Dell DVS Reference Architecture for Windows Server 2012 v1.5

Table of Contents

1 Introduction

1.1 Purpose of this document

This document describes:

1. Dell DVS Reference Architecture for Windows Server 2012, scaling from 50 to 600 VDI users or 260 -1040 RDSH sessions.
2. A VDI Experience Proof of Concept (POC) Solution, an entry level configuration supporting 10 VDI users.

This document addresses the architecture design, configuration and implementation considerations for the key components of the architecture required to deliver virtual desktops via RDS on Hyper-V 2012.

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 Physical Architecture Overview

The core architecture design consists of the Local Tier1 solution model. “Tier 1” in the DVS context defines from which disk source the VDI sessions execute. Local Tier1 applies to rack servers only while Shared Tier 1 can be rack or blade. Tier 2 storage is utilized for user profile/data and Management VM execution.

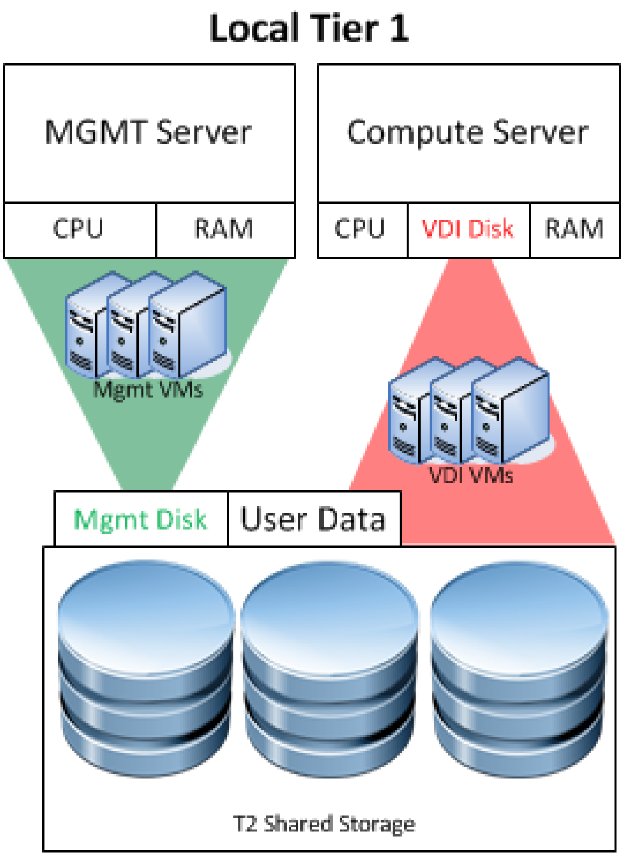

2.2 Layouts and Breakpoints

The solution architecture will follow the traditional DVS distributed design model consisting of 4 primary layers: Network, Compute, Management, and Storage. The Network and Storage layers can be optionally provided by the customer if suitable infrastructure is already in place. The Compute layer contains the hosts that serve the VDI sessions and the Management layer contains the components required to support the VDI broker infrastructure.

The following highlights the key layout and scaling elements of the solution.

2.3 Local Tier 1 – Solution Layers

Only a single high performance Force10 S55 48-port switch is required to get started in the Network layer. This switch will host all solution traffic consisting of 1Gb iSCSI and LAN sources. Additional switches can be added and stacked as required to provide High Availability for the Network layer.

The Compute layer consists of the server resources responsible for hosting the user sessions, whether shared via RDSH (formerly Terminal Services) or pooled via RDVH (see section 4.5.1 for a detailed explanation of each role). The RDVH role requires Hyper-V as well as hardware assisted virtualization so must be installed into the parent parition of the Hyper-V instance. The RDSH role is enabled within dedicated VMs on the same or dedicated hosts in the Compute layer.

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 RDS VMs necessary to support the infrastructure as well as a file server to host SMB shares for user Profile Disks or data.

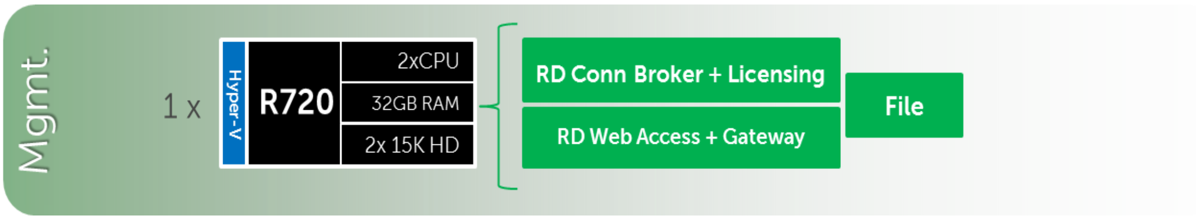

The Storage layer is made up by the capacity dense and performance capable Equallogic 4100E iSCSI array. 12TB is provided in base form that can scale as high as 36TB to suit your capacity requirements. A second 4100E can be added to group the two arrays to provided greater capacity or performance.

