- VDI Network:**
  - dvVMwareView (VLAN ID: --) containing Virtual Machines (4): MSSQL, NAS, VCENTER, VCS.
  - VDI Network-DVUplinks-109 containing dvUplink1 (1 NIC Adapter) with vmnic1 10.10.10.201.

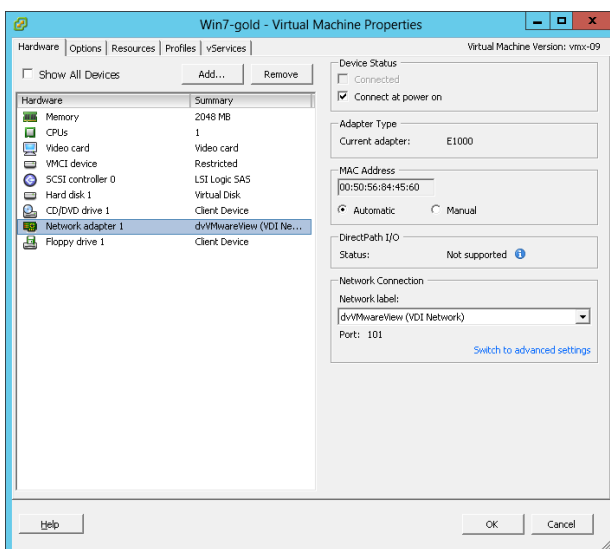
18) Now, we are going to remove the unused standard vSwitch.



19) The Standard Switch networking looks as follows:



20) Check the gold image to make sure the Network Connection has been migrated.

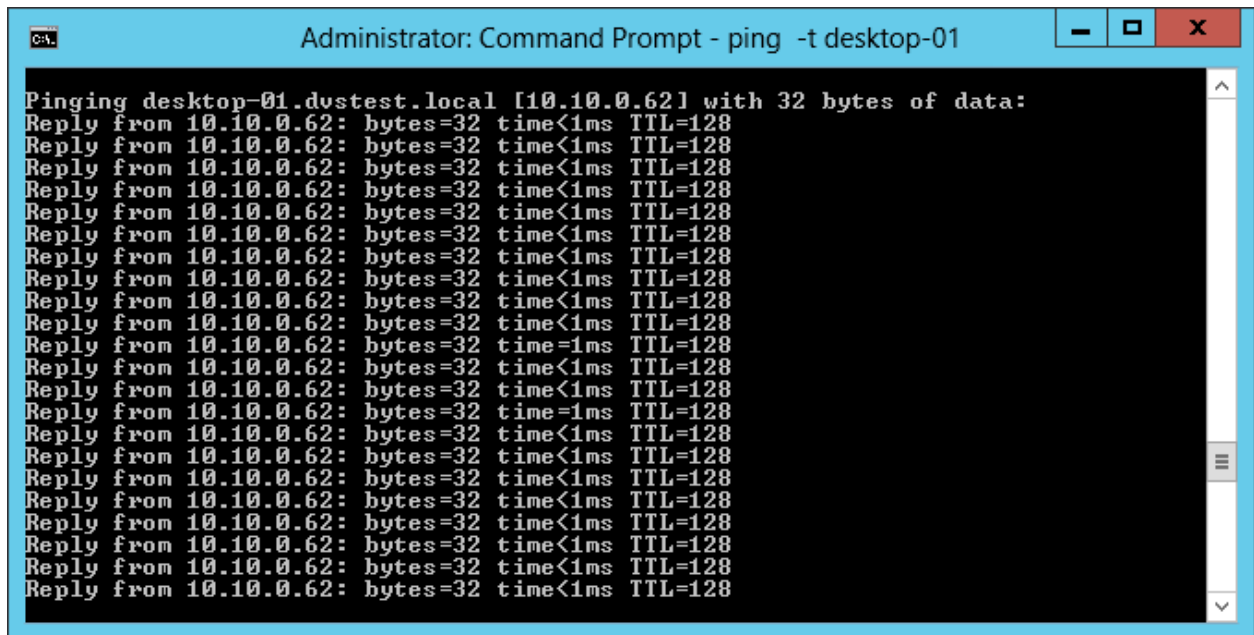


21) Take a new base snapshot for the gold image and deploy a new pool of Floating View Composer linked clones. Alternately, if there are existing VMs in the environment, their networking would have already been migrated. The new snapshot is very important though.

## Validation

After 10 VMs had been deployed using the DVS, the each VM was logged into using domain credentials to validate network connectivity.

In addition to logging into each VM, the VMs were each pinged from a management server that was on a separate ESXi server and vSwitch configuration. Each VM could ping and responded in less than 1ms as seen in Figure 5.



```
Administrator: Command Prompt - ping -t desktop-01
Pinging desktop-01.dvstest.local [10.10.0.62] with 32 bytes of data:
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time=1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
```

Figure 57

## Conclusion

The VMware vSphere Distributed Switch worked as expected with Horizon View 5.2 VMs and brings many features and enhancements that may be useful to customers in a variety of computing environments. In addition to the features and enhancement updates, customers and systems administrators may want to use Distributed Switches for ease of management across a large number of ESXi hosts and clusters.

## References

[http://www.vmware.com/files/pdf/vsphere\\_pricing.pdf](http://www.vmware.com/files/pdf/vsphere_pricing.pdf)

<http://www.vmware.com/products/datacenter-virtualization/vsphere/distributed-switch.html>

<http://www.vmware.com/files/pdf/techpaper/Whats-New-VMware-vSphere-51-Network-Technical-Whitepaper.pdf>

### 13.6.1 VMware vSphere Storage and Regular vMotion for Management Tiers

#### Executive Summary

VMware vSphere® vMotion presents a unified migration architecture that migrates live virtual machines, including their memory and storage, between vSphere hosts. With the introduction of VMware vSphere 5.1, vMotion also does not have any requirement for shared storage. Previously, VMs would have to be turned off to perform the same task. This shared-nothing live migration feature offers administrators significantly more simplicity and flexibility in managing and moving virtual machines across their virtual infrastructures compared to the traditional vMotion and Storage vMotion migration solutions.

A series of tests were conducted to investigate and validate the implications of vSphere 5.1 vMotion when used with VMware Horizon View 5.2 in various scenarios including the migration of the View Connection Server and a VMware View Composer VM. The testing was to ensure that the vMotion features worked as expected, but does not extensively study performance impact or interoperability with other products or features

Test results show the following:

- All tests worked as expected and migrated the appropriate resources as requested
- There were very minor interruptions in access to the VM during the migrations




## Technology Deep Dive

VMware vMotion and Storage vMotion are key, widely adopted technologies which enable the live migration of virtual machines on the vSphere platform. vMotion provides the ability to live migrate a virtual machine from one vSphere host to another host, with no perceivable impact to the end user. Storage vMotion technology provides the ability to live migrate the virtual disks belonging to a virtual machine across storage elements on the same host. Together, vMotion and Storage vMotion technologies enable critical datacenter workflows, including automated load-balancing with DRS and Storage DRS, hardware maintenance, and the permanent migration of workloads.

vSphere 5.1 vMotion transfers the entire execution state of a running virtual machine from the source VMware vSphere® ESXi™ host to the destination ESXi host over a high speed network. The execution state primarily consists of the following components:

- The virtual machine's virtual disks
- The virtual machine's physical memory
- The virtual device state, including the state of the CPU, network and disk adapters, SVGA , and so on
- External network connections



VMware vMotion moves live, running virtual machines from one host to another while maintaining continuous service availability.

Figure 58

## Prerequisites

This document assumes that you already have a working VMware® Virtual Center Cluster with VMware Horizon View 5.2 installed and configured and you have working knowledge of both.

Additionally the following components are required for VMware vMotion to work correctly:

- VMware ESXi servers must have compatible CPUs (or use EVC).
- VMware ESXi servers must have consistent network and network labels.
- The VMware vSphere Web Client must be installed.
- The hosts must be licensed for VMware vSphere vMotion.
  - VMware vSphere license must be Essentials Plus or higher for vMotion.
  - VMware vSphere license must be Standard or Higher for Storage vMotion.
- Virtual machine disks must be in persistent mode or be raw device mappings (RDMs).
- The destination host must have access to the destination storage.
- On each host, a VMkernel port group for vMotion must be configured.
  - This should be a port group on the same network subnet
- Ensure that virtual machines have access to the same subnets on source and destination hosts.
- Virtual machines that are attached to a virtual intranet with vMotion cannot be migrated, even if the destination host has a virtual intranet configured with the same network label.

## VMware vMotion Test Results

In the following procedure, the VMware vSphere 5.1 and Horizon View 5.2 environment have already been set up and configured using standard VMware vSwitches for the management and iSCSI network connectivity and VMware Distributed vSwitches (DVS) for the VMware View VMs. In all tests, Desktop-01 VM was used to test the vMotion capabilities as well as the View Connection Server VM. For brevity, the Desktop-01 VM was shown in the procedures within this document.

## Components Overview

### Hardware

- Dell PowerEdge R710 servers
- 1x EqualLogic PS6110XS
  - Firmware 6.0.4
- PowerConnect 6248 (Client switch)
- PowerConnect 8024 (iSCSI switch)

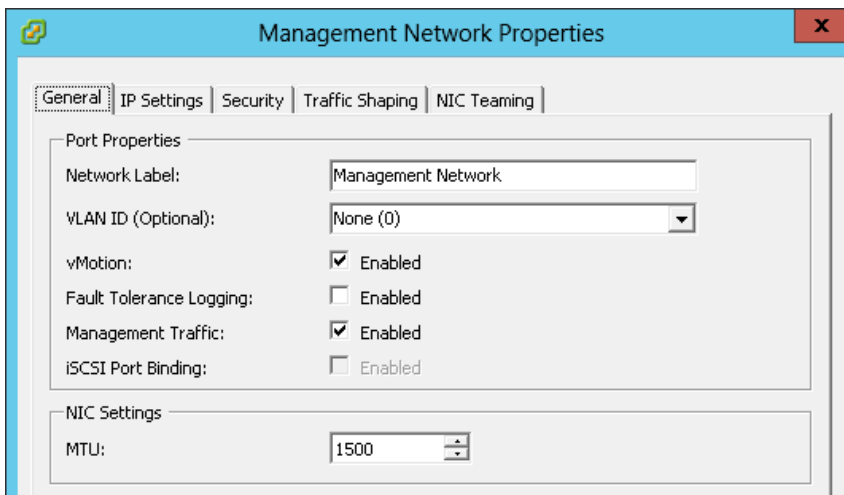
### Software

- VMware® ESXi 5.1 Update 1
- VMware® Virtual Center 5.1 Update 1
- VMware® Horizon View 5.2
- Microsoft SQL Server 2008 R2
- Microsoft Windows Server 2008 R2
- EqualLogic SANHQ 2.5

## VMware vMotion

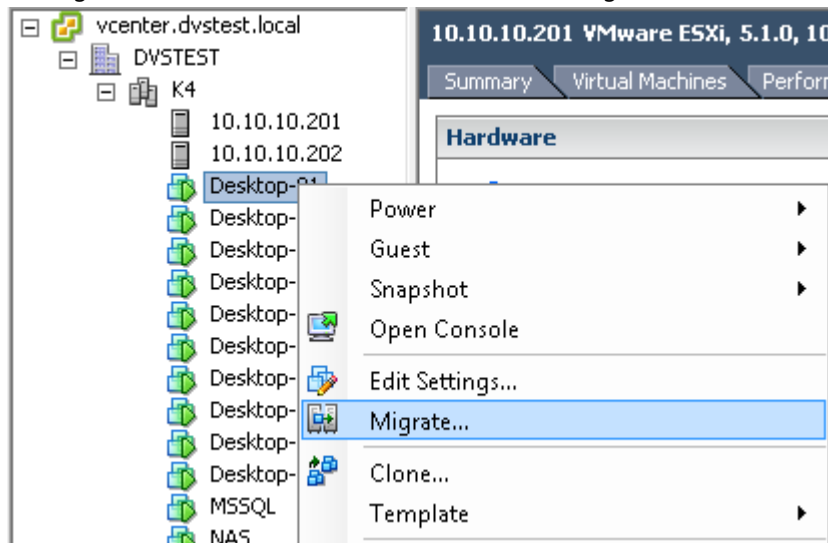
To test VMware vMotion, a VM was migrated from one ESXi server to another in the same VMware Virtual Center Cluster. The steps taken were as follows.

- 1) First, we need to ensure that VMware vMotion is enabled on a network. For this exercise, it was enabled on the Management network, but this is not typically a best practice. Best practices are documented extensively by VMware and vary depending on customer needs.



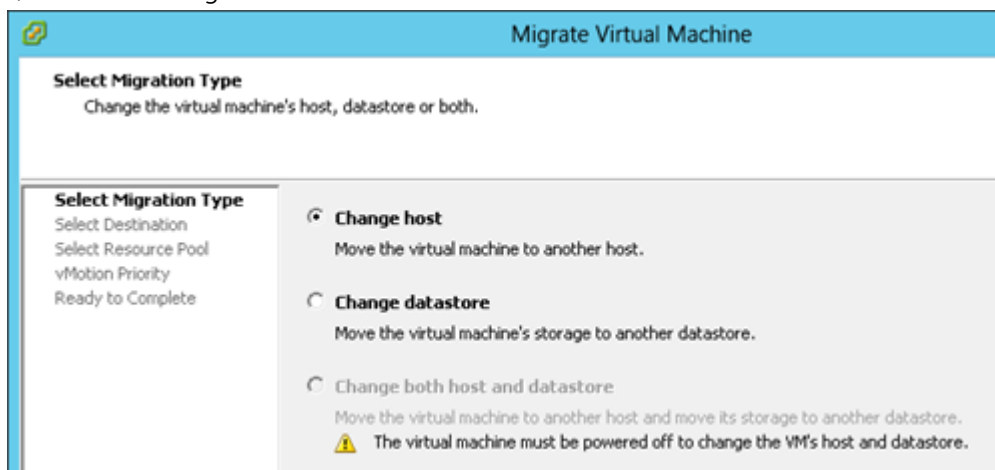
**Figure 59**

2) Right click on the VM to vMotion and click "Migrate".



**Figure 60**

3) Select "Change host" and click Next.



**Figure 61**

4) Expand the Cluster ("K4" below) and select the ESXi host that you wish to migrate the VM to. In the below example, the VM is being moved from 10.10.10.202 to 10.10.10.201. Ensure that the bottom window says "Validation succeeded" and click Next to continue.

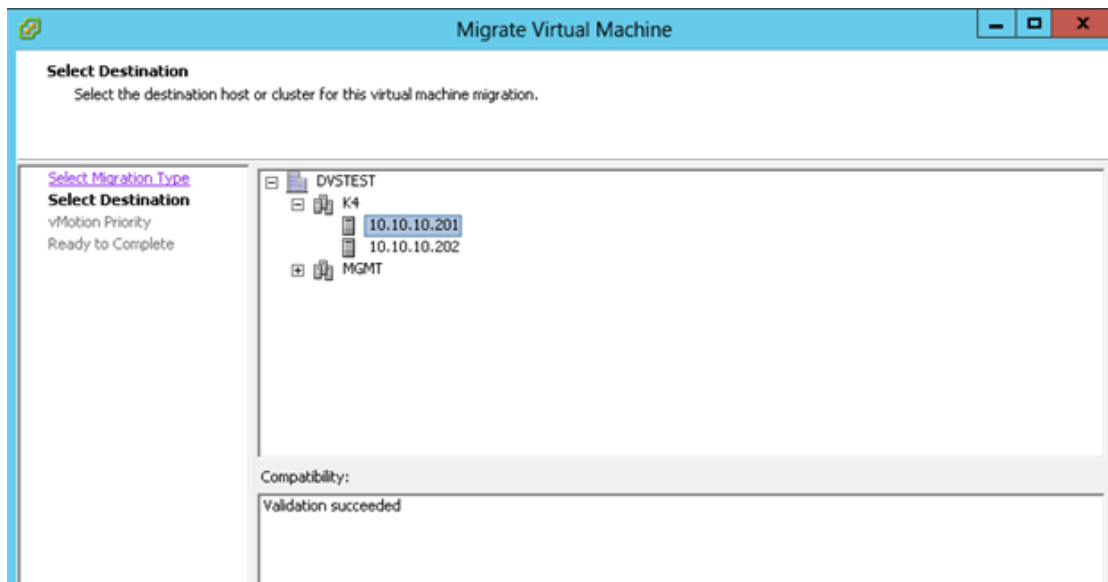


Figure 62

5) Select "High Priority" and click Next.

