Reference Architecture for SAN-Less Blades

Dell Wyse Datacenter for Citrix

Dell Wyse Solutions Engineering
March 2014

A Dell Reference Architecture

Revisions

Date	Description
March 2014	Initial release

Table of contents

1 Introduction

1.1 Purpose of this document

This document describes:

1. The Dell Wyse Datacenter Reference Architecture for Dell PowerEdge M1000E and Citrix XenDesktop using no default shared storage.

This document addresses the architecture design, configuration and implementation considerations for the key components of the architecture required to deliver virtual desktops via XenDesktop 7.1 using VMware vSphere 5.5.

1.2 Scope

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

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

2 Solution Architecture Overview

2.1 Introduction

The Dell Wyse Datacenter 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
- Endpoint Layer

These components have been integrated and tested to provide the optimal balance of high performance and lowest cost per user. Additionally, the Dell Wyse Datacenter Solution includes an approved extended list of optional components in the same categories. These components give IT departments the flexibility to custom tailor the solution for environments with unique VDI feature, scale or performance needs.

2.1.1 Physical Architecture Overview

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

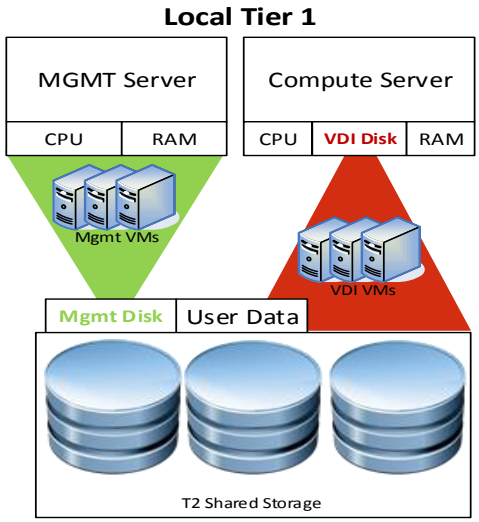

2.1.2 Dell Wyse Datacenter SAN-less – Solution Layers

The SAN-less solution uses the Dell PowerEdge M1000 including a pair of high-performance PowerEdge M I/O Aggregator (IOA) switches in the A Fabric. The IOAs handle all east-west traffic between the blades within the chassis without having to go top of rack (ToR). This integrated switch connects two 1Gb NICs to the A fabric of all blade hosts in the chassis. ToR switching is optional and best satisfied using a high performance Force10 switch.

The compute layer consists of the server resources responsible for hosting the XenDesktop user sessions hosted via VMware vSphere.

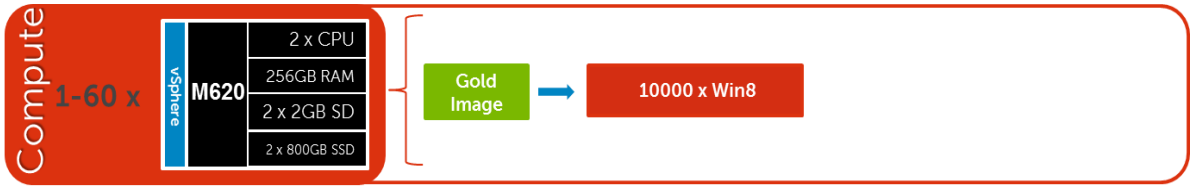

VDI management components are dedicated to their own layer so as to not negatively impact the user sessions running in the compute layer. This physical separation of resources provides clean, linear, and predictable scaling without the need to reconfigure or move resources within the solution as you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

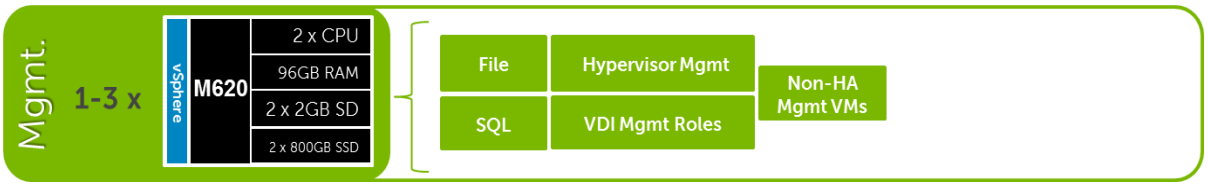

2.1.2.1 Local Tier 1 Blade – HA

In the Local Tier 1 architecture for blades, there is no need for a ToR switch unless iSCSI is present to provide mgmt. layer HA. The A Fabric IOA switches connect directly to the core or distribution network layer. Both Management and Compute servers connect to all VLANs in this model via a single vSwitch. The following diagram illustrates the server NIC to chassis switch connections, vSwitch assignments, as well as logical VLAN flow in relation to the core switch.

2.2 Local Tier 1 for Blade Servers

The SAN-less architecture shown below makes use of the Citrix Machine Creation Services (MCS) technology to provision the desktop pool from a gold master image residing locally on each compute host. Each blade is outfitted with a pair of high performance 800GB SSDs configured in RAID 1 to house either desktop or management VMs. This solution is available with as little as two blades, one mgmt and one compute, scaling to support thousands of users. Aside from the local SSDs, the blade hardware configuration is in line with the standard DWD recommendation.

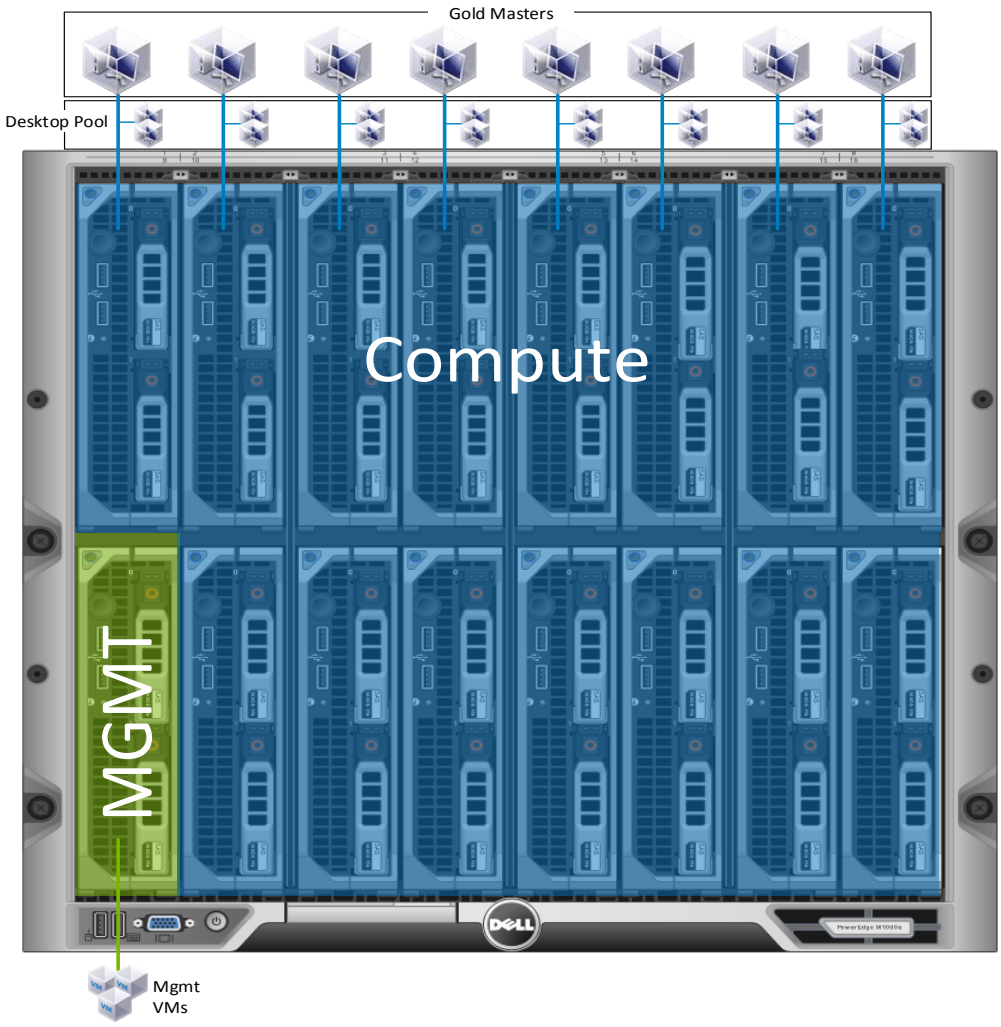

2.2.1.1 Local Tier 1 – Mgmt layer HA

The Local Tier 1 solution model for blade servers provides a high-performance configuration that does not require shared storage but Tier 2 is added to provide HA to the management layer infrastructure. User VDI sessions are hosted locally on SSDs in each blade server using Citrix MCS for desktop delivery. Only a pair of PowerEdge M I/O Aggregator switches is required in the A Fabric which supports all East—West communication between VMs within the chassis without going ToR. The B Fabric, ToR iSCSI switches and Tier 2 storage are optional.

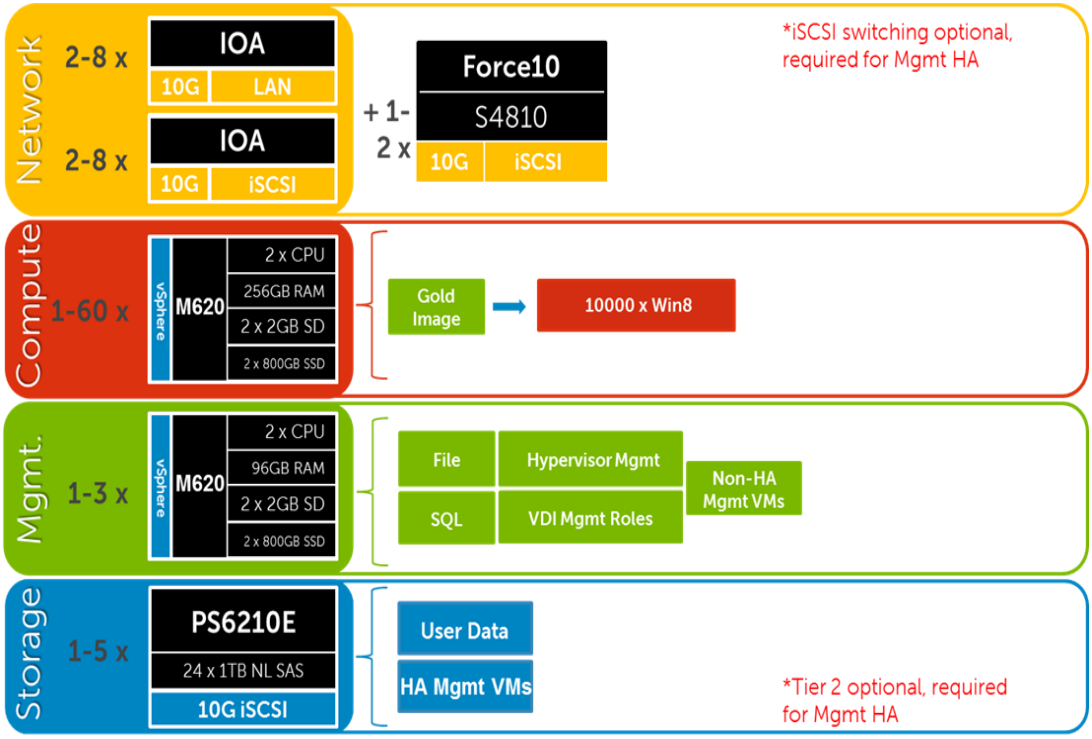

2.2.1.2 Local Tier 1 – Network Architecture

In the Local Tier 1 architecture for blades, there is no need for a ToR switch unless iSCSI is present for HA. The A Fabric IOA switches can connect directly to the core or distribution network layer. Both Management and Compute servers connect to all VLANs in this model via a single vSwitch. The following diagram illustrates the server NIC to chassis switch connections, vSwitch assignments, as well as logical VLAN flow in relation to the core switch.

