



DVS Enterprise

Reference Architecture Addendum for Accelerated Graphics on



XenDesktop

12/5/2013
Phase 5
Version 1.6

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

Copyright © 2013 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

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

Contents

1 Introduction	1
1.1 Purpose	1
1.2 Scope	1
1.3 What's New in This Release	1
2 Solution Architecture Overview.....	2
2.1 Introduction	2
2.2 Physical Architecture Overview	3
2.3 DVS Enterprise – Solution Layers	4
2.4 Local Tier 1	5
2.4.1 Local Tier 1 – Rack (iSCSI – EQL)	5
2.5 Shared Tier 1.....	7
2.5.1 Shared Tier 1 – Rack (iSCSI – EQL)	7
2.5.2 Shared Tier 1 – Rack (FC – CML).....	8
3 Solution Architecture for XenDesktop 7.....	10
3.1 XenDesktop with HDX 3D Pro.....	10
3.2 Graphics Compute Server Infrastructure.....	10
3.2.1 Local Tier 1 Rack	10
3.2.2 Shared Tier 1 Rack (iSCSI).....	11
3.2.3 Shared Tier 1 Rack (FC).....	12
4 Solution Performance and Testing.....	14
4.1 Load Generation and Monitoring	14
4.1.1 Login VSI – Login Consultants	14
4.1.2 Liquidware Labs Stratusphere UX	14
4.1.3 VMware vCenter	15
4.2 Testing and Validation	15
4.2.1 Testing Process	15
4.2.2 Pass-Through Graphics Results.....	15
4.2.3 Shared Graphics Results.....	26
Acknowledgements.....	29
About the Authors	30



1 Introduction

1.1 Purpose

This document describes:

1. Dell DVS Reference Architecture specifically as it pertains to applicable pass-through and shared graphic acceleration options.

This document addresses the design, configuration and implementation considerations for the key components of the architecture required to deliver graphically-accelerated virtual desktops via XenDesktop 7 on VMware vSphere 5. Please refer to the DVS Enterprise Reference Architecture for Citrix XenDesktop for more details on the broader solution stack and design considerations.

1.2 Scope

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

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

1.3 What's New in This Release

- Shared graphics (vSGA) using NVIDIA K1 Grid cards



2 Solution Architecture Overview

2.1 Introduction

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


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

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




2.2 Physical Architecture Overview

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



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




2.3 DVS Enterprise – Solution Layers

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




The Compute layer consists of the server resources responsible for hosting the XenDesktop user sessions, hosted via VMware vSphere for local or shared tier 1 solution models (local tier 1 pictured below).




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





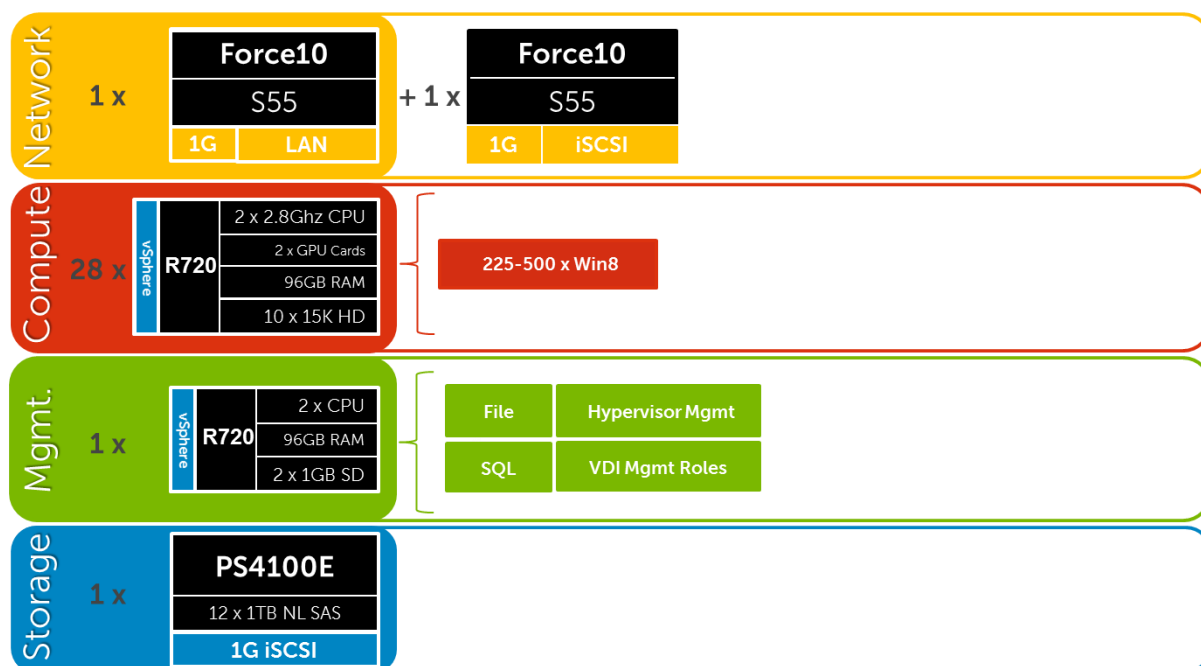
The Storage layer consists of options provided by EqualLogic for iSCSI and Compellent arrays for Fiber Channel to suit your Tier 1 and Tier 2 scaling and capacity needs.