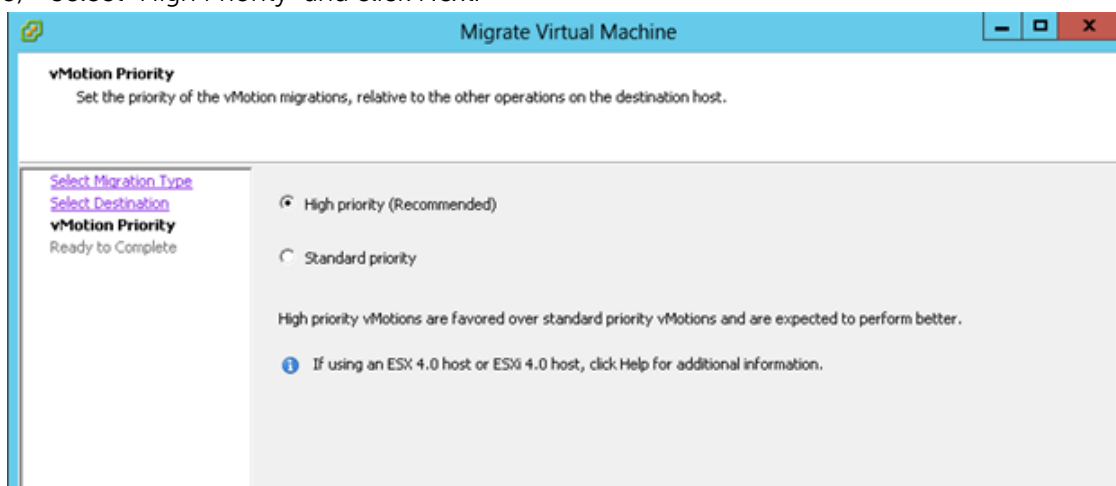
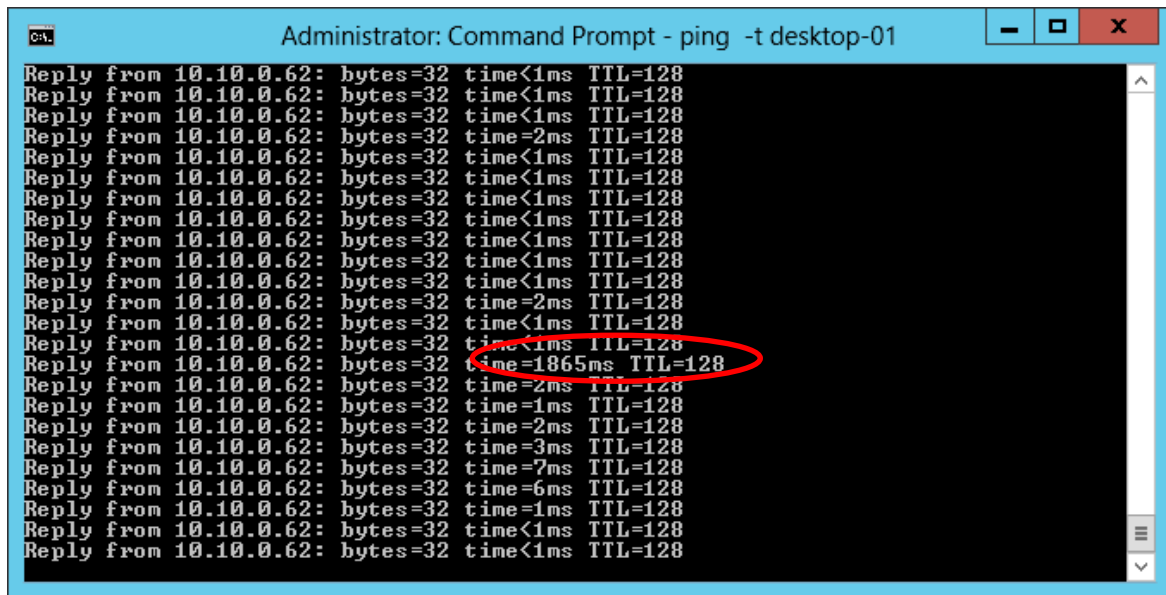


Figure 63

6) Click Finish

## VMware vMotion Test Results

During the test, the VM was pinged at the default interval. There was absolutely no disruption in service, although there was one spike in latency.



```
C:\> Administrator: Command Prompt - ping -t desktop-01
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time=2ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time=2ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time=1865ms TTL=128
Reply from 10.10.0.62: bytes=32 time=2ms TTL=128
Reply from 10.10.0.62: bytes=32 time=1ms TTL=128
Reply from 10.10.0.62: bytes=32 time=2ms TTL=128
Reply from 10.10.0.62: bytes=32 time=3ms TTL=128
Reply from 10.10.0.62: bytes=32 time=7ms TTL=128
Reply from 10.10.0.62: bytes=32 time=6ms TTL=128
Reply from 10.10.0.62: bytes=32 time=1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
```

Figure 64

## VMware Storage vMotion

To test VMware Storage vMotion, a VM was migrated from one datastore to another on the same VMware ESXi host using shared storage devices on an EqualLogic storage array. The steps taken were as follows.

- 1) Right click on the VM to vMotion and click "Migrate".




Figure 65

- 2) Select "Change datastore" and click Next.

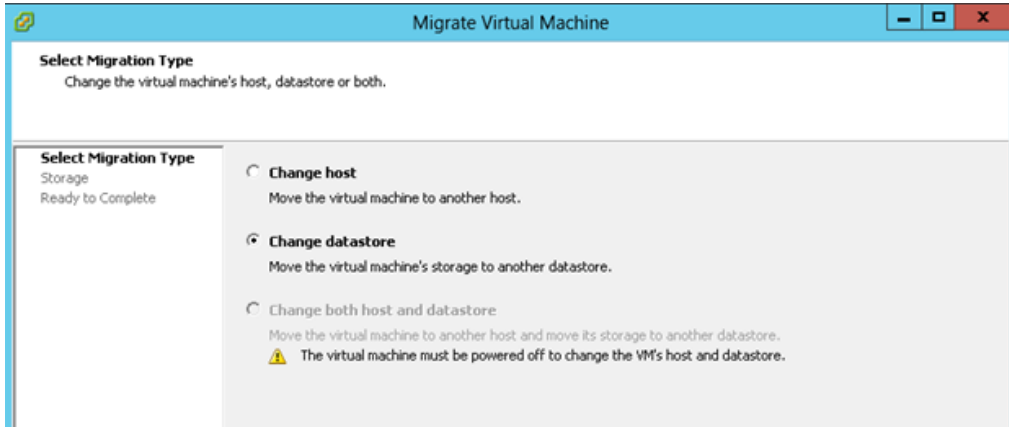


Figure 66

- 3) In this example, the VMs disk resources are being moved from vdi-1 to vdi-2 datastores. Select "vdi-2", ensure the bottom panel says "Validation succeeded" and click Next.

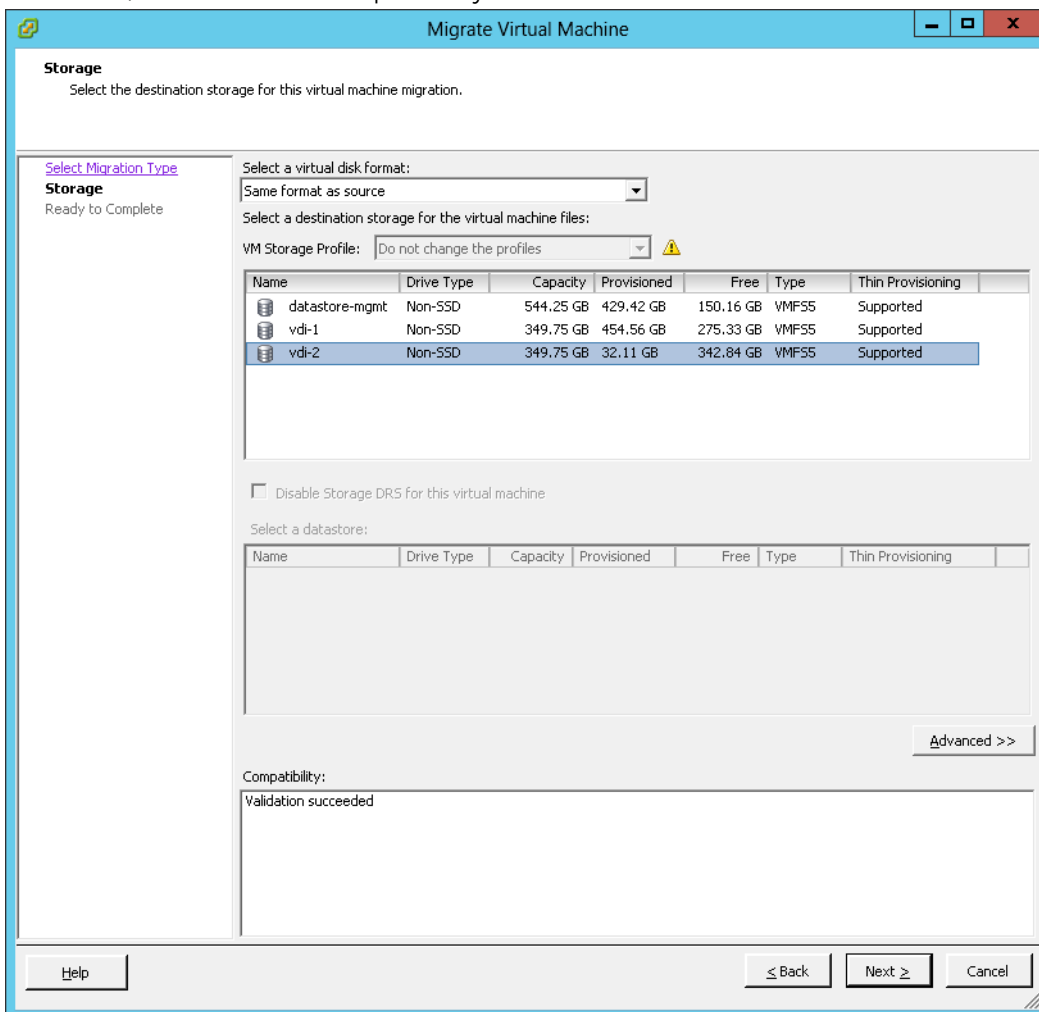


Figure 67

- 4) Click Finish and the VM will start migrating to the new datastore.

Recent Tasks							
Name	Target	Status	Details	Initiated by	Requested Start Time	Start Time	Completed Time
Relocate virtual machine	Desktop-01	30%	Copying Virtual Machine Files	DVSTEST\pic...	5/30/2013 12:05:22 ...	5/30/2013 12:05:22 ...	

Figure 68

- 5) The process will complete.

Recent Tasks							
Name	Target	Status	Details	Initiated by	Requested Start Time	Start Time	Completed Time
Relocate virtual machine	Desktop-01	Completed		DVSTEST\pic...	5/30/2013 12:05:22 ...	5/30/2013 12:05:22 ...	5/30/2013 12:12:48 ...

Figure 69

6) The VM is now located on "vdi-2" datastore.

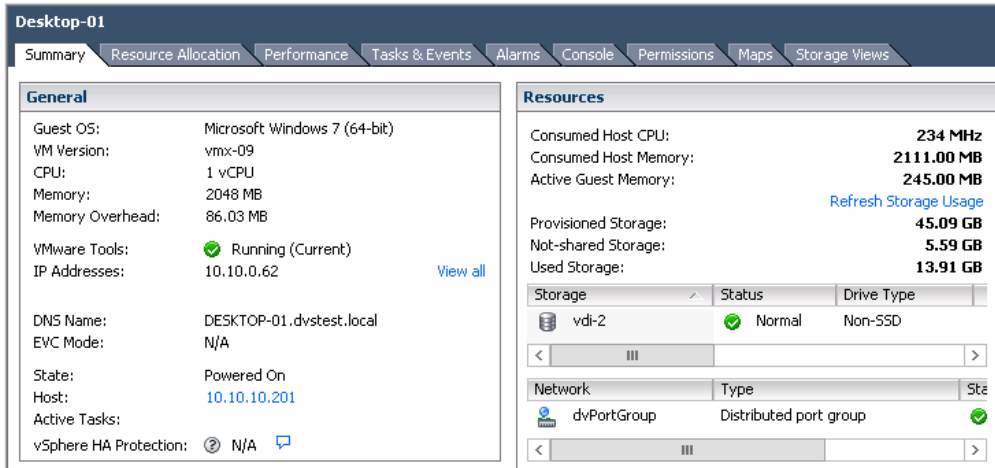
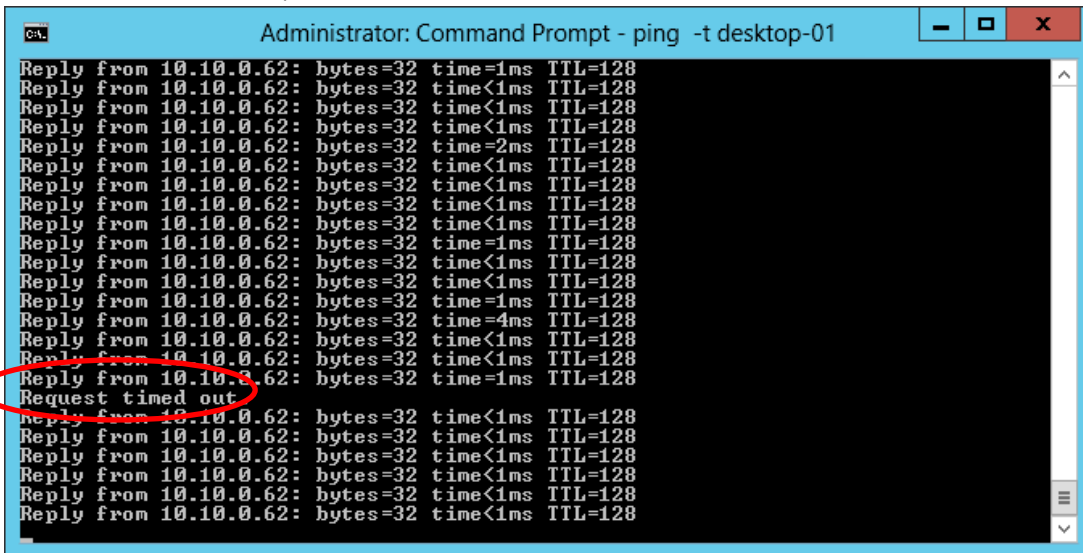


Figure 70


## VMware Storage vMotion Test Results

During the test, the VM was pinged at the default interval. There were two small disruptions in service.

One at about 70% complete:



And one ping drops when the process completes followed by a high ping of 11ms:



## VMware vMotion without Shared Storage

To test VMware Storage vMotion, a VM was migrated from one ESXi server to another as well as from one server's local disks to another in the same VMware Virtual Center Cluster. The steps taken were as follows.

**\*\*Note:** This procedure cannot be done using the standard VMware vSphere Client and must be done through the VMware vSphere web client.

- 1) The VM starts on the local datastore of host 10.10.10.201:

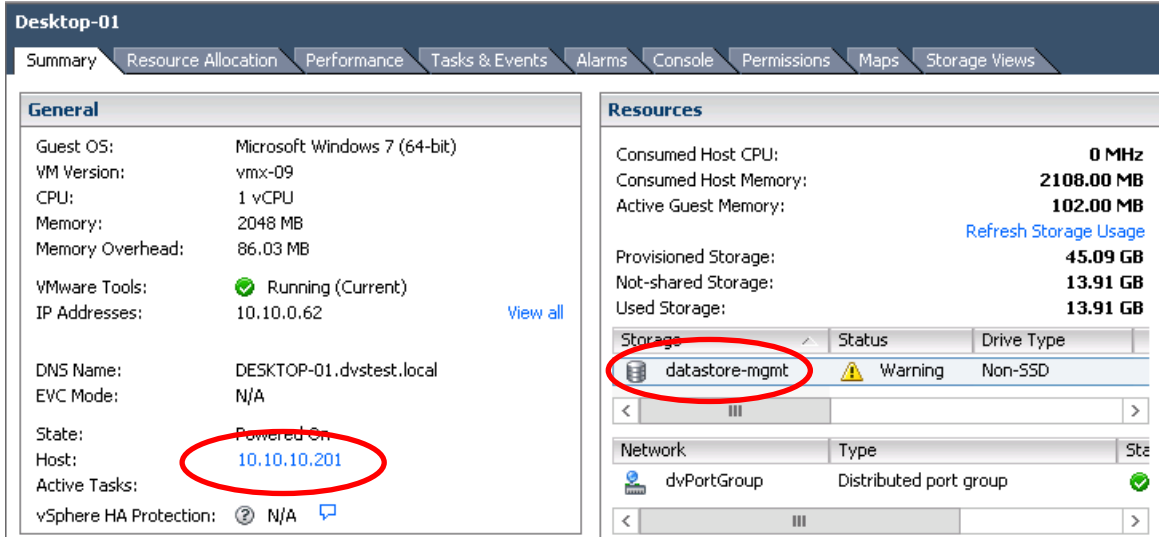


Figure 71

- 2) Using Internet Explorer, browse to [https://<virtual\\_center\\_ip\\_address>:9443](https://<virtual_center_ip_address>:9443) to access the





Figure 72




3) Log in with the appropriate credentials.

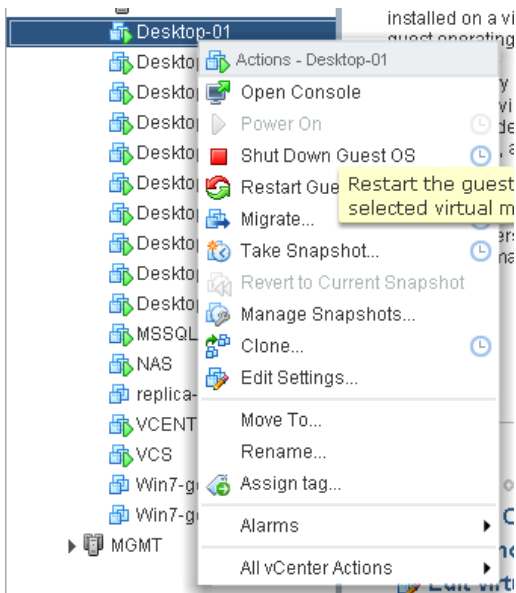


**Figure 73**

4) Click on vCenter, Hosts and Clusters, and then expand the arrows to see the VMs

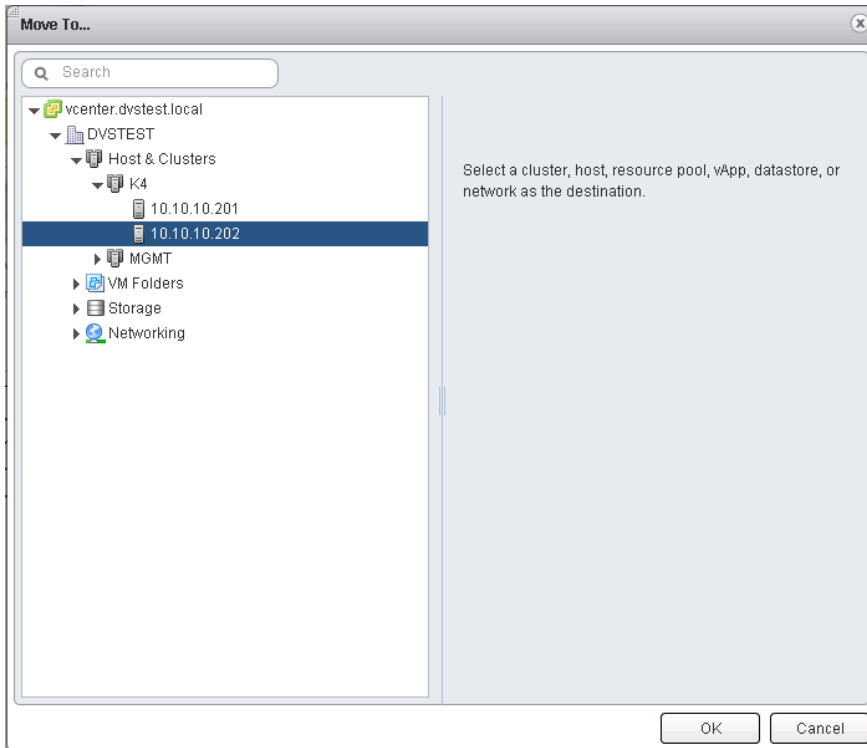


5) Right click on the VM that you wish to migrate and select "Move to".