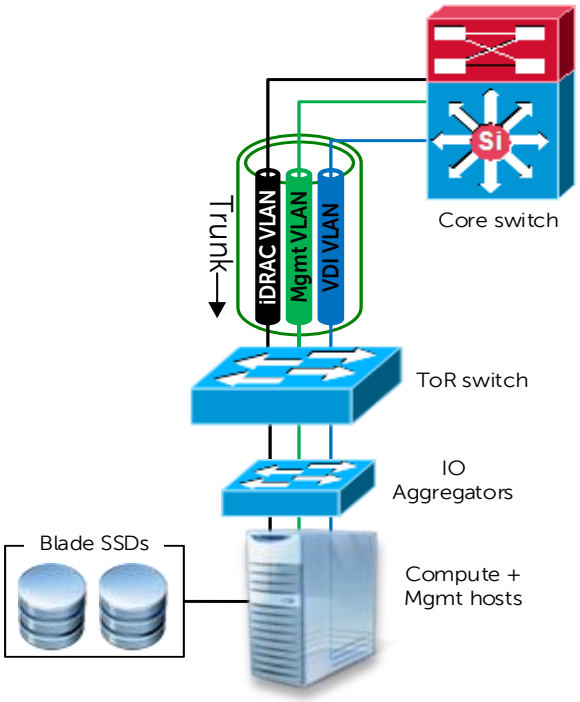

2.2.1.3 Local Tier 1 – Cabling (Blades – HA)

The following diagram depicts the LT1 blade solution including optional components for HA. The A Fabric, B Fabric and ToR switches are stacked, respectively.

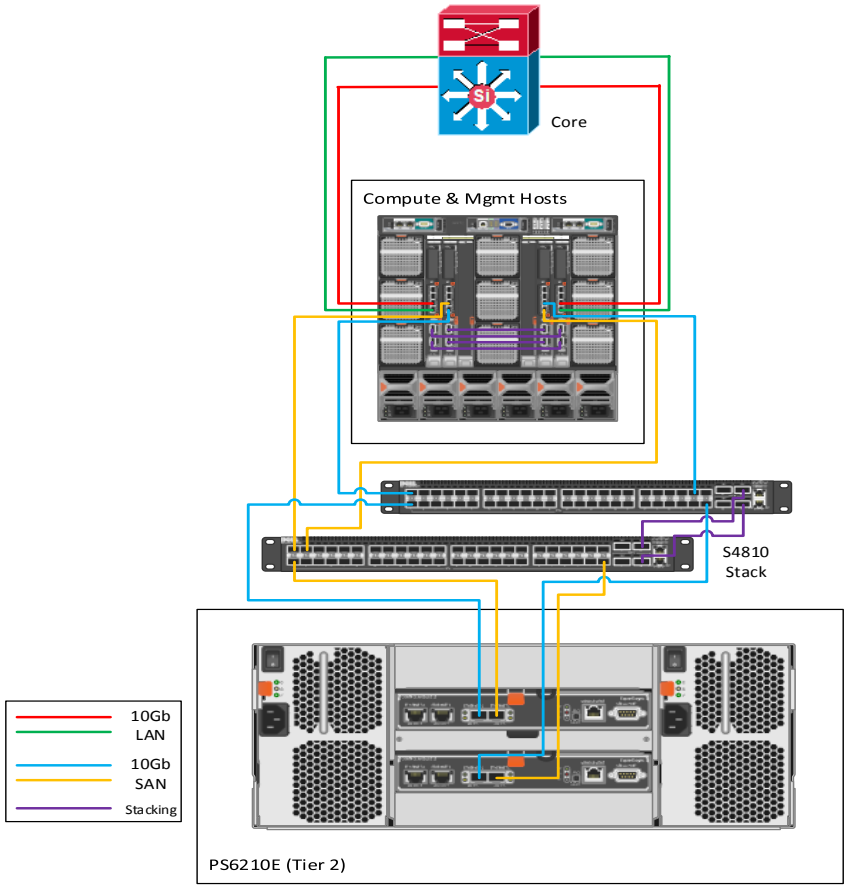

2.2.1.4 Local Tier 1 – Blade Scaling Guidance

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

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

3 Hardware Components

3.1 Network

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

Model	Features	Options	Uses
PowerEdge M I/O Aggregator (IOA)	Up to 32 x 10Gb ports + 4 x external SFP+ 2 x line rate fixed QSFP+ ports 2 optional FlexIO modules	2-port QSFP+ module in 4x10Gb mode	Blade switch for iSCSI in Shared Tier 1 blade solution, LAN + iSCSI in Local Tier 1 blade solution.
		4-port SFP+ 10Gb module	
		4-port 10GBASE-T copper module (one per IOA)	
		Stack up to 2 IOAs using QSFP+ ports	

Guidance:

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

For more information on the Dell IOA switch, please visit: LINK

3.2 Servers

3.2.1 PowerEdge M1000E

The Dell PowerEdge M-Series blade servers address the challenges of an evolving IT environment by delivering leading enterprise-class features and functionality. The M-Series delivers a unique array of options configured to meet the needs of your IT environment today and in the future. For more information please visit: Link

3.2.2 Server Configuration

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

3.1 Dell Wyse Cloud Clients

The following Dell Wyse end/ zero clients are the recommended choices for this solution.

3.1.1 ThinOS – T10D

T10D sets the standard for thin clients. Providing an exceptional user experience, the T10D features the incredibly fast Dell Wyse ThinOS, for environments in which security is critical—there's no attack surface to put your data at risk. Boot up in just seconds and log in securely to almost any network. The T10D delivers a superior Citrix VDI user experience, along with usability and management features found in premium thin clients. The T10D delivers outstanding performance based on its dual core system-on-a-chip (SoC) design, and a built-in media processor delivers smooth multimedia, bi-directional audio and Flash playback. Flexible mounting options let you position the T10D vertically or horizontally on your desk, on the wall or behind your display. Using about 7-watts of power in full operation, the T10D creates very little heat for a greener, more comfortable working environment. Link

3.1.2 ThinOS – D10D

Designed for knowledge workers and power users, the new Dell Wyse D10D is a high-performance thin client based on Dell Wyse ThinOS, the virus-immune firmware base designed for optimal thin client security, performance, and ease-of-use. Highly secure, compact and powerful, the D10D combines Dell Wyse ThinOS with a dual-core AMD 1.4 GHz processor and a revolutionary unified graphics engine for an outstanding user experience. The D10D addresses the performance challenges of processing-intensive applications like computer-aided design, multimedia, HD video and 3D modeling. Scalable enterprise-wide on-premise or cloud-based management provides simple deployment, patching and updates. Take a unit from box to productivity in minutes with auto configuration. Delivering outstanding processing speed and power, security and display performance, the D10D offers a unique combination of performance, efficiency, and affordability. For more information, please visit: Link

3.1.3 Windows Embedded 7 – Z90Q7

The Dell Wyse Z90Q7 is a super high-performance Windows Embedded Standard 7 thin client for virtual desktop environments. Featuring a quad-core AMD processor and an integrated graphics engine that significantly boost performance, the Z90Q7 achieves exceptional speed and power for the most demanding VDI and embedded Windows applications, rotational 3D graphics, 3D simulation and modeling, unified communications, and multi-screen HD multimedia. Take a unit from box to productivity in minutes. Just select the desired configuration and the Z90Q7 does the rest automatically—no need to reboot. Scale to tens of thousands of endpoints with Dell Wyse WDM software or leverage your existing Microsoft System Center Configuration Manager platform. The Z90Q7 is the thin

client for power users who need workstation-class performance on their desktop or within a desktop virtualization environment (x86 or x64). For more information, please visit: Link

3.1.4 Windows Embedded 8 – Z90Q8

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

3.1.5 Suse Linux – Z50D

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

3.1.6 Dell Wyse Zero – Xenith 2

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

3.1.7 Dell Wyse Zero – Xenith Pro 2

Dell Wyse Xenith Pro 2 is the next-generation zero client for Citrix HDX and Citrix XenDesktop, delivering ultimate performance, security and simplicity. With a powerful dual core AMD G-series CPU; Xenith Pro 2 is faster than competing devices. This additional computing horsepower allows dazzling HD multimedia delivery without overtaxing your server or network. Scalable enterprise-wide management provides simple deployment, patching and updates—your Citrix XenDesktop server configures it out-of-the-box to your preferences for plug-and-play speed and ease of use. Virus and malware immune, the

Xenith Pro 2 draws under 9 watts of power in full operation—that's less than any PC on the planet. For more information please visit: Link

3.1.8 Dell Wyse Cloud Connect

Designed to promote bring-your-own-device (BYOD) environments, Dell Wyse Cloud Connect allows you to securely access and share work and personal files, presentations, applications and other content from your business or your home. Managed through Dell Wyse Cloud Client Manager software-as-a-service (SaaS), IT managers can ensure that each Cloud Connect device is used by the appropriate person with the right permissions and access to the appropriate apps and content based on role, department and location. Slightly larger than a USB memory stick, Cloud Connect is an ultra-compact multimedia-capable device. Simply plug it into any available Mobile High-Definition Link (MHL) / HDMI port on a TV or monitor, attach a Bluetooth keyboard and mouse, and you're off and running. Easy to slip into your pocket or bag, it enables an HD-quality window to the cloud, great for meetings and presentations while on business travel, or for cruising the internet and checking email after a day of work. For more information, please visit: Link

