Dell | Product Group

Dell

Desktop Virtualization Solutions

Dell DVS Enterprise 6010

Solution Architecture (Rack & Blade)

Version 2

Citrix XenDesktop

February 2012

Note: The Dell ISS Solution has been renamed the Dell DVS Enterprise 6010. The Rack and Blade Configurations of more than 500 users are the same as the previous ISS Solution so this document has only been renamed. For configurations of 50 - 500 users, a new Reference Architecture document is now available.

The Entire Dell DVS Enterprise 6010 solution stack is now documented by the following documents:

- Dell DVS Enterprise 6010 Entry Configuration Reference Architecture for 50 - 500 Users (New Configuration) (Citrix and VMware versions available)
- Dell DVS Enterprise 6010 Rack and Blade Reference Architecture (over 500 users - previously named ISS) (Citrix and VMware versions available)
- Dell DVS Enterprise 6010 Solution Architecture (over 500 users - previously named ISS) (Citrix and VMware versions available)

Table of Contents

Table of Figures

1 Introduction

Dell's Desktop Virtualization Solutions offering is a comprehensive solution portfolio designed to enable customers to experience the benefits of virtual end user 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 VDI 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 user experience will be degraded from that of a traditional PC. If the architecture is oversized, then the cost per desktop/user will be economically undesirable for VDI adoption.

This white paper provides a detailed review of the Dell Integrated Solution Stack Enterprise Plus architecture and workload characterization testing done by the Dell Solutions Center. The solution is based on Citrix® XenDesktop® and VMware® vSphere® 4 ESXi with Dell™ PowerEdge™ blade servers, Dell EqualLogic™ storage, and Juniper network switches. Based on extensive engineering work in architectural design and scalability testing, the Dell Desktop Virtualization Solutions can be quickly and confidently deployed into production environments, thereby helping to eliminate much of the costly and time consuming trial and error often encountered during complex custom infrastructure design and implementation projects. This solution provides scalability and manageability with no single point-of-failure. Network architectures, storage configurations, and scalability of the Dell Desktop Virtualization Solutions are discussed in detail.

1.1 Audience and Scope

The intended audience for this white paper is IT administrators and managers who are planning to deploy a desktop virtualization solution. This paper provides a technical overview of the Dell DVS Integrated Solution Stack, and the engineering, validation and scalability testing which went into the components to provide a robust enterprise class desktop virtualization solution.

1.2 Executive Summary of Scalability Findings

The purpose of the testing detailed within this report was to enable Dell to validate and characterize the Dell Integrated Solution Stack for Citrix XenDesktop. This solution guide relates to the Enterprise Plus configuration which is based on Dell PowerEdge 11th generation Blades, Dell EqualLogic storage arrays and Dell/Juniper network switches.

Throughout the testing process the Dell Solutions Center focused on ensuring that the testing replicated production environments wherever possible rather than simply trying to achieve the maximum density of virtual desktops on a single server. The secondary objective was to ensure that the user experience was not degraded as the load increased. Finally the team deliberately took the approach that the desktop host servers should have available "headroom" at peak load to account for typical user behavior which cannot be replicated in a test environment. To ensure the desktop host servers maintained this headroom the CPU usage was monitored to ensure this load did not exceed 80% and that the consumed memory did not exceed 85% which provides the added benefit of ensuring best cost and performance balance.

To ensure the results gathered from the Liquidware Labs monitoring suite "Stratusphere UX" the team complemented all test runs with a subjective user session. This session was launched when the environment was nearing full load so the team was able to benchmark the worst possible performance on each test run. This subjective testing fully supported the information being gathered from Stratusphere UX which is detailed later in this report.

It is worth highlighting that there are numerous variables that can impact the density of virtual desktops that can be hosted on a given server configuration. Dell's approach to evaluating virtual desktop density is to conduct testing that mirrors actual real world production environments as closely as possible. To develop the density guidelines provided in this paper, Dell created very real world workload scenarios and configurations which included virus scan activity, boot and login storms to reflect typical production workloads.

The table below shows a summary of the optimal number of desktops hosted within the environment and a summary of the desktop and user data IOPS. The IOPS have been taken from the highest values in each of the reported test runs.

Workload	Full Bundle Density	Full Bundle Density (N+1)	Single Server Density (N+1)	IOPS per desktop	IOPS per User
Basic	960	880	74	7.29	0.83
Standard	720	660	55	11.1	0.97
Premium	504	462	39	13.9	0.59

The solution validation, workload characterization and resulting data enable Dell to provide a very robust and scalable Integrated Solution Stack which ensures customers' success through the Dell Desktop Virtualization Solutions assessment and design methodology. Typically customers are faced with many unknown design criteria to determine a solution which meets their technical and strategic infrastructure plans. This paper outlines the effort the Dell Desktop Virtualization Solutions team put into removing this burden to accelerate and simplify the implementation of a Desktop Virtualization Solution.

2 Integrated Solution Stack Overview

The Dell DVS Integrated Solution Stack is available in two bundles to best meet customer needs as part of the Dell Desktop Virtualization Solutions. This paper is focused on the Enterprise Plus Solution Stack.

Enterprise

- All inclusive, enterprise-class hardware, software and services
- Built on robust, industry-standard components.
- Scalable from 500-5,000 users

Enterprise Plus

- Higher datacenter-density design helps reduce rack space
- Blade architecture is efficient and easy-to-maintain
- Shared storage enables time-saving features like vMotion and improves image management efficiency.

	Enterprise	Enterprise Plus
Virtual PC	Yes	Yes
Redundancy	Yes	Yes
Recoverability	Yes	Yes
High DC Density	No	Yes
Dynamic Motion	No	Yes

Figure 1: Overview of Dell Desktop Virtualization Solutions

Juniper EX 4500 40 port top of rack switch

- Virtual chassis-Up to 10 Juniper EX4500-40T connected to create a logical device.
- Up to 480 10/100/1000BASE-T ports additional 20 GbE or ten 10 GbE uplink

Dell PowerEdge Blade Servers

- Industry leading energy efficient PowerEdge M1000e blade enclosure.
- 11th generation M610 blade server
  - Intel Xeon 5600 processor with QuickPath memory technology
  - Unified Server Configurator (USC)- simplifies provisioning functions such as system deployment, system updates, hardware configuration and diagnostics.

PowerConnect network switches

- PowerConnect M8024 chassis switch
- PowerConnect M6220 chassis switch

Dell EqualLogic PS series iSCSI Arrays

- PS6010 XVS SSD/SAS Hybrid array (Tier 1 DDVS storage)
  - Dual 10GbE controllers with a total of 4GB battery-backed memory
  - Eight (8) hot-plug Solid State Drives (SSD) drives and Eight (8) hot-plug SAS hard disk drives.
- PS6510E SATA array (Tier 2 and 3 DDVS storage)
  - Dual 10GbE controllers with a total of 4GB battery-backed memory
  - 7,200 RPM 1TB SATA II drives

Figure 2: Overview of Dell Enterprise Plus Integrated Solution Stack

PowerConnect J-EX4200 48 port top of rack switch

- Virtual chassis-Up to 10 Juniper EX4200 connected to create a single logical device.
- Supporting up to 480 10/100/1000BASE-T ports and a additional 20 GbE or ten 10 GbE uplink

Dell PowerEdge Blade Servers

Dell PowerEdge R710 Rack Servers

- 125% more memory capacity and more integrated I/O than the previous generation
- Unified Server Configurator (USC)- simplifies provisioning functions such as system deployment, system updates, hardware configuration and diagnostics.

Dell EqualLogic PS series iSCSI Arrays

- PS6500E SATA array (Tier 2 and 3 DDVS storage)
  - Dual 1 GbE controllers with a total of 4GB battery-backed memory
  - 7,200 RPM 1TB SATA II drives

Figure 3: Overview of Dell Enterprise Integrated Solution Stack

2.1 Dell PowerEdge Servers

2.1.1 Dell PowerEdge Blade Servers

Blade Modular Enclosure: The Dell PowerEdge M1000e is a high-density, energy-efficient blade chassis that supports up to sixteen half-height blade servers, or eight full-height blade servers, and six I/O modules. A high-speed passive mid-plane connects the server modules to the I/O modules, management, and power in the rear of the chassis. The enclosure includes a flip-out LCD screen (for local configuration), six hot-pluggable/redundant power supplies, and nine hot-pluggable N+1 redundant fan modules.

Blade Servers: Dell's latest virtualization-ready blade servers and enclosures— 11th generation Dell PowerEdge M610 blade server based on the Intel® *Nehalem* processors. The Dell PowerEdge M610 supports 12 DIMM slots and up to 192 GB of RAM. The M610 blade servers use Intel QuickPath technology to provide a high-speed link to the memory modules. The 11th generation blade servers come with the next generation system management tool: Dell Unified Server Configurator (USC). This helps customers reduce operating costs by simplifying deployment and management. The Dell USC supports diagnostics, firmware updates, and hardware configuration.

I/O Modules: The enclosure provides three redundant fabrics using six I/O modules. The modules are populated with two PowerConnect I/O Modules the M6220 and M8024.

The PowerConnect M6220 is one of Dell's most advanced switch offerings with advanced core switching capabilities specifically designed for the PowerEdge M1000e Modular Server Enclosure. This 20-port (16 internal, 4 external) Gigabit Ethernet Layer 3 switch is stackable and offers optional 10 Gigabit Ethernet uplinks and supports the latest version of the Internet Protocol—IPv6—enabling broader worldwide scalability. High performance stacking is supported for up to twelve systems, and advanced security and Quality of Service (QoS) features make this switch ideal for blade server applications such as Layer 3 routing, High Performance Cluster Computing (HPPC), and iSCSI storage.

The PowerConnect M8024 switch provides twenty four 10GbE ports to the Dell™ PowerEdge M1000e blade chassis. This valuable capability delivers 10GbE to each of the up-to 16 server blades. In addition to the 16 internal 10GbE ports, the M8024 delivers flexible external I/O connectivity choices, utilizing up-to two modular 10Gb Base-T, CX-4 and/or SFP+ uplinks. Each SFP+ module provides four 10GbE ports. When using two SFP+ modules, a total of eight external 10GbE ports are provided. Each CX-4 module provides three 10GbE ports. Two CX-4 modules yield a total of six external 10GbE ports. Two 10Gb Base-T modules provide four 10GbE ports. For even more flexibility, the M8024 supports mixed environments using combinations of SFP+, 10G Base-T, and CX-4 modules.

Chassis Management: The Dell PowerEdge M1000e has integrated management through a redundant Chassis Management Controller (CMC) module for enclosure management and integrated keyboard, video, and mouse (iKVM) modules. Through the CMC, the enclosure supports FlexAddress technology which enables the blade enclosure to lock the World Wide Names (WWN) of the Fibre Channel controllers and Media Access Control (MAC) addresses of the Ethernet controllers to specific blade slots. This enables seamless swapping or upgrading of blade servers with Ethernet and Fibre Channel controllers without affecting the LAN or SAN configuration.

2.1.2 Dell PowerEdge Embedded Server Management

Embedded Management with Dell's Lifecycle Controller: The Lifecycle Controller is the engine for advanced embedded management and is delivered as part of iDRAC Enterprise in 11th-generation Dell PowerEdge blade and rack servers. It includes 1 GB of managed and persistent storage that embeds systems management features directly on the server, thus eliminating the media-based delivery of system management tools and utilities previously needed for systems management. Embedded management includes:

- **Unified Server Configurator (USC)** aims at local 1-to-1 deployment via a graphical user interface (GUI) for operating system install, updates, configuration, and for performing diagnostics, on single, local servers. This eliminates the need for multiple option ROMs for hardware configuration.
- **Remote Services** are standards-based interfaces that enable consoles to integrate, for example, bare-metal provisioning and one-to-many OS deployments, for servers located remotely. Dell's Lifecycle Controller takes advantage of the capabilities of both USC and Remote Services to deliver significant advancement and simplification of server deployment.
- **Lifecycle Controller Serviceability** aims at simplifying server re-provisioning and/or replacing failed parts and thus reduces maintenance downtime.

2.2 Dell EqualLogic Storage

2.2.1 PS6010 XVS Hybrid Series iSCSI Arrays

The Dell EqualLogic PS6010 features two 10 Gigabit Ethernet iSCSI network ports per controller, fast processors, 2 GB of cache per controller, support for RAID 6, increased drive capacity, a monitoring application, and SAN Headquarters (SAN HQ), at no additional cost. The PS6010 delivers up to 10x the bandwidth per port and 2.5x the bandwidth per array compared to previous generation PS series arrays.

The PS6010XVS reinforces the EqualLogic principles of intelligence, automation and scalable performance, with the addition of optimized data response for your applications through intelligent data tiering. Using a combination of low-latency SSD and performance oriented SAS drives; the PS6010XVS automatically places data on the appropriate disk technology using historical I/O and performance patterns and helps optimize data response for your applications while maximizing efficiency.

The ability to sense and intelligently respond to data requirements is a valuable feature for applications like virtual desktops that can benefit from low-latency SSD during desktop initialization (i.e. Bootstorm) and often require cost effective, performance-oriented storage for user data files.

All PS Series storage arrays include SAN configuration features and capabilities that sense network connections, automatically build RAID sets, and conduct system health checks to help ensure that all components are fully functional. A PS6010XVS is designed to be installed, configured, and serving storage in less than one hour.

- Intelligent data placement across SSD and 15,000 RPM SAS helps deliver excellent user responsiveness for tiered workloads
- Balance cost and performance with a combination of SSD and 15,000 RPM SAS in a single array

2.2.2 Dell EqualLogic PS6510E Series iSCSI Arrays

The Dell EqualLogic PS6510E is a virtualized iSCSI SAN that combines intelligence and automation with fault tolerance to provide simplified administration, rapid deployment, enterprise performance and reliability, and seamless scalability.