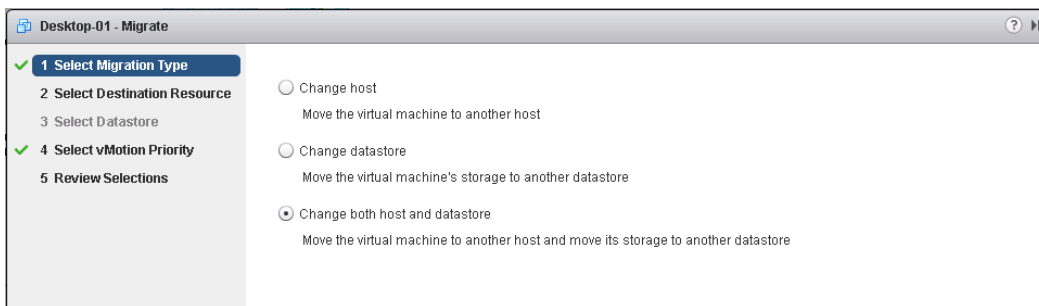
**Figure 74**

6) In this exercise, we are migrating the VM from 10.10.10.201 to 10.10.10.202. Click OK.



**Figure 75**

7) Select "Change both host and datastore" and click Next.



**Figure 76**

8) Select the appropriate resource, ensure it says "Compatibility checks succeeded" and click Next

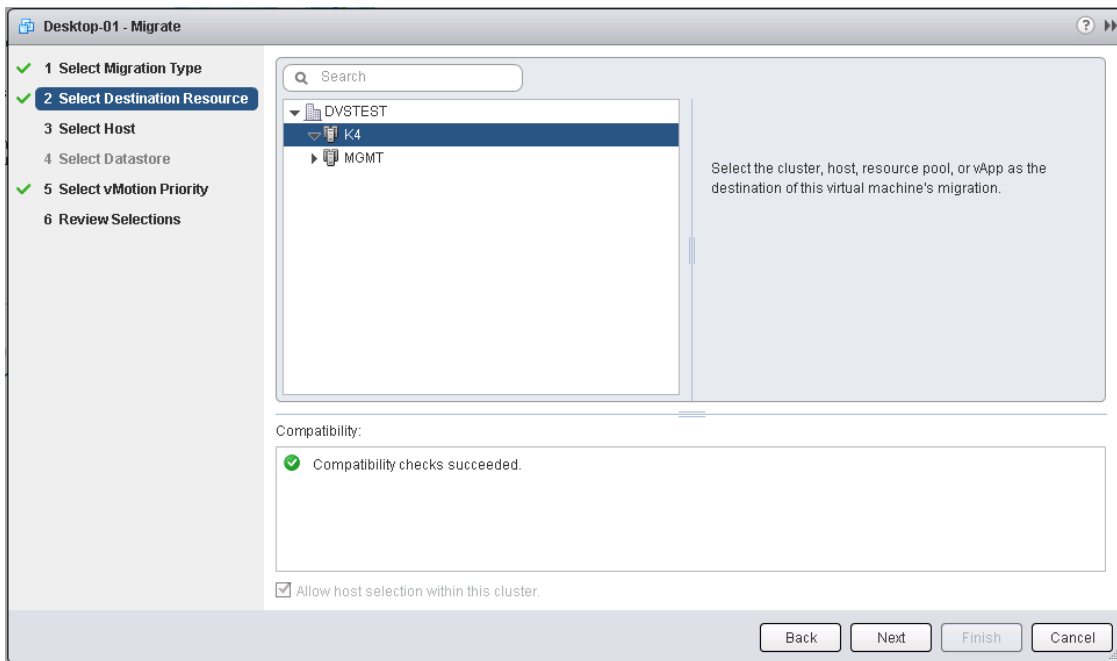


Figure 77

9) Select the appropriate ESXi host and click Next.

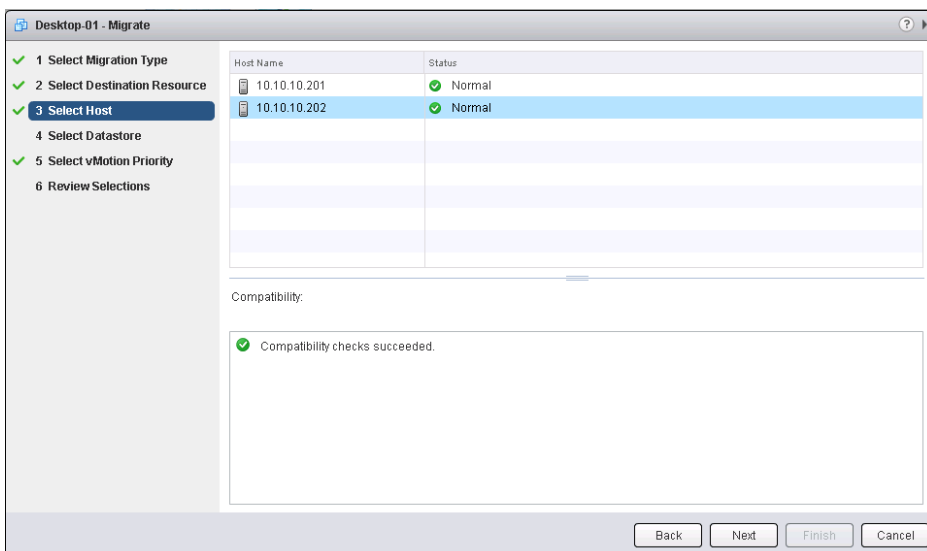


Figure 78

10) Select the appropriate datastore and click Next. Note that this is local storage on another server in the below screenshot.

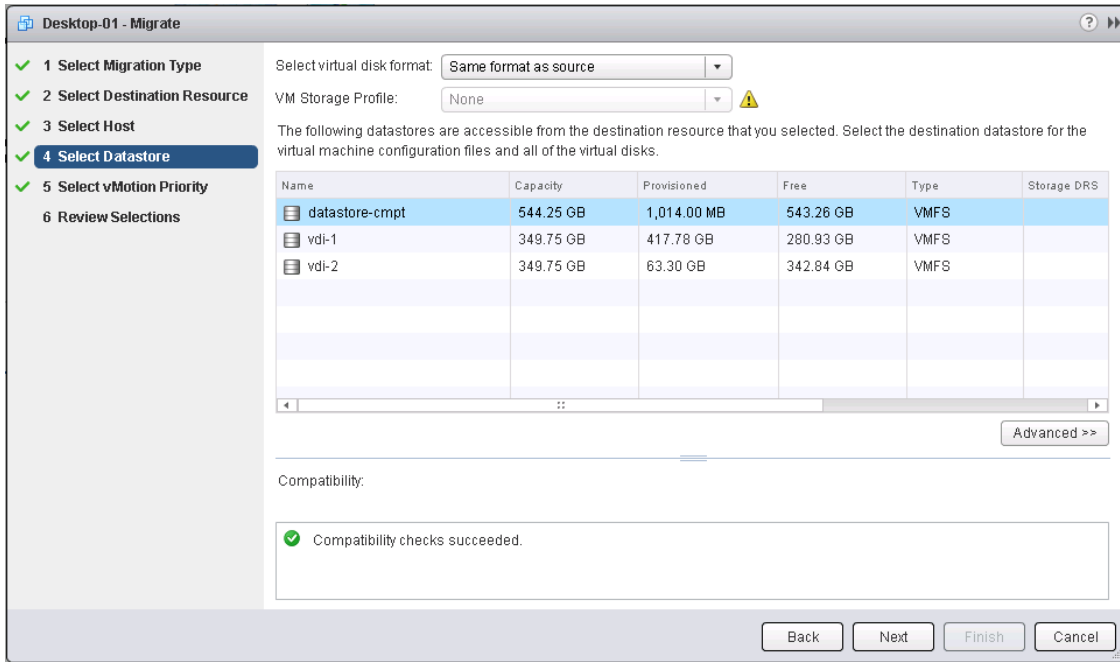


Figure 79

11) The free space on the destination datastore is shown before the vMotion starts.

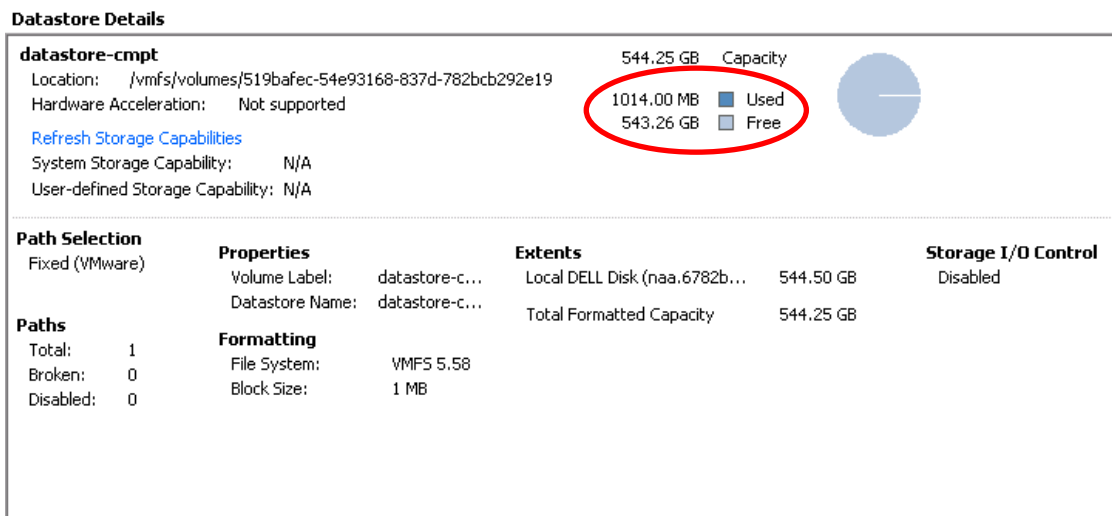


Figure 80

12) Select Next.

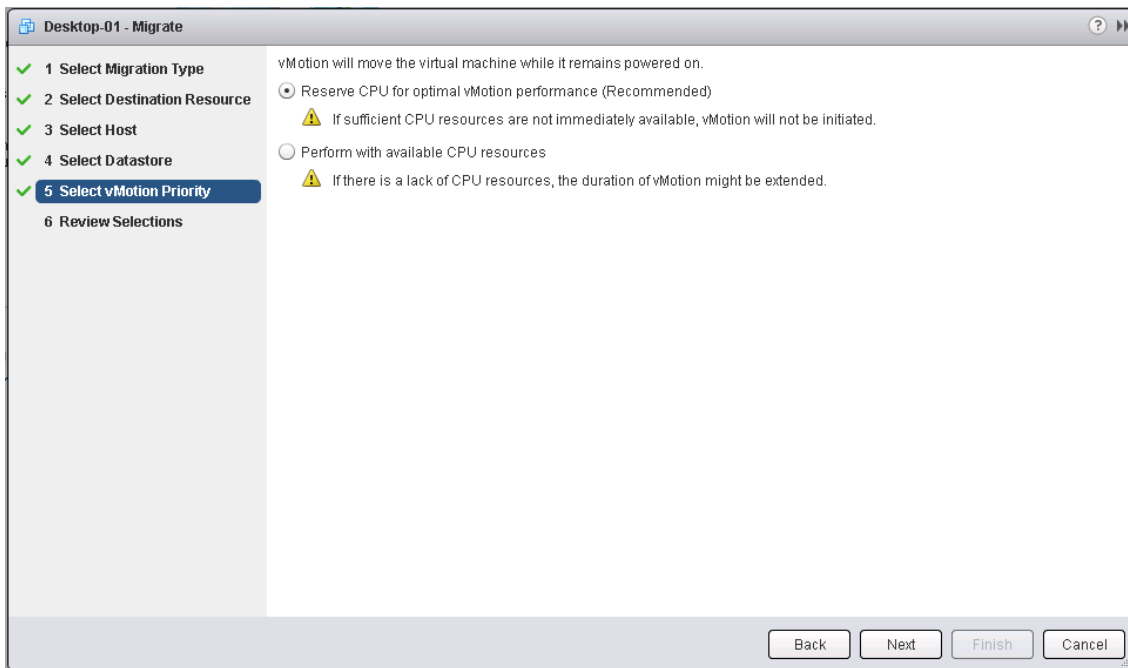


Figure 81

13) Review and click Finish.

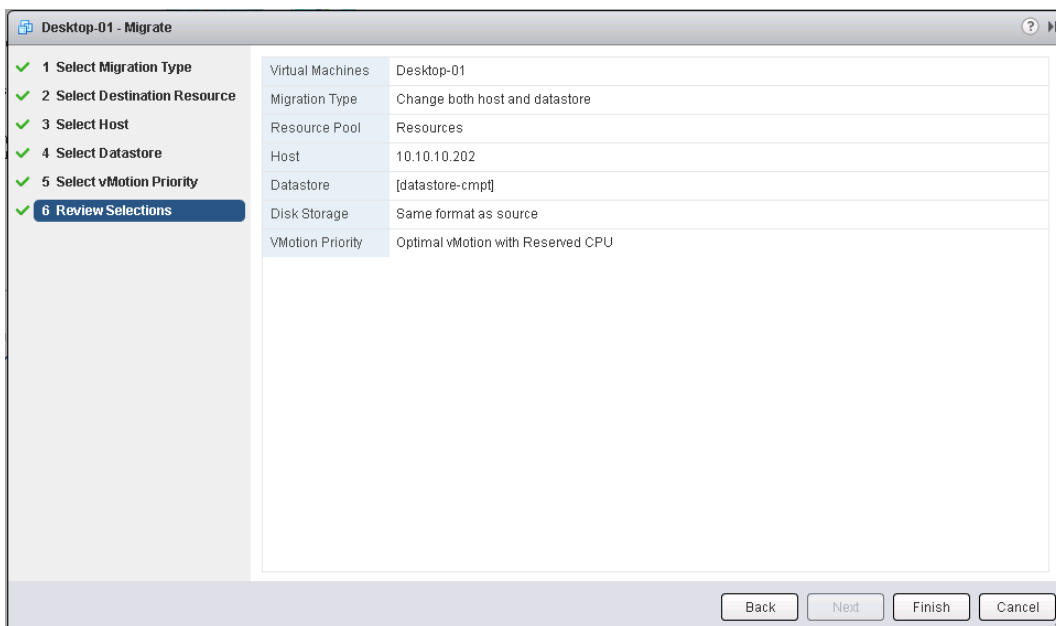


Figure 82

14) The progress can be monitored either through the vSphere Web or regular client.

Task Name	Target	Status	Initiator	Queued For	Start Time	Completion Time	Server
Relocate virtual machine	Desktop-01	56 %	DVSTEST\Nicholas...	0 ms	5/30/2013 12:50 PM		vcenter.dvstest.local
Move into resource pool	K4	Completed	DVSTEST\Nicholas...	0 ms	5/30/2013 12:46 PM	5/30/2013 12:46 PM	vcenter.dvstest.local

Figure 83

15) The process completed successfully.

Name	Target	Status	Details	Initiated by	Requested Start Time	Start Time	Completed Time
Relocate virtual machine	Desktop-01	Completed		DVSTEST1\nc...	5/30/2013 12:50:56 ...	5/30/2013 12:50:56 ...	5/30/2013 12:53:53 ...
Move into resource pool	K4	Completed		DVSTEST1\nc...	5/30/2013 12:46:22 ...	5/30/2013 12:46:22 ...	5/30/2013 12:46:22 ...

Figure 84

16) The VM is now on 10.10.10.202 and "datastore-cmpt" datastore.

**Desktop-01**

Summary | Resource Allocation | Performance | Tasks & Events | Alarms | Console | Permissions | Maps | Storage Views

**General**

Guest OS: Microsoft Windows 7 (64-bit)  
 VM Version: vmx-09  
 CPU: 1 vCPU  
 Memory: 2048 MB  
 Memory Overhead: 86.06 MB  
 VMware Tools: ✔ Running (Current)  
 IP Addresses: 10.10.0.62 [View all](#)

DNS Name: DESKTOP-01.dvstest.local  
 EVC Mode: N/A

State: **Powered On**  
 Host: **10.10.10.202**  
 Active Tasks:  
 vSphere HA Protection: ? N/A 🗨

**Resources**

Consumed Host CPU: **0 MHz**  
 Consumed Host Memory: **2103.00 MB**  
 Active Guest Memory: **102.00 MB** [Refresh Storage Usage](#)

Provisioned Storage: **45.09 GB**  
 Not-shared Storage: **13.93 GB**  
 Used Storage: **13.93 GB**

Storage	Status	Drive Type
datastore-cmpt	<span style="color: green;">✔</span> Normal	Non-SSD

Network

Network	Type	Sta
dvPortGroup	Distributed port group	<span style="color: green;">✔</span>

Figure 85

17) The datastore shows an additional ~15 GB of used space:

**Datastore Details**

**datastore-cmpt** 544.25 GB Capacity

Location: /vmfs/volumes/519bafec-54e93168-837d-782bcb292e18  
 Hardware Acceleration: Not supported  
[Refresh Storage Capabilities](#)

System Storage Capability: N/A  
 User-defined Storage Capability: N/A

15.01 GB Used  
 529.24 GB Free

Path Selection	Properties	Extents	Storage I/O Control
Fixed (VMware)	Volume Label: datastore-c... Datastore Name: datastore-c...	Local DELL Disk (naa.6782b... 544.50 GB Total Formatted Capacity 544.25 GB	Disabled

**Paths**

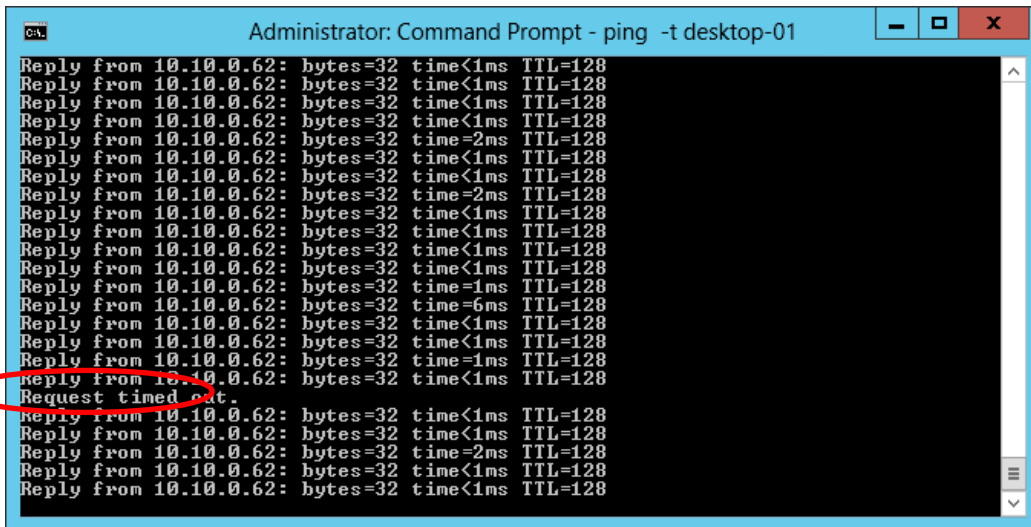
Paths	Formatting
Total: 1 Broken: 0 Disabled: 0	File System: VMFS 5.58 Block Size: 1 MB