3.1.9 Dell Venue 11 Pro

Meet the ultimate in productivity, connectivity and collaboration. Enjoy full laptop performance in an ultra-portable tablet that has unmatched flexibility for a business in motion. This dual purpose device works as a tablet when you're out in the field but also enables you to work on your desktop in the office thanks to an optional dock. For more information, please visit: Link

3.1.10 Dell Chromebook 11

The lightweight, easy-to-use Dell Chromebook 11 helps turn education into exploration - without the worries of safety or security. Priced to make 1:1 computing affordable today, Chromebook 11 is backed by Dell support services to make the most of your budget for years to come. The Chrome OS and Chrome browser get students online in an instant and loads web pages in seconds. A high-density battery supported by a 4th Gen Intel® processor provides up to 10 hours of power. Encourage creativity with the Chromebook 11 and its multimedia features that include an 11.6" screen, stereo sound and webcam. For more information, please visit: Link

4 Software Components

4.1 Citrix XenDesktop

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

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

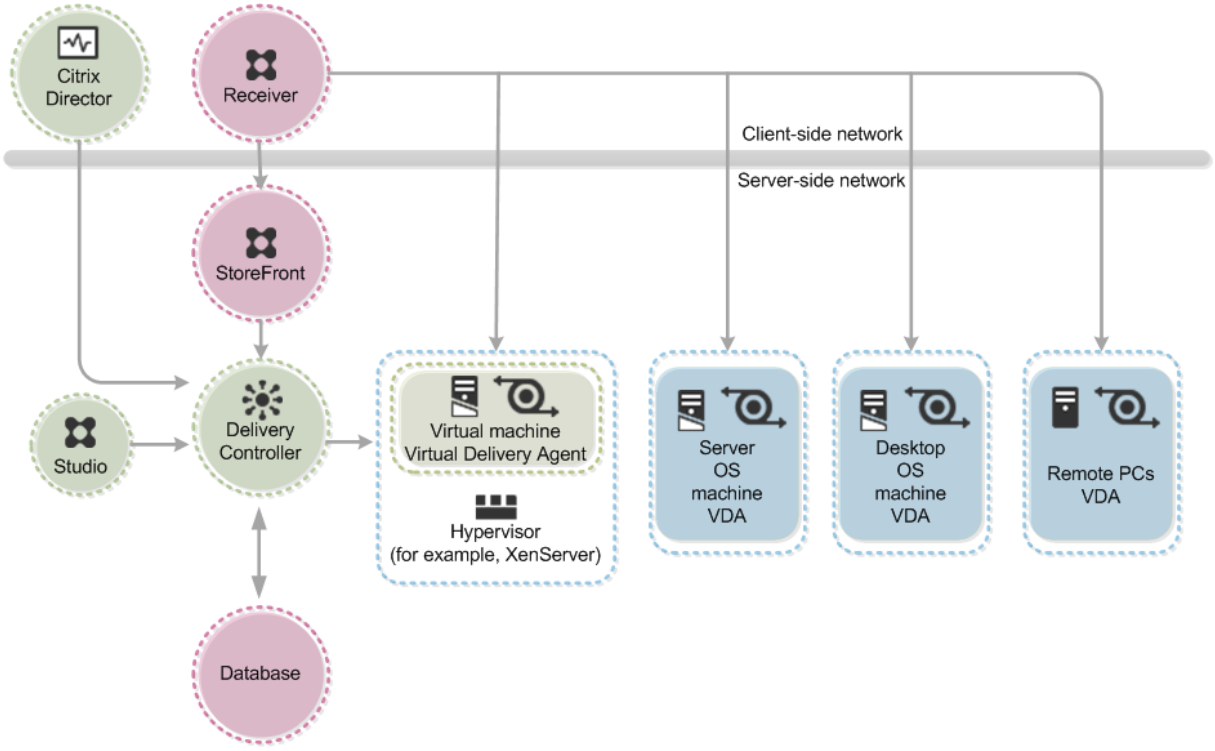

The core XenDesktop components include:

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

- **Director**
  - Director is a web-based tool that enables IT support teams to monitor an environment, troubleshoot issues before they become system-critical, and perform support tasks for end users. You can also view and interact with a user's sessions using Microsoft Remote Assistance.
- **Receiver**
  - Installed on user devices, Citrix Receiver provides users with quick, secure, self-service access to documents, applications, and desktops from any of the user's devices including smartphones, tablets, and PCs. Receiver provides on-demand access to Windows, Web, and Software as a Service (SaaS) applications.
- **Delivery Controller** (DC)
  - Installed on servers in the data center, the controller authenticates users, manages the assembly of users' virtual desktop environments, and brokers connections between users and their virtual desktops.
- **StoreFront**
  - StoreFront authenticates users to sites hosting resources and manages stores of desktops and applications that user's access.
- **License Server**
  - The Citrix License Server is an essential component at any Citrix-based solution. Every Citrix product environment must have at least one shared or dedicated license server. License servers are computers that are either partly or completely dedicated to storing and managing licenses. Citrix products request licenses from a license server when users attempt to connect.
- **Machine Creation Services** (MCS)
  - A collection of services that work together to create virtual servers and desktops from a master image on demand, optimizing storage utilization and providing a pristine virtual machine to users every time they log on. Machine Creation Services is fully integrated and administrated in Citrix Studio.
- **Virtual Delivery Agent** (VDA)
  - The Virtual Desktop Agent is a transparent plugin that is installed on every virtual desktop or XenApp host (RDSH) and enables the direct connection between the virtual desktop and users' endpoint devices.

4.1.1 Machine Creation Services (MCS)

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

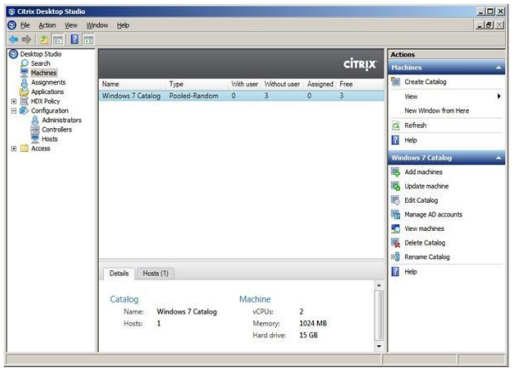

- **Random**: Virtual desktops are assigned randomly as users connect. When they logoff, the desktop is reset to its original state and made free for another user to login and use. Any changes made by the user are discarded at log off.
- **Static**: Virtual desktops are assigned to the same user every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

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

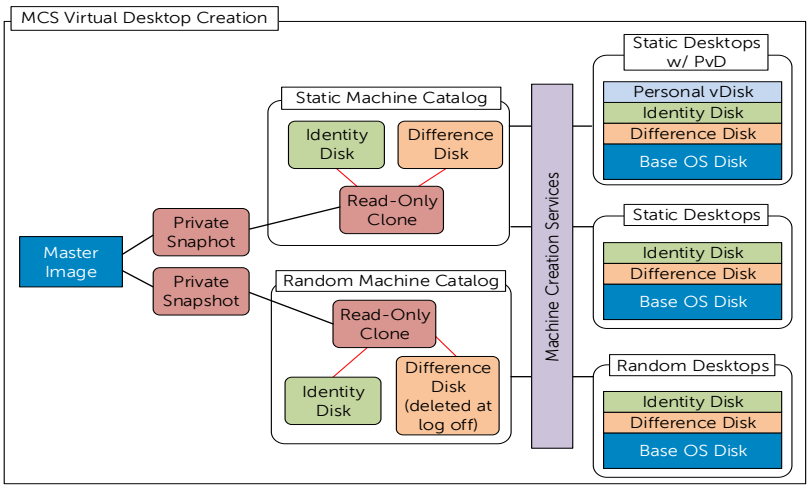

While traditionally used for small to medium sized XenDesktop deployments, MCS can bring along with it some substantial Tier 1 storage cost savings because of the snapshot/identity/difference disk methodology. The Tier 1 disk space requirements of the identity and difference disks when layered on top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

4.1.2 Citrix Personal vDisk Technology

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

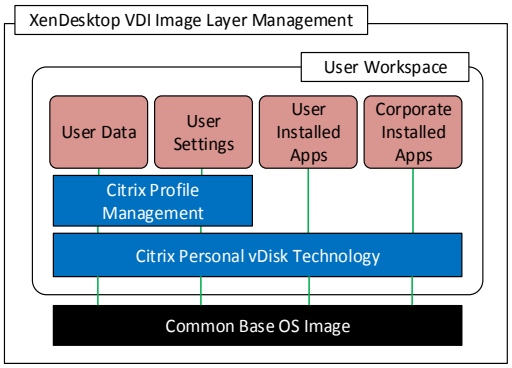

With Citrix Personal vDisk, each user receives personal storage in the form of a layered vDisk, which enables them to personalize and "persist" their desktop environment.

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

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

4.1.3 Citrix Profile Manager

Citrix Profile Management is a component of the XenDesktop suite which is used to manage user profiles and minimize many of the issues associated with traditional Windows Roaming profiles in an environment where users may have their user profile open on multiple devices at the same time. The profile management toolset has two components, the profile management agent which is installed on any device where the user profiles will be managed by the toolset, which will be the virtual desktops. The second component is a Group Policy Administrative Template, which is imported to a group policy which is assigned to an organizational unit within active directory which contains the devices upon which the user profiles will be managed.

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

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

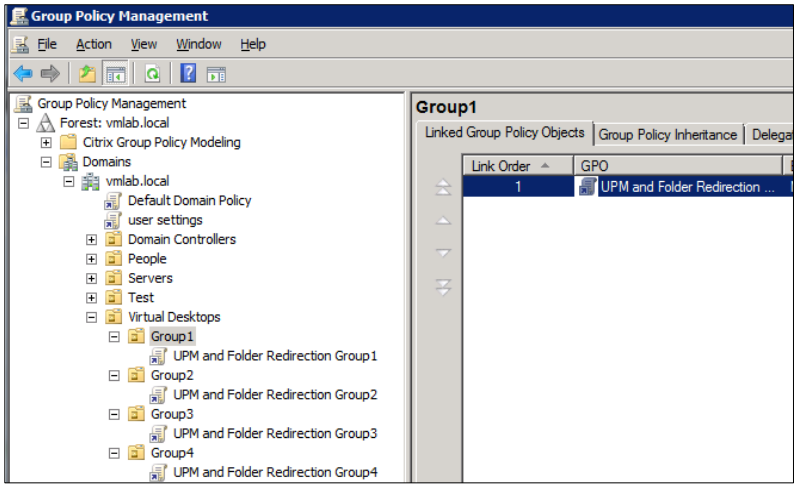

4.2 VDI Hypervisor Platforms

4.2.1 VMware vSphere 5.5

VMware vSphere 5.5 is the enterprise virtualization platform used for building VDI and cloud infrastructures. VMware vSphere 5.5 includes three major layers: Virtualization, Management and Interface. The Virtualization layer includes infrastructure and application services. The Management layer is central for configuring, provisioning and managing virtualized environments. The Interface layer includes the vSphere client and the vSphere web client.