The PS6510E delivers the lowest cost-per-gigabyte in the PS6010 Series with up to 96TB of SATA storage in a high-density 4U chassis. A PS Series SAN built with PS6510E arrays can scale up to 1.5PB under a single, intuitive management interface.

The PS6510 line features a redesigned 10GbE dual-port controller delivering up to 10x the bandwidth per port and 2.5x the bandwidth per array compared to previous generations.

The Dell EqualLogic SAN devices provide the following capabilities:

Reliability: Dell EqualLogic PS6000 Series arrays have hot-swappable redundant components, a choice of RAID types, and hot-spare disks. They also include the Auto-Stat Disk Monitoring System (ADMS) which proactively scans disk drives in the background to help detect media anomalies and correct them.

Scalability: As each array is added to the storage group, the storage capacity and performance, in terms of both bandwidth and IOPS, are increased. This increased capacity can be utilized without downtime. Thin-provisioning permits predefinition of a group of Virtual Machine File System (VMFS) volumes which are larger than the physical space, allowing for physical space to be added and utilized without downtime, when necessary.

Self-Managing Arrays: The arrays offer many self-managing features such as automatic load balancing and storage tiering. A single storage pool can have different models that offer a range of capacity and performance parameters. In addition, different arrays in a storage pool can be configured with different RAID levels, and volumes will automatically be migrated between RAID levels based on performance data and usage patterns. All data and volume movement can be performed online with zero downtime.

Top-Tier Data Protection Software: Advanced data protection features such as Auto Replication and Auto-Snapshot Manager is standard with EqualLogic arrays. The Auto-Snapshot Manager integrates with VMware vCenter and VMware's native snapshot technology to provide intelligent and consistent snapshots.

SAN Headquarters (HQ): SAN HQ provides centralized access to detailed performance and event data, reporting and intelligent alerts from the EqualLogic storage array groups. SAN HQ uses Simple Network Management Protocol (SNMP) to collect performance, alarm, and health status data from dozens of array groups that could be spread across multiple locations around the world.

3 Integrated Solution Stack Bundles

The Dell Integrated Solution Stack provides a simplified and scalable desktop virtualization solution designed with two ISS bundles to address the majority of customer needs and use cases for Desktop Virtualization. These bundles are Enterprise and Enterprise Plus and provide a prescriptive and scalable architecture; the bundles have limited configuration options to provide predictive scalability and Dell Pro-Support for Solutions. Whether the customer prefers a Managed Solution from Dell or chooses to manage the solution in-house, the Dell ISS bundles are consistent and can be leveraged as the horizontal platform. The ISS bundles include the following key features:

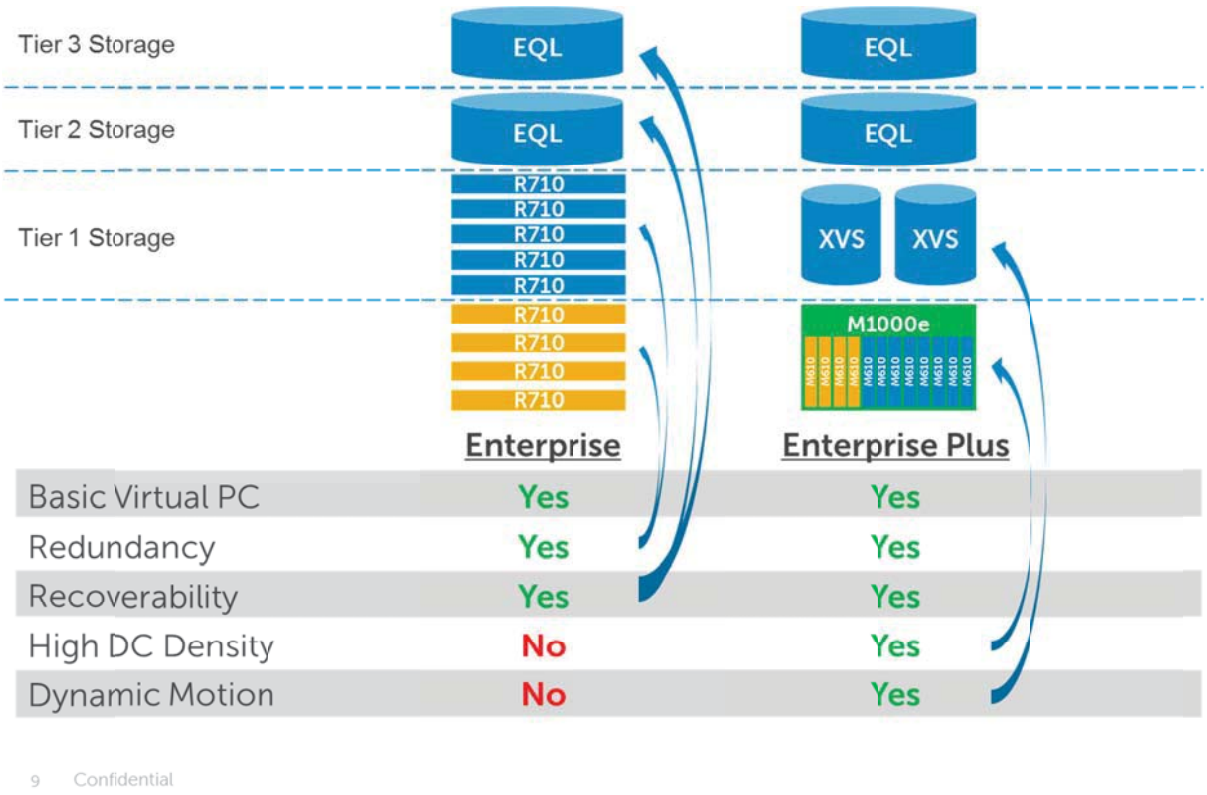

Figure 4: Dell Integrated Solution Stack Bundles

To provide the highest service levels for the most critical components of the solution, the management stack and database services have been designed into dedicated and highly available solutions fully supported by Dell, VMware, Citrix and Microsoft. This also provides an added benefit of increased user densities and centralized management services for multiple Integrated Solution Stacks.

Figure 5: XenDesktop Resource Pools

4 Characterization and Scalability Testing

The Dell Desktop Virtualization Solutions Group developed a methodology to load test the Enterprise Plus Integrated Solution Stacks to determine linear scalability with acceptable performance of both the infrastructure and the end user desktop environment. The primary goal of the testing was to provide a real world workload characterization including anti-virus software, login and boot storms and not simply to fit as many user workloads as possible onto a single server.

4.1 Workload Generation and Assessment Methodology

4.1.1 Workload Generation

4.1.1.1 Login VSI – Login Consultants

The Login VSI tool (developed by Login Consultants) was used to generate the desktop workloads. Login VSI is a widely-used 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. The test script starts 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). These launchers are managed via a Microsoft Management Console which is used to configure whether the sessions are launched in parallel (sessions are created from each launcher in a round robin mechanism) or sequential (all sessions to be connected from each launcher are connected before the next launcher is used).

4.1.1.2 Liquidware Labs Stratusphere UX

Liquidware Labs™ 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 reports, graphs and summary tables of key information and performance metrics. In addition the Stratusphere Hub generated a magic quadrant style scatter plot showing the machine and I/O 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;

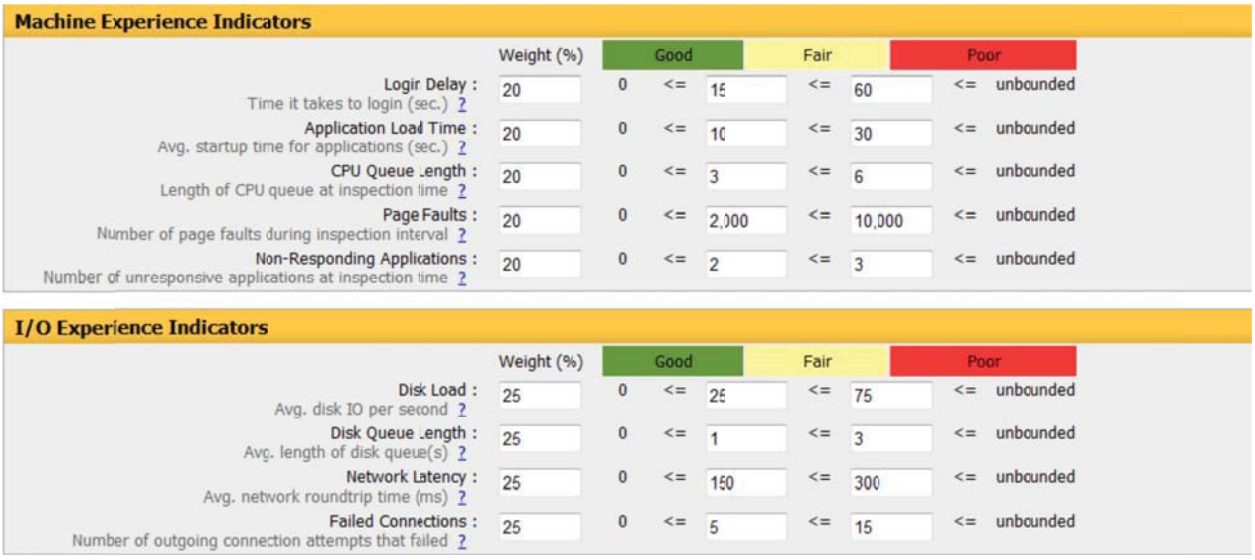

Figure 6: Stratusphere UX Metrics

Since each metric has a specified weight these results for each metric are combined to provide an overall UX score. **This is a key metric in the solution development process to define maximum load characteristics for the Integrated Solution Stack.**

4.1.2 Assessment and Monitoring Tools

The tools outlined below were used to monitor different aspects of the environment during testing. The key infrastructure areas assessed were host CPU and memory utilization, VDI session performance and end user experience, storage I/O peaks and averages, storage space and latency along with network performance impact. The following tools were utilized to gather the key data.

4.1.2.1 EqualLogic SAN HQ

EqualLogic SANHQ was used for monitoring the Dell EqualLogic storage units in each Integrated Solution Stack. SAN HQ has been used to provide IOPS data of the SAN. This has allowed the team to understand the IOPS required by each tier of the solution for each user workload type.

4.1.2.2 VMware vCenter

VMware vCenter was used to gather key data (CPU, Memory and Network usage) from each of the desktop hosts during each test run. This data was exported for each host and then consolidated to show data from all hosts. The report did not and does not typically include specific performance metrics for the two infrastructure host servers. The servers were monitored during testing and were seen to be performing at such a level that in the event one of these hosts should fail the remaining host would be able to deliver all services. The data along with other key metrics is utilized within the Dell Desktop Solutions Group sizing tools to adequately size a customer environment as part of the Dell Blueprint Assessment process.

4.1.2.3 MRTG

The final monitoring tools used within the environment is MRTG, this tool is used to report the performance of each network interface on the switches. MRTG has been configured to monitor all ports on the PowerConnect J series switches and blade interconnects. It is worth noting that one of the limitations of MRTG is that its report format is on a port by port basis as a HTML graph. For the purposes of this report the decision was made that rather than including a large number of graphs on a port by port basis the team have been continuously monitoring the switch interfaces with MRTG and found that no port has been seen to exceed 10-15% utilization, this is supported by the Network utilization graphs from the ESX hosts shown later within this report.

4.2 Load Generation Methodology & Workloads

4.2.1 Workload Generation Methodology

The testing undertaken by the Dell Solutions Center team was intended to provide a set of results that were as closely aligned with an actual production environment as possible. In addition multiple runs for each workload were completed to allow the team to understand and record any performance differences or anomalies within the environment.

The following testing scenarios were completed for ***each user workload***;

Pipe-Clean Run

- The pipe-clean run is used to ensure all components within the environment are functioning correctly and results are being gathered as expected

Runs 1, 2 & 3 – Full bundle, boot on demand

- All hosts are available and hosting desktop sessions, the desktop group is configured to maintain 20% of the group powered on and available for users to connect, as a result the environment will power desktops on while users are connecting and performing the workload tasks

Run 4 – Full bundle, pre-booted

- All hosts are available, all desktops are powered on and ready for users to connect, the desktops are left to idle for 20 minutes after the last desktop registers with the connection brokers. This idle period allows VMware ESXi to have optimized memory usage through Transparent Page Sharing. The environment is only running user tasks during the test.

Runs 5, 6 & 7 – Single server

- Single Desktop host testing, this is performed when the desktops are only available on a single host, the temporary WriteCache data for these desktops is held on a single volume on the EqualLogic PS6010XVS arrays and the image is streamed from a single Citrix Provisioning Server.

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

Parallel

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

Sequential

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

All test runs which involved a fully populated Integrated Solution Stack were conducted using the Login VSI "Parallel Launch" mode, all sessions were launched over an hour to try and represent the typical 9 a.m. logon storm. Once the last user session was connected, the sessions were left to run for 15 minutes prior to the sessions being instructed to logout at the end of the current task sequence, this allowed 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, and just as with the full bundle test runs sessions, single server test runs were left to run for 15 minutes after the last user connected prior to the sessions being instructed to log out.

4.2.2 Workloads

When designing a Desktop Virtualization Infrastructure it is critical to classify and assess the user workloads. The Dell Desktop Virtualization Solution methodology includes a Dell blueprint process designed to assess and categorize a customer's environment according to the workloads defined in this section. In the Dell Desktop Virtualization solution this will map directly to the SLA levels we offer in our Integrated Solution Stack. There are three desktop workload classifications, each of which is bound by specific metrics and capabilities. XenDesktop comes with a few extra features that affect the manner in which Dell could deploy the Integrated Solution Stack. In particular, this has to do with the concept of "presented apps" or "terminal Services" offerings. At this time presented apps or terminal services are only available in Dell's Custom Desktop Virtualization Solutions and were not configured or tested. Figure 7: Dell DVS virtual desktop user workloads below provides an overview of the workloads characterized and defined for the Dell Integrated Solution Stack.