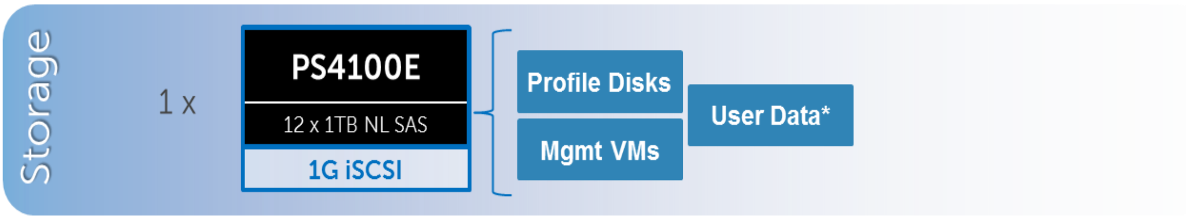

2.4 Dell DVS 10-Seat Trial Kit

To get up and running as quickly as possible with pooled VDI, Dell is offering an extremely affordable solution capable of supporting 10 concurrent users for a minimal investment. This architecture leverages an inexpensive single server platform intended to demonstrate the capabilities of VDI for a small environment or focused POC/ trial. Networking is provided optionally in this solution and all VDI roles/ sessions are hosted on a single server.

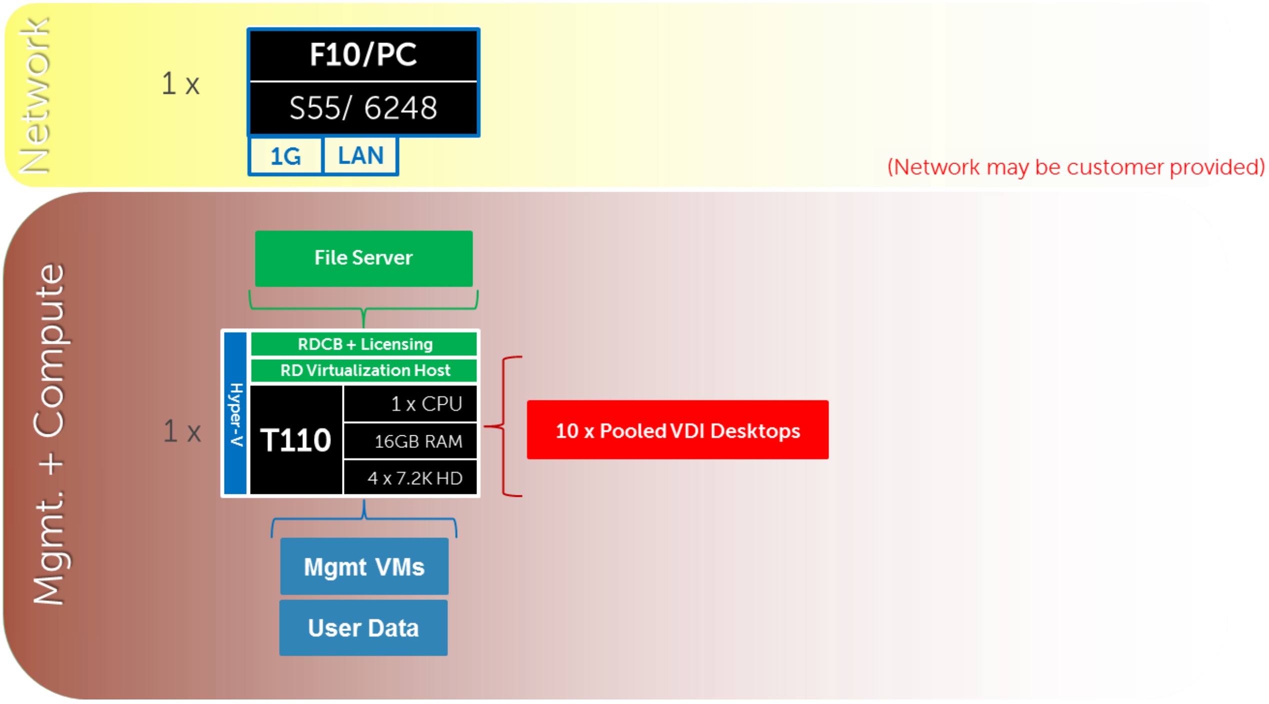

For more information on the 10-user POC offering, please see Appendix A.