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

5 Solution Architecture for XenDesktop 7

5.1 Compute Server Architecture

5.1.1 Local Tier 1 Blade

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

Local Tier 1 Compute Host – PowerEdge M620
2 x Intel Xeon E5-2690v2 Processor (3GHz)
256GB Memory (16 x 16GB DIMMs @ 1600Mhz)
VMware **vSphere** on 2 x 2GB internal SD
PERC H310 RAID Controller – RAID 1
2 x 800GB Value MLC SSD
Broadcom 57810-k 10Gb DP KR NDC (LAN)
iDRAC7 Enterprise w/ vFlash, 8GB SD

Local Tier 1 Management Host – PowerEdge M620
2 x Intel Xeon E5-2670v2 Processor (2.5GHz)
96GB Memory (6 x 16GB DIMMs @ 1600Mhz)
VMware **vSphere** on 2 x 2GB internal SD
PERC H310 RAID Controller – RAID 1
2 x 800GB Value MLC SSD
Broadcom 57810-k 10Gb DP KR NDC (**iSCSI-HA**)
iDRAC7 Enterprise w/ vFlash, 8GB SD

5.2 Management Server Infrastructure

The Management role requirements for the base solution are summarized below. Use data disks for role-specific application files such as data, logs and IIS web files in the Management volume. Present Tier 2 volumes with a special purpose (called out above) in the format specified below:

Role	vCPU	RAM (GB)	NIC	OS + Data vDisk (GB)	Tier 2 Volume (GB)
DDC	4	8	1	40 + 5	-
Storefront	1	2	1	40 + 5	-
Citrix License Server	1	2	1	40 + 5	-
vCenter	2	8	1	40 + 5	50 (VMDK/ VHDX)
SQL Server	2	8	1	40 + 5	200 (VMDK/ VHDX)
TOTALS	10	28	5	225	250