User Workload	VM Memory Allocation	VM Memory Reservation	OS Image
Basic	1GB	0.5GB	Shared
Standard	1.5GB	1GB	Shared
Premium	2.5GB	1.5GB	Shared plus Optional Profile Virtualization to provide persistent image as necessary.

Figure 7: Dell DVS virtual desktop user workloads

The tasks undertaken by the users in each of the workloads are outlined below;

Basic Workload

This workload emulates a typical task worker

- The basic workload is very light in comparison to Standard
- Only 2 apps are open simultaneously
- Only apps used are IE, Word and Outlook
- Idle time total is about 1:45 minutes

Standard Workload

This workload emulates a medium knowledge worker using Office, Internet Explorer and Bullzip PDF. Once a session has been started the medium workload will repeat every 12 minutes. During each loop the response time is measured every 2 minutes. The medium workload opens up to 5 apps simultaneously. The type rate is 160ms for each character. Approximately 2 minutes of idle time is included to simulate real-world users.

Each loop will open and use:

- Outlook 2007, browse 10 messages.
- Internet Explorer, one instance is left open (BBC.co.uk), one instance is browsing Wired.com, Lonelyplanet.com and heavy flash app gettheglass.com.
- Word 2007, one instance to measure response time, one instance to review and edit document.
- Bullzip PDF Printer & Acrobat Reader, the word document is printed and reviewed to PDF.
- Excel 2007, a very large randomized sheet is opened.
- PowerPoint 2007, a presentation is reviewed and edited.
- 7-zip: using the command line version the output of the session is zipped.

Premium Workload

The Premium workload is based on the standard workload; the differences in comparison to the Standard workload are:

- Type rate is 130ms per character.
- Idle time total is only 40 seconds.
- The heavy workload opens up to 8 apps simultaneously

In order to ensure the environment is as representative of a production environment as possible Anti-Virus was deployed to all servers and virtual desktops. The environment utilized McAfee® VirusScan Enterprise 8.7i, managed centrally by the ePolicy Orchestrator which controls agent deployment, definition updates and policies.

Since the Virtual desktops will be based on a single read-only image streamed from the provisioning servers there are a number of exclusions that will be configured to optimize performance of the On-Access scanning per McAfee and Citrix best practices.

4.3 Scalability Testing Results

4.3.1 Summary

The table below shows the desktop densities and storage IOPS on a per desktop basis and user basis (based on file data for the user profile and home directory storage). The IOPS values are an average based on the peak usage points for full bundle testing for the boot on demand test runs. As noted in the executive summary of this paper, the testing was developed to provide characterization of the infrastructure based on real world workloads.

Workload	Full Bundle Density	Full Bundle Density (N+1)	Single Server Density	Single Server Density (N+1)	IOPS per desktop (Tier 1)	IOPS per User (Tier 2)
Basic	960	**880**	80	**73**	7.29 (7000 total)	0.83 (800 total)
Standard	720	**660**	60	**55**	11.1 (8000 total)	0.97 (700 total)
Premium	504	**462**	42	**38.5**	13.9 (7000 total)	0.59 (300 total)

Figure 8: Results summary by workload

In the table above the Full Bundle densities refer to the number of desktops that were running (while providing an acceptable user experience) on all 12 Desktop host servers. The Full Bundle (N+1) density figures are calculated based on the assumption that the environment will be configured to allow for the failure of a single host with no impact to user experience. In a production environment it is recommended that the N+ 1 vSphere HA Cluster design be implemented to ensure the solution can continue to function in the event that a single server fails or requires maintenance again, with minimal impact to the end user experience.

The "IOPS per desktop" values presented relate to the temporary data that is written to the Citrix Provisioning Services WriteCache while each desktop is running, while the "IOPS per user" values relate to the file activity to the user profile store and home directory. The environment uses folder redirection to minimize the amount of user data stored in the user profile to ensure that the user logon times are not impacted by quantity of data stored in these folders. These values have been taken when the environment demonstrated peak load and are presented as the average IOPS per user and per desktop at peak load.

The major reason that the Premium workload demonstrates a lower "IOPS per user" IOPS value is due to the fact that there were issues with sessions hanging due to the increased workload. Upon investigation the Dell Solution Center team determined that since the Outlook Personal File was being redirected to the user's home drive it was causing instability in the Outlook response times. It was decided that in order to resolve this issue the Microsoft Outlook Personal file (.pst) folder would be moved back to a location on the local machine rather than being held on the user's home directory. This resolved many of the issues related to user sessions hanging during the test runs; as a result the "IOPS per Desktop" includes the traffic to the user .pst files in the Premium figures.

(***Note: Outlook .pst files are used since the VSI Login tools is not currently written to support a full Microsoft Exchange environment)***

The Dell Integrated Solution Stack provides an optional Tier 3 storage solution for data protection and replication. Some of the test runs indicated that storage latency was negatively impacted by the replication schedule. To assess this situation a number of test runs were executed in which all volumes except those hosting the desktop WriteCache data were replicated to a second EqualLogic PS6510E array every 30 minutes. The 30 minute replication schedule was chosen to ensure that replication took place a number of times during each test run. The inclusion of this data enables Dell Solutions Architects and customers to understand the effect of replication on the environment during working hours. Since Tier 3 storage is optional and customer data protection policy will dictate configuration it will be critical to carefully design the schedule to allow minimum impact to production activity. Dell provides a recommended replication design policy as a starting point which minimizes production impact during peak periods of I/O activity as part of the Dell Integrated Solution Stack offerings.

When sizing an environment one key area of consideration is the storage used for the desktop WriteCache, Figure 9: WriteCache Read: Write Ratio - Full Bundle Boot on demand shows the IOPS for this write cache data and shows the average read write ratio for the each of the workloads during the Boot on demand runs.

Workload	Average Read IOPS	Average Write IOPS	Read%	Write %
Basic	11.43	403.95	**3%**	**97%**
Standard	49.59	391.44	**11%**	**89%**
Premium	40.56	286.93	**12%**	**88%**

Figure 9: WriteCache Read: Write Ratio - Full Bundle Boot on demand

The Utility Cluster Performance for the full bundle testing at peak load on each workload is summarized in Figure 10: Utility Cluster Summary by Workload.

As expected, the CPU and Memory required by each of the cluster nodes decreases as the number of desktops and users decreases rather than by the workload of the user. This is due to the fact that each of these services is impacted by either the number of desktop power actions (SQL Services are impacted from the XenDesktop Controllers, Provisioning Services and VMware vCenter), or the number of users accessing data on the file shares hosted by the File services node.

Workload	File Service Node CPU	File Services Node Committed Memory	SQL Services Node CPU	SQL Service Node Committed Memory
Basic	45% Typical (Peaking to 65%)	1.86 GB	35% Typical (Peaking to 100%)	7.8 GB
Standard	35% Typical (Peaking to 65%)	1.66 GB	35% Typical (Peaking to 100%)	5.87 GB
Premium	25% Typical (Peaking to 50%)	1.64 GB	20% Typical (Peaking to 100%)	5.14 GB

Figure 10: Utility Cluster Summary by Workload

4.3.2 Desktop Host Performance Results

The Desktop Host performance graphs shown in this section of the report demonstrate a summary of key performance metrics within each desktop host during each test run. During testing the Dell Solutions Center team found that if memory was significantly over allocated this resulted in the ESX hosts needing to start "swapping" memory.The result of this was a jump in the IOPS on the storage arrays and a jump in CPU usage as the host was trying to "recover" available memory from each of the virtual desktops. The net effect of this was that user experience was degraded. The target for CPU utilization was a peak of between 80 and 85% this allows the host additional headroom for real-world spikes in user activity in production while maximizing the desktop density on a given host. The Dell Integrated Solution Stacks provide an optimal server configuration to ensure memory and CPU are utilized up to 85% which provides the best cost/performance benefit.

4.3.2.1 Full Bundle Pre-Booted Basic Users

Pre-Booting virtual desktops prior to typical usage patterns provides an optimal end user experience and is a recommended infrastructure resource management strategy. The Solution Architect will work with the customer to ensure the desktop pool boot policy meets the customer requirements based on Dell Blueprint and customer analysis.

The following graph shows the pre-booted test run, the first section of the graph (07:05 a.m. to 08:05 a.m.) shows the virtual desktops booting up, then the 20 minute idle period where the hosts are steady at around 15% CPU usage and from around 08:30 a.m. CPU usage starts to climb as users start to connect to the virtual desktops. It is worth noting that as with the IOPS results the peak values are not significantly different, except that during the period when users are connected the host demonstrates a smoother rate of change compared with the runs when the desktops are booted on demand. User sessions logoff at 09:50 a.m.

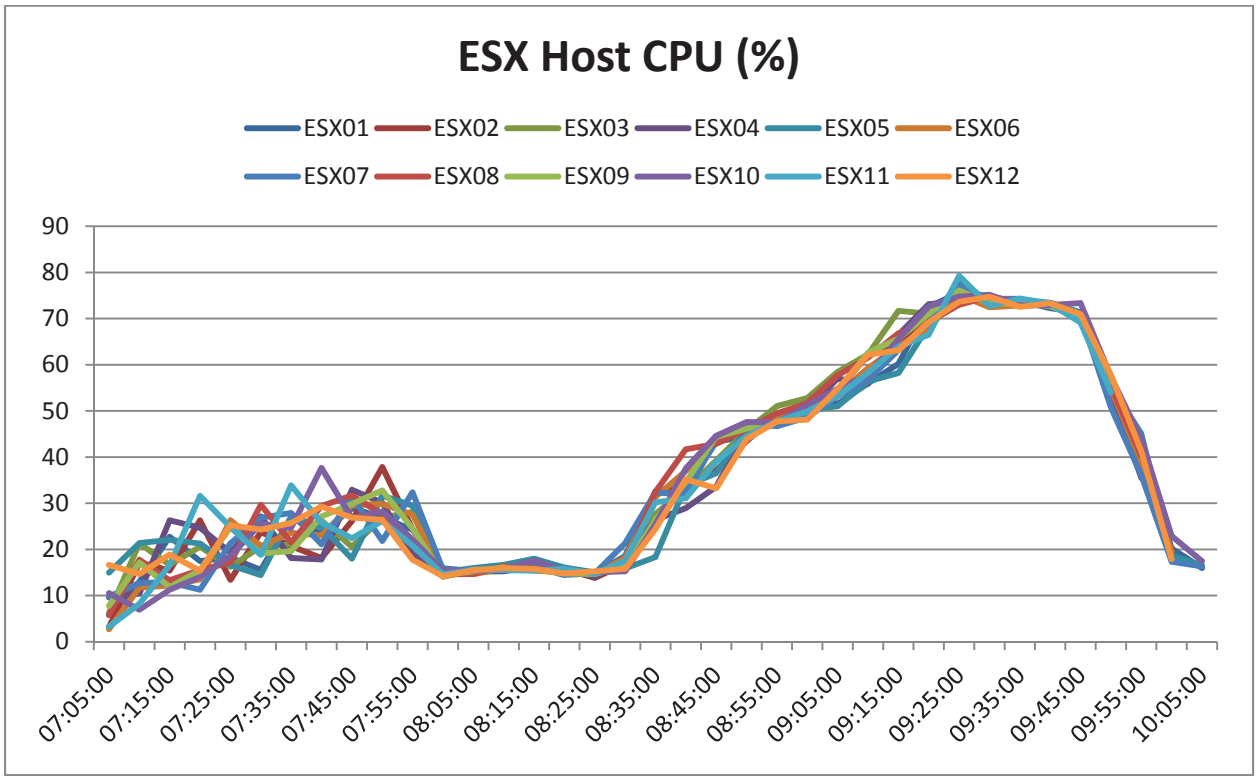

At the same stage of the test run the hosts were all consuming approximately 85GB of RAM which brings the hosts to approximately 89% of installed RAM being consumed, at this level there is no swap file in use therefore the need to swap memory to disk is having no effect on CPU or storage IOPS.

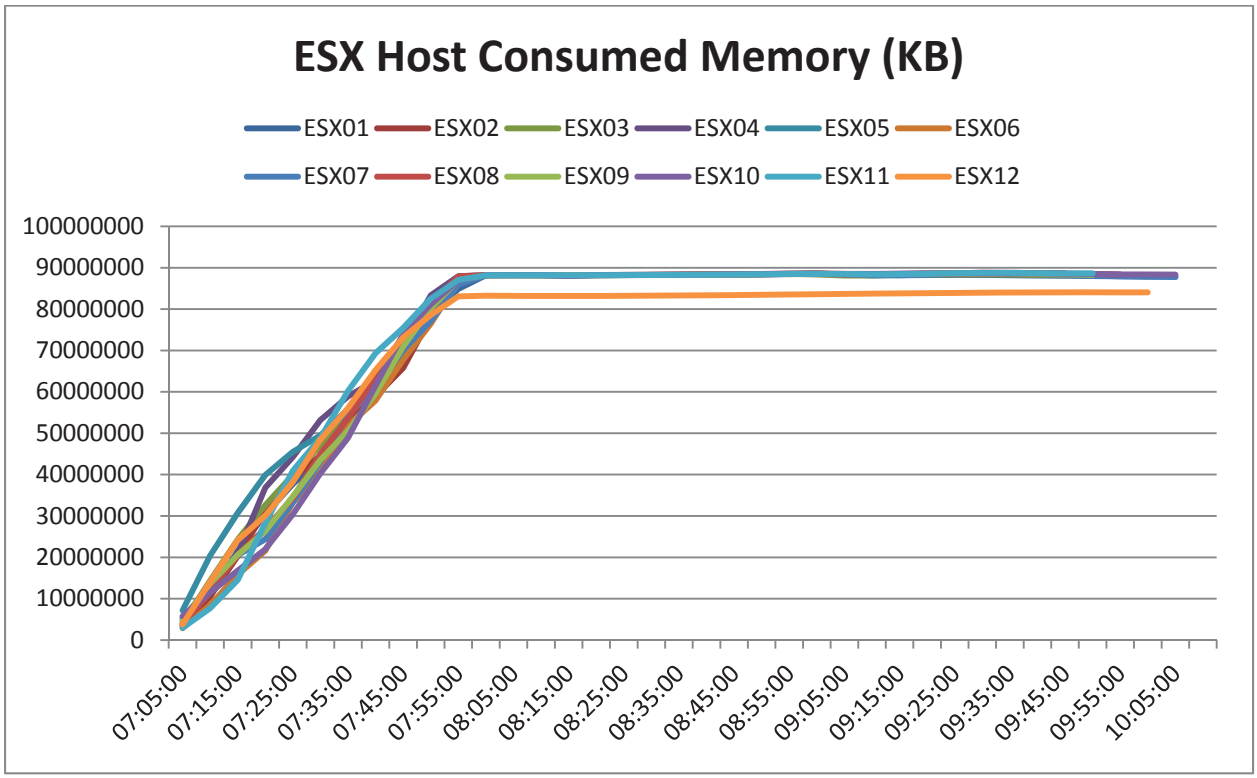

The EqualLogic PS6010XVS arrays are primarily used to host the Citrix Provisioning Server vDisks and the Write Cache data for the virtual desktops. The following graphs show the overall IOPS usage for these two arrays throughout the test run. As these show there is very little read activity on the arrays the bulk of the traffic is write related (approximately 97% write) and IOPS remained below 20ms.