2.5 Local Tier 1 – 50 User/ Pilot

For small scale deployments or pilot efforts intended to familiarize yourself with the Dell enterprise solution architecture, we offer a 50 user/ pilot solution completely scalable to the maximum supported configuration. The architecture for the 50 user pilot follows a non-distributed model with all VDI and Management functions running on a single host. If additional scaling is desired, you can grow into a larger distributed architecture seamlessly with no loss on your initial investment.

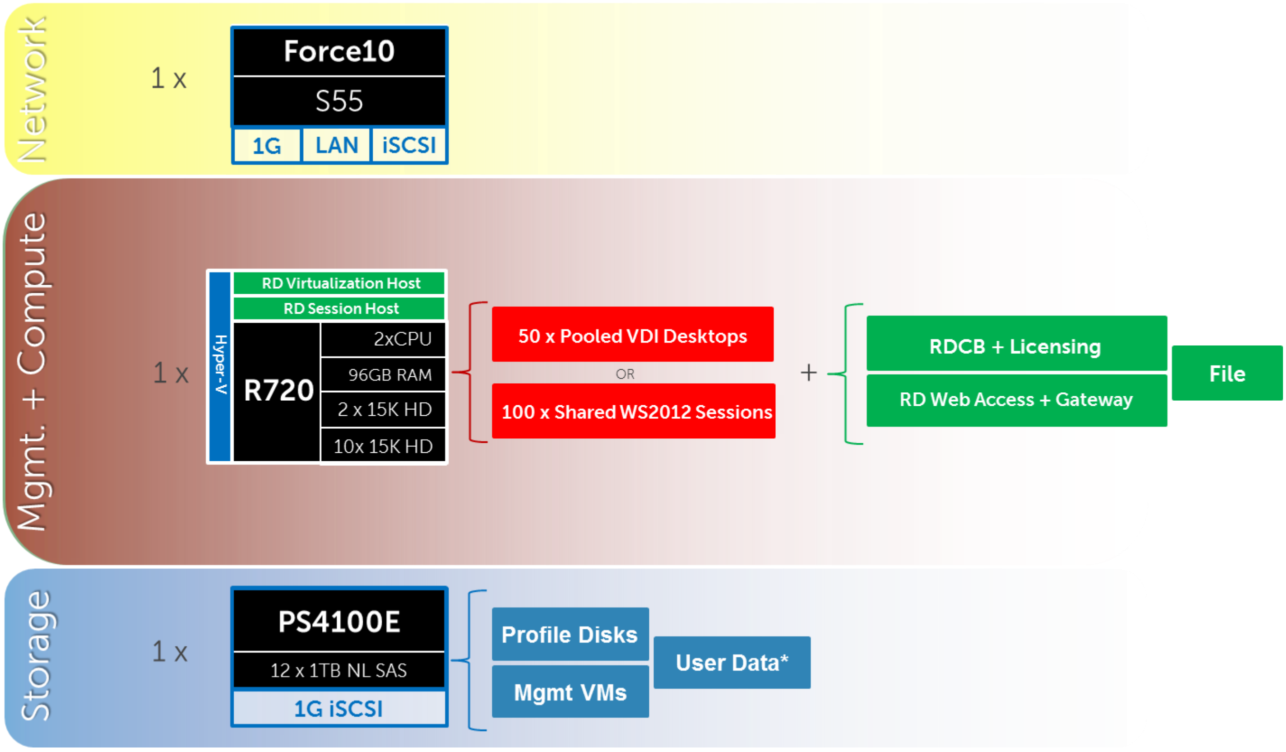

2.6 Local Tier 1 – Combined

As a logical entry point to the distributed RDS solution stack, a combined architecture is offered to host both the RD Virtualization Host (RDVH) and RD Session Host (RDSH) roles within the same physical Compute host while separating the Management layer. This will enable users requiring either shared RDP or pooled VDI sessions to be hosted on the same physical server. The value of this solution is a minimum infrastructure investment with maximum VDI flexibility easily tailored to shared and pooled user types. Horizontal scaling is achieved simply by adding additional Compute hosts. Additional information on the hardware components can be found in section 3 below.