Figure 86

## VMware vMotion without Shared Storage Test Results

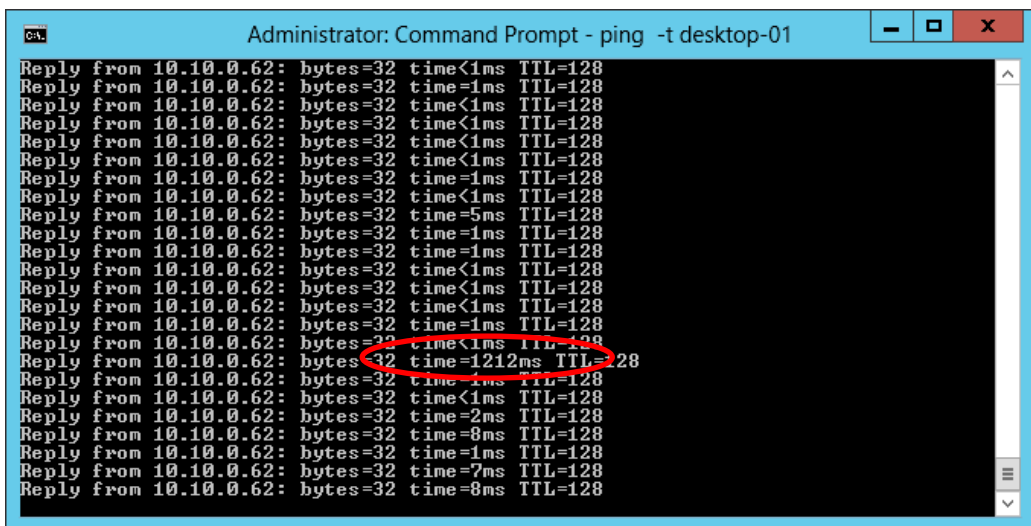
During the test, the VM was pinged at the default interval. There was only one small disruption in service followed by a

One at about 77% complete:



```
Administrator: Command Prompt - ping -t desktop-01
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time=2ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time=2ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time=1ms TTL=128
Reply from 10.10.0.62: bytes=32 time=6ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time=1ms TTL=128
Request timed out.
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time=2ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
```

And one high ping as the process completed:



```
Administrator: Command Prompt - ping -t desktop-01
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time=1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time=1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time=5ms TTL=128
Reply from 10.10.0.62: bytes=32 time=1ms TTL=128
Reply from 10.10.0.62: bytes=32 time=1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time=1ms TTL=128
Reply from 10.10.0.62: bytes=32 time=1ms TTL=128
Reply from 10.10.0.62: bytes=32 time=1212ms TTL=128
Reply from 10.10.0.62: bytes=32 time=1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time=2ms TTL=128
Reply from 10.10.0.62: bytes=32 time=8ms TTL=128
Reply from 10.10.0.62: bytes=32 time=1ms TTL=128
Reply from 10.10.0.62: bytes=32 time=7ms TTL=128
Reply from 10.10.0.62: bytes=32 time=8ms TTL=128
```

## VMware vMotion Cold Migration

To test VMware vMotion, a VM was migrated from one ESXi server to another as well as from one server's local disks to another in the same VMware Virtual Center Cluster using the "Change both host and datastore" option in the VMware vSphere client. In order to accomplish this, the virtual machine must be powered off to change both the host and datastore when using the vSphere client. The steps taken were as follows.

- 1) If the VM was created by VMware View, place the VM in maintenance mode through VMware Horizon View Administrator. Otherwise, when you shut down the VM, View will boot it up again.

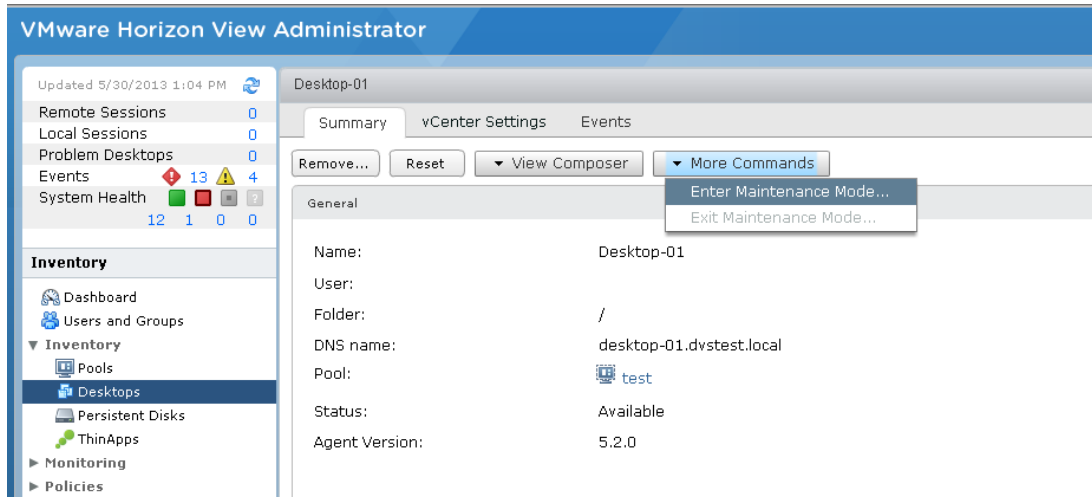


Figure 87

- 2) Select OK.

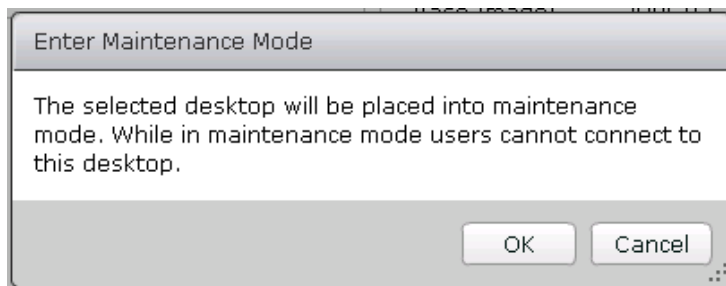


Figure 88

- 3) Shut down the VM through the vSphere client.

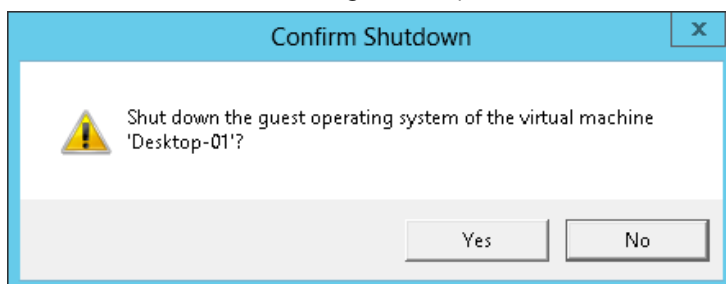


Figure 89



4) Once the host is shut off, right click the VM and select "Migrate".

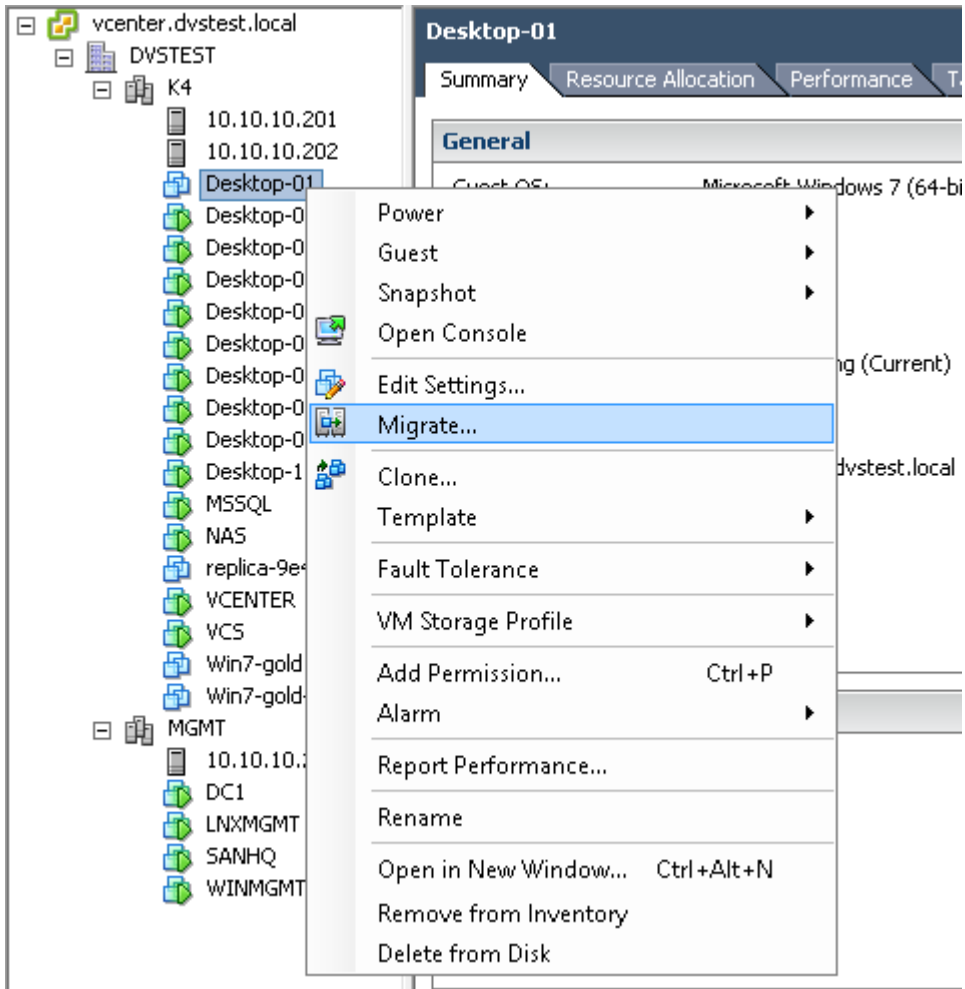


Figure 90

5) Select "Change both host and datastore" and click Next.

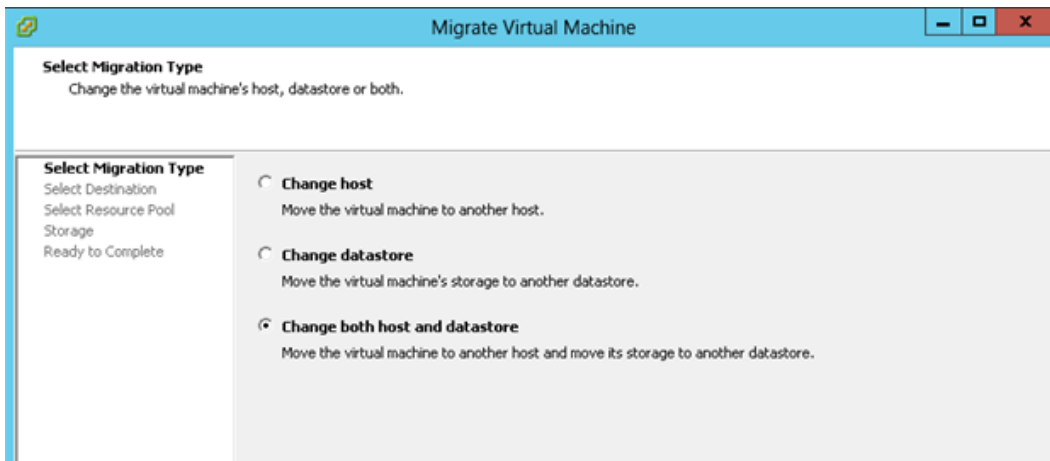


Figure 91

6) In this exercise, we are migrating from 10.10.10.202 to 10.10.10.201. Click Next.

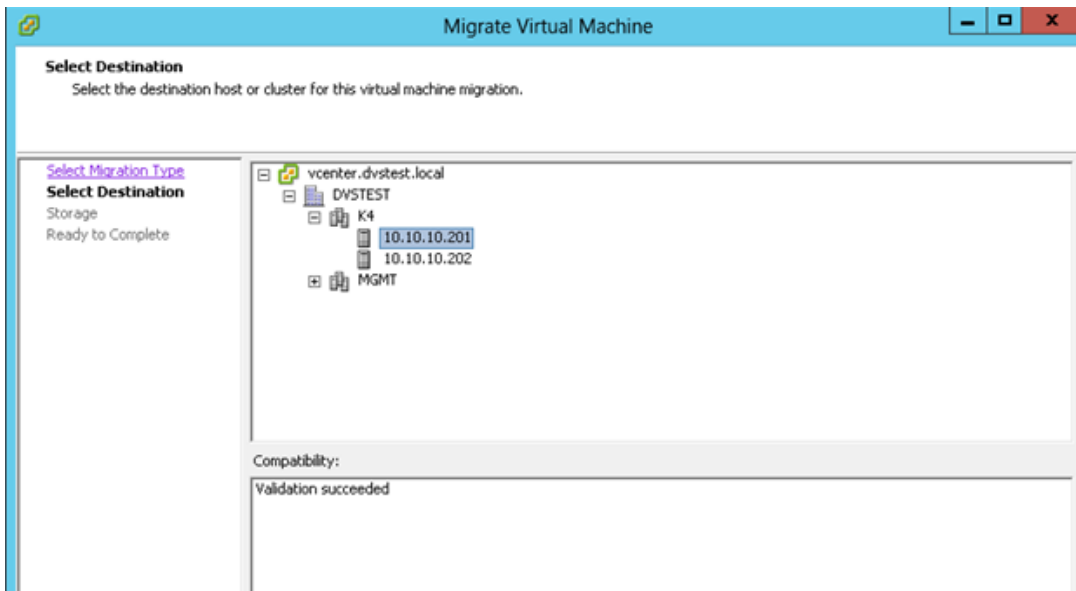


Figure 92

7) Select the local datastore on that host and click Next:

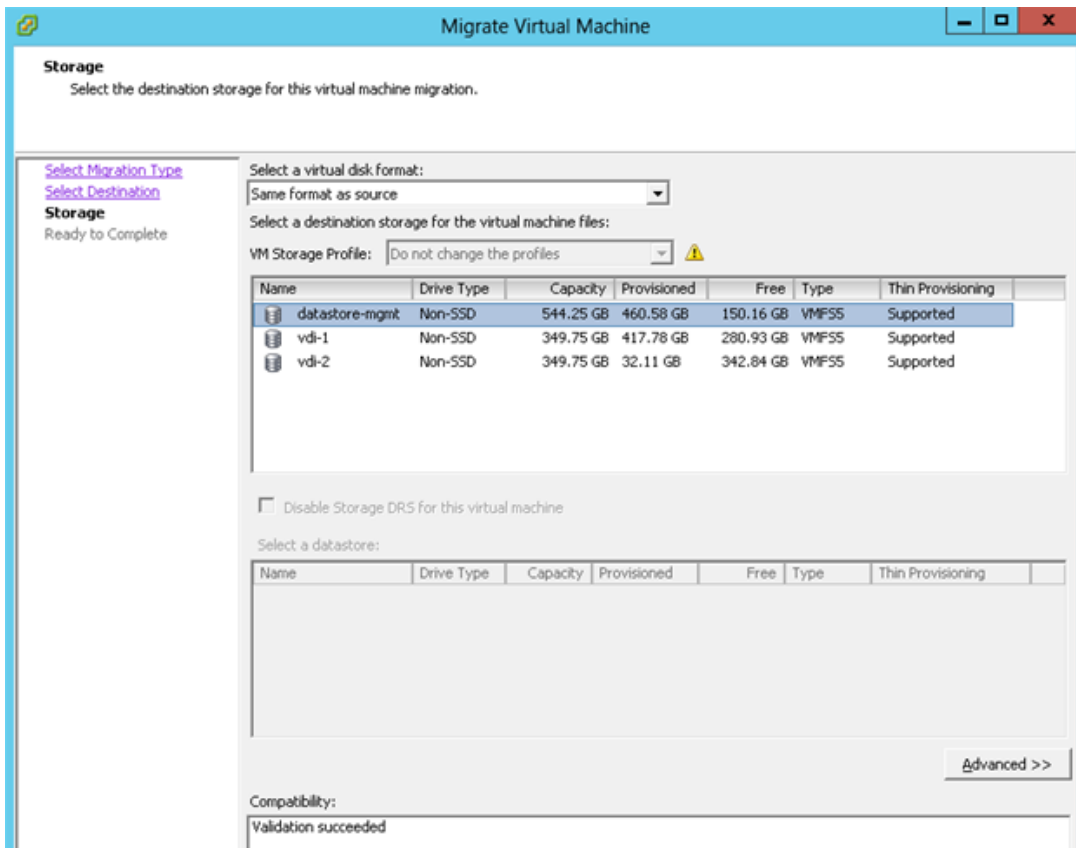


Figure 93

8) Click Finish.

9) The machine is migrated and progress can be monitored in "Recent Tasks".

Name	Target	Status	Details	Initiated by	vCenter Serv...	Requested Start
Relocate virtual machine	Desktop-01	31%	Copying Virtual Machine files	DVSTEST\nicholas_busick	vcenter...	5/30/2013 1:10:1
Initiate guest OS shutdown	Desktop-01	Completed		DVSTEST\nicholas_busick	vcenter...	5/30/2013 1:06:5
Power On virtual machine	Desktop-01	Completed		DVSTEST\nicholas...	vcenter...	5/30/2013 1:06:4

Figure 94

10) The VM is now located on 10.10.10.201 and the local datastore.