2.4 Local Tier 1


2.4.1 Local Tier 1 – Rack (iSCSI – EQL)

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

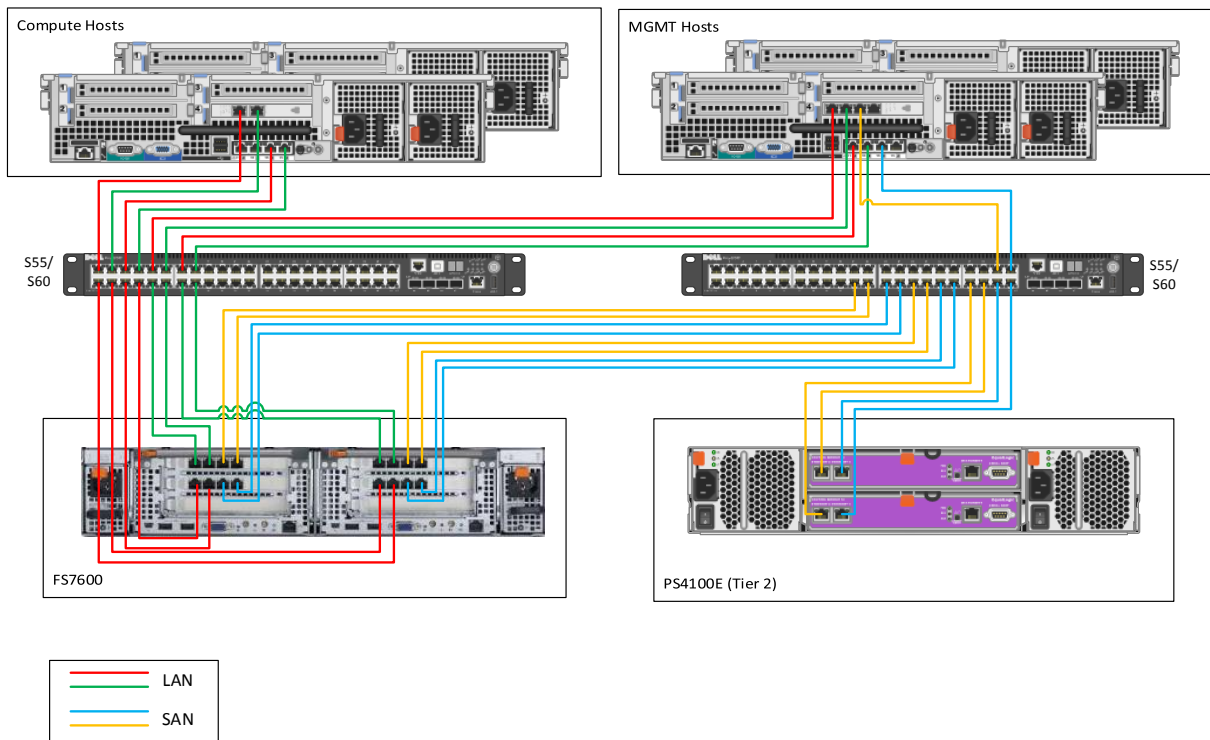


2.4.1.1 Local Tier 1 – Network Architecture (iSCSI)

In the local tier 1 architecture, a single Force10 switch can be shared among all network connections for both Management and Compute, up to 1000 users. Over 1000 users DVS recommends separating the network fabrics to isolate iSCSI and LAN traffic as well as making each switch stack redundant. Only the Management servers connect to iSCSI storage in this model. All Top of Rack (ToR) traffic has been designed to be layer 2/ switched locally, with all layer 3/ routable VLANs trunked from a core or distribution switch. The following diagrams illustrate the logical data flow in relation to the core switch.