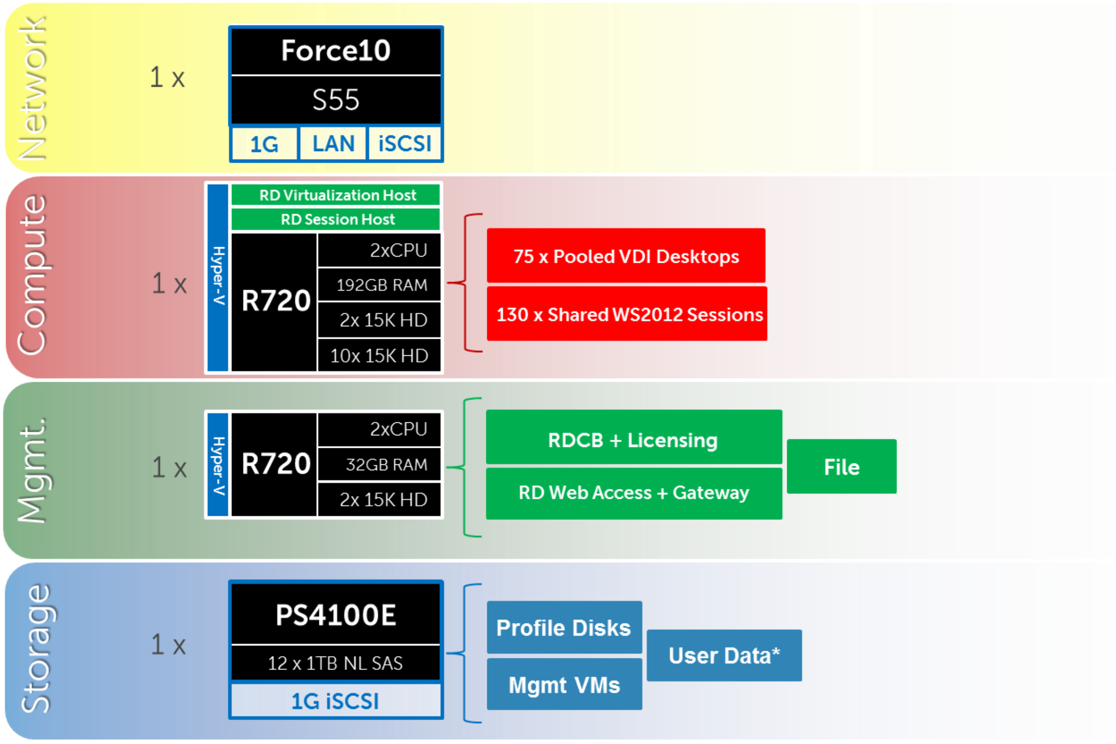

2.7 Local Tier 1 – Base

In the base distributed architecture the RDVH or RDSH roles are assigned to a dedicated Compute host. This architecture can support either a single RDVH or RDSH Compute host or one of each. This solution provides maximum Compute host user density for each broker model and allows clean linear upward scaling. You'll notice that the hardware spec is slightly different for the two Compute host types, giving additional RAM to the virtualization host. This of course can be adjusted to suit your specific needs. Additional information on the hardware components can be found in section 3 below.