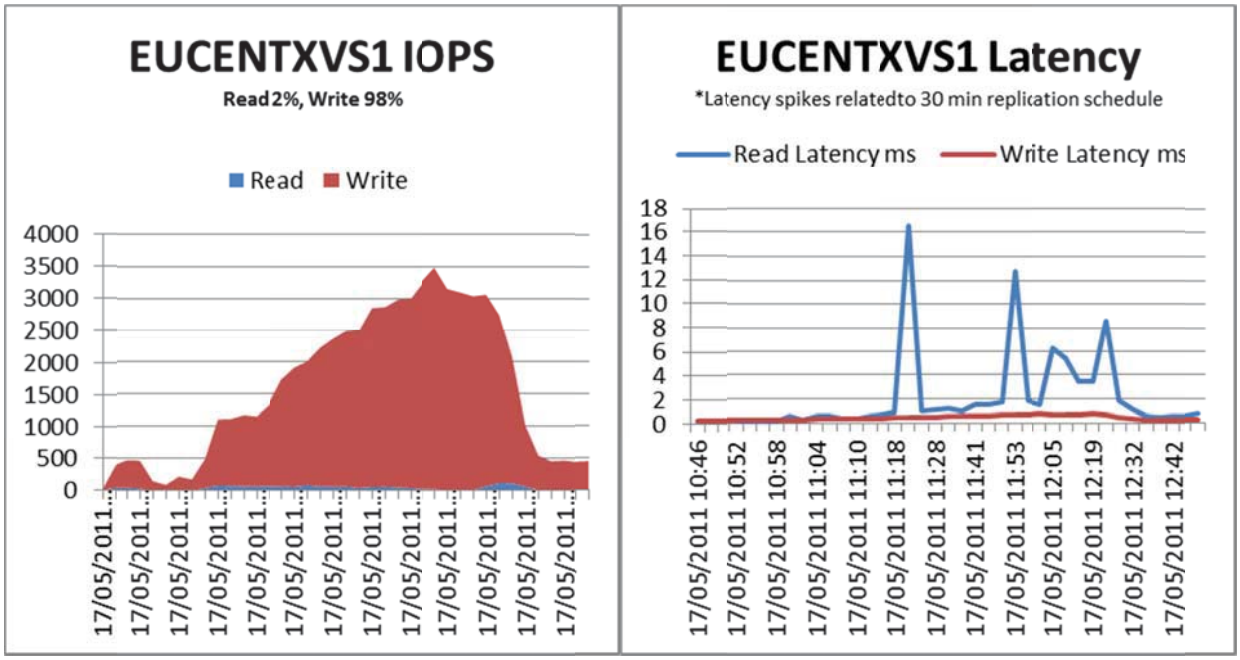

The file Share IOPS shown in the following graph and table is a subset of the PS6510E IOPS within Tier 2.

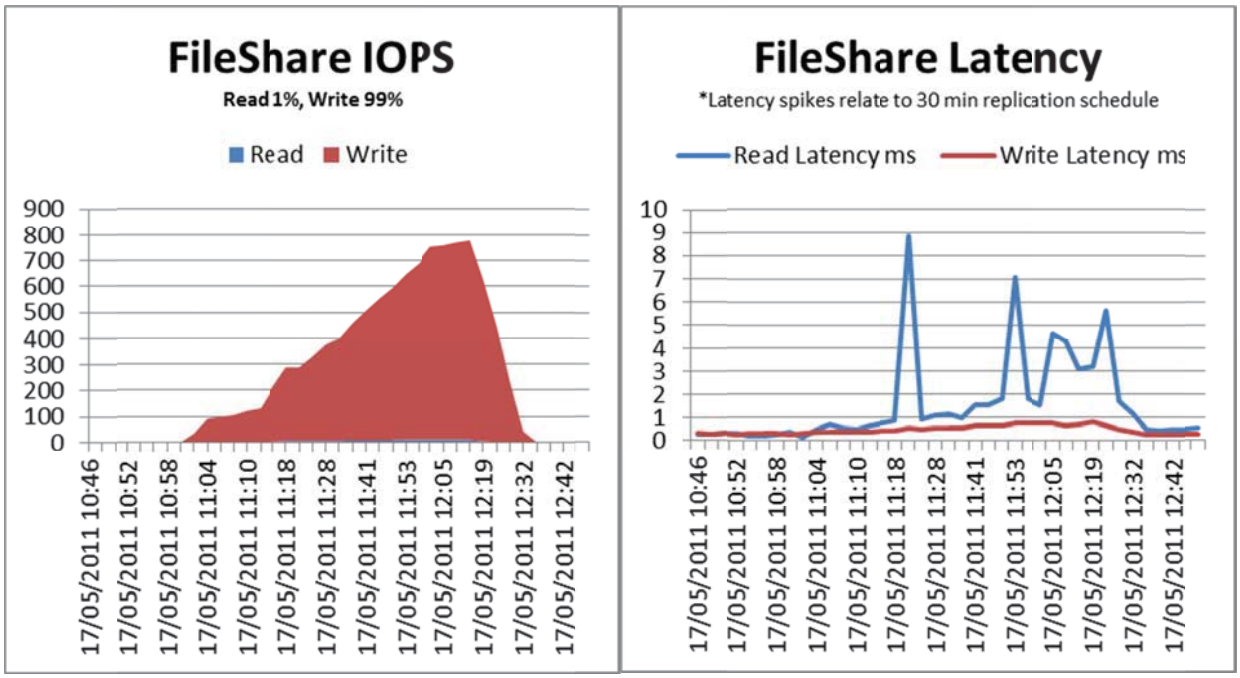

The following network usage graphs show the combined receive and transmit rates for all physical network interfaces in the desktop hosts. With a maximum receive rate of just under 30MBps (234Mbps) it is clear that no single interface in the hosts is in danger of becoming saturated. The transmit rate is always below 4MBps (31Mbps) so even with thee vales combined it is clear that each host is running at less than 25% of a single 1Gbps interface (the hosts actually had 2x1Gbps and 1x10Gbps interfaces).

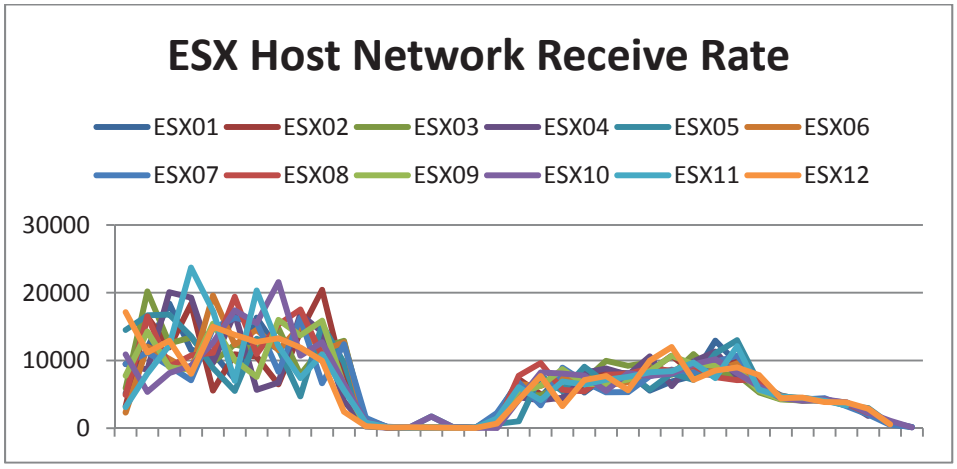

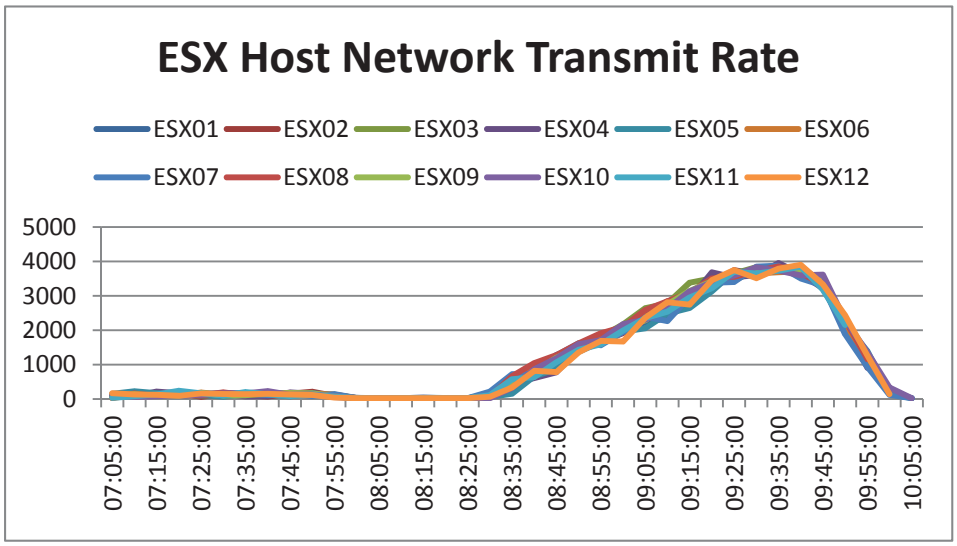

See the Appendix for additional results graph.

4.4 Scaling Out the Integrated Solution Stack

To provide a scalable and predictive solution stack Dell developed scalability and design criteria based on the extensive testing and characterization within the Dell Solutions Center lab. Based on this analysis the following design criteria have been established to scale the ISS Enterprise Plus. This drives a pod-based design criteria which allows modular scalability and requires certain design constraints as outlined below to ensure positive end user experience as it relates to infrastructure capacities.

ISS Enterprise Plus Base Pod:

- One M1000e blade enclosure
- Two M610 management servers in a VMware vSphere 4.1 cluster.
- Two M610 physical utility cluster servers in a Microsoft 2008 R2 mirror cluster.
- Five M610 virtual machine host servers in a VMware vSphere 4.1 cluster.
- Two Tier 1 EqualLogic PS 6010XVS storage array (Scalability depends on workload)
- One Tier 2 EqualLogic PS 6510E storage array (Scalability depends on workload)
- Two top of rack PowerConnect-J EX 4500 switches supports up to 2 pods.
- Pod Max: The pod can support up to twelve virtualization hosts. At this point additional storage, networking and management and utility cluster hosts are required to maintain service levels.

Once the pod design criteria are met then a new ISS base pod is required starting with management servers and networking to scale. The quantity of users within the ISS pod will be determined based by the mix of user types. Please work with the Dell Virtualization Specialist and Solutions Architect to determine user scalability per pod based on the results of a Dell Blueprint Assessment.

5 Integrated Solution Stack Architecture

This section describes the high-level solution architecture.

5.1.1 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 / down to support business needs.
- Resilient – The solution must be able to withstand the failure of a single component.

5.1.2 Scalable

The architecture is designed to provide a scalable platform:

- The components can be scaled horizontally by adding additional physical and virtual servers to the server pools

- The architecture has been designed to eliminate bandwidth and performance bottlenecks as much as possible
- The architecture has been designed to allow future horizontal and vertical scaling with the objective of reducing the future cost of ownership of the infrastructure; The table below summarises the scalability options for the components of the infrastructure:

Component	Horizontal scalability	Vertical scalability
Virtual Desktop Host servers	Additional hosts and clusters added as necessary	Additional RAM Higher specification servers
Provisioning servers	Additional servers added to the Provisioning Server farm	Additional network and I/O capacity added to the servers
XenDesktop Controllers	Additional servers added to the XenDesktop Site	
Web Interface Servers	Additional servers added to the web Interface environment	Additional virtual machine resources (RAM and CPU)
Load Balancers	Additional virtual appliances to split the Web Interface and XML Broker traffic	Upgrade license to support increased network throughput and/or increase virtual machine resources (RAM and CPU)
Database Services	Migrate databases to a dedicated SQL servers and increase the number of cluster nodes	Additional RAM and CPU for the cluster nodes
File Services	Split user profiles and home directories between multiple file servers in the cluster (will require additional cluster nodes)	Additional RAM and CPU for the cluster nodes
VMware vCenter	Deploy additional servers and use linked mode to optimise management	Additional RAM and CPU for the cluster nodes (limited to 200 hosts/2000 VM's per server)

ISS Scaling Strategy

5.1.3 Resilient

Dell's approach in creating the design for the solution is based on the following strategies:

- Reducing or eliminating single points of failure where possible.
- Using standardised hardware and virtualising the operating systems of the infrastructure components to allow the migration of services between servers in the event of server failures (whether physical or virtual servers).
- The scope of this design is to provide standard image mode virtual desktops to the user community; however, the infrastructure is scalable to handle increased requirements in the future.