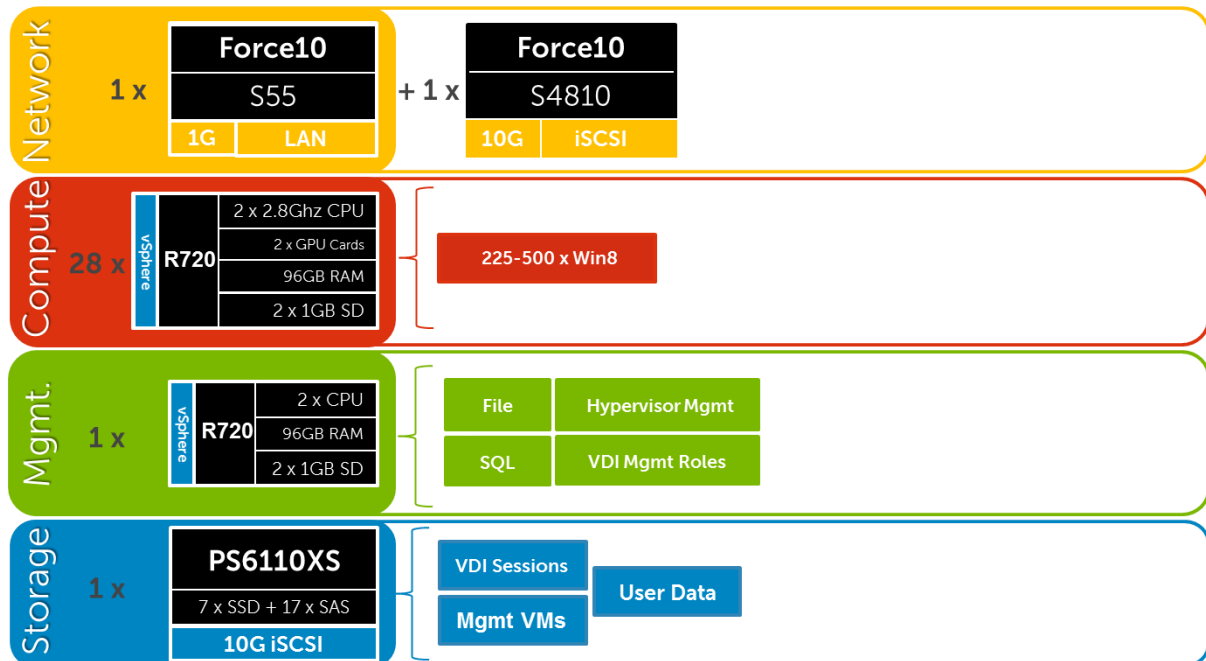
2.4.1.2 Local Tier 1 Cabling (Rack – HA)



2.5 Shared Tier 1


2.5.1 Shared Tier 1 – Rack (iSCSI – EQL)

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




2.5.1.1 Shared Tier 1 Rack – Network Architecture (iSCSI)

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

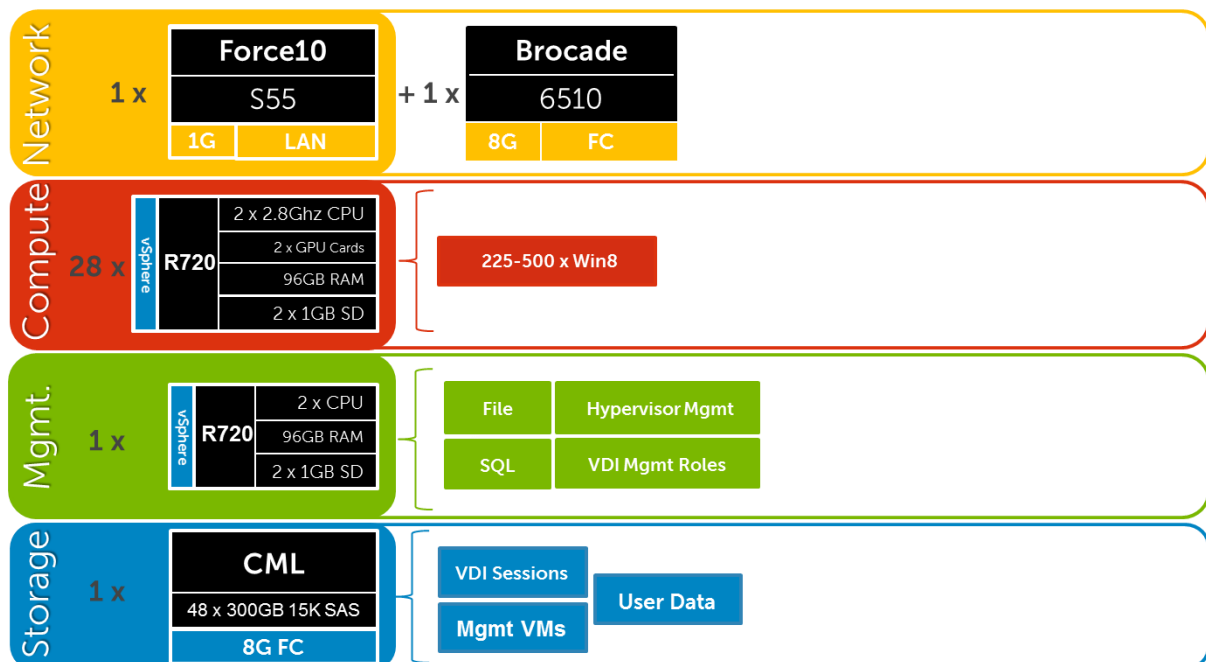


2.5.1.2 Shared Tier 1 Cabling – (Rack – EQL)




2.5.2 Shared Tier 1 – Rack (FC – CML)

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




2.5.2.1 Shared Tier 1 Rack – Network Architecture (FC)

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



2.5.2.2 Shared Tier 1 Cabling (Rack – CML)



3 Solution Architecture for XenDesktop 7

3.1 XenDesktop with HDX 3D Pro

XenDesktop with HDX 3D Pro is a desktop and app virtualization solution that supports high-end designers and engineers of 3D professional graphics applications and provides cost-effective support to viewers and editors of 3D data. With XenDesktop, you can deliver a persistent user experience and leverage other virtualization benefits such as single-image management and improved data security.