**Desktop-01**

Summary | Resource Allocation | Performance | Tasks & Events | Alarms | Console | Permissions | Maps | Storage Views

**General**

Guest OS: Microsoft Windows 7 (64-bit)  
VM Version: vmx-09  
CPU: 1 vCPU  
Memory: 2048 MB  
Memory Overhead: 152.97 MB  
VMware Tools: Not running (Current)  
IP Addresses: 10.10.0.62 [View all](#)

DNS Name: DESKTOP-01.dvstest.local  
EVC Mode: N/A

State: Powered Off  
Host: 10.10.10.201  
Active Tasks:  
vSphere HA Protection: N/A

**Resources**

Consumed Host CPU: 0 MHz  
Consumed Host Memory: 0.00 MB  
Active Guest Memory: 0.00 MB [Refresh Storage Usage](#)

Provisioned Storage: 45.11 GB  
Not-shared Storage: 11.65 GB  
Used Storage: 11.65 GB

Storage	Status	Drive Type
datastore-mgmt	Normal	Non-SSD

**Network**

Network	Type	Sta
dvPortGroup	Distributed port group	

**Commands**

- Power On
- Edit Settings
- Migrate
- Clone to New Virtual Machine
- Convert to Template

**VM Storage Profiles** [Refresh](#)

VM Storage Profiles:  
Profiles Compliance:

Figure 95

11) Exit maintenance mode on the VM using VMware Horizon View Administrator.

**VMware Horizon View Administrator**

Updated 5/30/2013 3:24 PM

Remote Sessions: 0  
Local Sessions: 0  
Problem Desktops: 0  
Events: 12 (Warning), 4 (Error)  
System Health: 12 (Warning), 1 (Error), 0 (Info), 0 (Info)

**Inventory**

- Dashboard
- Users and Groups
- Inventory
  - Pools
  - Desktops**
  - Persistent Disks
  - ThinApps
- Monitoring
- Policies
- View Configuration

**Desktop-01**

Summary | vCenter Settings | Events

Remove... | Reset | View Composer | More Commands

**General**

Name: Desktop-01  
User:  
Folder: /  
DNS name: desktop-01.dvstest.local  
Pool: test  
Status: Maintenance mode  
Agent Version: Unknown

More Commands menu:  
Enter Maintenance Mode...  
Exit Maintenance Mode...

Figure 96

12) Click OK.

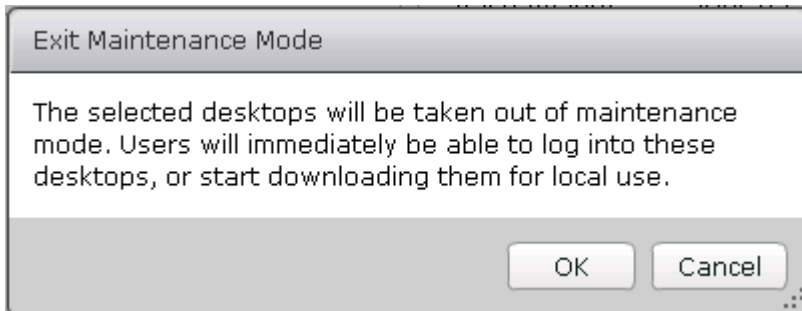


Figure 97

13) The VM will automatically power on.

Recent Tasks				
Name	Target	Status	Details	Initiated by
Power On virtual machine	Desktop-01	Completed		DVSTEST1\svc.view

Figure 98

14) The VM is now located on 10.10.10.201 and its local datastore.

Desktop-01

Summary Resource Allocation Performance Tasks & Events Alarms Console Permissions Maps Storage Views

**General**

Guest OS:  
VM Version: vmx-09  
CPU: 1 vCPU  
Memory: 2048 MB  
Memory Overhead: 86.06 MB  
VMware Tools: Running (Current)  
IP Addresses: 10.10.0.62

DNS Name:  
EVC Mode: N/A

State: Powered On  
Host: 10.10.10.201

Active Tasks:  
vSphere HA Protection: N/A

**Commands**

- Shut Down Guest
- Suspend
- Restart Guest
- Edit Settings
- Open Console
- Migrate
- Clone to New Virtual Machine

**Annotations**

Notes: Edit

**Resources**

Consumed Host CPU: 0 MHz  
Consumed Host Memory: 0.00 MB  
Active Guest Memory: 0.00 MB  
[Refresh Storage Usage](#)

Provisioned Storage: 45.01 GB  
Not-shared Storage: 13.86 GB  
Used Storage: 13.86 GB

Storage	Status	Drive Type
datastore-mgmt	Normal	Non-SSD

**Network**

dvPortGroup	Type
dvPortGroup	Distributed port group

**VM Storage Profiles**

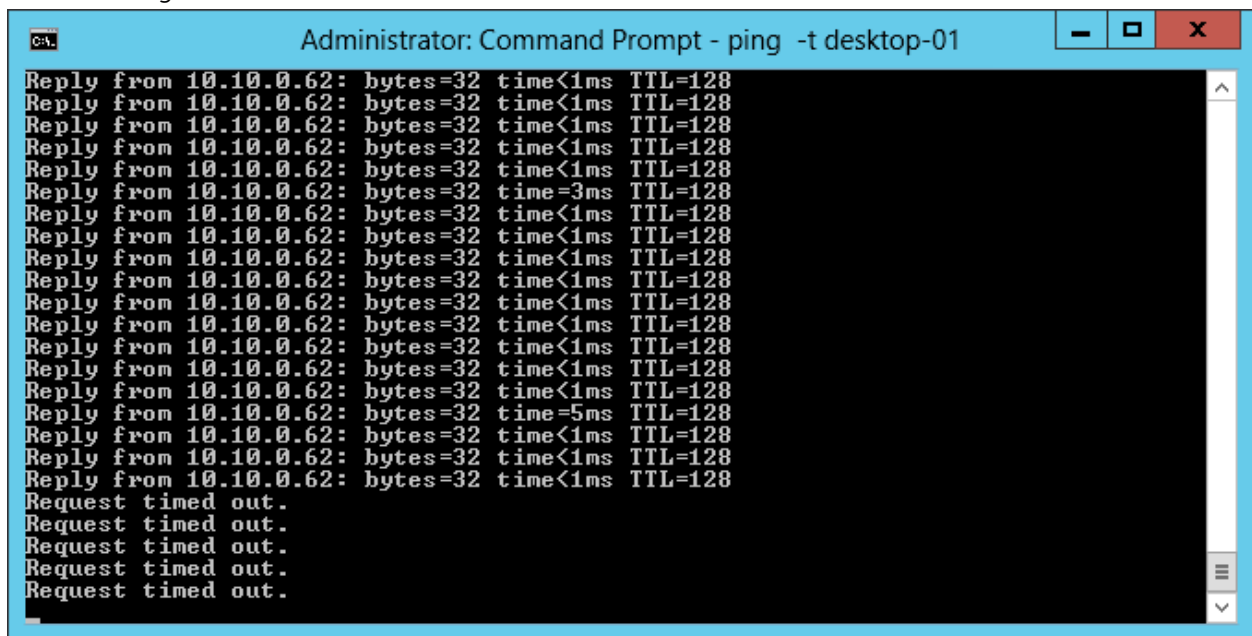
VM Storage Profiles: [Refresh](#)  
Profiles Compliance:

Figure 99

## VMware vMotion Cold Migration Test Results

The VM will be down for the entire vMotion process. This could be anywhere between minutes and hours depending on a number of factors such as disk space used within the VM, network speed, processing power of ESXi servers, and how busy the environment currently is.

The following screenshot shows that the VM will not be available as soon as it is shut down:



```
Administrator: Command Prompt - ping -t desktop-01
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time=3ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time=5ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Reply from 10.10.0.62: bytes=32 time<1ms TTL=128
Request timed out.
Request timed out.
Request timed out.
Request timed out.
Request timed out.
```

## Conclusion

VMware vSphere 5.1 brings important updates to vMotion capabilities that allow for much more flexible virtualization environments. The updates allow administrators to perform maintenance and configuration changes with a minimal amount of user disruption and downtime.

The testing done in this paper found that VMware vMotion worked as advertised and provided very high availability for the VMs that were migrated from one piece of hardware to the next. There were no notable issues or concerns with using this technology in combination with VMware Horizon View 5.2.

## References

<http://www.vmware.com/files/pdf/techpaper/VMware-vSphere51-vMotion-Perf.pdf>

<http://blogs.vmware.com/performance/2012/08/vmware-vsphere-5-1-vmotion-architecture-performance-and-best-practices.html>

### 13.7 AS 50, 200, 800, and 1000 (IOA and MSL)

Discussion on the merits and pieces of the AS Stacks (this section will be updated post RTS) along with description information

## 13.8 Win8 characterization and testing

### Executive Summary

This document is the test architecture and results document for testing conducted on the impact of Windows 8 on VMware Horizon View 5.2. It contains the testing methodology and results. Based on the functional details of Windows 8, the impact on Compute host local Tier1 storage and sizing for basic, standard and premium users will be determined as a result of this testing.

### VMware Horizon View 5.2

Horizon View 5.2 leverages a new vSphere capability that implements a new disk format for VMs on VMFS that allow for grain reduction size & more efficiently utilization of allocated blocks by filling it with real data. Unused space is reclaimed and View Composer desktops stay small. It also supports Windows 8 virtual desktops as guest OS and Client. Horizon View Storage Accelerator optimizes storage load by caching common image blocks when reading virtual desktop images. Space Efficient Disks continuously reduce the storage needed per desktop. Both these technologies improve storage capacity and utilization, thereby reducing costs of additional hardware.

### Configuration

Below is the Solution Software Configuration. For additional configuration information see the reference architecture document.

<b>Solution Configuration – Software Components</b>	<b>Description/Version</b>
VMware vCenter	Version 5.1 U1
VMware Horizon View	Version 5.2
Hypervisor	VMware ESXi 5.1 U1
Microsoft SQL Server	Version 2008 R2
Windows 8 Enterprise (32 bit)	VDI Clients for characterization tests Login VSI launchers
Windows Server 2008 R2 SP1	VMs for hosting VMware View, vCenter Server, MSSQL server and other infrastructure VMs.

Below is the Solution Hardware Configuration. For additional configuration information see the reference architecture document.

<b>Solution Configuration - Hardware Components:</b>		<b>Description</b>
<b>VMware Compute Host</b>	1 x Dell PowerEdge R720 Server: <ul style="list-style-type: none"><li>• ESXi 5.1</li><li>• Intel-(R)-Xeon- (R) CPU E5-2690 @ 2.9 GHZ</li><li>• 196 GB @ 1600 MHZ</li><li>• 10 x 146GB 15K SAS internal disk drives</li><li>• Broadcom BCM5720 1 GbE NIC</li><li>• PERC H710P RAID Controller</li></ul>	For ESXi Environment 10 x 146 GB drives will be configured on RAID 10  This host will be hosting the Windows 8 VDI desktops.

<b>VMware Management Host</b>	<p>1 x Dell PowerEdge R710 Server:</p> <ul style="list-style-type: none"> <li>• VMware ESXi 5.1</li> <li>• 2 x HexCore Intel® Xeon® X5670 2.9 GHz Processors</li> <li>• 96 GB RAM @ 1333 MHZ</li> <li>• 8 x 146GB 15K SAS internal disk drives</li> <li>• 1 x Broadcom 5709 1GbE NIC, Quad-Port</li> </ul>	<p>For ESXi Environment 8 x 146 GB drives will be configured on RAID 10</p> <p>Each VM will host the following workloads on Windows server 2008 R2 SP1:</p> <ul style="list-style-type: none"> <li>• VMware vCenter</li> <li>• VMware View Connection Server</li> <li>• Microsoft SQL server (View Connection Server and vCenter Database)</li> <li>• File Server</li> </ul>
<b>Test Management Server (Login VSI Infrastructure Mgmt Server)</b>	<p>1x Dell PowerEdge R710 servers:</p> <ul style="list-style-type: none"> <li>• VMware ESXi 5.1</li> <li>• 2 x HexCore Intel® Xeon® X5670 2.9 GHz Processors</li> <li>• 96 GB RAM @ 1333 MHZ</li> <li>• 2 x 146GB 15K SAS internal disk drives</li> <li>• 1 x Broadcom 5709 1GbE NIC, Quad-Port</li> <li>• Login VSI version 3.7</li> </ul>	<p>VMware ESXi 5.1 will be installed on the Test Mgmt R710 server.</p> <p>Test Infrastructure VMs will be created on this server.</p> <p>The VMs are loaded with the Windows 2008 R2 SP1 operating system.</p> <p>Each VM will host the following workloads:</p> <ul style="list-style-type: none"> <li>• Domain Controller (test1.com)</li> <li>• Stratusphere Hub (5.3.1)</li> <li>• McAfee</li> </ul>
<b>Test Launcher Server [VDI Workload Generator]</b>	<p>1x Dell PowerEdge R710 servers:</p> <ul style="list-style-type: none"> <li>• VMware ESXi 5.1</li> <li>• Intel-(R)-Xeon- (R) CPU X5670 @ 2.9 GHz</li> <li>• 96 GB RAM @ 1333 MHZ</li> <li>• 2 x 146GB 15K SAS internal disk drives</li> <li>• Broadcom 5709 1GbE NIC</li> <li>• Login VSI version 3.7</li> </ul>	<p>VMware ESXi 5.1 will be installed on the Test Mgmt R710 server.</p> <p>Login VSI launchers will be created on this server that will initiate VMware View sessions from each VM to simulate VDI workload.</p> <p>The VMs are loaded with the Windows 8 Enterprise 32bit operating system.</p>
<b>Network</b>	<p>1 x Dell PowerConnect 6248 1Gb Ethernet Switch</p>	<p>Deployment Stack and Test Stacks configured on same physical switch.</p>
<b>Performance Monitoring</b>	<p>VMware Virtual Center 5.1</p>	<p>Performance data will be captured from VSphere client.</p>

## Test results and analysis

The primary focus of the tests is to determine the maximum number of desktops that can be deployed with Windows 8 using VMware Horizon View 5.2 without compromising performance. All tests included a single Management server along with a single Compute Host, for hosting virtual desktops.

To determine the density data all tests were conducted on the compute host. The virtual desktops created using VMware Horizon View are placed on local Tier 1 storage.

The primary objective of the testing are:

- Determine the CPU, Memory, Disk Latency and Network impact of integrating Windows 8 on VMware Horizon View 5.2 on to the VDI stack using the Basic, Standard and Premium task worker.
- Determine the performance impact of Windows 8 on the local disk during peak I/O activity such as boot storms and login storms.

- 
- 

### ***Test: Basic Run (115 Users)***

The following validation was done for 115 Basic users on R720 host with 2.9 GHz processor and 10 146GB 15K disks. Validation was performed using DVS standard testing methodology using Login VSI to manage the VMware View linked clones and stratosphere UX and ESXi to deliver performance results. The CPU usage for this test reached 88.8% thus confirming that the server with this configuration can support up to 115 basic users.

The graphs below show the CPU, Consumed Memory, Local disk IOPS, Disk Latency, Network and VDI UX scatter plot results from this validation.




Figure 1

As seen from the above graph maximum CPU utilization was approximately 88.8%. After all users started logging in, CPU usage was spiked from 9 to 90 and became stable at 88% once all users logged in and dropped as the users logged off.






Figure 2

A Spike is evident after all VMs start logging in however active memory is less than 50% of available memory during the peak of validation activity.




Figure 3

As seen from the above graph, overall network performance was good.




Figure 4

A Spike is evident after all VMs start logging in however the read/write activity was fluctuating from 12:25 PM to 1:35 PM and became stable after all users finished logging in from 1:40 PM to 2:35 PM. Total Peak IOPS measured during the login state was 1906 giving an IOPS value of 16.57 per user and total peak IOPS measured during the steady state (after all users logged in) was 744 giving an IOPS value of 6.46 per user.




Figure 5

As seen from the above graph, overall disk latency is below 20ms which is in acceptable range.




Figure 6

From the above graph we can see that all 115 sessions registered perfect performance which is within the acceptable tolerance for this test.

**Test: Standard Run (95 Users)**

The following validation was done for 95 Standard users on R720 host with 2.9 GHz processor and 10 146GB 15K disks. Validation was performed using DVS standard testing methodology using Login VSI to manage the VMware View linked clones and stratosphere UX and ESXi to deliver performance results. The CPU usage for this test reached 88.71% thus confirming that the server with this configuration can support up to 95 standard users.

The graphs below show the CPU, Consumed Memory, Local disk IOPS, Disk Latency, Network and VDI UX scatter plot results from this validation.




Figure 1

As seen from the above graph maximum CPU utilization was approximately 87.71%. After all users

started logging in, CPU usage was spiked from 8 to 92 and became stable between 82-87% once all users logged in and dropped as the users logged off.




Figure 2

A Spike is evident after all VMs start logging in however active memory is less than 50% of available memory during the peak of validation activity.




Figure 3

As seen from the above graph, overall network performance was good.




Figure 4

A Spike is evident after all VMs start logging in however the read/write activity was fluctuating from 1:55 PM to 2:55 PM and became stable after all users finished logging in from 2:56 PM to 2:45 PM. Total Peak IOPS measured during the login state was 972 giving an IOPS value of 10.23 per user and total peak IOPS measured during the steady state (after all users logged in) was 578 giving an IOPS value of 6.08 per user.




Figure 5

As seen from the above graph, overall disk latency is below 20ms which is in acceptable range.




Figure 6

From the above graph we can see that all 95 sessions registered perfect performance which is within the acceptable tolerance for this test.

***Test: Premium Run (80 Users)***

The following validation was done for 80 Premium users on R720 host with 2.9 GHz processor and 10 146GB 15K disks. Validation was performed using DVS standard testing methodology using Login VSI to manage the VMware View linked clones and stratosphere UX and ESXi to deliver performance results. The CPU usage for this test reached 85.77% thus confirming that the server with this configuration can support up to 80 standard users.

The graphs below show the CPU, Consumed Memory, Local disk IOPS, Disk Latency, Network and VDI UX scatter plot results from this validation.




Figure 1

As seen from the above graph maximum CPU utilization was approximately 85.77% After all users started logging in, CPU usage was spiked from 8 to 89 and became stable between 82-85% once all users logged in and dropped as the users logged off.




Figure 2

A Spike is evident after all VMs start logging in however active memory is less than 50% of available memory during the peak of validation activity.




Figure 3

As seen from the above graph, overall network performance was good.




Figure 4

A Spike is evident after all VMs start logging in however the read/write activity was fluctuating from 9:25 AM to 10:15 AM and became stable after all users finished logging in from 10:16 AM to 11:05 AM. Total Peak IOPS measured during the login state was 1747 giving an IOPS value of 21.83 per user and total peak IOPS measured during the steady state (after all users logged in) was 560 giving an IOPS value of 7 per user.