5.1.4 Physical Architecture Overview

The diagram below provides a summary of the Dell Integrated Solution Stack logical design deployed for the Enterprise Plus bundle. This diagram outlines the logical design of the compute resource pools for the virtual desktops, management and cluster host pools including the storage tiers.

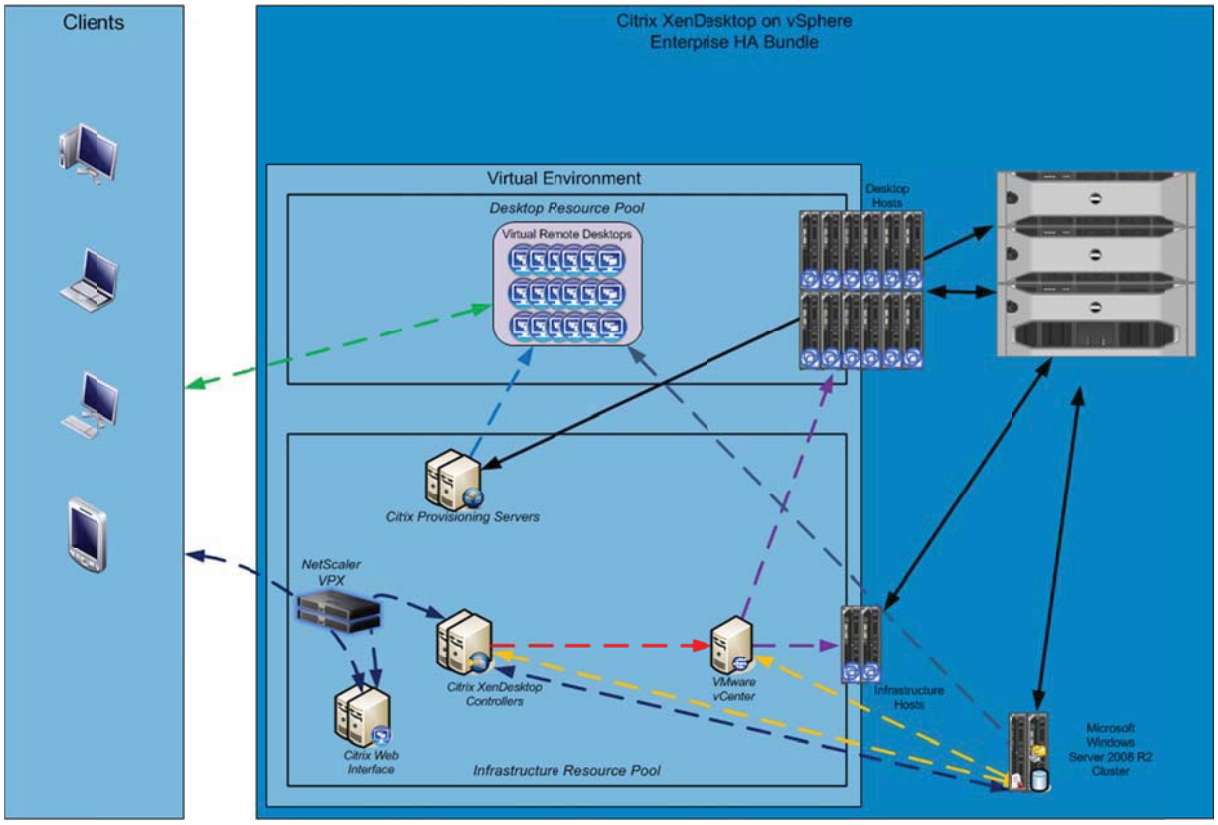

Figure 11: ISS Logical Design

All XenDesktop Virtual Desktops will be run as virtual instances hosted VMware vSphere infrastructure. In addition to the Virtual Desktop hosts there will be a pair of vSphere host servers that will provide compute resources for infrastructure services such as web servers, broker services, provisioning services and load balancing services within the management servers. The solution will also leverage a physical Microsoft Windows Server 2008 R2 Cluster to host file and database services within the environment.

As shown in Figure the Virtual Infrastructure will be configured into two (3) resource pools. The first pool will deliver infrastructure services, the second will deliver resources for the Virtual Desktops and the third will provide file and database services.

5.1.5 ISS Host Configuration

ISS Full Bundle Server Configuration	
	Enterprise Plus
Virtual Desktop Host	(12) x M610 PowerEdge Servers
Hypervisor	VMware ESXi 4.1 U1 Standard
Management Host	(2) x M610 PowerEdge Servers
Cluster Host	(2) x M610 PowerEdge Servers
Storage Tier1	EqualLogic PS 6010 XVS SSD/SAS hybrid array
Storage Tier 2 & 3	EqualLogic PS 6510 E SATA array
Network	PowerConnect J EX 4500 with virtual chassis
Chassis Configuration	
I/O module for A1	PowerConnect M6220
I/O module for B1	PowerConnect M8024
I/O module for C1	None
I/O module for A2	PowerConnect M6220
I/O module for B2	PowerConnect M8024
I/O module for C2	None
Management	Redundant Chassis Management Controller (CMC)
KVM	Integrated Avocent keyboard, video, and mouse (iKVM) switch

Virtual Desktop Host Configuration	
Server Model	M610 PowerEdge Blade
Processor	(2) x Intel Xeon E5670, 2.93Ghz, 12M Cache
Memory	96 GB (12 x 8 GB, RDIMMS 1333 MHz)
Add-in Network Card	Broadcom 57711 Dual Port 10 GbE I/O Card for M-Series Blades
On-board Network Card	Broadcom 5709 2 port 1 GbE with w/ TOE.
Local storage and controller	Diskless configuration
Hypervisor	VMware ESXi 4.1 U1 Standard Edition
Management Host Configuration	
Server Model	M610 PowerEdge Blade
Processor	(2) x Intel Xeon E5670, 2.93Ghz, 12M Cache
Memory	96 GB (12 x 8 GB, RDIMMS 1333 MHz)
Add-in Network Card	Broadcom 57711 Dual Port 10 GbE I/O Card for M-Series Blades
On-board Network Card	Broadcom 5709 2 port 1 GbE with w/ TOE.
Local storage and controller	Diskless configuration
Hypervisor	VMware ESXi 4.1 U1 Standard Edition
Utility Cluster Host Configuration	
Server Model	M610 PowerEdge Blade
Processor	(2) x Intel Xeon E5670, 2.93Ghz, 12M Cache
Memory	24GB (12 x 2 GB, RDIMMS 1333 MHz)
Add-in Network Card	Broadcom 57711 Dual Port 10 GbE I/O Card
On-board Network Card	Broadcom 5709 2 port 1 GbE with w/ TOE.
Local storage and controller	Perc 6 with 256 MB Cache (2) 146 GB 15k SAS
Operating System	Microsoft Server 2008 R2 64 bit

5.2 Citrix XenDesktop 5 Configuration

All Virtual Desktops will be run as virtual instances hosted on VMware vSphere infrastructure. In addition to the Virtual Desktop hosts there will be a pair of management host servers that will provide compute resources for infrastructure services such as, web servers, broker services, provisioning services and load balancing services. Wherever possible these services will be provided by virtual servers. See Figure for logical design.

The solution will also leverage a physical Microsoft Windows Server 2008 R2 Utility Cluster to host file and database services. This provides no single point of failure for the core database and file services.

The Virtual Infrastructure will be configured into two (2) vSphere clusters. The first pool will deliver Infrastructure Management services while the second will deliver resources for the Virtual Desktops themselves.

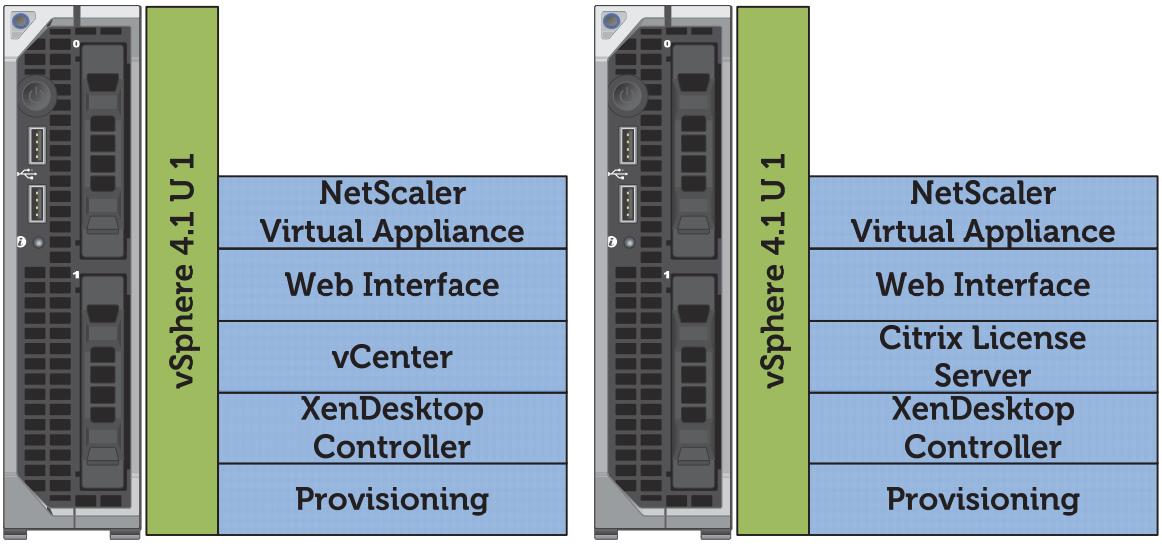

Figure 12: Management host cluster

Citrix XenDesktop Software Components:

- XenDesktop Controller 5.0
- Citrix License Server 11.6.1
- Citrix Web Interface 5.4
- Citrix Provisioning Server 5.6 Service Pack 1
- Citrix Netscaler VPX 200
- Citrix Profile Manager

5.2.1 Citrix License Server 11.6.1

The Citrix License server component requires minimal system resources and no hardware requirements above those required by the operating system used to host the service. The license service will run in Windows Server 2008 R2 virtual machine. Resilience for this service will be provided by VMware High Availability features which will be configured in the Management Host cluster.

5.2.2 Citrix Web Interface 5.4

As with the Citrix License server detailed above the resource requirements for Web interface are minimal and will run in a Windows Server 2008 R2 virtual machine as well. The Web Interface virtual machine will reside on the Management Host cluster.

5.2.3 Provisioning Server 5.6 Service Pack 1

The provisioning service streams a single desktop image (vDisk) to create multiple virtual desktops on one or more servers in a data center. This facility greatly reduces the amount of storage required compared to other methods of creating virtual desktops. The more target devices using the same vDisk image, the less vDisks need to be created; making vDisk management easier.

5.2.4 XenDesktop Controller

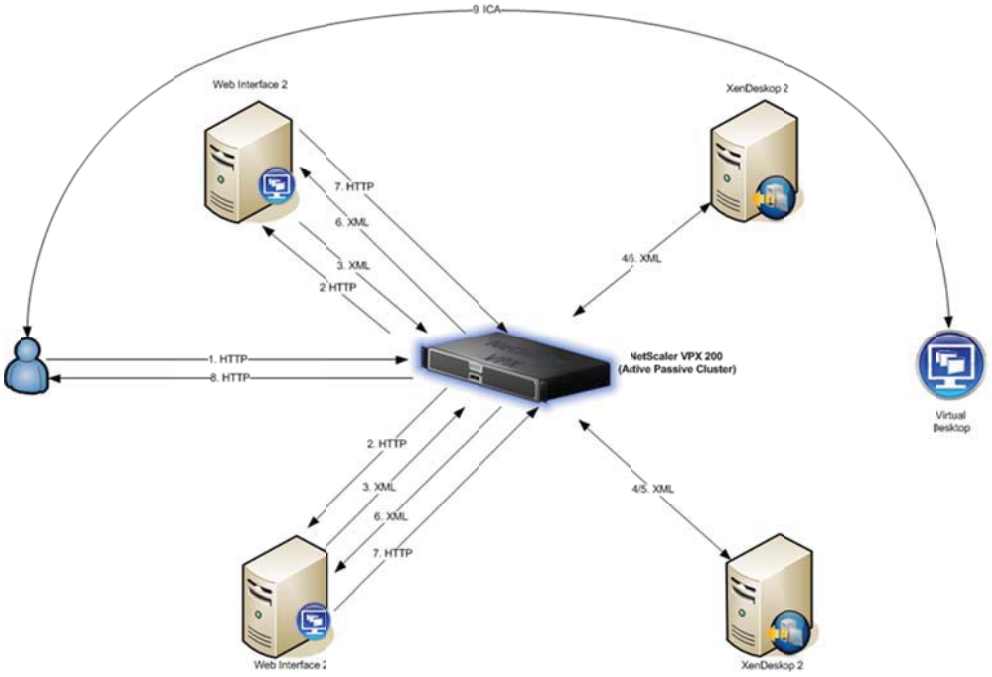

Figure 13: XenDesktop Controller Overview

The XenDesktop Controllers manage the virtual desktops and the sessions running on them. These servers broker connection requests from client endpoint devices and provide the connection information via the Web interface server to the endpoint device to allow a direct ICA connection to the Virtual Desktop Agent running within the virtual machine. The Controller also manages licensing and the database that contains the configuration and session information for the site. A site is a group of controllers that are administered and operate as a single entity.

5.2.5 Citrix Profile Manager

Citrix Profile Management is a component of the XenDesktop suite which is used to manage user profiles an minimize many of the issues associated with traditional Windows Roaming profiles in an environment where users may have their profile open on multiple devices at the same time.