Use HDX 3D technologies with:

- Computer-aided design, manufacturing, and engineering (CAD/CAM/CAE) applications
- Geographical information system (GIS) software
- Picture Archiving Communication System (PACS) workstations for medical imaging
- Latest OpenGL, DirectX, CUDA and CL versions supported
- Latest NVIDIA Grid cards

Comparing the available options:

Name	Description	Considerations
SVGA (Super Video Graphics Array)	VMware WDDM (Windows Display Driver Model) 1.1-compliant driver	Software rendering for 2D and 3D. Live vMotion is supported.
vSGA (Virtual <u>Shared</u> Graphics Acceleration)	Multiple virtual machines leverage physical GPUs installed locally in the ESXi hosts to provide hardware-accelerated 3D graphics to multiple virtual desktops	512MB max video memory per VM, half from RAM, half from GPU. GPU resources are assigned first come, first serve.
vDGA (Virtual <u>Dedicated</u> Graphics Acceleration)	Graphics acceleration capability provided by VMware ESXi for high-end workstation graphics where a discrete GPU is needed	More suited to persistent desktop use case. Live vMotion is not supported (cold only). Intel VT-d required.

For more information, please visit: [LINK](#)

3.2 Graphics Compute Server Infrastructure

3.2.1 Local Tier 1 Rack

In the Local Tier 1 model, VDI sessions execute on local storage on each compute server. Due to the local disk requirement in the compute layer, this model supports rack servers only. In this model, only the management server hosts access iSCSI storage to support the solution's management role VMs. Because of this, the compute and management servers are configured with different add-on NICs to support their pertinent network fabric connection requirements. Refer to section 2 for cabling implications. Due to the reduced densities on GPU-enabled compute hosts,



recommended RAM has been reduced. The management server host does not require local disk to host the management VMs.

Local Tier 1 GFX Compute Host – PowerEdge R720
2 x Intel Xeon E5-2680v2 Processor (2.8Ghz)
96GB Memory (6 x 16GB DIMMs @ 1600Mhz)
VMware vSphere on internal 2GB Dual SD
10 x 300GB SAS 6Gbps 15k Disks (VDI)
PERC H710 Integrated RAID Controller – RAID10
Broadcom 5720 1Gb QP NDC (LAN)
Broadcom 5720 1Gb DP NIC (LAN)
iDRAC7 Enterprise w/ vFlash, 8GB SD
2 x 110W PSUs

Local Tier 1 Management Host – PowerEdge R720
2 x Intel Xeon E5-2670v2 Processor (2.5Ghz)
96GB Memory (6 x 16GB DIMMs @ 1600Mhz)
VMware vSphere on internal 2GB Dual SD
Broadcom 5720 1Gb QP NDC (LAN/iSCSI)
Broadcom 5719 1Gb QP NIC (LAN/iSCSI)
iDRAC7 Enterprise w/ vFlash, 8GB SD
2 x 110W PSUs

3.2.2 Shared Tier 1 Rack (iSCSI)

In the Shared Tier 1 model, VDI sessions execute on shared storage so there is no need for local disk on each server. To provide server-level network redundancy using the fewest physical NICs possible, both the compute and management servers use a split QP NDC: 2 x 10Gb ports for iSCSI, 2 x 1Gb ports for LAN. 2 additional DP NICs (2 x 1Gb + 2 x 10Gb) provide slot and connection redundancy for both network fabrics. All configuration options otherwise are identical to the Local Tier 1 host.

Shared Tier 1 Compute Host – PowerEdge R720
2 x Intel Xeon E5-2680v2 Processor (2.8GHz)
96GB Memory (6 x 16GB DIMMs @ 1600Mhz)
VMware vSphere on 2 x 1GB internal SD
Broadcom 57800 2 x 10Gb SFP+ + 2 x 1Gb NDC
1 x Broadcom 5720 1Gb DP NIC (LAN)
1 x Intel X520 2 x 10Gb SFP+ DP NIC (iSCSI)
iDRAC7 Enterprise w/ vFlash, 8GB SD
2 x 110W PSUs



Shared Tier 1 Management Host – PowerEdge R720
2 x Intel Xeon E5-2670v2 Processor (2.5GHz)
96GB Memory (6 x 16GB DIMMs @ 1600Mhz)
VMware vSphere on 2 x 1GB internal SD
Broadcom 57800 2 x 10Gb SFP+ + 2 x 1Gb NDC
1 x Broadcom 5720 1Gb DP NIC (LAN)
1 x Intel X520 2 x 10Gb SFP+ DP NIC (iSCSI)
iDRAC7 Enterprise w/ vFlash, 8GB SD
2 x 110W PSUs