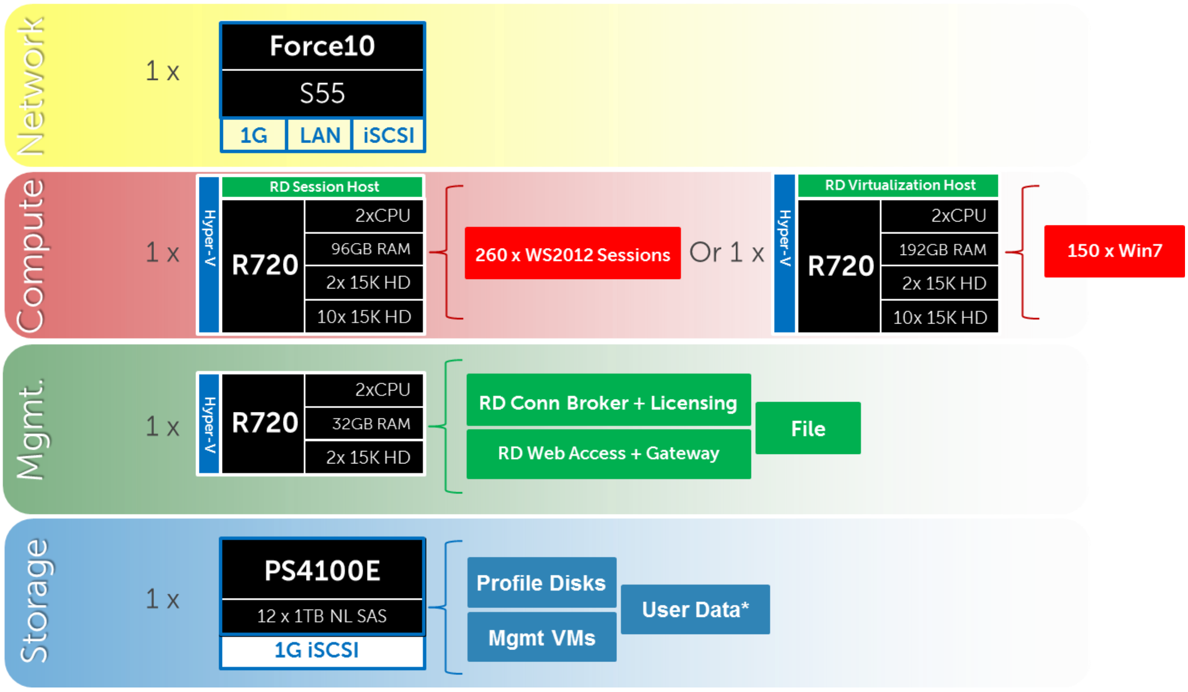

2.7.1 Fully Expanded

The fully expanded architecture provides linear upward scale for both the RDVH and RDSH roles optimized for 600 pooled VDI sessions or over 1000 shared. See Appendix for test results. This solution supports up to 4 Compute hosts of any combination running either RDVH or RDSH roles to meet the needs of the enterprise. Additional information on the hardware components can be found in section 3 below. The overall infrastructure is capable of supporting a much higher number of sessions but additional management infrastructure will be required. The native RDS management infrastructure was designed to support in the neighborhood of 500 sessions.