In order to further optimize the profile management, folders within the user profile that are used to store data will be redirected the users' home drive via the file services. The folder redirection will be managed via group policy objects within Active Directory through the Profile Management tools.

5.2.6 Netscaler VPX 200

The Netscaler platform has been incorporated in the Dell Integrated Solution Stacks to provide load balancing capabilities for the Citrix Web Interface and XML services at this stage, however, since these appliances have the ability to provide secure remote access to the Citrix XenDesktop environment their inclusion in the solution provides the ability for the solution to be upgraded in the future if this functionality is required.

As shown in Figure 13: XenDesktop Controller Overview the Netscaler appliances will be configured as a High availability pair, this means that the appliance will operate as an active/Passive cluster. The appliances will provide load balancing services for the following XenDesktop5 components;

- Citrix Web Interface
- XenDesktop XML service
- XenDesktop Desktop Director
- Provisioning Services TFTP service

5.2.7 Microsoft Cluster Services

A Physical Windows "Utility" Cluster will be deployed. The Microsoft Cluster will be deployed to provide a resilient architecture for the SQL server databases and file services required within the environment.

During the design process an alternate option of a virtualized file server protected by VMware High Availability was considered however, it was felt that since the file server would be responsible for hosting the user data and user profiles it is recommended that the file shares are deployed as a clustered resource. A secondary option of a virtualised cluster was also considered, however, a virtualised Microsoft cluster hosted on iSCSI storage is currently unsupported by VMware on vSphere 4.1 and would only be supported by Microsoft for its customers with a Premier agreement; this resulted in the cluster becoming a physical instance.

Since the cluster is required for file services Dell decided to use the passive node to host the Microsoft SQL databases required by the Citrix and VMware components within the architecture. The primary reason for this decision was ensure the solution made efficient use of the hardware and software licensing. The alternative would have been to deploy the SQL databases as mirrored databases on virtual SQL servers, however, to achieve automated failover of the databases the solution would have required an additional three (3) virtual Windows servers for the Primary, Secondary and witness roles in the database mirrors.

5.3 VMware vSphere 4.1 Configuration

The solution is based on VMware vSphere 4.1 Update 1 Standard Edition. Each virtualization host will run ESXi on the Dell embedded SD memory card in a stateless host configuration. A Virtual Center virtual machine will be utilized in the stack running Windows 2008 R2 residing on one of the Virtual Infrastructure Management Servers.

Note: vSphere 5 Standard Edition licensing will be downgraded to vSphere 4 until Citrix XenDesktop fully supports vSphere 5. Licenses will be procured according to the vSphere EULA in a per CPU model vs. the new vRAM entitlement model in vSphere 5. Future upgrades to vSphere 5 by the customer will require additional vmware licensing according to vSphere 5 EULA.

This database for vCenter will reside on the Utility cluster running under Microsoft 2008 R2 Clustered Services. Also, under this same cluster resource we will run a separate database for VMware Update Manager which will be used for host updates.

The Utility cluster will provide failover and ease of maintenance for SQL services. The cluster will help mitigate against hardware failure and also simplify maintenance operations for both the hardware and the software.

A "gold image" VM will be generated for each base Operating System required within the solution. Customization Scripts will be added to vCenter to ensure all cloning activity from these VM's produce VM's with new SID (VMware customization scripts are used to ensure each cloned machine is properly sys-prepped to ensure there are no duplicated SID's) and also licensing, naming convention and some default network settings etc. are part of the Dell ISS Implementation Solution.

Within vCenter, there will be a single Datacenter containing 2 HA clusters. These 2 clusters will contain the Virtual Desktops and the Infrastructure VM's respectively.

Each of the 2 vSphere Standard Edition clusters will have the same HA cluster settings;

- HA Enabled
- DRS Disabled (Not licensed)
- Power Management (Disabled)
- HA settings (Enabled, most conservative)
- HA tolerates One Host Failure
- Admission Control Enabled.

HA will ensure that when a single host fails, its guest VM's are powered up elsewhere in the cluster. It will also provide placement for all new VM's powered up within the cluster i.e. it will pick the host with the most resource available.

6 Network Architecture

This section provides an overview of the Dell Desktop Virtualization Solution Stack network design.

6.1 Understanding the Dell Blade Network Architecture in Dell ISS Enterprise Plus

The Dell M1000e blade chassis has three separate fabrics referred to as A, B, and C[1]. Each fabric has two I/O modules, for a total of six I/O modules in the chassis. The I/O modules are A1, A2, B1, B2, C1, and C2. Each I/O module can be an Ethernet physical switch or an Ethernet pass-through[2]. In this solution, fabric B contains redundant PowerConnect M8024 10 Gigabit I/O module and fabric A contains redundant PowerConnect M6220 10 Gigabit I/O module.

Each half-height blade server has a dual-port onboard 1Gb NIC and two optional dual-port mezzanine I/O cards. See Figure 4 for the PowerEdge M610 I/O connectivity mapping.

---

[1] http://support.dell.com/support/edocs/systems/pem/en/HOM/HTML/about.htm#wp1212100

[2] Fabrics B and C support Fibre Channel and InfiniBand, which are outside the scope of this solution.

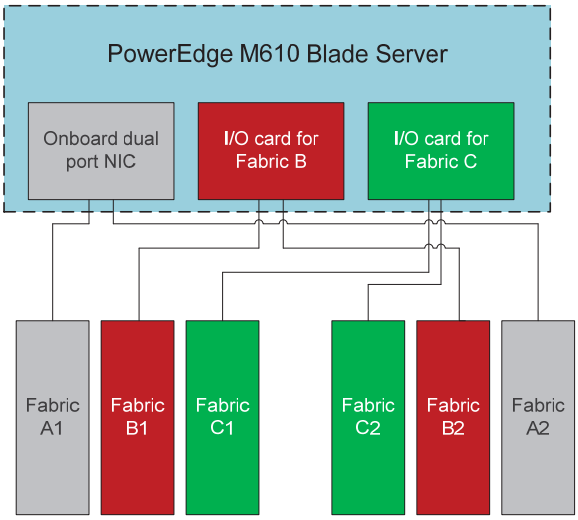

Figure 4: Adapter and I/O Modules connection in Chassis for M610 Blade Server

The Dell ISS Enterprise Plus Hosts in this solution leverages Broadcom 57711 dual-port 10 Gb Ethernet mezzanine I/O cards to connect both Fabrics B1 and B2. PowerEdge M610 blade servers feature two onboard 1 Gb Broadcom NetXtreme II™ 5709 NICs, which connect to Fabric A1 and A2. Fabric C is available for future expansion. The M1000e Chassis I/O modules are uplinked to the top of rack Juniper EX 4500 as outlined in Figure 5.

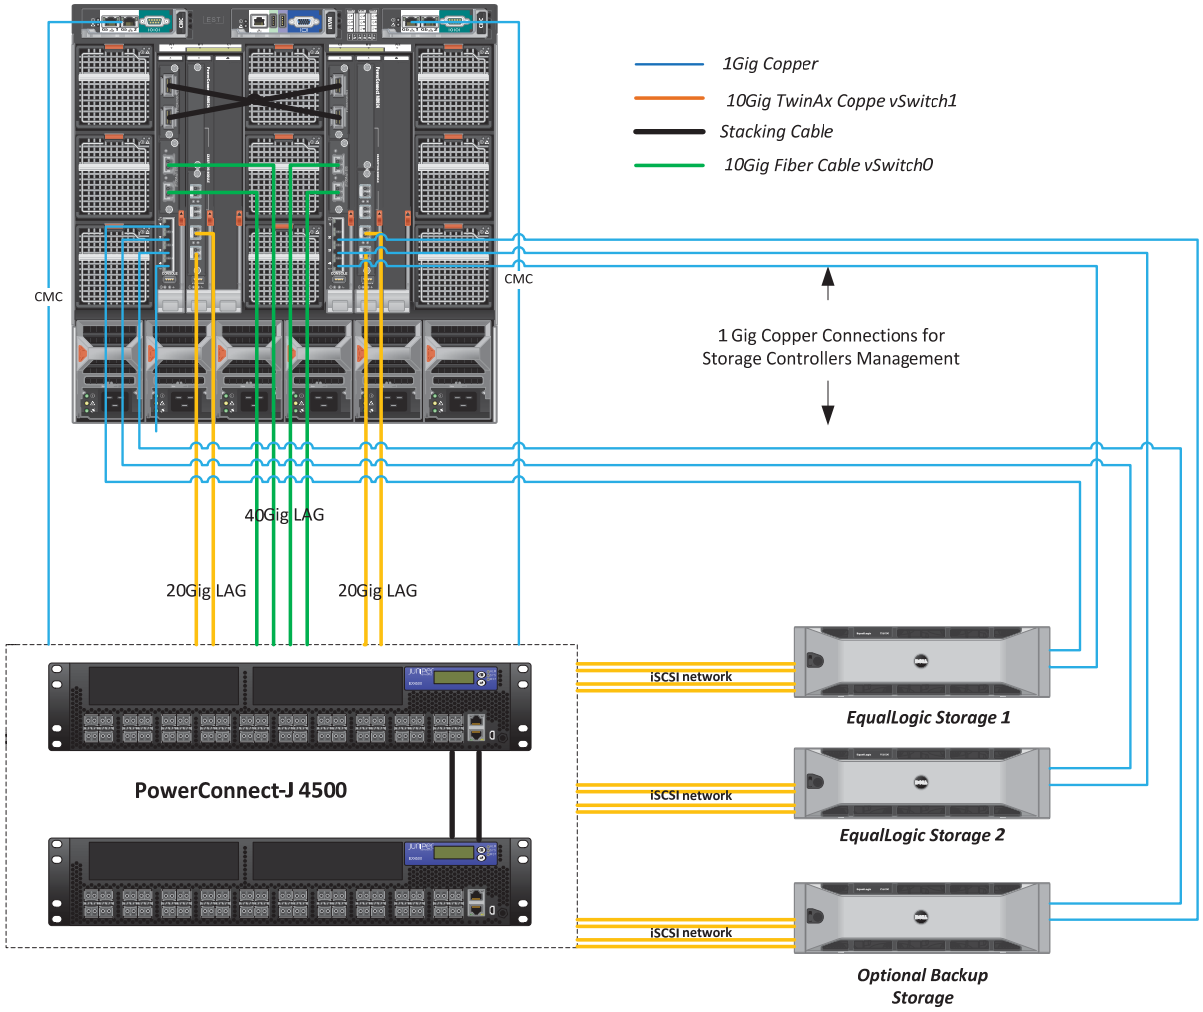

Figure 5: Dell ISS Enterprise Plus Network Design

6.2 vSphere Network Overview

vSphere network traffic is typically comprised of four types: virtual machine traffic, management traffic, vMotion traffic, and iSCSI traffic. Two separate physical networks are created to support different traffic types (see Dell ISS Enterprise Plus Network Design).

The vSphere virtual network is designed into two logical Virtual Switches (vSwitch). The first is vSwitch0 which supports Service Console, vMotion and iSCSI. The second is vSwitch1 which supports the desktop and vCenter production networks.

The Dell Blade chassis network for the ISS Enterprise Plus has vSwitch0 traffic on Fabric B and vSwitch1 on Fabric A. (See Figure 5)

Each network is segmented into vLAN's for network segmentation and security. These are iSCSI, Management and Production Network traffic and assigned to vSphere Port Groups within each vSwitch. Network traffic is tagged with the respective VLAN ID for each traffic type in the virtual switch.

The Dell PowerEdge R710 has redundant dual 1 GbE uplinks for vSwitch0 and vSwitch1.

Design Principles for Network Architecture: The following design principles are used to develop the network architecture to support the Dell Integrated Solution Stacks:

- **Redundancy:** Both LAN and iSCSI SAN have redundant I/O modules. Redundancy of the network adapters is achieved through NIC teaming at the virtual switch. There is no single point-of-failure.
- **iSCSI SAN physical isolation:** The iSCSI SAN should be physically separated from the LAN to avoid any network congestion.
- **Logical isolation of vMotion using VLAN:** vMotion traffic is unencrypted. It is important to logically isolate the VMotion traffic using VLANs. *Not applicable in XenDesktop Dell ISS*.
- **Optimal performance:** Load balancing is used to achieve the highest throughput possible.

Based on these principles, industry best practices from Citrix, VMware and Dell for load balancing, failover, and VLAN configuration are provided with the Dell Integrated Solution Stacks.

This section assumes a certain understanding of networking concepts in a virtualized environment. See *http://www.vmware.com/files/pdf/virtual_networking_concepts.pdf* for an overview of VMware's networking capabilities.

6.3 iSCSI SAN Architecture

In this section we discuss the iSCSI SAN configuration for the Dell ISS Enterprise Plus and Enterprise solutions.

6.3.1 Inter Switch Links

It is important to understand the ISL requirements for switches connected to the EqualLogic storage. ISL requirements vary based on I/O required. For the Dell Integrated Solution Stacks ISL's apply to the Enterprise Plus solution stack. Dell has designed a 40 GbE ISL to support iSCSI network traffic along with vMotion and Management traffic.

6.3.2 Storage Connectivity

Each controller on the Dell EqualLogic PS Series array has four 10 GbE ports (two active and two passive connections). The two active ports are split between two different physical switches to provide redundancy. There are redundant 1 GbE management ports as well. The passive ports are connected in a similar fashion. Illustrated in Figure 5

6.3.3 Virtual Switch Design

A virtual switch (vSwitch0) is created and two 10 Gigabit physical network adapter ports are connected as uplinks to the virtual switch. These network adapter ports are in turn connected to PowerConnect M8024 switch I/O modules in Fabrics B1 and B2.