3.2.3 Shared Tier 1 Rack (FC)

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

Shared Tier 1 Compute Host – PowerEdge R720
2 x Intel Xeon E5-2680v2 Processor (2.8GHz)
96GB Memory (6 x 16GB DIMMs @ 1600Mhz)
VMware vSphere on 2 x 1GB internal SD
Broadcom 5720 1Gb QP NDC (LAN)
Broadcom 5720 1Gb DP NIC (LAN)
2 x QLogic 2562 8Gb DP FC HBA
iDRAC7 Enterprise w/ vFlash, 8GB SD
2 x 110W PSUs

Local Tier 1 Management Host – PowerEdge R720
2 x Intel Xeon E5-2670v2 Processor (2.5GHz)
96GB Memory (6 x 16GB DIMMs @ 1600Mhz)
VMware vSphere on 2 x 1GB internal SD
Broadcom 5720 1Gb QP NDC (LAN)
Broadcom 5720 1Gb DP NIC (LAN)
2 x QLogic 2562 8Gb DP FC HBA
iDRAC7 Enterprise w/ vFlash, 8GB SD
2 x 110W PSUs



In the above configurations, the R720-based DVS Enterprise Solution can support the following user counts per server (PVS).

Local/ Shared Tier 1 Rack Densities		
Mode	NVidia K1	NVidia K2
vSGA (Shared)	18	20
vDGA (Pass-Through)	8	4



4 Solution Performance and Testing

4.1 Load Generation and Monitoring

4.1.1 Login VSI – Login Consultants

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

4.1.2 Liquidware Labs Stratusphere UX

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

- Active Directory
- DNS
- DHCP
- Anti-Virus

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

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

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



4.1.3 VMware vCenter

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

4.2 Testing and Validation

4.2.1 Testing Process

The purpose of the single server testing is to validate the architectural assumptions made around the server stack. Each user load is tested against 4 runs. A pilot run to validate that the infrastructure is functioning and valid data can be captured and 3 subsequent runs allowing correlation of data. Summary of the test results will be listed out in the below mentioned tabular format.

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

Login VSI has two modes for launching user’s sessions;

- Parallel
 - Sessions are launched from multiple launcher hosts in a round robin fashion; this mode is recommended by Login Consultants when running tests against multiple host servers. In parallel mode the VSI console is configured to launch a number of sessions over a specified time period (specified in seconds)
- Sequential
 - Sessions are launched from each launcher host in sequence, sessions are only started from a second host once all sessions have been launched on the first host, and this is repeated for each launcher host. Sequential launching is recommended by Login Consultants when testing a single desktop host server. The VSI console is configured to launch a specified number of session at a specified interval specified in seconds

All test runs which involved the 6 desktop hosts were conducted using the Login VSI “Parallel Launch” mode, all sessions were launched over an hour to try and represent the typical 9am logon storm. Once the last user session has connected, the sessions are left to run for 15 minutes prior to the sessions being instructed to logout at the end of the current task sequence, this allows every user to complete a minimum of two task sequences within the run before logging out. The single server test runs were configured to launch user sessions every 60 seconds, as with the full bundle test runs sessions were left to run for 15 minutes after the last user connected prior to the sessions being instructed to log out.

4.2.2 Pass-Through Graphics Results

These results describe validation efforts undertaken on vSphere using NVidia K1 and K2 Grid cards to study their behavior when used with graphic-intensive applications. 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 modelling 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.


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**).

Fixed Frame-Rate Video Component

When considering End User Experience 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 (European standard) is transmitted at 25 FPS. This combination of facts is used to derive a further workload. A movie clip at 30 fps (Trailer from Hobbit movie) has been created. In the pass-through environment with n (4 or 8) virtual desktops, this clip was played on 1 of the desktops. The actual frame-rate delivered by the protocol was then monitored to investigate if it remained above 25 i.e. still at broadcast-level.

The following graphs show the measurement of the FPS while this video was running:





It is clear from these results that both Grid cards are able to sustain the frame rate at above 30 fps. However the Dell Wyse Z90D7 is not capable of sustaining the frame rate during the playback. It can be seen that the thin client is not capable of supporting the processing required to maintain the frame rate, and consequently the end user experience will suffer. It seems that the superior CPU and memory in the OptiPlex 7010 enhances the playback.

Further tests will be run, with the appropriate benchmark companion workload running on the other n-1 desktops to ensure that the performance can be maintained when the system is more fully loaded. It is expected that with additional workload on the server the frame rate may drop especially for the K1 card. This is supported by results observed in the subsequent section on SPECviewperf tests.

Subjective End-User Experience

The type of endpoint can have a significant impact on the user experience regardless of the GPU being used. Although we were able to connect to desktops and execute graphics applications with thin clients such as the Wyse Z90D7, there was a noticeable lag in the video being displayed as well as when interacting with the session. In contrast, performance and response times on thick endpoints like the OptiPlex 980 and 7010 were good. Additionally, as expected, the K2 desktops provided the best experience. When using thick endpoints, graphics applications responded well on the K2 desktops with smooth playback/display while occasional jitter/lag was observed on the K1 desktops. In conclusion, the best performance and end user experience was observed with OptiPlex endpoints connecting to desktops using K2 GPUs.