Figure 5

As seen from the above graph, overall disk latency is below 20ms which is in acceptable range.




Figure 6

From the above graph we can see that all 80 sessions registered perfect performance which is within the acceptable tolerance for this test.

### Results Summary Table

The summary table below shows the desktop densities and storage IOPS on a per server basis for Basic, Standard and Premium users.

Workload	Server	CPU	Login State	Login	Steady State	Steady State
----------	--------	-----	-------------	-------	--------------	--------------

	Density		IOPS	State IOPS per User	IOPS	IOPS per User
<b>Basic</b>	115	88.8%	1906	16.57	744	6.46
<b>Standard</b>	95	88.71%	972	10.23	578	6.08
<b>Premium</b>	80	85.77%	1747	21.83	560	7

Sustained IOPS per user can be calculated by taking the sustained IOPS on the ESXi Server local drive and dividing it by number of users. For basic workload, per user IOPS will be  $744/115 = 6.46$  per user. Similarly we can calculate it for standard and premium workload.

## Conclusion

Based on testing conducted the impact on storage was found to be minimal. While evaluating the desktop density per server, server CPU for windows 8 is very high compared to Windows 7 which is reducing the user density for each workload. Hence we can conclude on the basis of this testing that 115 Basic, 95 Standard and 80 Premium user sessions can be supported on the R720 platform with less than 90% CPU utilization and minimal impact on performance.

## 13.9 Compellent 6.3 code update

Based on testing customers should take advantage of the performance improvements as shown below. As seen in graphs below almost linear scaling occurs out to ~8000 user sessions and the 6.3 code shows huge improvements over the previous 6.1.2 codebase.

### Test Configuration

#### **Storage Center:**

- 2 - SC8000 Controllers (SCOS 6.3.1)
- 16 – 8 Gbps FC Front End Ports
- 16 – 6 Gbps SAS Back end ports
- 20 – 200 GB SSD drives (19 Active + 1 Hot Spare)
- 496 – 400 GB 10k Drives (491 Active + 5 Hot Spares)

#### **Servers:**

- 48 – PowerEdge M620 Blades with 320 GB RAM w/ Dual 8-Core Intel Xeon 2.70GHz Processors (2-8Gbps FC ports)
- 6 – ESXi Clusters with 8 Hosts each

#### **Software and Operating Systems:**


- Hypervisor: VMware ESXi 5.0
- Connection Broker: VMware View 5.1 w/Composer for Linked Clones
- Guests: Windows7 1vCPU & 1024 MB RAM

## Load Generation:

VMware View Planner 2.1

Load Profile – 20 Second think time (~3.6-4.5 IOPS per guest)

### VDI Scale Validation - SC8000 Steady State Latency - 6.1.2 vs. 6.3.1



## 13.10 Converged Networking

### Executive Summary

As enterprise IT organizations support increasingly complicated and diverse networking requirements and business needs, new solutions are needed to keep cost, complexity, and management overhead to a minimum. Simpler approaches are being sought to solve these problems without compromising availability, performance, or user experience. Dell's Active Infrastructure provides a pre-integrated converged networking solution that allows today's fast paced IT organizations to focus on providing value to the business and end user while also reducing cost, complexity, and management overhead in the networking and overall computing infrastructure.

### Introduction

Dell's Active Infrastructure is a family of converged infrastructure offerings that combine servers, storage, networking, and infrastructure management into an integrated system that provides general purpose virtual resource pools for applications and private clouds. These systems blend intuitive infrastructure management, an open architecture, flexible delivery models and a unified support model to allow IT to rapidly respond to dynamic business needs, maximize efficiency and strengthen IT service quality. Active Infrastructure includes vStart, as well as Active System, a new converged infrastructure offering from Dell.

One of the key attributes of the DVS Enterprise - Active System 800v is the convergence of SAN and LAN over the same network infrastructure. LAN and iSCSI SAN traffic share the same physical connections from servers to storage. The converged network is designed using Data Center Bridging

(IEEE 802.1) and Data Center Bridging Exchange (IEEE 802.1AB) technologies and features. The converged network design drastically reduces cost and complexity by reducing the components and physical connections and the associated efforts in deploying, configuring, and managing the infrastructure.

### **Traditional Networking Approach**

Traditionally, IT managers and administrators have maintained multiple networks to meet different connectivity requirements. Typically, this included separate networks for low-latency server clustering, a network for client-server connectivity and management and finally, a network for storage connectivity. The clustering network typically was based on InfiniBand in high-performance environments. The management and client-server connectivity made use of IP networks over Ethernet and the storage network was based on Fibre Channel or more recently iSCSI. This model requires increased investment in data center infrastructure including the number of switches, switch ports, cables, and interface and mezzanine cards required to connect every server with its corresponding networks. The resulting network lacked the level of flexibility that is needed in today's dynamic data centers.

### **Dell Active Infrastructure Approach**

With Dell Active Infrastructure, IT managers and administrators enjoy the benefits of a flexible networking solution that consolidates networking resources, incorporates advanced levels of Ethernet functionality, and reduces complexity and overall maintenance and administration of networking resources. This model reduced the number of switches, switch ports, cables, and interface and mezzanine cards required to connect every server with its corresponding networks. The resulting network is incredibly flexible and enables IT managers and administrators to dynamically change to support today's data center requirements.

A key component of Dell Active Systems offerings, Dell Force10 data center networking switches enable converged, intelligent, and scalable switching capabilities for 10GbE and 40GbE fabrics. Force10 switches comply with Dell Virtual Network Architecture (VNA) standards-based interoperability, adaptive workload intelligence, and efficient scalability, and enable organizations to deploy a single physical network infrastructure that provides deterministic performance and QoS. Force10 switches are specifically designed to work seamlessly with Dell storage area network (SAN) arrays and the underlying Dell Fluid Data architecture to help deliver automated, optimized performance.

### **High Level Solution**

Dell pre-integrated Active Systems allow IT organization to deploy new data center platforms modularly, on a just-in-time basis. Each system leverages Dell's extensive enterprise experience paired with joint engineering with leading virtualization providers, consistent platform reference architectures, an efficient merge center, and comes complete with Dell deployment services.

IT organizations can significantly accelerate the rollout of new IT services, while ensuring predictable and consistent deployment of each additional platform. Each Active System is centrally managed by Active System Manager. The Dell Active System portfolio spans multiple scalable starting points, and easily integrates into your existing data center fabric.

One of the key components of the Active Systems converged networking solution are Dell Force10 switches. Force10 switches are available as either discrete components or embedded into a converged, chassis-based architecture. They offer an efficient modular design with built-in Ethernet stacking capability for cost-effective scalability. In addition to the capability to stack these switches across chassis, their switching functionality also allows for enhancing efficient east-west traffic flow

while optimizing north-south uplink capacity. Other features include flexible connectivity options through FlexIO modular switch technology that help deliver support for a variety of modules and cabling standards. These switches also support and integrate IP storage through a single 10GbE connection for enhanced efficiency. A performance-optimized design helps provide high performance and low latency for dynamic workloads and demanding east-west traffic flow. Additional features include access control lists (ACLs), internal traffic distribution and prioritization, high port density, and power and cooling optimization.

### **Converged Network Features**

Leveraging the Fluid Data architecture and Virtual Network Architecture (VNA) standards, Dell Force10 data center networking switches and Dell EqualLogic PS Series storage systems offer features such as storage and network virtualization, automatic configuration, and services orchestration. Compliance with the DCB-standard extensions also enables Dell data center fabric solutions to work better together in a range of topologies such as scale-up, scale-out, and data center-in-a-box architectures. Dell storage and networking can also be architected to support a range of specialized data center topologies that leverage advances such as converged 10GbE connectivity efficiently and cost-effectively.

### **Storage Performance**

Ensuring performance, responsiveness, and reliability of enterprise storage is an important consideration for organizations. As a result, many organizations rely on separate, dedicated networks to meet service requirements.

Dell Force10 data center networking switches are designed to deliver seamless, automated performance and QoS for the full range of Dell storage offerings. On-board intelligence enables Force10 switches to automatically detect and identify Dell storage systems on the network using the Link Layer Discovery Protocol (LLDP). The network is then configured automatically for the optimized data flow requirements of a particular storage system. For verification, Dell has tested and certified Force10 interoperability with Dell EqualLogic storage in a variety of data center application environments. Force10 switches are optimized to meet the demands of a range of different network and storage topologies, including scale-out and scale-up architectures and converged, chassis-based, data center-in-a-box infrastructures.

### **Scale-up and Scale-out Architecture**

Organizations that depend on business agility often look for ways to maximize return on investment. One way to help accomplish this goal for their IT environments is to consolidate different types of network traffic onto a single 10GbE fabric without sacrificing performance and control. A scale-up or scale-out architecture based on Dell EqualLogic storage arrays and Dell Force10 data center networking switches enables consolidation of both iSCSI and LAN traffic of storage data on a single 10GbE fabric.

EqualLogic PS Series storage features a frameless peer storage architecture with multipath I/O technology that is designed to scale performance and capacity together to meet increasing demand. EqualLogic PS Series iSCSI arrays include features such as storage virtualization, automated data tiering, and dynamic load balancing to enhance data center efficiency and scalability.

Force10 switches for data center networking leverage DCB-standard extensions to help ensure that an EqualLogic PS Series–based architecture delivers the intended performance over a converged 10GbE fabric. This converged network scale-out architecture not only helps organizations leverage cost, efficiency, and flexibility benefits of 10GbE, but also helps enhance storage efficiency and streamline management and operations.

## Solution Details

The DVS Enterprise - Active System 800v design is based upon a converged network. All LAN and iSCSI traffic within the solution share the same physical connections. The following section describes the converged network architecture of DVS Enterprise - Active System 800v.

### Converged Networking Architecture

Connectivity between hypervisor hosts and converged network switches: The compute cluster hypervisor hosts, PowerEdge M620 blade servers, connect to the Force10 S4810 switches through the PowerEdge M I/O Aggregator I/O Modules in the PowerEdge M1000e blade chassis. The management cluster hypervisor hosts, PowerEdge R620 rack servers, directly connect to the Force10 S4810 switches.

### Components

Converged networking within the Active System 800v solution consists of Dell provided hardware and industry standard protocols. The hardware consists of Dell Force10 S4810 switches, Dell PowerEdge M I/O Aggregator modules, Broadcom 57810-k Dual port 10GbE KR Blade NDCs, and EqualLogic PS6110 iSCSI SAN arrays. The combination of these systems enables DVS Enterprise to utilize these technologies, features, and capabilities to support the converged network architecture. The protocols supported by the converged networking architecture are discussed further in this section.

### Data Center Bridging (DCB)

A collection of standards that defines a unified 802.3 Ethernet media interface. DCB enables a lossless Ethernet fabric for converging multiple networks into a single common converged network infrastructure while ensuring QoS and reliability (see Figure 1). Some of the standards defined within DCB (PFC, ETS, and DCBX) are discussed in more detail below.




Figure 100

### Priority Flow Control (PFC)

This capability provides zero packet loss under congestion by providing a link level flow control mechanism that can be controlled independently for each priority.

### Enhanced Transmission Selection (ETS)

This capability provides a framework and mechanism for bandwidth management for different traffic types by assigning bandwidth to different frame priorities.

## Data Center Bridging Exchange (DCBX)

This functionality is used for conveying the capabilities and configuration of the above features between neighbors to ensure consistent configuration across the network.

## How Converged Networking Components Work Together

This section goes into detail about different converged networking components in the DVS Enterprise -Active System 800v solution work together to provide to provide a comprehensive solution.

### Connectivity between the two converged network switches

The two Force10 S4810 switches are connected using Inter Switch Links (ISLs) using two 40 Gbps QSFP+ links. Virtual Link Trunking (VLT) is configured between the two Force10 S4810 switches. This design eliminates the need for Spanning Tree-based networks; and also provides redundancy as well as active-active full bandwidth utilization on all links.

Connectivity between the converged network switches and iSCSI storage arrays: Each EqualLogic PS6110 array in DVS Enterprise - Active System 800v uses two controllers. The 10Gb SFP+ port on each EqualLogic controller is connected to the Force10 S4810 switches. This dual controller configuration provides high availability and load balancing.

Figure 2 below illustrates the resultant logical converged network connectivity within the DVS Enterprise - Active System 800v solution.




Figure 101

## Converged Network Configuration

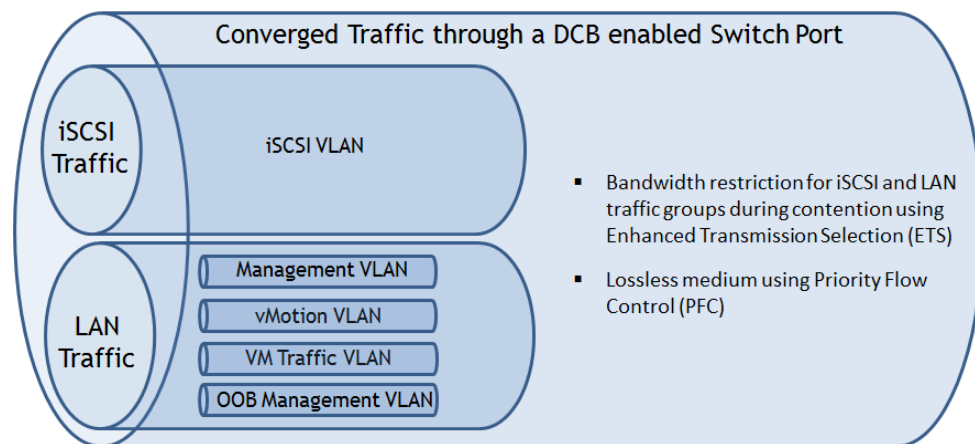
Within the DVS Enterprise - Active System 800v environment, DCB settings are configured within the Force10 S4810 switches. Utilizing the DCBX protocol, these settings are then automatically

propagated to the PowerEdge M I/O Aggregator modules. Additionally, the DCB settings are propagated to the network end nodes, including the Broadcom Network Adapters in PowerEdge R620 rack servers, the Broadcom NDCs in the PowerEdge M620 blade servers, and the EqualLogic PS6110 storage controllers. The DCB settings are not propagated to the Force10 S55 out-of-band management switch and the associated out-of-band management ports but the out-of-band management traffic going to the core from Force10 S55 switch traverses through the Force10 S4810 switches. When the out-of-band management traffic traverses through the Force10 S4810 switches, it obeys the DCB settings.

DCB technologies enable each switch-port and each network device-port in the converged network to simultaneously carry multiple traffic classes, while guaranteeing performance and QoS. In case of DVS Enterprise - Active System 800v, DCB settings are used for the two traffic classes: (i) Traffic class for iSCSI traffic, and (ii) Traffic class for all non-iSCSI traffic, which, in the case of DVS Enterprise - Active System 800v, are different LAN traffic types. DCB ETS settings are configured to assign bandwidth limits to the two traffic classes. These bandwidth limitations are effective during periods of contention between the two traffic classes. The iSCSI traffic class is also configured with Priority Flow Control (PFC), which guarantees lossless iSCSI traffic.

The Broadcom Network Adapters and the Broadcom NDCs support DCB and DCBX. This capability, along with iSCSI hardware offload, allows DVS Enterprise - Active System 800v solution to include an end-to-end converged network design, without requiring support from the VMware vSphere hypervisor for DCB.

Figure 3 below provides a conceptual view of converged traffic with Data Center Bridging in DVS Enterprise - Active System 800v.



**Figure 102 - Conceptual View of Converged Traffic Using DCB**



## Virtual Link Trunking (VLT) for S4810s

Inside each DVS Enterprise - Active System 800v, a Virtual Link Trunking interconnect (VLTi) is configured between the two Force10 S4810 switches using the Virtual Link Trunking (VLT) technology. VLT peer LAGs are configured between the PowerEdge M I/O Aggregator modules and Force10 S4810 switches, and also between the Force10 S4810 switch and the Force10 S4810 switches.

Virtual Link Trunking technology allows a server or bridge to uplink a single trunk into more than one Force10 S4810 switch, and to remain unaware of the fact that the single trunk is connected to two different switches. The switches, a VLT-pair, make themselves appear as a single switch for a connecting bridge or server. Both links from the bridge network can actively forward and receive traffic. VLT provides a replacement for Spanning Tree-based networks by providing both redundancy and active-active full bandwidth utilization.

Major benefits of VLT technology are:

1. Dual control plane on the access side that lends resiliency.
2. Full utilization of the active LAG interfaces.
3. Rack-level maintenance is hitless and one switch can be kept active at all times.

Note that the two switches can also be stacked together. However, this is not recommended, as this configuration will incur downtime during firmware updates of the switch or failure of stack links.

## NIC Partitioning (NPAR) Configuration

In DVS Enterprise - Active System 800v, each port of the Broadcom 57810-k Dual port 10GbE KR Blade NDCs in the PowerEdge M620 blade servers, and the Broadcom 57810 Dual Port 10Gb Network Adapters in PowerEdge R620 rack servers is partitioned into four ports using NPAR to obtain a total of eight I/O ports on each server. As detailed in the subsequent sections, one partition each on every physical I/O port is assigned to management traffic, vMotion traffic, VM traffic and iSCSI traffic.

The Broadcom NDC and the Broadcom Network Adapter allow setting a maximum bandwidth limitation to each partition. Setting maximum bandwidth at 100 will prevent the artificial capping of any individual traffic type during periods of non-contention. NPAR maximum bandwidth settings may be modified to limit the maximum bandwidth available to a specific traffic type, regardless of contention.

The Broadcom NDC and the Broadcom Network Adapter also allow setting relative bandwidth assignments for each partition. While utilizing NPAR in conjunction with Data Center Bridging (DCB) and Data Center Bridging Exchange (DCBX), the relative bandwidth settings of the partitions are not enforced. Due this fact, the relative bandwidth capability of the Broadcom NDCs and the Broadcom Network Adapters are not utilized in DVS Enterprise - Active System 800v.

## iSCSI hardware offload

In DVS Enterprise - Active System 800v, iSCSI hardware offload functionality is used in the Broadcom 57810-k Dual port 10GbE KR Blade NDCs in the PowerEdge M620 blade servers, and also in the Broadcom 57810 Dual Port 10Gb Network Adapters in the PowerEdge R620 rack servers. The iSCSI offload protocol is enabled on one of the partitions on each port of the NDC or the Network Adapter.