Two VMkernel interface is created and associated with each physical network interface. The VMkernel interfaces are used by the software iSCSI initiator to connect to the storage arrays. This is illustrated in Figure . For detailed discussion regarding configuring VMware iSCSI software initiator, refer to the Technical Report: *Configuring VMware vSphere software iSCSI with Dell EqualLogic PS Series Storage*.

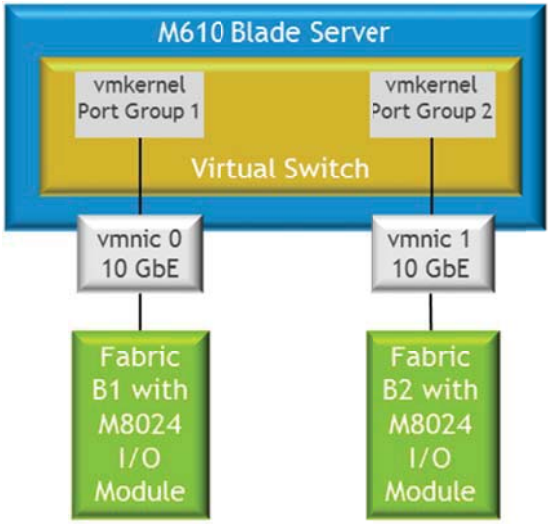

Figure 16: iSCSI Host Network Configuration

6.3.4 Jumbo Frames

Jumbo frames are enabled in the iSCSI network to improve performance. Jumbo frames are enabled on the VMkernel interfaces, virtual switches, and PowerConnect switches.

6.3.5 SAN Load Balancing

Multipathing is a technique that allows more than one physical path to be used to transfer data between a host and an external storage device. The Dell EqualLogic MPIO plugin will be installed to handle load balancing. VMware vSphere 4.1 offers many new and advanced enhancements to the software iSCSI initiator beyond basic iSCSI SAN connectivity. The most significant of these enhancements is the API support for third party mutipathing plugins. This provides a framework that enables the EqualLogic MEM to utilize VMware's Round Robin multipathing and to more intelligently route and efficiently load balance iSCSI traffic across multiple NICs.

The EqualLogic MEM offers:

- Ease of install
- Increased bandwidth
- Reduced network latency
- Automatic load balancing across multiple active paths
- Automatic connection management
- Automatic failure detection and failover
- Multiple connections to a single iSCSI target

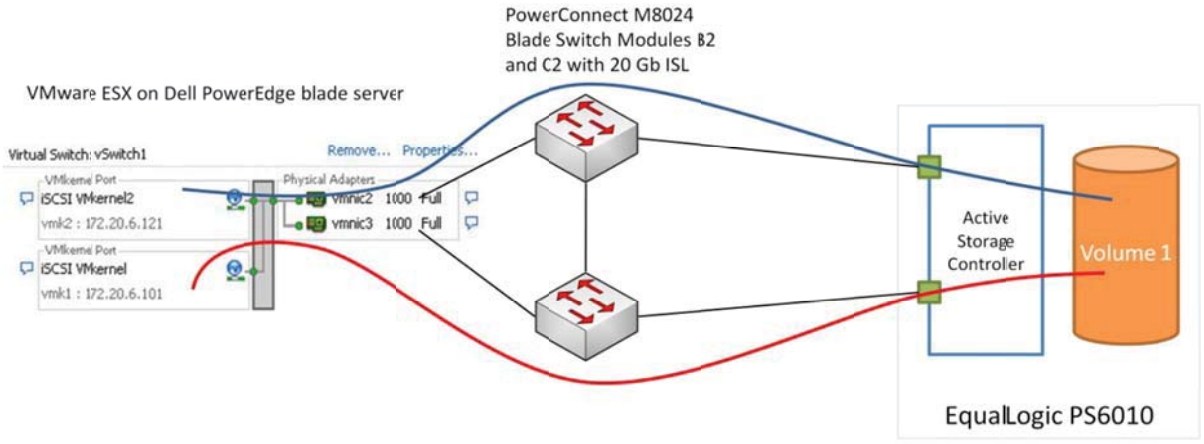

Figure 6: Multipathing using Round Robin

The Dell EqualLogic SAN automatically distributes iSCSI connections between the available two network interfaces in each controller. It also automatically distributes volumes between different storage devices in the storage pool. No additional configuration is required.

For a detailed discussion of iSCSI connections, see the VMware ESX/ESXi 4.0 "iSCSI SAN Configuration Guide" for a detailed description of how to configure the iSCSI connection.

For more details on the EqualLogic Multipathing Extension Module for vSphere see Configuring and Installing MEM for vSphere

7 Storage Configuration

The Dell Integrated Solution Stack storage architecture is segmented into three tiers to maximize performance and optimize costs. Within each Integrated Solution Stack the storage architecture varies to provide ideal cost, functionality and performance balance. The end user VDI image is segmented to take advantage the storage tiers based on I/O and Capacity requirements in each tier.

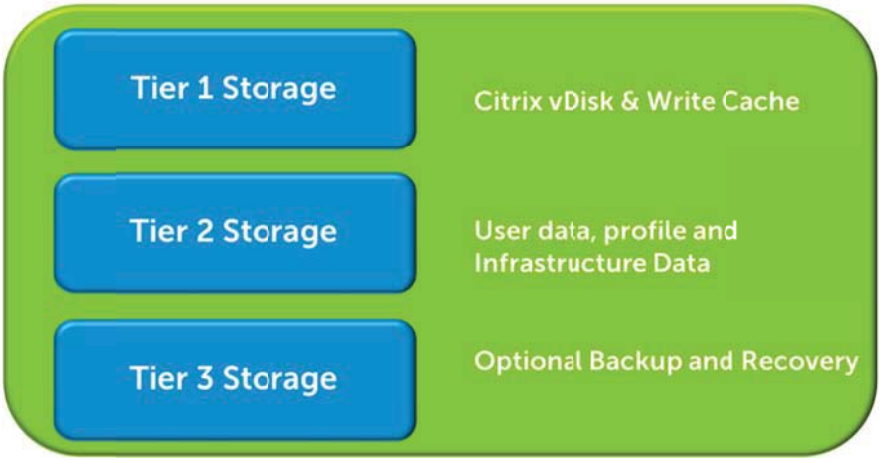

Figure 7: End User Disk and Data Segmentation

The ISS Enterprise Plus provides a 10 GbE based EqualLogic storage solution leveraging the PS 6010XVS SSD and SAS Hybrid Array for Tier 1 storage and the PS 6510E SATA array for Tier 2 and 3.

- Tier 1 storage is configured with RAID 6
- Tier 2 storage is configured with RAID 50
- Tier 3 is customer optional and configuration depends on replication requirements.

The configurations of the storage volumes are based on common best practices for space, I/O, block size and RAID configurations defined by Dell, VMware and Citrix. The recommended and tested configuration for the Dell ISS Enterprise Plus volumes were implemented as outlined below.

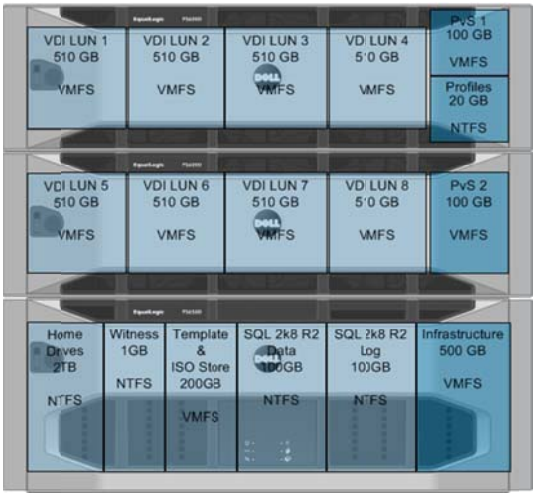

The three EqualLogic arrays (2x PS6010XVS and 1x PS6510E) were configured as a single Group and each were a member of a dedicated Storage Pool. Storage pools are created to ensure that the arrays do not try to distribute load across the units which is not currently supported on the XVS series arrays.

The workload characterization performed by the Dell Data Center Solution Center quantified the density of users which could run within the platforms before performance degraded beyond acceptable latency levels within the storage tiers. It is expected that latency above 20ms would provide a poor end user experience. The results of the Scalability analysis provided the Dell Desktop Virtualization Solutions Group with the necessary metrics to define a scalable POD architecture to ensure workloads would not over-utilize the infrastructure thus maintaining an acceptable or improved end user experience when moving users to a virtual desktop model.

See (section 2.2.1) for further details on the arrays.

See (4.3.2.1 Full Bundle Pre-Booted Basic Users) for test result data.

7.1 Systems Management

The following tools and software can be used to manage the Dell Integrated Solution Stack:

Deployment and change management using Lifecycle Controller: Dell PowerEdge M610 blade servers come with USC. This helps reduce operating costs by simplifying deployment and management. Key features include: Diagnostics, self-update (UEFI, Driver Pack update), firmware updates (BIOS, NIC FW, RAID Controllers), and hardware configuration.

Out-of-band CMC and iDRAC: The CMC provides a single, secure interface to manage the inventory, configuration, monitoring, and alerting for chassis components (iKVM, CMC), I/O modules, servers, and iDRAC. It also provides excellent real-time power management, monitoring, and alerting capabilities. The Dell chassis provides users with system-level power limiting, slot-based prioritization, and dynamic power engagement functionalities. The iDRAC on each server provides the flexibility to remotely manage the server through Console redirection and Virtual CD-ROM/DVD/Floppy/Flash capabilities.

Dell Management Console: The Dell Management Console (DMC), powered by Symantec, is a systems management solution designed to provide a central console offering features from basic hardware management to advanced enterprise functionality. The DMC creates a foundation for more advanced management functionality and provides a single view into the deployment, inventory, monitoring, and updating of your IT infrastructure - including servers, storage, desktops, notebooks, network devices, printers, and other non-Dell devices.

7.2 Storage Management

Dell EqualLogic storage arrays provide a rich set of management features that are available at no additional cost and come with easy-to-use management tools. For more information on features and management capabilities, see www.dell.com/equallogic.

8 Appendix

8.1 Full Bundle Boot on Demand Basic User

The intent of this document was to outline the methodology Dell Desktop Virtualization Solutions Group went through to define, validate and characterize workloads within the Integrated Solution Stacks. The set of data and reports created from the testing on each workload is immense and is best summarized as done in section 4. As a comparative to the Pre-Booted Basic user results provided of the Basic user workload we have provided the Boot on Demand. The primary difference with regard to resource utilization was the duration of the ramp into optimal resource utilization. The boot on demand scenario is a more gradual one timed with the login rate of users, whereas the pre-boot is very sharp as all users sessions are started at once to be ready for the start of business.

With 80 desktops running on each host the following graph shows that peak CPU usage is just below the desired 80% level.

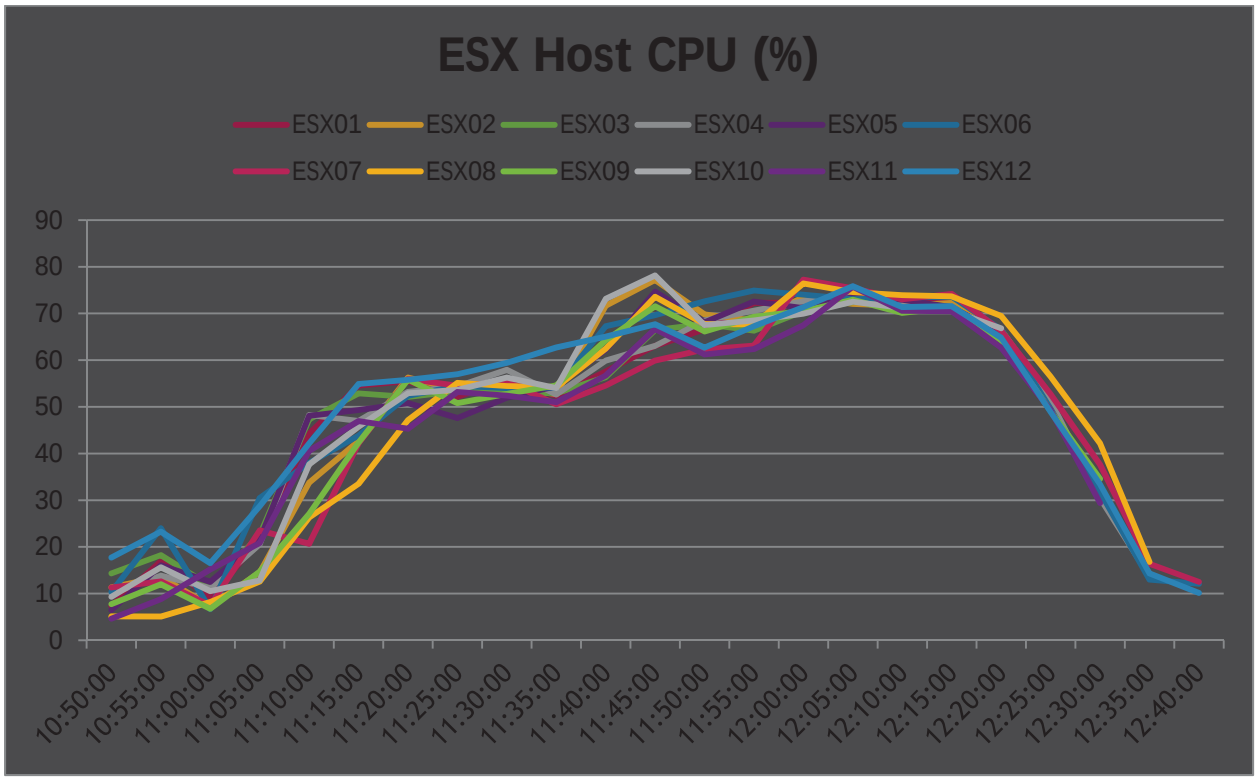

At the same stage of the test run the hosts were all consuming approximately 85GB of RAM which brings the hosts to approximately 89% of installed RAM being consumed, at this level there is no swap file in use therefore the need to swap memory to disk is having no effect on CPU or storage IOPS