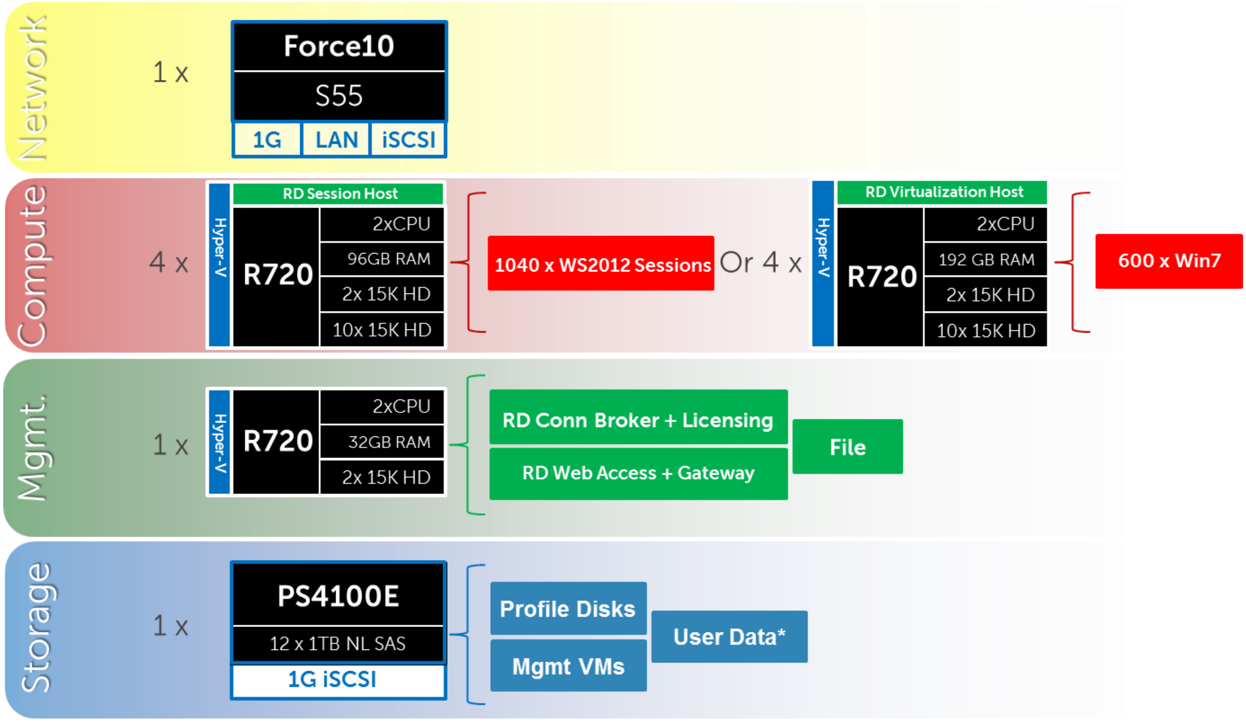

2.7.2 High Availability

High availability (HA) is currently offered to protect all layers of the solution architecture. An additional ToR switch is added to the Network layer and stacked to provide redundancy, additional Compute and Mgmt hosts are added to their respective layers, and Hyper-V clustering is introduced in the Management layer. Please see section 4.6 for more detailed information regarding HA in this architecture.

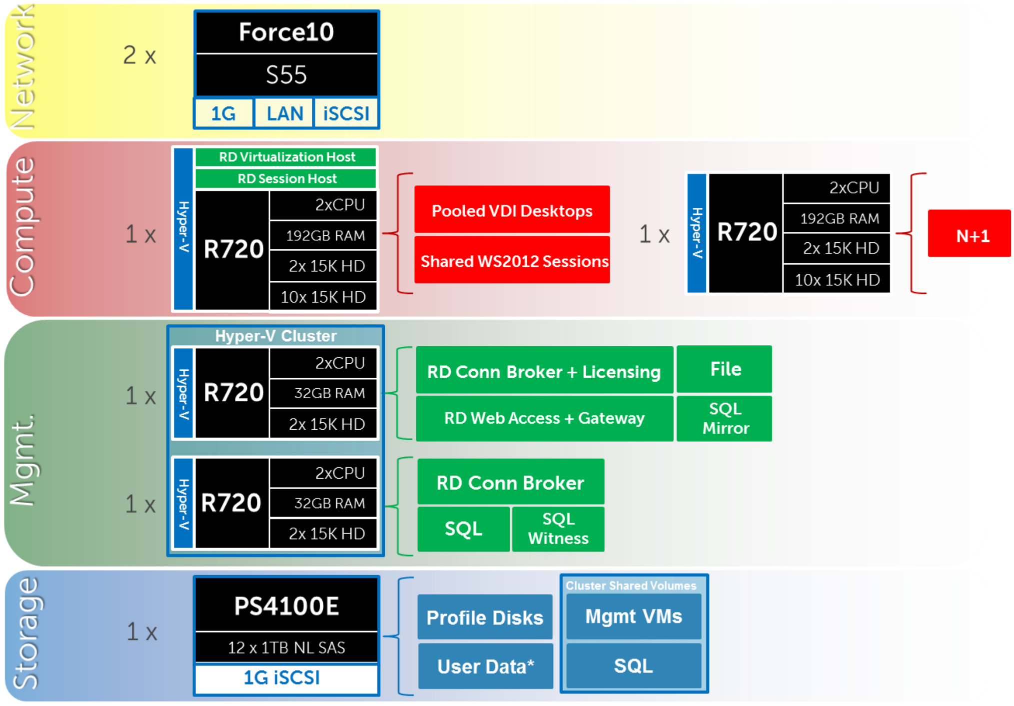

2.7.3 Solution Density Summary

Design Scale	Management Hosts	Compute Hosts	RDSH Sessions		RDVH Sessions	HA
10 User POC	0	1	0		10	-
50 User / Pilot	0	1	100	*or*	50	-
Combined	1	1	130	*and*	75	+ 1 Compute +1 Mgmt
Base	1	1	260	*or*	150	+ 1 Compute +1 Mgmt
Expanded	1	2	520	*or*	300	+ 1 Compute +1 Mgmt
Fully Expanded	1	4	1040	*or*	600	+ 1 Compute +1 Mgmt

2.8 Cabling Diagrams

2.8.1 Local Tier 1 – Network Architecture

In the Local Tier 1 architecture, a single Force10 S55 switch can be shared among all network connections for both Management and Compute layer components, for the upper limit of 500 pooled VDI sessions. Over 1000 users DVS recommends separating the network fabrics to isolate iSCSI and LAN traffic as well as making each network stack redundant. Only the Management servers connect to iSCSI storage in this model. All ToR traffic has been designed to be layer 2/ switched locally, with all layer 3/ routable VLANs trunked from a core or distribution switch. The following diagrams illustrate the logical data flow in relation to the core switch.