With iSCSI hardware offload, all iSCSI sessions are terminated on the Broadcom NDC or on the Broadcom Network Adapter.

### Traffic isolation using VLANs

Within the converged network, the LAN traffic is separated into four unique VLANs; one VLAN each for management, vMotion, VM traffic, and out-of-band management. The iSCSI traffic also uses a unique VM. Network traffic is tagged with the respective VLAN ID for each traffic type in the virtual switch. Routing between the management and out-of-band management VLANs is required to be configured in the core or the Force10 S4810 switches. Additionally, the Force10 S4810 switch ports that connect to the blade servers are configured in VLAN trunk mode to pass traffic with different VLANs on a given physical port. Table 4, below, provides an overview of different traffic types segregated by VLANs in the DVS Enterprise - Active System 800v, and which edge devices they are associated with.

Traffic Type (VLAN segregation)	Description	Associated Network Device
Management	vSphere management traffic and DVS Enterprise - Active System 800v management services	Broadcom NDC and Broadcom Network Adapter
vMotion	VMware vMotion traffic	Broadcom NDC and Broadcom Network Adapter
VM	LAN traffic generated by compute cluster VMs	Broadcom NDC and Broadcom Network Adapter
iSCSI	iSCSI SAN traffic	Broadcom NDC and Broadcom Network Adapter
Out-of-Band Management	Out-of-Band Management traffic	iDRAC, CMC, and EqualLogic Management Ports

Figure 103

Hypervisor network configuration for LAN and iSCSI SAN traffic: VMware ESXi hypervisor is configured for the LAN and iSCSI SAN traffic that are associated with the blade servers. LAN traffic in DVS Enterprise - Active System 800v solution is categorized into four traffic types: VM traffic, management traffic, vMotion traffic, and Out-of-Band (OOB) management traffic. OOB management traffic is associated with CMC, iDRAC, and EqualLogic SAN management traffic. VM traffic, management traffic, and vMotion traffic are associated with the blade servers in the compute cluster and the rack servers in the management servers. Similarly, iSCSI SAN traffic is also associated with the blade servers and the rack servers. On each hypervisor host within the compute cluster and the management cluster, a virtual switch is created for each of the three LAN traffic types associated with the blade and the rack servers, and also for the iSCSI traffic.

On the compute cluster hosts (the PowerEdge M620 blade servers), one vSwitch each is created for VM traffic, vSphere management traffic, vMotion traffic, and iSCSI traffic. Two partitions, one from each physical network port, are connected as uplinks to each of the virtual switches. This creates a team of two network ports, enabling NIC failover and load balancing for each vSwitch. On the management cluster hosts (the PowerEdge R620 rack servers), one vSwitch each is created for management traffic, vMotion traffic, and iSCSI traffic. In this case, all VMs are management VMs, so the VM traffic and the vSphere management traffic are on the same management VLAN. Due to this fact, the VM traffic port group and the vSphere management traffic port group are on the same vSwitch.

The resultant compute cluster and management cluster hypervisor host configuration is illustrated in

Figure 5.

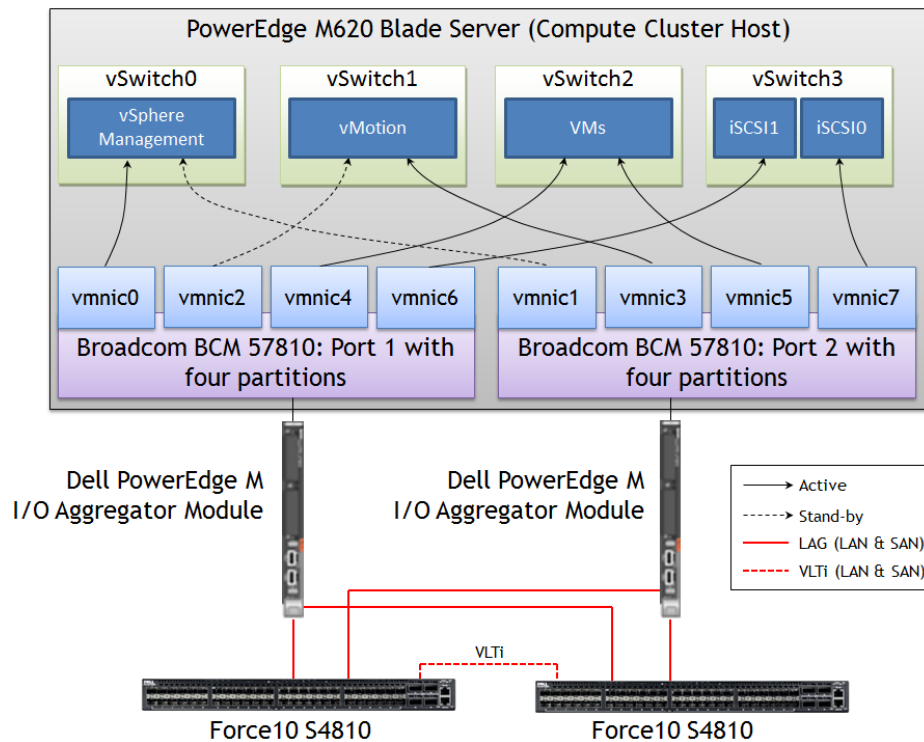


Figure 104 - vSwitch and NPAR Configuration for the Hypervisor Hosts

### Load Balancing and Failover

This solution uses *Route based on the originating virtual switch port ID* configuration at the vSwitch for load balancing the LAN traffic. Any given virtual network adapter will use only one physical adapter port at any given time. In other words, if a VM has only one virtual NIC, it will use only one physical adapter port at any given time. The reason for choosing this option is that it is easy to configure and provides load balancing across VMs, especially in the case of a large number of VMs.

### Uplinks

There are several options to uplink the Force10 switches to the core network. Selecting the uplink option depends on your core network and your requirements. One simple option is to create multiple uplinks on each switch and connect them to the core network switches. Uplink LAGs can then be created from the Force10 S4810 switches to the core network.

### Business Benefits

There are a multitude of business benefits from installing and utilizing converged networking in VDI environments. Some of them are measurable by using metrics such as support desk calls, while others are not, such as ease of use.

Some of the largest benefits can be summarized by the following points:

- Reduced CapEx (purchasing of hardware, licenses, etc)
- Reduced OpEx (simplifies operations and management required)
- Scalable & modular integrated system that delivers faster time to value
- Up to 50% reduction in infrastructure including adapters, switches, cables, and management ports over a traditional approach
- Consolidation of server access network elements, reducing potential points of failure and overall

- capital and operational costs for data centers
- Single point of contact for technical support

## Summary

Dell's Active Infrastructure is a family of converged infrastructure offerings that combine servers, storage, networking and infrastructure management into an integrated system that provides general purpose virtual resource pools for applications and private clouds. It provides a pre-integrated converged networking solution that allows today's fast paced IT organizations to focus on providing value to the business and end user while also reducing cost, complexity, and management overhead in their overall computing infrastructure.

## 13.11 Horizon Workspace

### Executive Summary

IT organizations today are faced with an increasingly mobile workforce that need access to data, applications, and communication methods to effectively collaborate, communicate and be productive. Traditional approaches typically involve incremental solutions that are difficult to manage, troubleshoot, and don't scale as users increase and mobile devices become more diverse.

VMware Horizon Workspace provides a complete solution that provides end users with a simple and seamless computing experience to access applications, data, and desktops from a variety of computing devices, both personal and business. It enables IT administrators to easily manage a diverse array of mobile devices in a unified and scalable manner, quickly deploy new applications as business requirements demand, and provide the security required to keep corporate data and intellectual property safe and secure.

### Introduction

In the past, IT organizations did not have to worry about providing a robust user experience for remote users. Mobile devices were expensive and impractical for large scale use. Internet connections were slow and frustrating. Largely, remote users were provided with a laptop that had a pre-installed operating system with all of the required software. Virtual Private Networks (VPNs) were used to connect from remote locations into the corporate network and access shared resources and files. Collaboration was difficult as communication was limited to email and phone calls.

Today, with the advent of new technologies, lower prices, and the availability of a diverse portfolio of mobile products, the corporate workforce is becoming increasingly mobile and global, IT administrators are struggling to provide computing environments where users are connected to the applications, data, and people they need, while also providing security for both the corporation and the users. IT organizations who do not react to the changing IT landscape risk users seeking alternative ways of collaborating and communicating, which could potentially expose private information and cause many security risks. The issues discussed can be tackled in a number of ways, which are discussed in the following sections.

### Traditional Approach

The traditional approach to addressing an increasingly mobile workforce has been to implement incremental solutions that address individual components and aspects of the computing experience. Typically, separate management applications had to be installed with separate administrative consoles for computing aspects such as user applications, files, communication, security, and device management. This approach was difficult to manage and not scalable as users increased. Mobile devices were typically provided strictly by the IT organization and only subsets of the employees were

provided with these.

Personal computing devices were out of the question because of the level of cross-over between personal and business data. If the device was lost or stolen, unsecured data on the device could be compromised and result in data breaches. Mobile employees typically needed to carry multiple devices to stay connected while on the go. It wasn't uncommon for someone to carry two phones and two laptops to keep connected to the applications, email, data, and people they needed in their personal and professional lives.

The lack of good mobile solutions and management capabilities caused IT organizations to be unable to provide alternative solutions for employees at all outside of the corporate network. This had a multitude of affects ranging from reduced productivity to security concerns. Users would find ways of circumventing the system and using public ways of communicating, storing files, and applications. This means that that information is completely out of the control of the IT administrators and a huge security risk for the company. Additionally, many organizations would fall out of compliance with regulatory bodies in their respective countries.

### **VMware Horizon Workspace Approach**

VMware Horizon Workspace takes a drastically different approach to managing an increasingly mobile and global workforce. It embraces the aspects of modern mobile computing and empowers employees to increase productivity. Employees are able to collaborate, access their data and applications, and communicate no matter where they are or what device they are using. IT administrators are able to easily manage a large variety of devices and applications to provide a consistent and robust user experience while maintaining the security and control that they need to ensure that corporate data is safe and secure.

Employees who want to bring their own devices to the workplace, which has been commonly referred to as "Bring your own device" or BYOD, is increasingly common and actually encouraged. This allows them the flexibility to use the devices they are most familiar and comfortable with. Personal computing devices are seamlessly managed and secured using policies and controls through the Horizon Workspace. A diverse portfolio of applications and devices can be easily managed and maintained while keeping costs low. Devices that are lost or stolen are much more secure as the data is encrypted and can be remotely wiped. Users prefer to carry a single mobile device that allows them access to both personal and corporate data and applications. Horizon Workspace provides this by providing a workspace or "sandbox" within the employee's device that is isolated, encrypted, and protected. Personal and corporate data is kept completely separate and the corporate workspace is completely enterprise-owned and controlled.

### **High Level Solution**

VMware Horizon Workspace provides a simplified end-user experience for an increasingly remote and global workforce. It also streamlines IT management for end-user mobility by combining applications and data into a single aggregated workspace. This workspace contains the data and applications necessary for employees to be productive, regardless of where they are based: at the office, at home, or on the go. For the administrator, the result is fewer management points and easier access. End users gain freedom of mobility through anytime, anywhere access.

### **VMware Horizon Workspace Features**

VMware Horizon Workspace enables enterprises to securely customize and manage an employee's corporate workspace in isolation from an employee's personal environment on smartphones, keeping personal and business assets separate. This solution provides greater security for corporate data while lowering the liability risks of bring-your-own-device (BYOD) programs.

Horizon Workspace is architected for a dynamic, multi-device workforce. IT administrators can allocate applications to users and groups of users—instead of to their devices—and users can self-provision rather than contact the helpdesk every time they need a new application. The result is that IT administrators can deliver a workspace with the right applications and data, on any device, on the fly. The company can then easily add new devices, new users or new applications to a user group without needing to reconfigure the devices or endpoints. Because it is centrally managed, IT administrators have a single point from which to apply user policies, making the entire solution always secure.

Typical users of VMware Horizon Workspace are IT administrators, help desk specialists, and end users. IT administrators can use the management points to control access to applications, files, and data. Help desk specialists can quickly manage resources relating to permissions and access control. Users get the benefit of a common experience while being able to access their corporate files and applications from anywhere.

Figure 1 below shows a high level overview of how users accessing resources from different devices, both local and mobile to the corporate network, are provided a common experience.

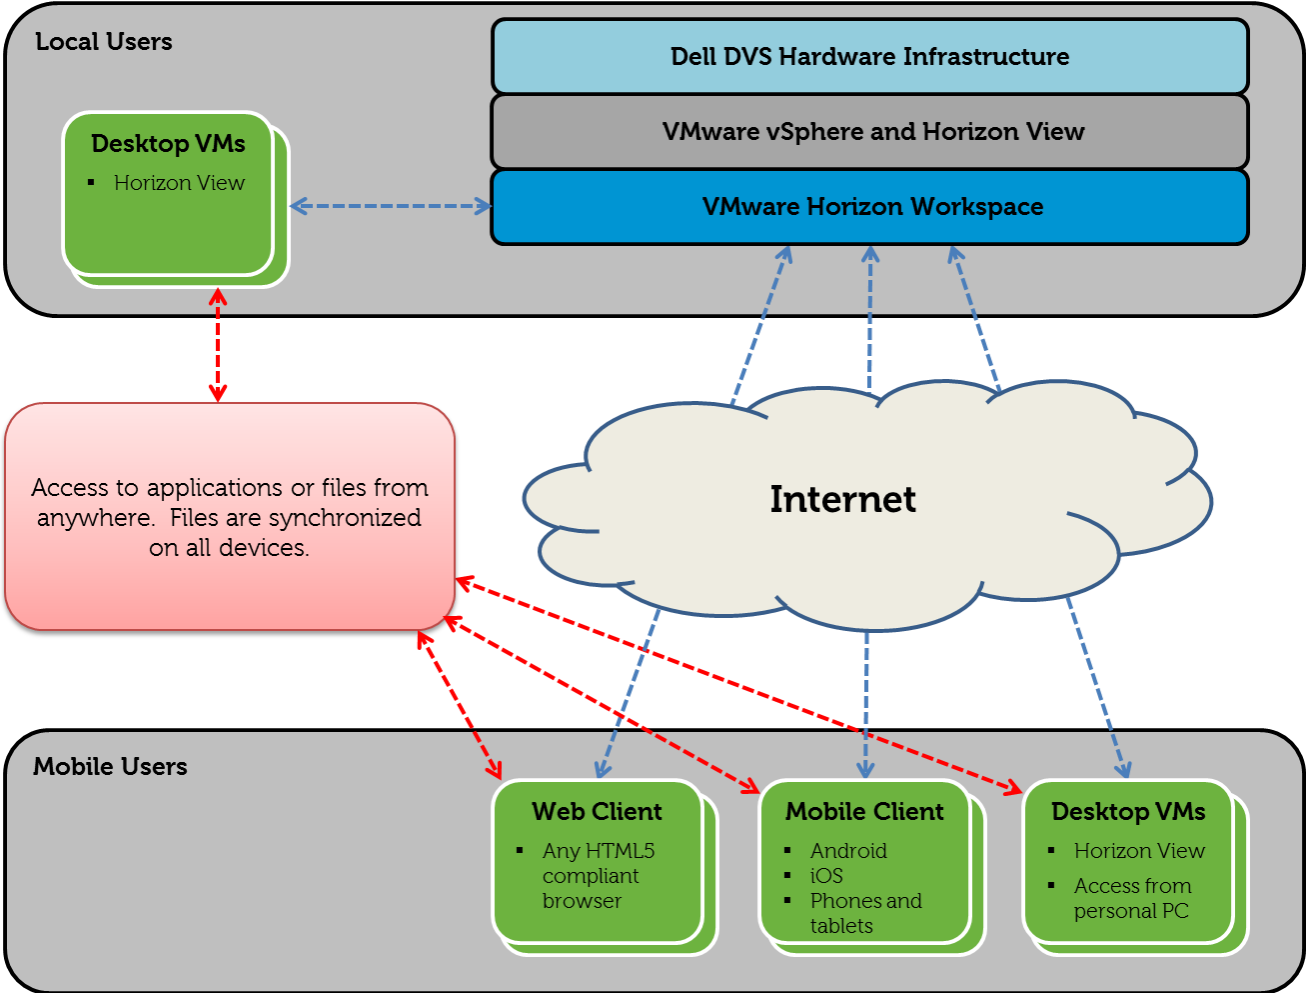


Figure 105

**Solution Details**

**VMware Horizon Workspace Architecture**

VMware Horizon Workspace is an integrated workspace that proves end users with access to their apps, data, and desktops from any of their devices while enabling IT administrators to easily manage entitlement and policy controls.

VMware Horizon Workspace is composed of the following components:

- VMware Horizon Workspace vApp
  - Configurator Virtual Appliance
  - Connector Virtual Appliance
  - Gateway Virtual Appliance
  - Management Virtual Appliance
  - Data Virtual Appliance
- VMware Horizon Workspace Client Agent
  - Web Client
  - Client for Windows
  - Client for Mac
  - Client for Android
  - Client for iOS

The purpose of each component is outlined below.

### **VMware Workspace vApp Server Component**

Horizon Workspace is a multi-virtual machine vApp, distributed as an Open Virtualization Archive (OVA) file. The VMware Horizon Workspace vApp can be directly deployed to VMware vCenter. The VMware Workspace vApp consists of 5 virtual appliances: the Configurator, Connector, Gateway, Management, and Data Virtual Appliance.

#### **Configurator Virtual Appliance**

The Configurator appliance provides a console and Web-based user interface for the centralized configuration of all other appliances in the vApp. It also provides central control of the network, gateway, vCenter, and SMTP settings of all appliances in the vApp.

#### **Connector Virtual Appliance**

The Connector appliance provides local user authentication as well as Active Directory binding and synchronization services. An IT administrator can define the directory replication schedule and also synchronize VMware Horizon View and VMware ThinApp pools and repositories for provisioning to end users.

#### **Gateway Virtual Appliance**

The Gateway appliance enables a single endpoint for all access to Horizon Workspace. As the central aggregation point for all user connections, the Gateway appliance routes requests to the appropriate destination and proxies requests on behalf of user connections.

#### **Management Virtual Appliance**

The Management appliance provides a Web-based Workspace administrative interface that allows an IT administrator to configure the application catalog, manage user entitlements, and configure groups and reporting for all the systems in the Workspace vApp.

#### **Data Virtual Appliance**

The Data appliance provides the datastore for user files, controls file sharing policy for local and mobile users, provides the file preview functionality, and serves the end-user Web interface for VMware Horizon Workspace.