Graphics-Specific Performance Analysis Results

Viewperf Benchmark

SPECviewperf is a widely used benchmark in the workstation domain for benchmarking graphics performance and on this basis it has been chosen as the appropriate benchmark to use for the pass-through use case. It was used for both K1 and K2 testing. The tests run in SPECviewperf are defined by viewsets. The viewsets chosen for the K1 and K2 pass-through testing are CATIA, EnSight, Maya, Pro/ENGINEER, SolidWorks, Teamcenter Visualization Mockup, Siemens NX and Lightwave. For a description of these viewsets please see <http://www.spec.org/gwpg/gpc.static/vp11info.html>.


The SPECviewperf workloads were run against both the Dell OptiPlex 7010 and the Dell Wyse Z90D7.



A number of tests were run against K1 grid card as follows

- Viewperf test run on a single VM against both endpoints with no companion tests running
- Viewperf test run against 2 OptiPlex endpoints with Heaven companion workload on remaining VMs
- Viewperf test run against Wyse Z90D7 with Solidworks eDrawing companion workload on remaining VMs
- Viewperf test run against Wyse Z90D7 endpoint with AutoCAD companion workload on remaining VMs

The SPECviewperf results recorded are presented in the following graph




A number of tests were run against K2 grid card as follows

- Viewperf test run on a single VM against both endpoints with no companion tests running. An additional test was run against OptiPlex 980
- Viewperf test run against OptiPlex 7010 endpoint with Heaven companion workload on remaining VMs
- Viewperf test run against Wyse Z90D7 and OptiPlex 7010 endpoints with Solidworks eDrawing companion workload on remaining VMs
- Viewperf test run against Wyse Z90D7 and OptiPlex 7010 endpoints with AutoCAD companion workload on remaining VMs



The SPECviewperf results recorded are presented in the following graph:



SPECviewperf reports results in frames per second. The most obvious result is that the K2 card gives a consistently higher score than the K1 card. This is to be expected as the NVidia GRID K2 card is positioned a high-end graphics card with performance similar to the Quadro K5000 card while the NVidia K1 is positioned as a mid-range graphics card equivalent to a Quadro K600.

It is also noticeable that the addition workload when companion tests are running affects the results for the K1 more than the K2.

For the K2 card these results appear in line with some of the results obtained on workstations submitted to the SPECviewperf website results. Please see a summary of submitted results for various workloads across multiple workstations (including 2 Dell models) at <http://www.spec.org/gwpg/gpc.data/vp11/summary.html>.

Results for Dell Precision Workstation R5500 with NVIDIA Quadro 6000

Company/Product	Catia	Ensign	Lightwave	Maya	ProE	Snx	SW	Tcvis
Dell Precision Workstation R5500 with NVIDIA Quadro 6000	50.68	58.23	59.33	108.77	9.95	60.14	58.27	47.8

These results were submitted by Dell

Improved results may be expected against workstations with higher spec. processors using the K5000 video card. It appears that the K2 GRID card is capable of competing against high end workstations but the K1 GRID card should be reserved for medium use cases.

In addition to the Viewperf metrics recorded the Frame rate per second was recorded from Citrix XenDesktop as described below.













It is first necessary to enable debug logging using the HDX3D cmd line tool using the command shown below:


`"C:\Program Files\Citrix\ICAService\HDX3DConfigCmdLineX64.exe" debug_logging 1`

Next, start DebugView – a copy comes with the HDX 3d Pro health check tool and this should be used. Once this tool has been started, it is possible to filter on “FPS” and save this frame-rate information to a log file.

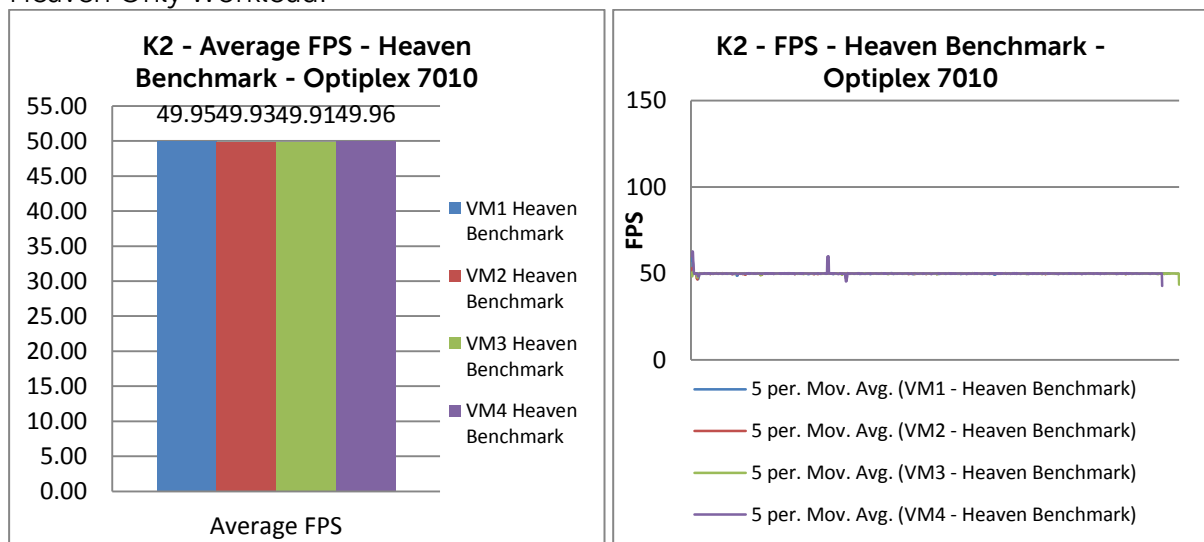
The graphs recorded for frame rate on some of the tests are reported below:







Heaven Only Workload:



An interesting result here is that on the K1 tests the companion workload maintained a reasonably high frame rate even though the Viewperf frame rate was low.

In addition to specific graphics results obtained results were gathered for the host performance during various workloads. This was done to ensure that the host was not stressed and the recorded results fall within the thresholds below.


Resource Utilization Parameter	Threshold
Compute Host CPU	85%
Computer Host Memory	85%
Network Throughput	85%
Tier 1 Storage Latency	20ms



Results were gathered using the VSphere Performance Reporting tool. The results shown below were gathered when running SPECviewperf tests on the K2 card. These represent the most load seen on the R720 host during the validation.

AUTOCAD & Viewperf Workload


The following graphs are gathered when the system is running the SPECviewperf test in conjunction with 3 companion AUTCAD tests.



eDrawings & Viewperf Workload

The following graphs are gathered when the system is running the SPECviewperf test in conjunction with 3 companion eDrawings workloads.

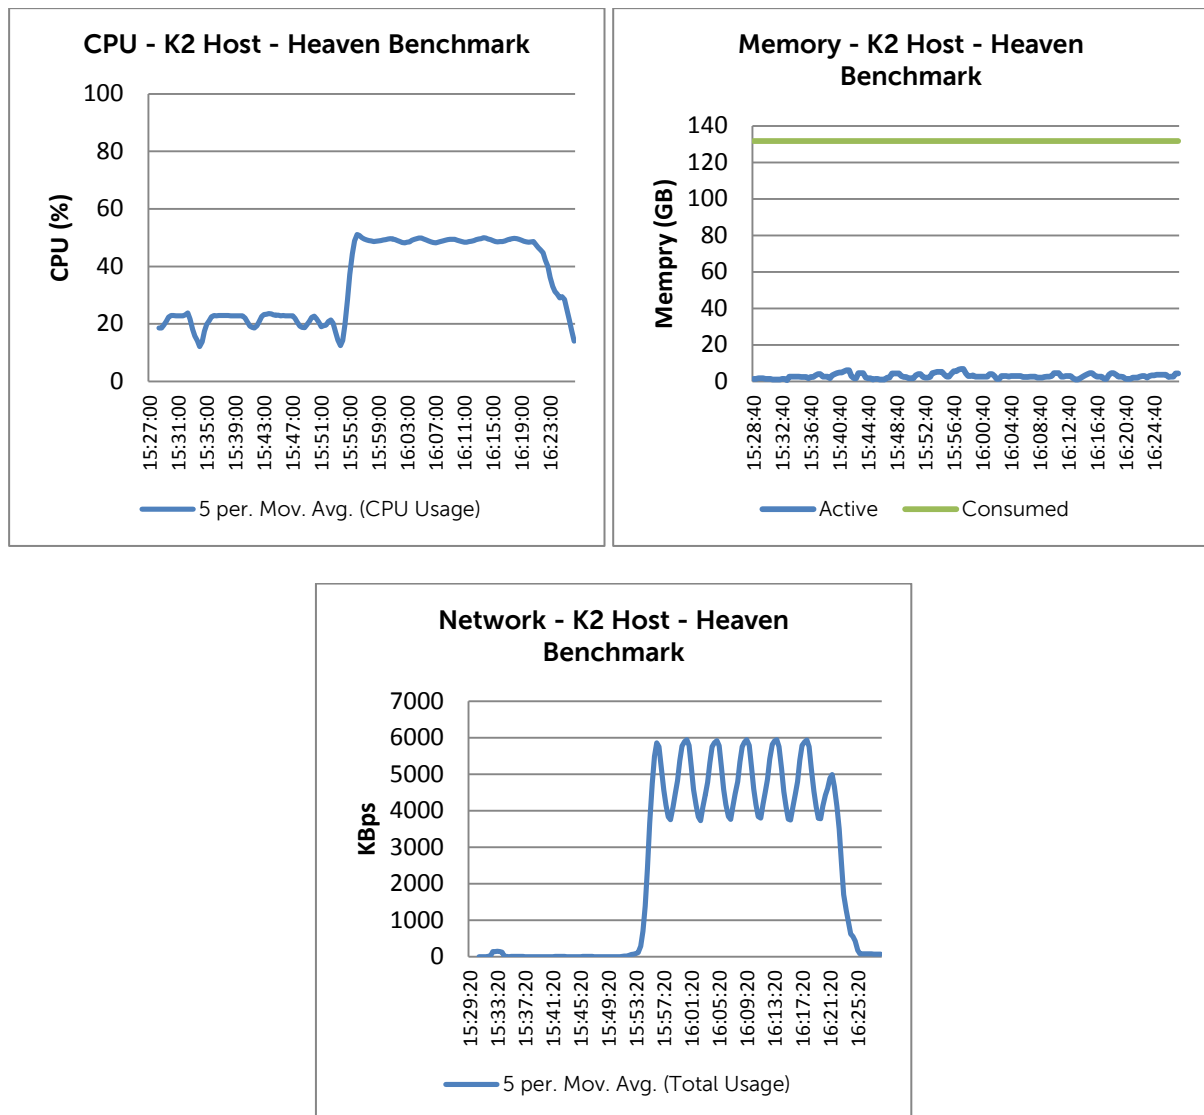




Heaven Benchmark – No Viewperf

In the following case the Heaven workload was running on all 4 VMs





It can be seen that in all cases the performance of the host remains well within the thresholds defined. Maximum CPU is approximately 50% and network usage does not go above 4 Mbps. Active and consumed memory are well within the capabilities of the host and no ballooning is evident.

Although only one set of results are included for the Datastore this shows that the local disk latency did not go above 1 millisecond. Similar results are observed in the raw data for other tests.

Conclusion

The results from the pass-through tests indicate that the R720 host is capable of handling the various graphics workloads. It should be noted that the specification for the processor is reduced from the normal 2.9GHz DVS spec. to 2.6 GHz because of power/thermal limitations.

The K2 card earns its place as the superior offering over the K1, as expected, and it seems to match the performance of many of the workstation results on the SPECviewperf website. However, there are probably newer higher specification workstations that show superior performance. The K1 cards performs well for the mid-size market segment it is positioned for.



4.2.3 Shared Graphics Results

This section describes validation efforts undertaken on vSphere using NVIDIA Grid K1 cards to study their behavior when used in shared mode for virtual desktops with graphics-intensive applications.