5.2.1 SQL Databases

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

- Citrix XenDesktop and
- vCenter

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

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

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

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

5.2.2 DNS

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

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

5.2.2.1 DNS for SQL

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

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

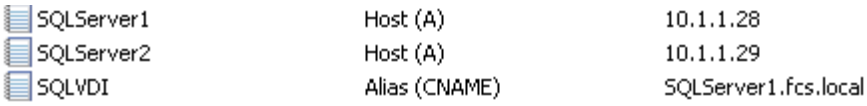

5.3 Virtual Networking

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

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

5.3.1 vSphere – Non HA

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

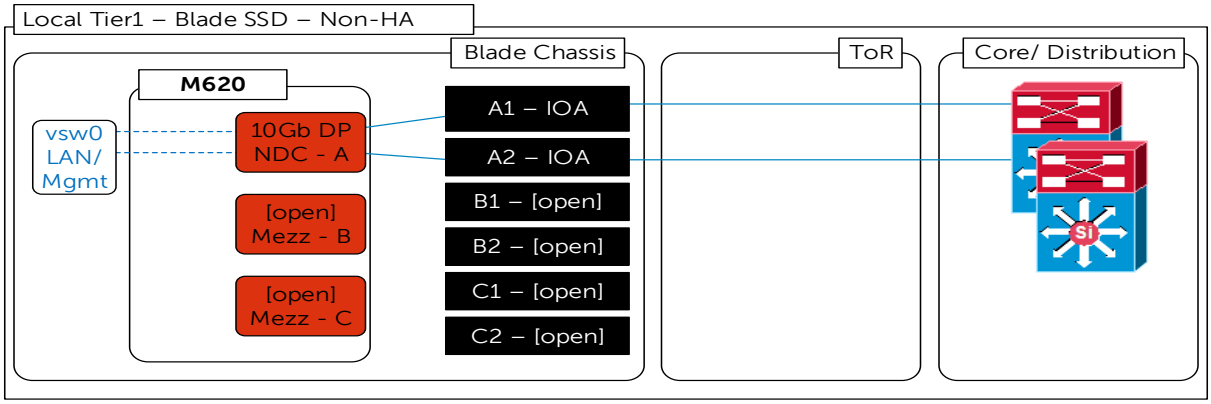

5.3.2 vSphere – HA

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

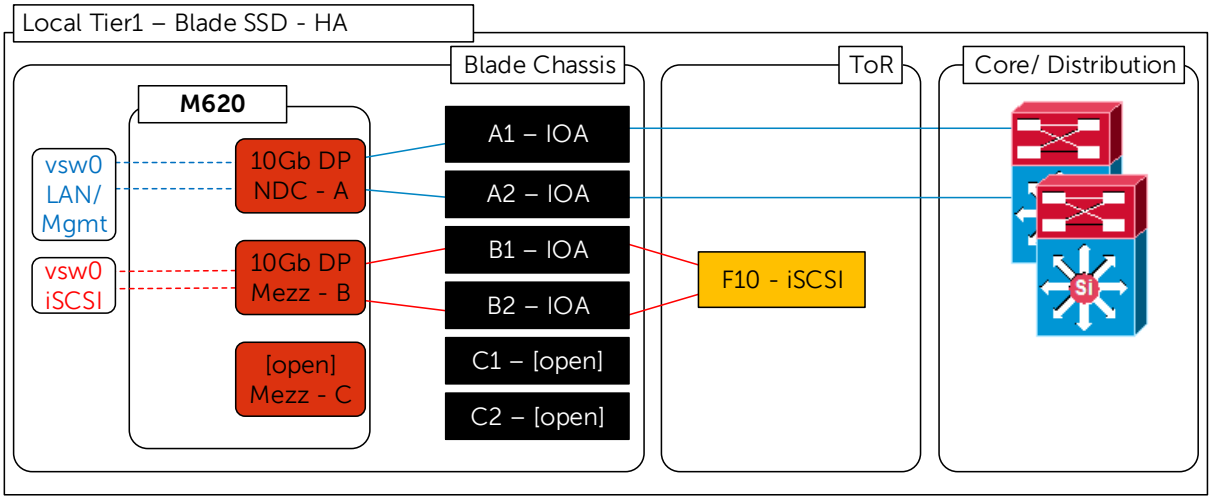

5.4 XenDesktop Communication Flow

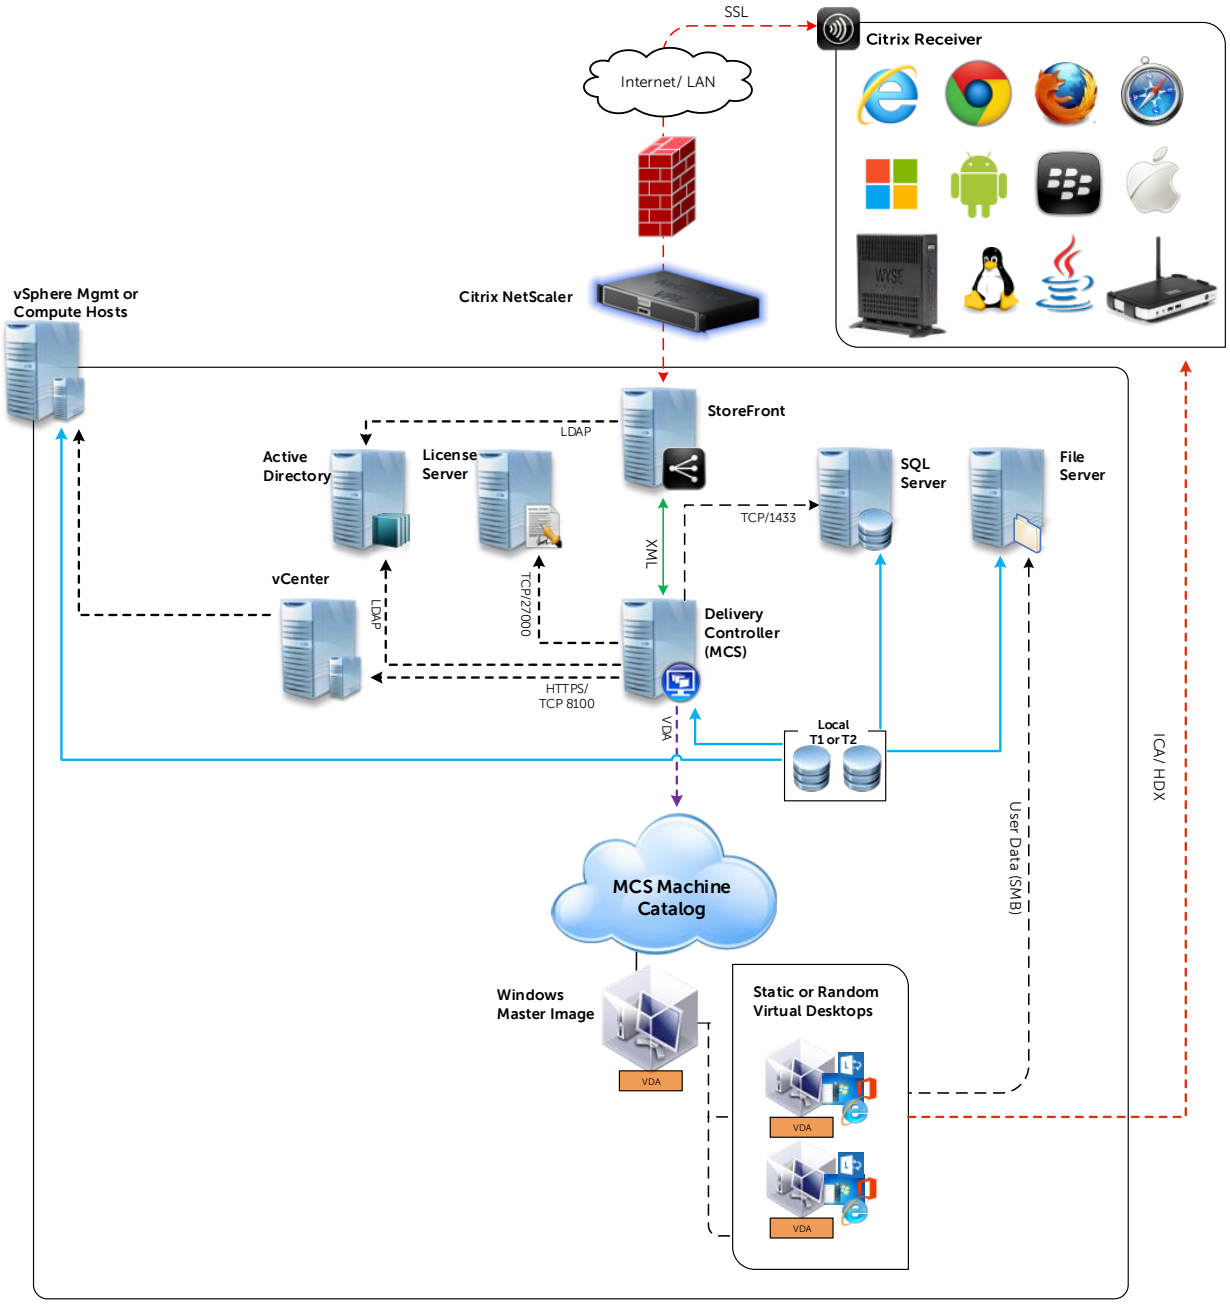

6 Solution Performance and Testing

6.1 Load Generation and Monitoring

6.1.1 Login VSI 4 – Login Consultants

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

6.1.2 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:

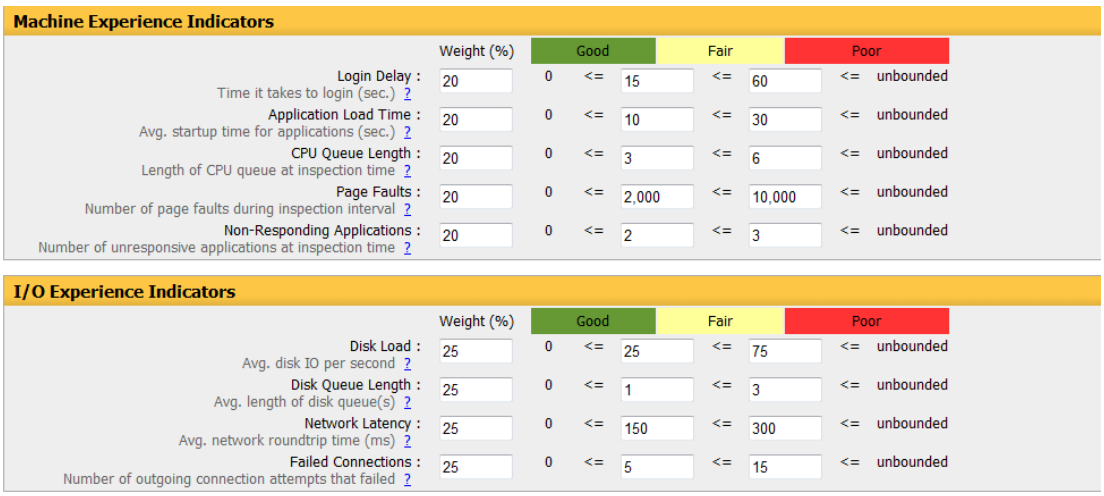

Machine Experience Indicators

	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

6.1.3 VMware vCenter

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

6.2 Performance Analysis Methodology

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

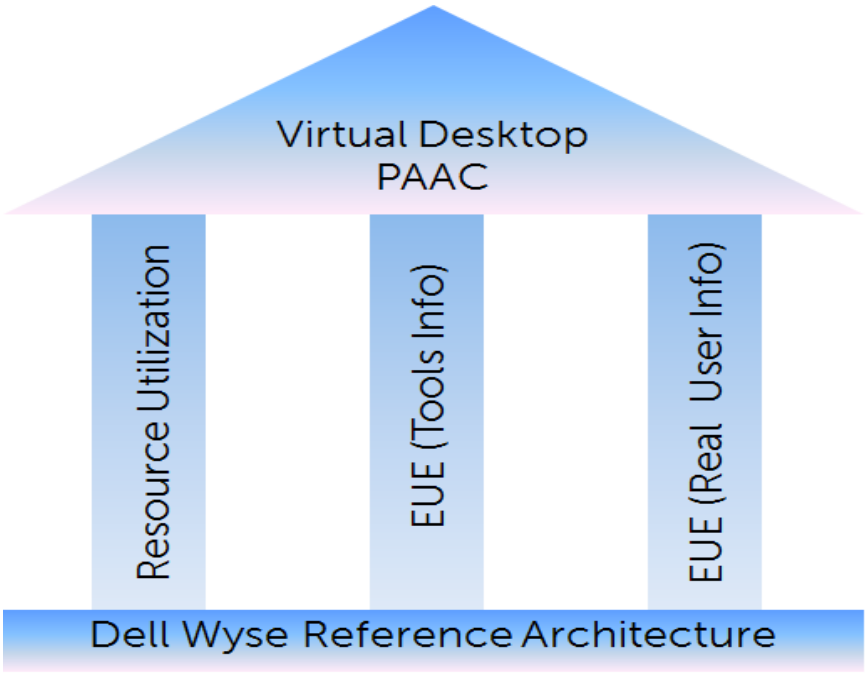

***Each of the pillars shown in the above diagram is discussed in more detail below.**

6.2.1 Resource Utilization

Poor end-user experience is one of the main risk factors when implementing desktop virtualization but the root cause for poor end-user experience is resource contention – hardware resources at some point in the solution have been exhausted, thus causing the poor end-user experience. In order to ensure that this has not happened (and that it is not close to happening), PAAC on Dell Wyse Datacenter solutions monitors the relevant resource utilization parameters and applies relatively conservative thresholds as shown in the table below. As discussed above, these thresholds are carefully selected to deliver an optimal combination of good end-user experience and cost-per-user,

while also providing burst capacity for seasonal / intermittent spikes in usage. These thresholds are used to decide the number of virtual desktops (density) that are hosted by a specific hardware environment (i.e. combination of server, storage and networking) that forms the basis for a Dell Wyse Datacenter RA.

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

6.2.2 EUE (Tools Info)

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

6.2.3 EUE (Real User Info)

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

6.2.4 Dell Wyse Datacenter Workloads and Profiles

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

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

6.2.4.1 Dell Wyse Datacenter Profiles

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

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

6.2.4.2 Dell Wyse Datacenter Workloads

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

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

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

- For single-server / single-host testing (typically carried out to determine the virtual desktop capacity of a specific physical server), users are logged in every 30 seconds.
- For multi-host / full solution testing, users are logged in over a period of 1-hour, to replicate the normal login storm in an enterprise environment.
- All desktops are pre-booted in advance of logins commencing.

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

6.2.5 Login VSI 3 versus Login VSI 4

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

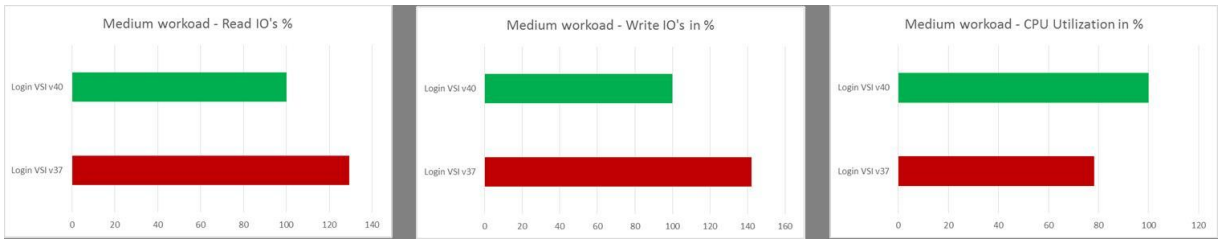

6.2.5.1 Login VSI Light Workload

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

6.2.5.2 Login VSI Medium Workload

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

- Once a session has been started the workload will repeat (loop) every 48 minutes.

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

Each loop opens and uses:

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

6.2.5.3 Login VSI Heavy Workload

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

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

6.3 Testing and Validation

6.3.1 Testing Process

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

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

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

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

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

6.3.2 Local Tier 1 for Blades Test Results

6.3.2.1 Configuration

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

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

In the solution validated the management host was a separate M620 with 2 * 15K 300GB disks in a RAID 1 configuration connected by a PERC H310. If the desire exists to keep all blades in the same configuration, the management functions could be hosted by any blade. Each node is identically configured in the same way as the compute node above. Based on the latencies observed during the Login VSI tests on the compute nodes the solution should scale well to accommodate up to 16 blades. In addition the Virtual Center appliance was deployed with sufficient capacity for up to 20 hosts and 2000 VMs.

The M1000E blade chassis was fitted with a Dell MIO aggregator. This supported the external communication from the chassis for VDI and Active directory access, and provided support for all the internal communication between the various VMs hosted on the blades in the chassis (East—West without going ToR). A simple management configuration was implemented using 3 management VMs only, a vSphere Virtual Center appliance, Citrix XenDesktop DDC, and Citrix Storefront server and license server combined. SQL Server 2008 Express was installed on the DDC to support XenDesktop

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

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

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

6.3.2.2 Test Results

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

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

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

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

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

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

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

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

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

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

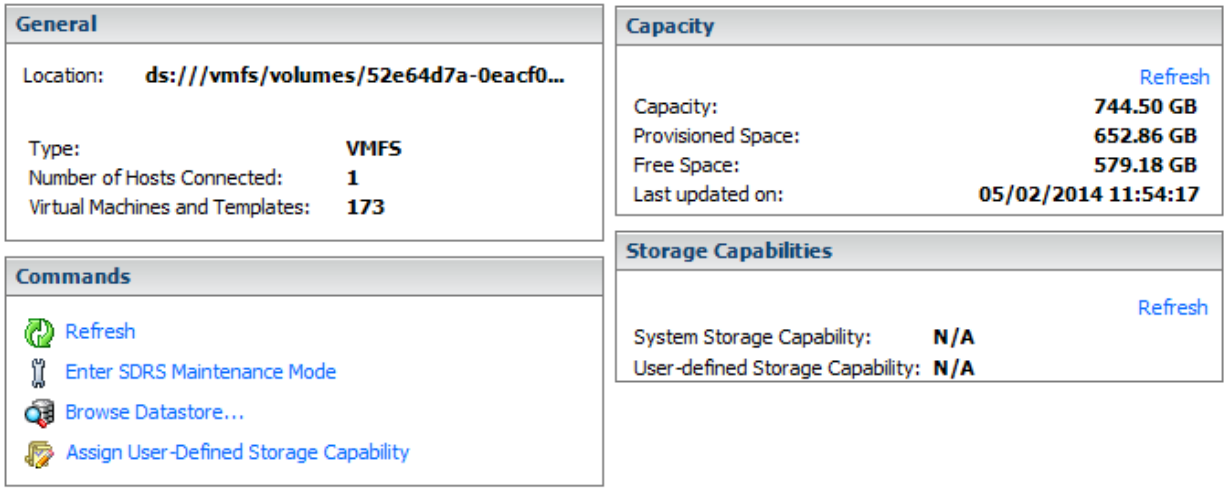

When all the VMs are powered on the swap files reduce the free space to 224GB as follows:

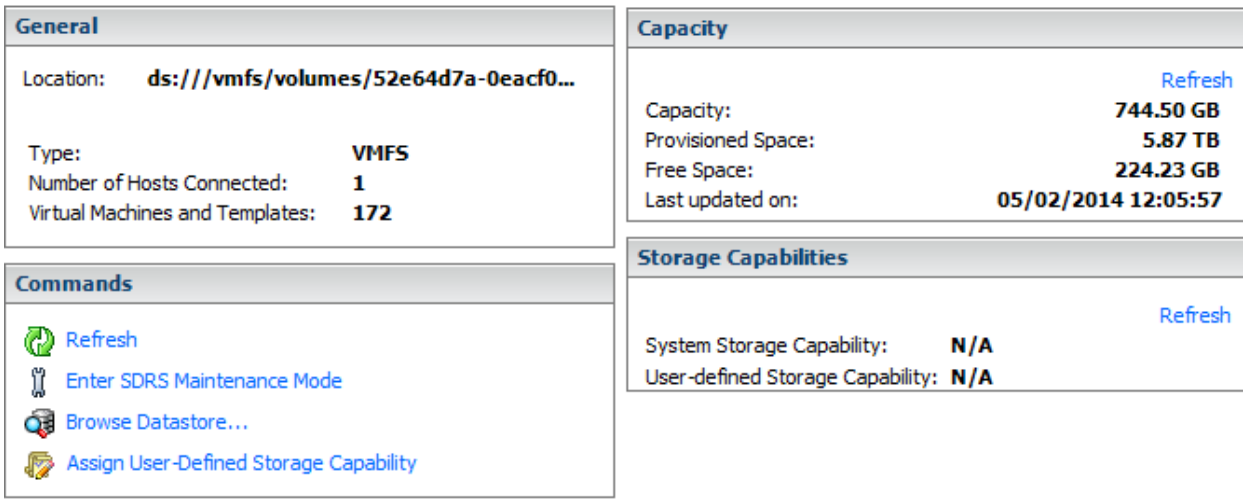

The SQL Express data file usage was 112MB. Adding a further 15 hosts should not increase this to more than 1.75GB. With the allocation used in the tests there was 13.7GB free on the DDC so there is room for expansion of the database as further hosts and VMs are added. Space usage on the Management Datastore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

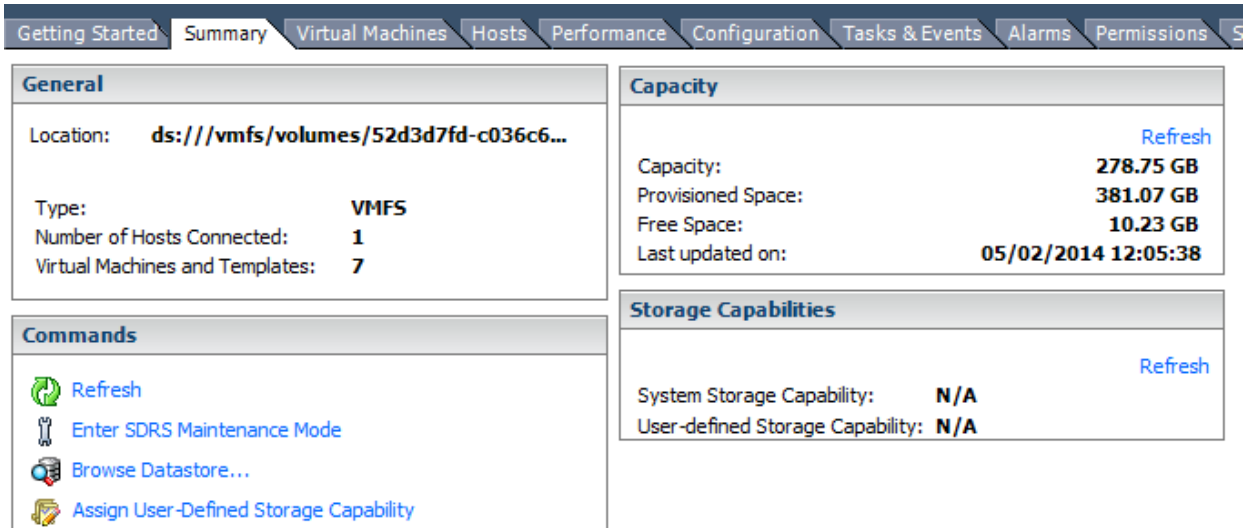

Smoke further tests were also performed to investigate the capability of the local SSD volumes within the VDI solutions:-