## **VMware Horizon Workspace User Client Component**

Users can access Horizon Workspace with Horizon Web client (an agentless client), Windows client, Mac client, Android client, or iOS client. Each client provides users with access to the Horizon Workspace user interface, but access to applications, desktops, and data varies depending on the client.

### **VMware Horizon Workspace Web Client**

The Horizon Workspace Web Client is an agentless client. It is the default client used when users access Horizon Workspace with a browser. Using the Horizon Workspace Web Client, users can access their Horizon Workspace Data, Horizon View Desktops and Horizon Workspace Web Applications.

### **VMware Horizon Workspace Client for Windows**

When Horizon Workspace Client for Windows is installed on users' Windows systems, they can access their Horizon Workspace Data and Windows applications (captured as ThinApp packages) locally. When this client is installed, a user's personal and shared folders and files are synchronized between their system and Horizon Workspace.

### **VMware Horizon Workspace Client for Mac**

When Horizon Workspace Client for Mac is installed on users' Apple Mac OS X systems, they can access their Horizon Workspace Data locally. When this client is installed, users' personal and shared folders and files are synchronized between their system and Horizon Workspace.

### **VMware Horizon Workspace Client for Android**

When Horizon Workspace Client for Android is installed on users' Android devices, they can access their Data and Web applications. They can also install mobile applications that administrators have curated from Google Play.

### **VMware Horizon Workspace Client for iOS**

When Horizon Workspace Client for iOS is installed on users' iOS devices, they can access their Data and Web Applications. They can also install mobile applications that you have curated from the Apple App Store.

Additionally, if the deployment is configured to access Horizon View desktops, iPad users can view their entitled desktops using Horizon View Client for iOS.

## **How VMware Horizon Workspace Components Work Together**

The VMware Horizon Workspace Gateway Virtual Appliance acts as a proxy and routes queries from the client to the appropriate Virtual Appliance based on what the user is requesting access to. It acts as a single endpoint and resides between client and all other Virtual Appliance communication.

The Configurator Virtual Appliance communicates with all of the other Appliances in the vApp and provides configuration for each.

The Connector Virtual Appliance provides authentication and synchronization services to validate user connections and permissions.

The Management Virtual Appliance is used to manage general user permissions and activities.



The Data Virtual Appliance is used to store user data and synchronizes data to the different devices that an end user uses.

The different Virtual Appliances each have specific functions and work together to provide all of the aspects needed to provide a robust and consistent computing environment for the end user.

Figure 2 below outlines the basic relationship between the different virtual appliances within the VMware Horizon Workspace vApp, the external services, and the client devices.

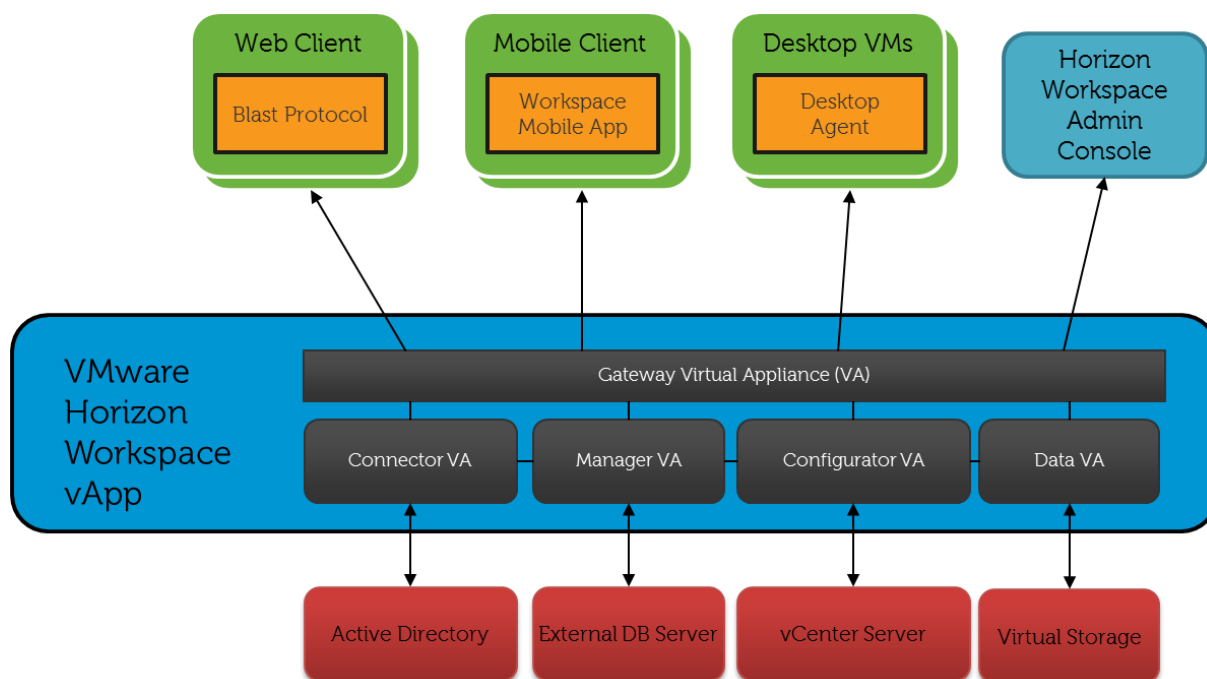


Figure 106

## Licensing

Licensing of VMware Horizon Workspace is as follows:

- Horizon Workspace and its components are licensed by per named user
- Available in quantities of 10 or 100 licenses.

## Business Benefits

There are a multitude of business benefits from installing and utilizing VMware vCenter Operations Manager for View in VDI environments. Some of them are measurable by using metrics such as support desk calls, while others are not, such as ease of use.

Some of the largest benefits can be summarized by the following points:

- Provides end users with the convenience of a self-service workspace
- Allows access to corporate application stores from any web-connected device
- Delivers simplified IT governance and visibility into end users' activity
- End users get a simple, secure login experience
- Enterprise-grade security that lowers risks for IT organizations and businesses
- Dramatically lowers cost of ownership of software-as-a-service (SaaS) and Web-based applications

## **Summary**

VMware Horizon Workspace provides a complete solution that provides end users with a simple and seamless computing experience to access to their apps, data, and desktops from a variety of devices, both personal and business. It enables IT administrators to easily manage a diverse array of mobile devices in a unified and scalable manner, quickly deploy new applications as business requirements demand, and provide the security required to keep corporate data and intellectual property safe and secure.

## 13.12 Dell Mobile Clinical Computing - AlwaysOn Point of Care

---

### Dell and VMware Solution Architecture - Secure, Reliable and Continuous Access to EMR and Patient Care Applications

Digitized medical records can reduce medical errors, improve patient safety and produce better clinical outcomes—but if the caregiver cannot access Electronic Medical Record (EMR) and clinical information systems efficiently or are required to change workstations to access particular applications, productivity suffers and user frustration can increase. Furthermore, security and HIPAA compliance requirements may stipulate that users provide login credentials for each application that is accessed, then exit from each application and log off the system when finished. These processes are cumbersome and time consuming, and the log off process may be overlooked altogether creating potential security and data protection issues.

When electronic devices replace paper charts and physician prescription pads having quick secure and reliable access to the various applications from any device in clinical workspaces becomes critical. This is why modernization of the point-of-care desktop has become an urgent priority for hospital IT professionals as well as hospital caregivers. Desktops and patient care applications must be immediately accessible and available to caregivers regardless of workstation, device or location.

Dell Mobile Clinical Computing (MCC) and VMware have a solution. The Dell Mobile Clinical Computing – VMware AlwaysOn Point of Care™ solution provides an innovative virtualization architecture solution for the specific application access and stringent security requirements of medical organizations. By virtualizing desktops with VMware Horizon View 5.2™ and hosting them on VMware vSphere™, a key component of Dell's Mobile Clinical Computing architecture, and using this validated architectural design, healthcare organizations can now have unparalleled desktop and application reliability and availability – and secure rapid access from virtually any device. The caregiver can rapidly and securely access their desktop and application workspace by tapping a proximity card, swiping their finger on a biometric device or entering their username and password and they will be immediately presented with their clinical desktop to resume caring for patients.

The Dell MCC - AlwaysOn Point of Care offering allows clinicians and staff to achieve the rapid, secure access and continuous level of availability they demand. In a typical IT environment only the production environment is backed up. But by having an active-active desktop environment running identical desktop images, even if there is a failure at the primary site, end users can promptly access their desktops and applications. If a healthcare provider's infrastructure is compromised through a natural disaster or other unexpected outage, caregivers—who are many times among the first responders—can be assured they can reach their clinical desktops and applications where and when they are needed the most.

This new Dell and VMware architectural design features continuous monitoring capabilities, as well as load balancing with constant data replication across sites to ensure that if the primary site is down VMware Horizon View 5.2 will seamlessly failover to the secondary site so the caregiver experiences minimal disruption. As a result, IT can now deliver non-stop point of care desktops with all applications and data readily available where and when they are needed most.

### 13.12.1 Benefits

- Validated Dell and VMware solution architecture
- Rapid secure non-stop access to clinical desktops and Electronic Medical Record applications

- Provides constant monitoring, load balancing and data replication features for optimizing performance and uptime
- Single sign on access, session mobility, highly reliable
- End-to-end Dell and VMware Horizon View 5.2 implementation with Dell Mobile Clinical computing Infrastructure
- High performance, secure, reliable Dell Wyse thin clients

### **13.12.2 Solution Elements**

Dell and VMware have jointly validated the Dell Mobile Clinical Computing - AlwaysOn Point of Care offering. This purpose built architecture integrates Dell's Mobile Clinical Computing solution with VMware and additional Dell technology partner solutions to meet the uniquely stringent need for security, data protection, flexible scaling, and multiple single sign on mechanisms for health care organizations.

### **13.12.3 Imprivata OneSign**

Imprivata OneSign® provides Single Sign On and strong authentication, permitting users to access all workstations and applications they are authorized to use. By configuring and linking multiple instances of virtual appliances at both sites with fault tolerance, during the site failover, desktop agents can continuously look up the next working instance without disrupting the workflow. Users easily connect to virtual desktops and applications via Imprivata OneSign single sign on from the access point all the way through to the EMR application.

### **13.12.4 Dell and VMware Solutions for secure, reliable and continuous access to EMR and patient care applications**

Dell and VMware have created the Dell Mobile Clinical Computing - AlwaysOn Point of Care of Care offering - an innovative virtualization architecture solution for the specific application access and stringent security requirements of medical organizations. It integrates technology and solutions from Dell, Dell Wyse, VMware and our thriving ecosystem of technology partners.

The solution leverages and integrates Dell's Mobile Clinical Computing and a highly available VMware virtualization infrastructure, with single sign on and secure high performance thin clients validated to health care industry requirements. Together, Dell, VMware, and our partners can deliver rapid and secure access to high performance continuous desktops, critical medical applications and records. More than that, we deliver true peace of mind for medical professionals—and the highest levels of patient care.

For additional information about how the Dell Mobile Clinical Computing with VMware AlwaysOn Point of Care offer is built and validated, please consult the Reference Architecture at: [www.dell.com/virtualdesktops](http://www.dell.com/virtualdesktops) or <http://www.vmware.com/solutions/industry/healthcare/point-of-care.html> or contact your Dell PartnerDirect Reseller [www.dell.com/partnerdirect](http://www.dell.com/partnerdirect)

## 13.13 Mobile Secure Desktop

---

Dell and VMware Solution Architecture - Mobile, Secure Access to Applications and Data

Today's workforce is no longer tethered to traditional stationary desktops. New devices have proliferated in organizations of all sizes. Employees are increasingly mobile, and more than 60% of enterprise firms and 85% of SMB organizations are looking to initiate Bring Your Own Device (BYOD) programs. But while end users are eager to embrace "bring your own device" trends, IT departments—faced with tight budgets—are challenged with how to best support and manage these new devices while protecting corporate data as it is accessed across networks and locations.

Dell and VMware have a solution. The Dell Desktop Virtualization Solutions (DVS) Enterprise - Mobile Secure Desktop architecture provides an innovative way for IT to support device diversity and BYOD initiatives by improving user access and mobility, streamlining application updates, enhancing data security, and delivering the highest-fidelity user experience. This solution virtualizes desktops and hosts them on VMware® vSphere™, a key component of VMware Horizon View 5.2™, and uses Dell's DVS Enterprise infrastructure so organizations can now have unparalleled desktop and application access across devices and locations.

And because this solution ties desktop environments to user identities instead of devices, end users are free to access their data and applications from any qualified device, whether in the office or halfway around the world. This solution addresses three key requirements including:

### 13.13.1 Mobility

The Mobile Secure Desktop solution is built on Dell DVS Enterprise and VMware Horizon View 5.2. It places desktops in the datacenter and provides secure access to applications and data from a multitude of client devices including PC workstations, thin clients and mobile device. This enables BYOD support with persistence for true session mobility across devices—so users can access the same desktop from different devices. With persona management and optional support for user-installed applications, the Mobile Secure Desktop solution provides a personalized user experience across devices and sessions. In addition to providing session persistence across devices, VMware Horizon View 5.2 uses PCoIP protocol to deliver the best desktop user experience from any device.

### 13.13.2 Security

By integrating support for two-factor authentication (RSA SecurID, RADIUS authentication), the DVS Enterprise - Mobile Secure Desktop solution emphasizes data and application security. In addition to providing the right level of access to the right resources, it also simplifies patch management and update management. Since all the desktops are hosted in a centralized infrastructure, this solution streamlines deployment of updates and patches to desktops helping ensure that no vulnerabilities exist in the environment due to outdated, unpatched or orphaned systems. This solution also integrates VMware vShield™ to provide superior security for the environment.

### **13.13.3 Management**

One of the key challenges facing organizations today is the ability to monitor and manage the entire end user compute environment including; desktops, access policies and service levels. The DVS Enterprise - Mobile Secure Desktop solution with optionally integrated vCenter Operations Manager (vCOPs), provides an integrated dashboard with intelligent response on all desktop- related events. This helps IT admins to provide the right amount of intervention and guidance when the demands on the virtual infrastructure appear to exceed an expected range of behavior. The solution can also include vCenter Configuration Manager (vCM) for importing suggested configurations and to meet required regulatory compliance standards.

A critical aspect to mobile secure desktops and true BYOD support is securely managing and enabling corporate access. Dell Wyse project Stratus offers simple, secure, cloud-based management for today's dynamic IT environment. The integration of Dell Wyse Stratus into the Mobile Secure Desktop solution provides IT administrators with an intelligent and dynamic cloud-based console to securely manage and enable corporate access to any device, including smartphones, tablets, thin clients, zero clients, and PC's – regardless of whether the device is owned by the company or by an individual employee.

### **13.13.4 Solution Elements**

Dell and VMware have jointly validated the Dell DVS Enterprise - Mobile Secure Desktop offering. This purpose built architecture integrates Dell's DVS Enterprise solution with VMware and additional Dell technology partner solutions to meet the need for flexible, mobile and secure end user computing in today's dynamic business environment.

### **13.13.5 Compliance - VMware vCenter Configuration Manager**

A key requirement for many organizations is managing compliance to various government and industry regulations. vCenter Configuration Manager (vCM) automates critical configuration and compliance management tasks including configuration data collection, configuration change execution, configuration reporting, change auditing, and compliance assessment.

### **13.13.6 Cortado ThinPrint**

Most of the use cases supported by this solution have a location-aware printing requirement. Cortado ThinPrint software, OEM by VMware, provides the ability to take advantage of location-aware printing from a wide range of devices.

### **13.13.7 Imprivata OneSign**

Imprivata OneSign® provides Single Sign On and strong authentication, permitting users to access all workstations and applications they are authorized to use. By configuring and linking multiple instances of virtual appliances at both sites with fault tolerance, during the site failover, desktop agents can continuously look up the next working instance without disrupting the workflow. Users easily connect to virtual desktops and applications via Imprivata OneSign single sign on from the access point all the way through to the EMR application.

## Dell and VMware solutions for mobile, secure access to applications and data

Dell and VMware have created the Dell DVS Enterprise - Mobile Secure Desktop solution – a validated architecture that integrates technology from Dell, VMware and the Dell technology partner ecosystem. The solution leverages Dell server, storage and networking infrastructure, mobile, wireless and wired networks, VMware Horizon View 5.2 and vSphere, vShield security services, management and monitoring components to protect data, monitor the infrastructure and secure access for virtual any end point device. This solution is optimized for organizations looking to drive higher levels of productivity by improving end-user access across devices and locations, reduce costs by streamlining desktop and application management, and enhance security and compliance.

The solution leverages and integrates Dell's DVS Enterprise end to end desktop virtualization infrastructure and highly available VMware virtualization platforms, with single sign on, personal management, and additional high performance industry recognized solutions. Together, Dell, VMware, and our partners deliver rapid and secure access to high performance desktops, business applications and corporate data with the Dell DVS Enterprise - Mobile Secure Desktop solution.

For additional information about how the Dell DVS Enterprise - Mobile Secure Desktop solution is built and validated, please consult the Dell DVS Enterprise - Mobile Secure Desktop Reference Architecture at [www.dell.com/virtualdesktops](http://www.dell.com/virtualdesktops) or <http://www.vmware.com/solutions/desktop/mobile-secure-desktop/overview.html> or contact your Dell PartnerDirect Reseller [www.dell.com/partnerdirect](http://www.dell.com/partnerdirect)

## 14 Reference

---

VMware references:

- [VMware vSphere Edition Comparisons](#)
- [VMware vSphere Availability Guide](#)
- [VMware vSphere 5.1 documentation](#)

Dell PowerEdge References:

- [Dell PowerEdge M1000e Technical Guide](#)
- [Dell PowerEdge M I/O Aggregator Configuration Quick Reference](#)

Dell EqualLogic references:

- [EqualLogic Technical Content](#)
- [Dell EqualLogic PS Series Architecture Whitepaper](#)
- [Configuring iSCSI Connectivity with VMware vSphere 5 and Dell EqualLogic PS Series Storage](#)
- [Configuring and Installing the EqualLogic Multipathing Extension Module for VMware vSphere 5.1, 5.0 and 4.1 and PS Series SANs](#)
- [How to Select the Correct RAID for an EqualLogic SAN](#)
- [Using Tiered Storage in a PS Series SAN](#)
- [Monitoring your PS Series SAN with SAN HQ](#)

Dell Management reference:

[Dell Management Plug-In for VMware vCenter references – Solution Brief](#)