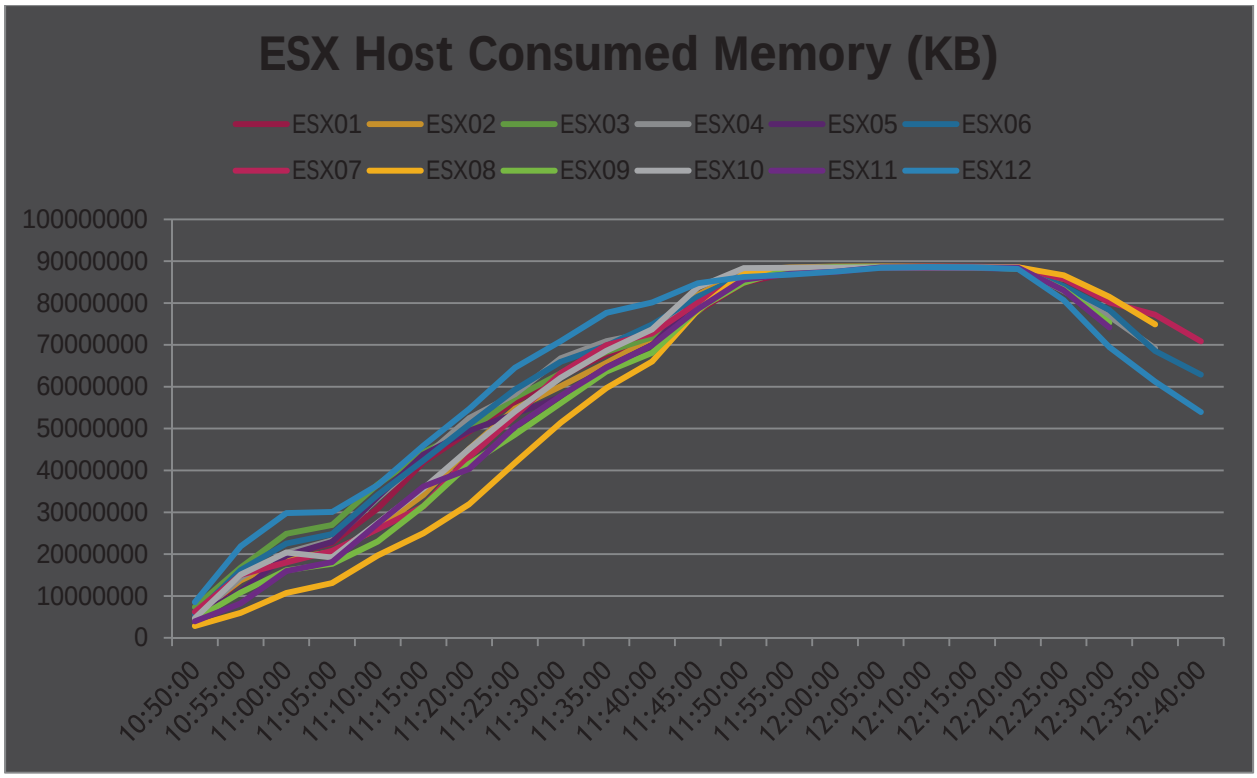

The following network usage graphs show the combined receive and transmit rates for all physical network interfaces in the desktop hosts. With a maximum receive rate of just under 30MBps (234Mbps) it is clear that no single interface in the hosts is in danger of becoming saturated. The transmit rate is always below 4MBps (31Mbps) so even with the vales combined it is clear that each host is running at less than 25% of a single 1Gbps interface (the hosts actually had 2x1Gbps and 1x10Gbps interfaces). As such network is not considered to be a limiting factor within the environment

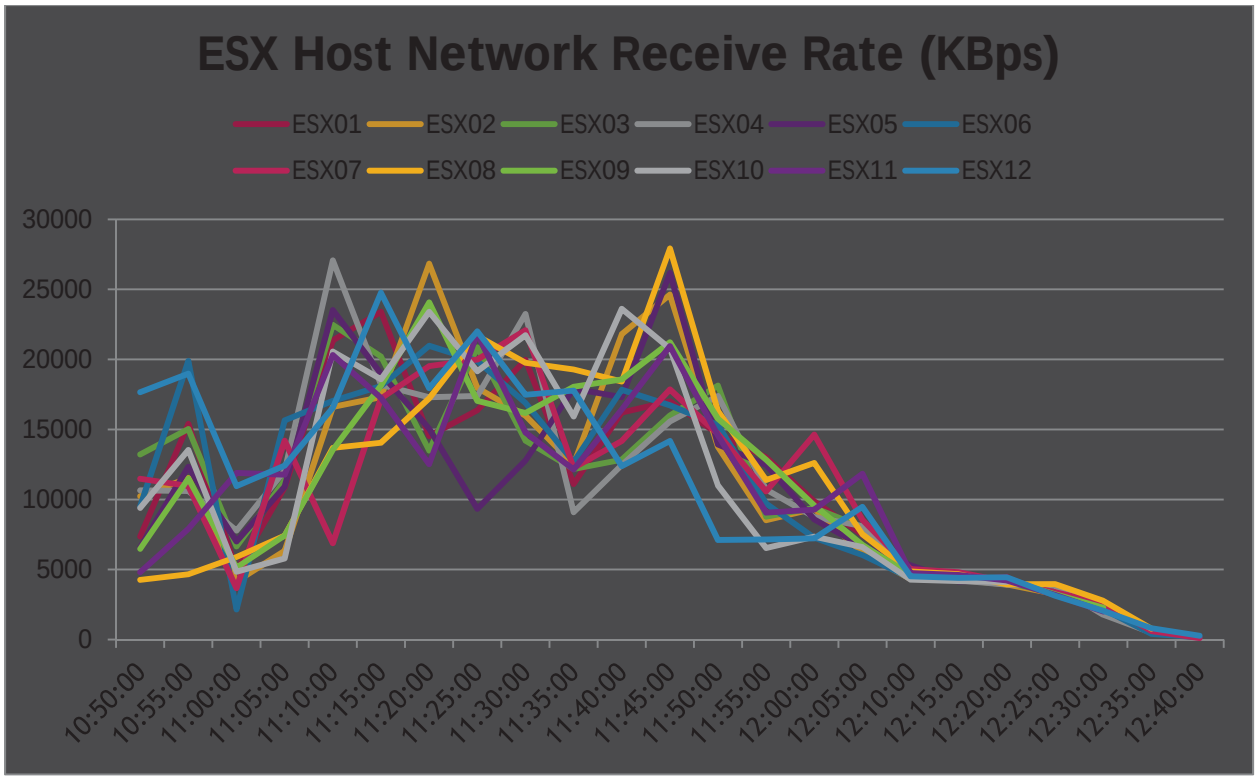
ESX Host Network Receive Rate (KBps)
ESX01 ESX02 ESX03 ESX04 ESX05 ESX06
ESX07 ESX08 ESX09 ESX10 ESX11 ESX12
30000
25000
20000
15000
10000
5000
0
10:50:00 10:55:00 11:00:00 11:05:00 11:10:00 11:15:00 11:20:00 11:25:00 11:30:00 11:35:00 11:40:00 11:45:00 11:50:00 11:55:00 12:00:00 12:05:00 12:10:00 12:15:00 12:20:00 12:25:00 12:30:00 12:35:00 12:40:00

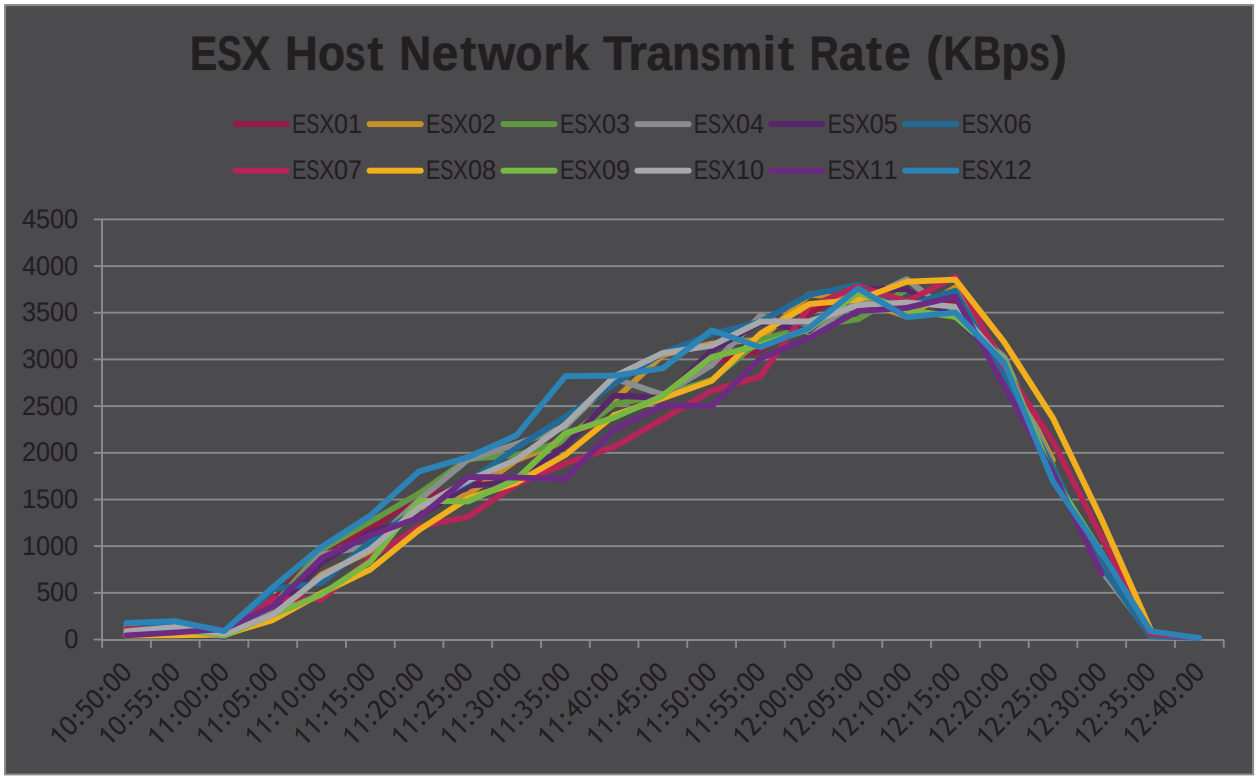
ESX Host Network Transmit Rate (KBps)
ESX01 ESX02 ESX03 ESX04 ESX05 ESX06
ESX07 ESX08 ESX09 ESX10 ESX11 ESX12
4500
4000
3500
3000
2500
2000
1500
1000
500
0
10:50:00 10:55:00 11:00:00 11:05:00 11:10:00 11:15:00 11:20:00 11:25:00 11:30:00 11:35:00 11:40:00 11:45:00 11:50:00 11:55:00 12:00:00 12:05:00 12:10:00 12:15:00 12:20:00 12:25:00 12:30:00 12:35:00 12:40:00

9 References

Citrix Documentation

- XenDesktop eDocs
  http://support.citrix.com/proddocs/topic/xendesktop/cds-library-wrapper-rho.html
- Netscaler VPX 200
  http://support.citrix.com/proddocs/topic/netscaler/ns-gen-nsvpx-wrapper-con-93.html

VMware Documentation

- ESXi Configuration Guide
  http://www.vmware.com/pdf/vsphere4/r40_u1/vsp_40_u1_esxi_server_config.pdf
- VMware Virtual Networking Concepts
  http://www.vmware.com/files/pdf/virtual_networking_concepts.pdf
- What's new in vSphere Networking
  http://www.vmware.com/files/pdf/VMW_09Q1_WP_vSphereNetworking_P8_R1.pdf

EqualLogic Documentation

- Configuring VMware vSphere software iSCSI with Dell EqualLogic PS Series Storage
  http://www.equallogic.com/resourcecenter/assetview.aspx?id=8453
- PS Series Array Network Connection and Performance Guidelines
  http://www.equallogic.com/resourcecenter/assetview.aspx?id=5229
- Configuring VMware vSphere Software iSCSI with Dell EqualLogic PS Series Storage
  http://www.equallogic.com/resourcecenter/assetview.aspx?id=8453

PowerEdge Blade Server

- Blade Server Manual http://support.dell.com/support/edocs/systems/pem/en/index.htm
- Integrating Blade Solutions with EqualLogic SANs
  http://www.dell.com/downloads/global/partnerdirect/apj/Integrating_Blades_to_EqualLogic_SAN.pdf

10 About the Authors

The Dell Desktop Virtualization Solutions Group is responsible for developing end-to-end solutions for End User Computing. This group is a collection of Architects, Engineers and Marketing resources with a great deal of domain and field experiences in the data center and desktop virtualization space.

Julian Tym is a Solution architect within the Dell Solution Centers team with 12 years' experience working in the enterprise IT sector. Julian has specialized for the last 9 years in designing, validating and deploying alternate End User Computing solutions, including Desktop Virtualization, Application Virtualization, Server Based computing and User profile management solutions. Julian's current role in the Dell Solution Center group involves solution development, solution design & validation and customer engagements (briefings, workshops, architecture design sessions and proof of concepts).

John Kelly is a Systems Engineer within the Dell Solution Center team with 10 years' experience working in the enterprise IT sector. John has successfully delivered enterprise-scale OS migration, server virtualization and desktop virtualization projects in the educational, local government and financial sectors. John's current role in the Dell Solution Center group involves overall responsibility for the implementation and strategic direction of testing and validation within the desktop virtualisation domain.

Todd Day is a Solution Architect within the Dell Product group with 14 years of enterprise solutions experience. In this role at Dell, he works as part of the Desktop Virtualization Solutions Group, developing alternative approaches to traditional desktop architectures and solutions.