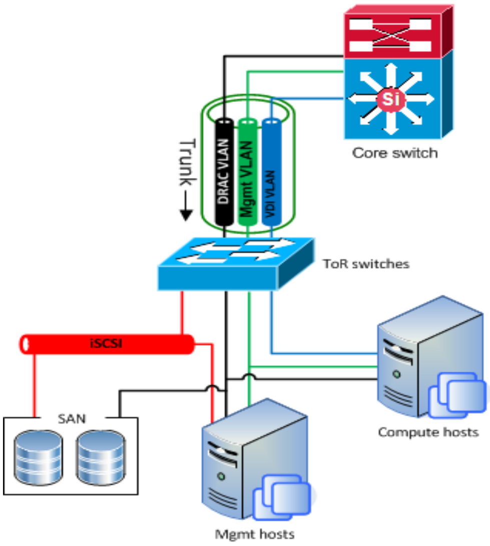

2.8.2 Local Tier 1 Cabling

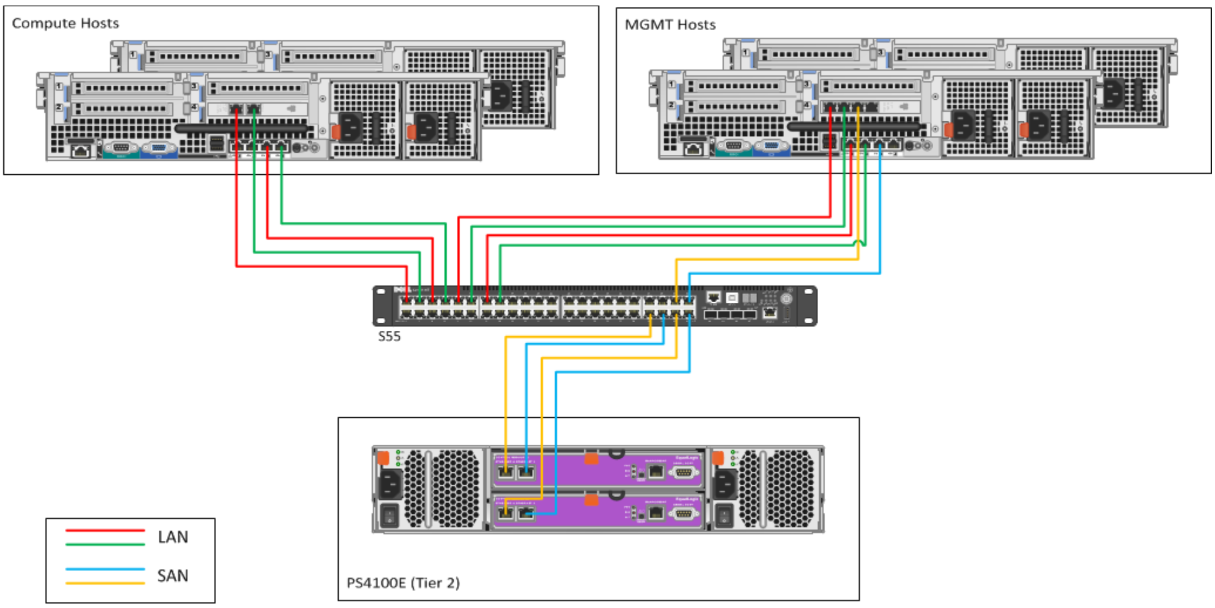

3 Hardware Components

3.1 Network

The following sections contain the core network components for the DVS local and shared Tier 1 solutions. General cabling guidance to consider in all cases is that TwinAx is very cost effective for short 10Gb runs and for longer runs fiber with SFPs should be used.

3.1.1 Force10 S55 (ToR Switch)

Model	Features	Options	Uses
Force10 S55	44 x BaseT (10/100/1000) + 4 x SFP	Redundant PSUs 4 x 1Gb SFP ports the support copper or fiber 12Gb or 24Gb stacking (up to 8 switches) 2 x modular slots for 10Gb uplinks or stacking modules	ToR switch for LAN and iSCSI in Local Tier 1 solution