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

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

6.3.2.3 VMware vSphere 5.5 Test Charts

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

6.3.2.4 MCS Standard User Workload (170 Users)

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

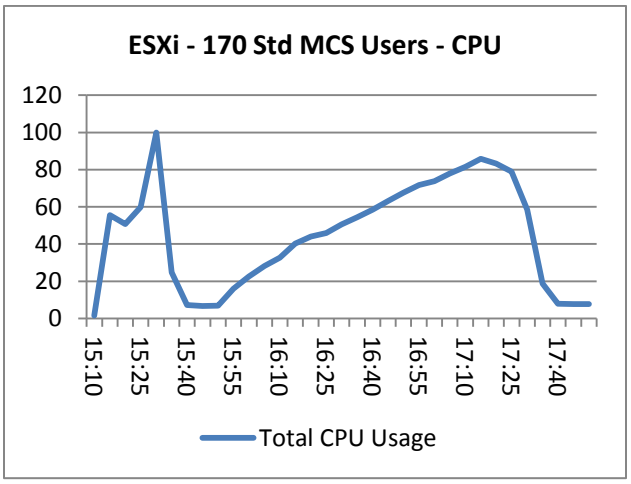

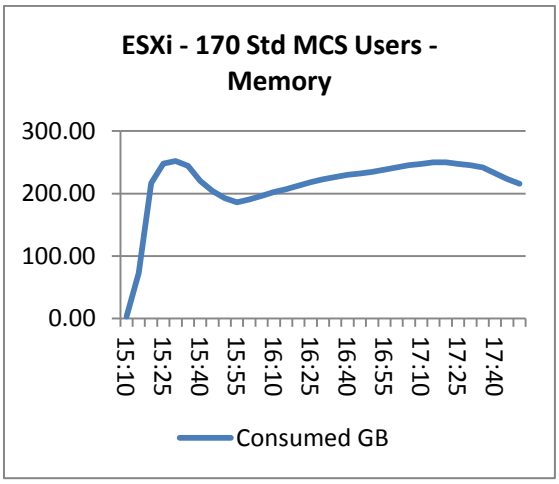

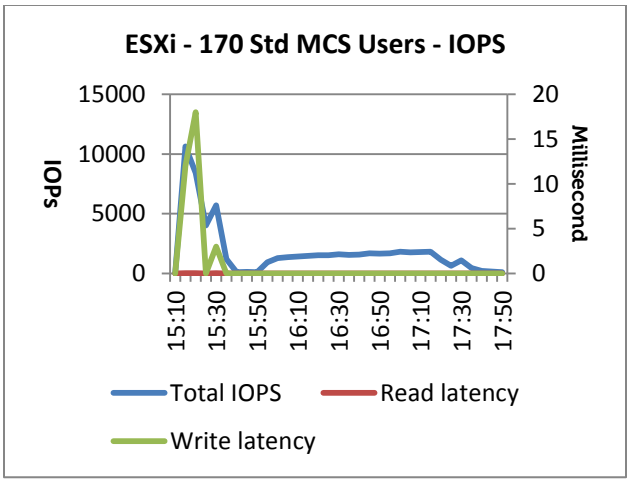

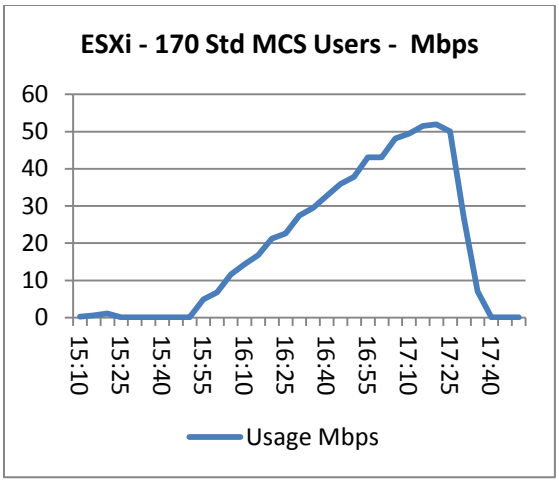

6.3.2.5 MCS Enhanced User Workload (110 users)

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

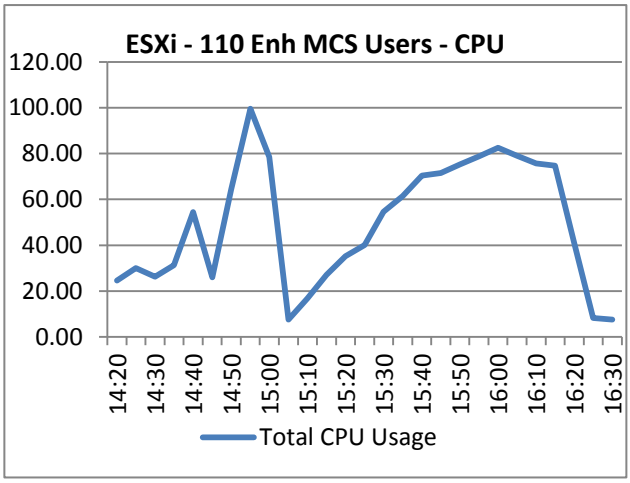

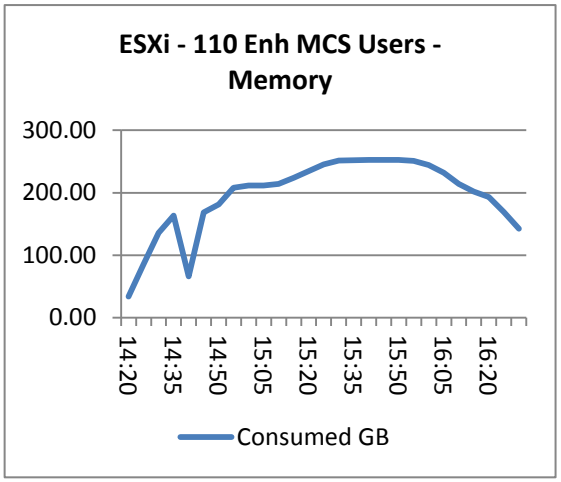

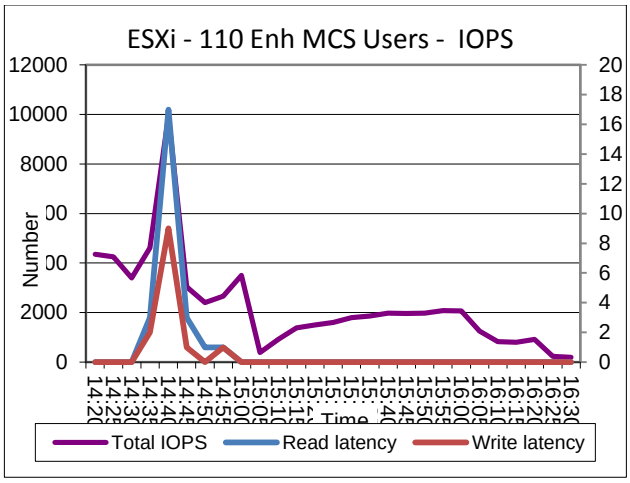

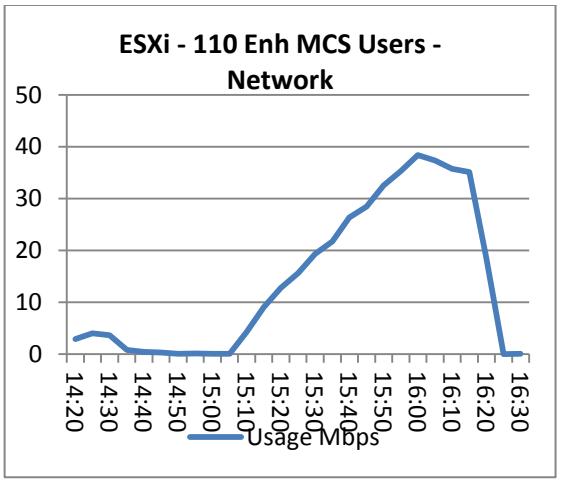

6.3.2.6 MCS Professional User Workload (90 users)

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

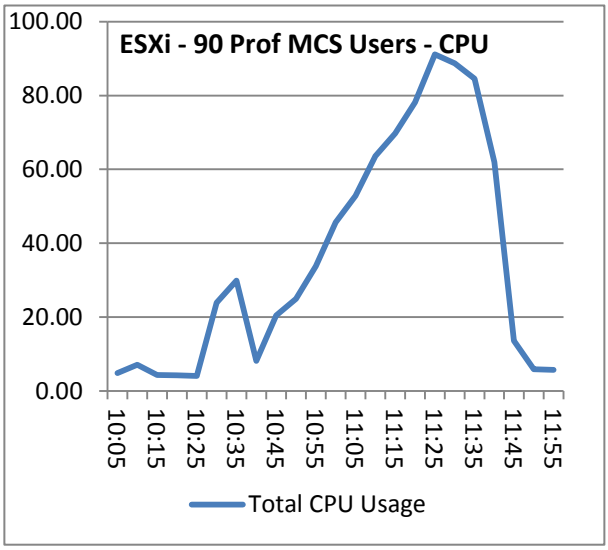

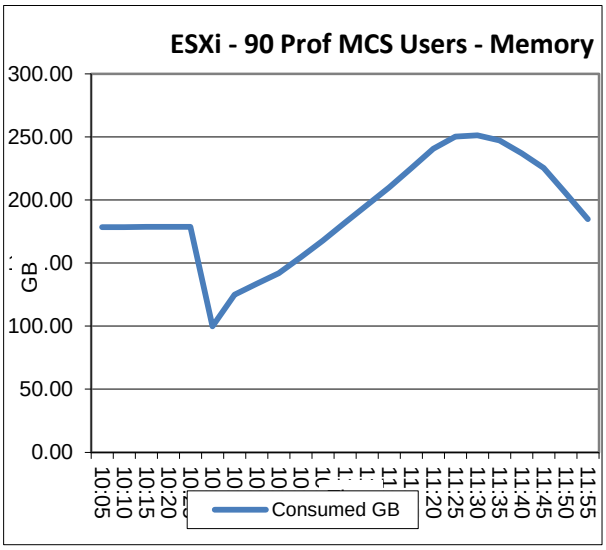

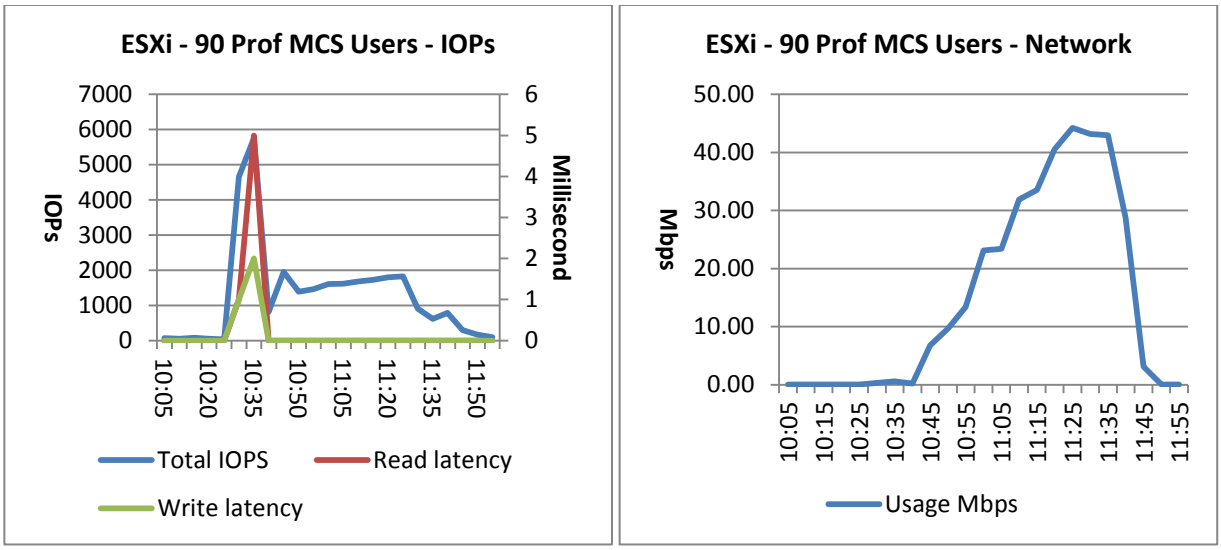

6.3.2.7 MCS Professional User Workload (90 users) with 0% preboot

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

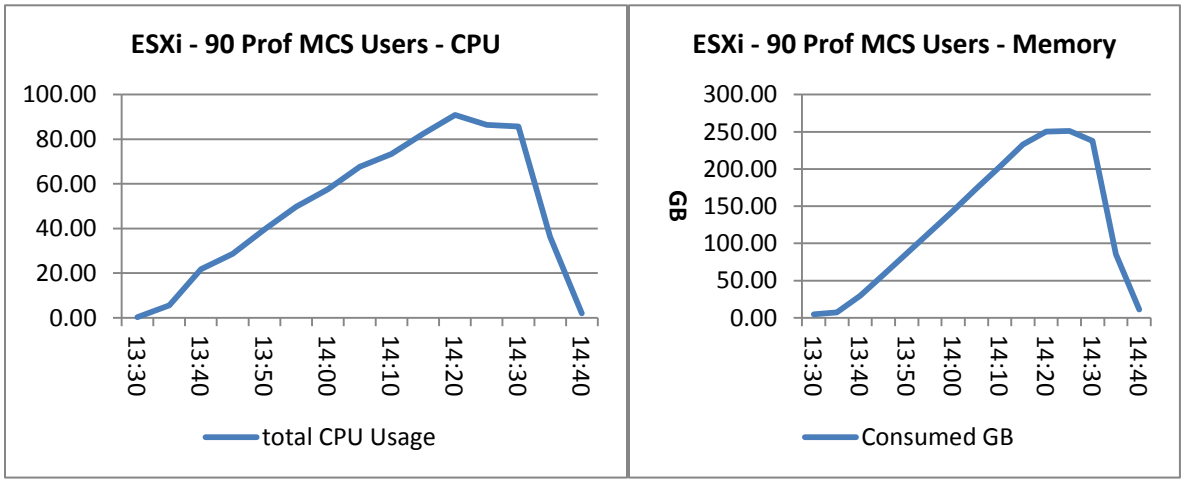

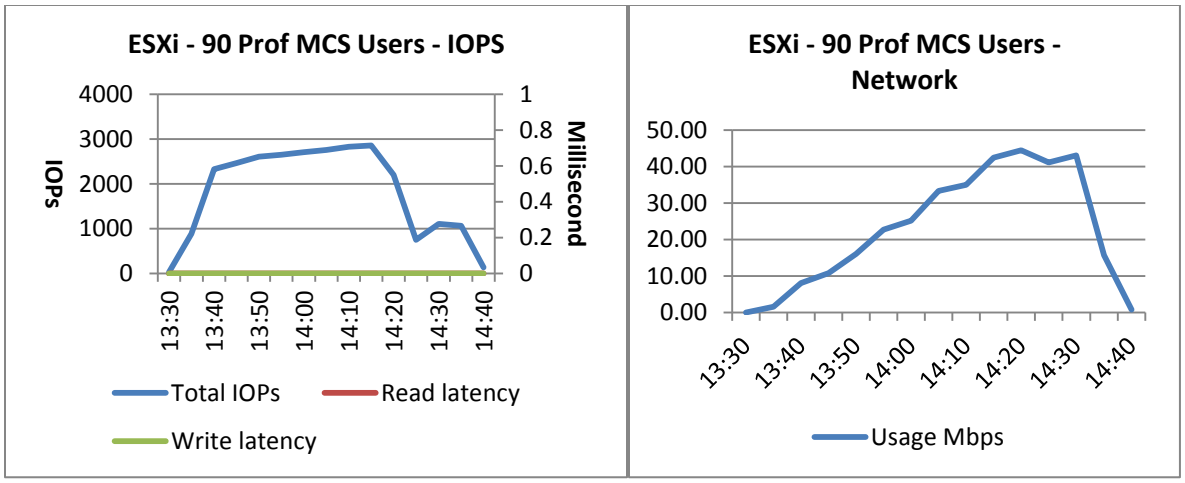

6.3.2.8 IOMeter test against SSD disks

These graphs show IOMETER results against the SSD disks using the following IOMETER profile. This was performed on a single VM on the compute host using a 20GB disk configured on the SSD volume. The test was run to gather information that would indicate the overall capability of the SSD volume when hosting VDI VMs. The access specifications was configured at proposed by Jim Moyle of Atlantis at the following location
http://blog.atlantiscomputing.com/2013/08/how-to-use-iometer-to-simulate-a-desktop-workload