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 modelling 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 shard mode (vSGA). VMware vSGA provides the ability for multiple virtual machines to leverage physical GPUs installed locally in the ESXi hosts to provide hardware-accelerated 3D graphics. i.e. a VM does not have direct access to a GPU on the K1 cards but multiple VM share a GPU through 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 K1 Grid cards in shared mode. All of the following results were gathered with two K1 cards installed in the server.


Test Results Summary

This validation was performed for XenDesktop 7 virtual machines running on an R720 host, 96GB of RAM and dual 2.5 GHz processors. Validation was performed using DVS standard testing methodology using LoginVSI load generation tool for VDI benchmarking that simulates production user workloads. The “multimedia workload” profile in LoginVSI was used for testing the graphics capabilities. In addition, the Fishbowl HTML5 test was also used. As a result of this testing, the optimal supported configuration proved to be:

- Maximum of **18** XenDesktop 7 virtual machines per physical R720 host
- Each XenDesktop 7 virtual machine configured with four 2x vCPUs
- Each XenDesktop 7 virtual machine configured with 4GB RAM

Workload	Server Density	CPU	Login State IOPS	Login State IOPS per User	Steady State IOPS	Steady State IOPS per User
Multimedia Workload & Fishbowl HTML5 Test	18	83.3% - 88.5%	276	15.3	145	8







This virtual server-based configuration was able to sustain a user density of 18 sessions per physical host server depending on the type of user workload chosen as well as the file access and scanning configuration of the anti-virus tools.


Server Resource Performance (18 Users)

The graphs below show the CPU, Consumed Memory, Local disk IOPS, Disk IOPS, Network and VDI UX scatter plot results from this validation. The CPU usage for this test reached 88.5% thus confirming that the server with this configuration can support up to 18 Standard users. Memory was configured 4GB per VM. The following graphs show CPU, memory, local disk IOPS, network and VDI UX scatter plot results.





Stratusphere UX indicates all desktops had a good user experience. All desktops were in the upper right corner of the upper right quadrant indicating that 18 standard provisioned users can be supported with good performance.



Acknowledgements

Thanks to Gus Chavira, Principal Architect for VMware solutions at Dell, for his expertise and assistance on graphic acceleration using VMware vSphere.



About the Authors

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

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

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

Jerry Van Blaricom is a Sr. Systems Engineer in the Desktop Virtualization Solutions Group at Dell. Jerry has extensive experience with the design and implementation of a broad range of enterprise systems and is focused on making Dell's virtualization offerings consistently best in class.