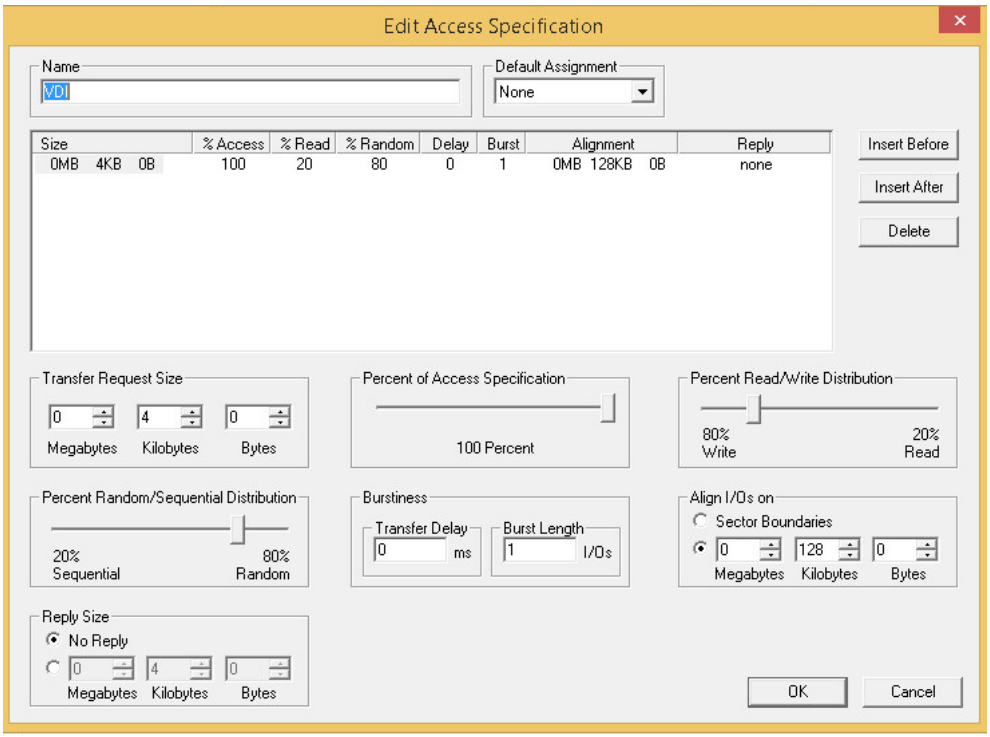

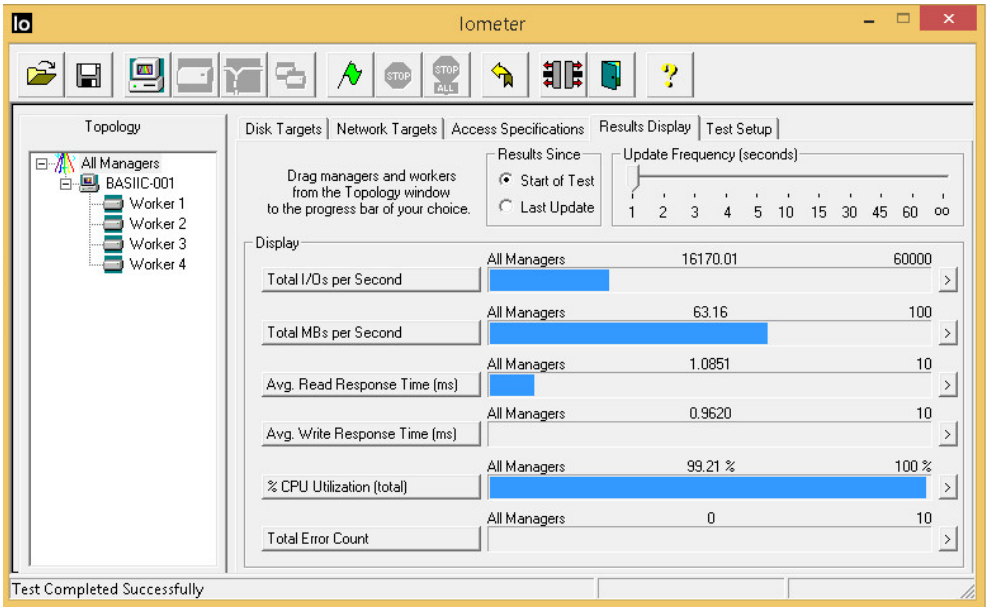

It can be seen that at 16000 IOPs with 80% write the latency remains at approx. 1 millisecond. While IOMETER test scenarios are very subjective and it is arguable whether the test specification used above truly represents a VDI workload, it does represent a reasonable approximation and the result of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

Number of Desktops per Server	Approx. Max IOPS per Desktop
170 Standard	95 IOPS
110 Enhanced	145 IOPS
90 Premium	180 IOPS

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

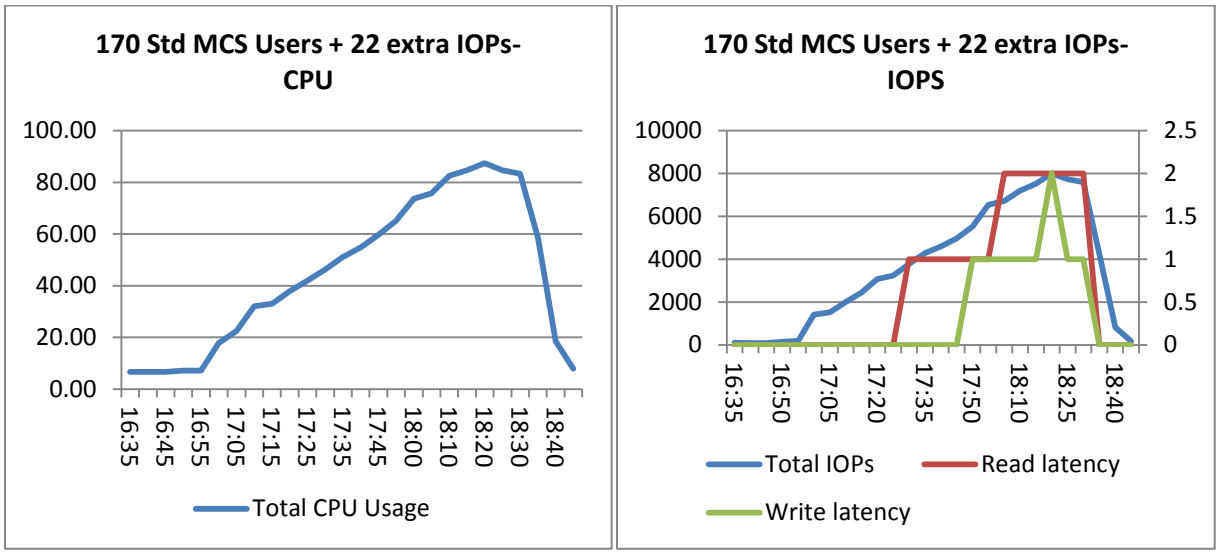

The login script launched an automated Iometer workload in each VDI session to generate extra IOPs within the session. This generated sustained IOPs of approximately 8000. As can be seen from the graphs above the latency of the local tier 1 storage on the SSD volumes did not rise above 2 milliseconds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

6.3.2.10 MCS Standard User Workload (170 Users) Logon storm

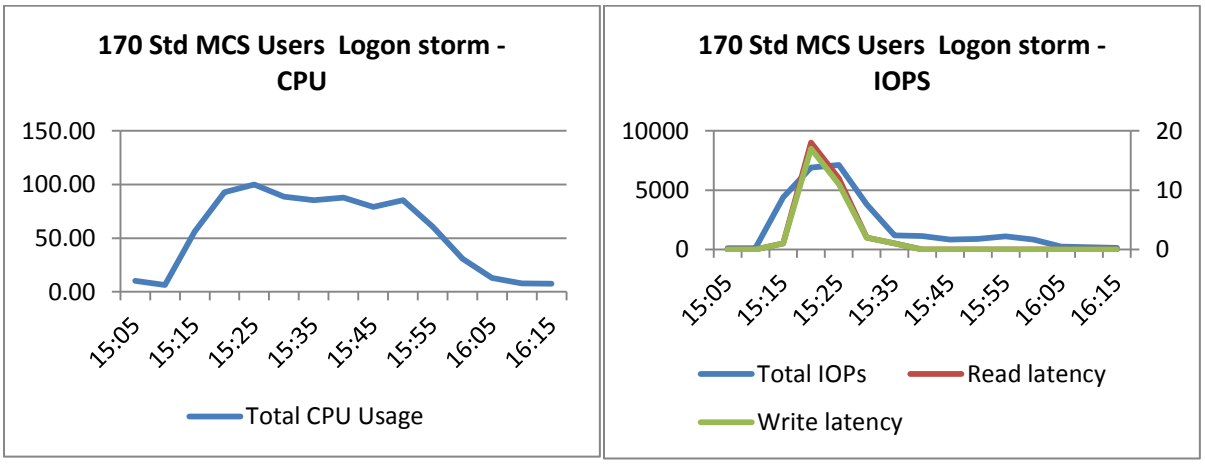

During the logon period of 10 minutes the Read and Write latency on the SSD volume rose to 17 milliseconds but quickly dropped to 0 once the sessions reached steady state. A latency of 17 seconds is still below a value that could affect VDI performance. The shortened logon interval produced a peak CPU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

About the Authors

Peter Fine is the Sr. Principal Engineering Architect for the 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.

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

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