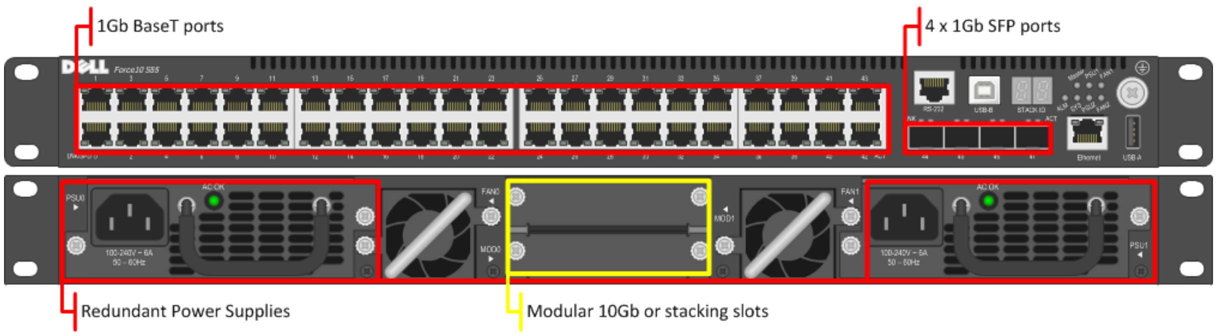

Guidance:

- 10Gb uplinks to a core or distribution switch are the preferred design choice using the rear 10Gb uplink modules. If 10Gb to a core or distribution switch is unavailable the front 4 x 1Gb SFP ports can be used.
- The front 4 SFP ports can support copper cabling and can be upgraded to optical if a longer run is needed.

For more information on the S55 switch and Dell Force10 networking, please visit: LINK

3.1.1.1 Force10 S55 Stacking

The Top of Rack switch in the Network layer can be optionally stacked with a second switch, if greater port count or redundancy is desired. Each switch will need a stacking module plugged into a rear bay and connected with a stacking cable. Switch stacks greater than 2 should be cabled in a ring configuration with the last switch in the stack cabled back to the first. Uplinks should be configured on all switches in the stack back to the core to provide redundancy and failure protection.

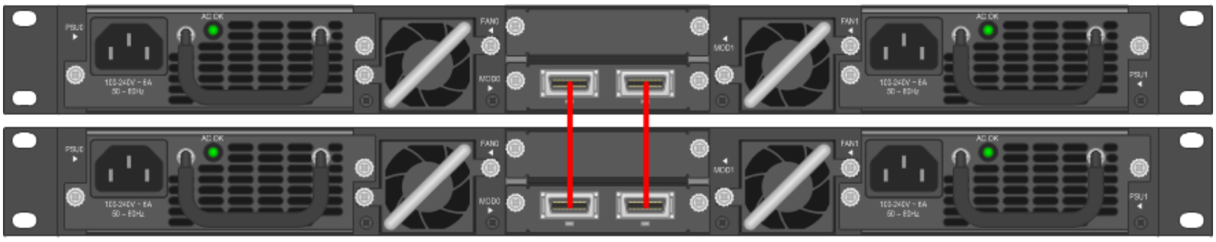

Please reference the following Force10 whitepaper for specifics on stacking best practices and configuration: LINK

3.2 Servers

3.2.1 Local Tier 1 Rack

The server platform for the Windows Server 2012 RDS enterprise solution is the best-in-class Dell PowerEdge R720. This dual socket CPU platform runs the fastest Intel Xeon E5-2600 family of processors, can host up to 768GB RAM, and supports up to 16 2.5" SAS disks. Uncompromising performance and scalability in a 2U form factor.

In the local Tier 1 model, VDI sessions execute on the local storage of each Compute server. Due to the local disk requirement in the Compute layer, this model supports rack servers only. In this model only the Management server hosts access shared storage to support the solution's Management role VMs. Because of this, the Compute and Management servers are configured with different add-on NICs to support their pertinent network fabric connection requirements. Refer to section 2.3.1 for cabling implications. The management server host has reduced RAM, CPU and fewer disks, since its VMs execute on shared Tier 2 storage. Both servers require a pair of hard disks configured in RAID1 to support the host operating system.

Local Tier 1 Compute Host – PowerEdge R720
2 x Intel Xeon E5-2690 Processor (2.9Ghz)
192GB Memory (12 x 16GB DIMMs @ 1600Mhz) (**RDVH**)
Or 96GB Memory (12 x 8GB DIMMs @ 1600Mhz) (**RDSH**)
Microsoft Windows Server 2012 Hyper-V
2 x 146GB SAS 6Gbps 15k Disks (OS)
10 x 146GB SAS 6Gbps 15k Disks (VDI)
PERC H710 Integrated 1GB RAID Controller – RAID10
Broadcom 5720 1Gb QP NDC (LAN)
Broadcom 5720 1Gb DP NIC (LAN)
iDRAC7 Enterprise w/ vFlash, 8GB SD
2 x 750W PSUs

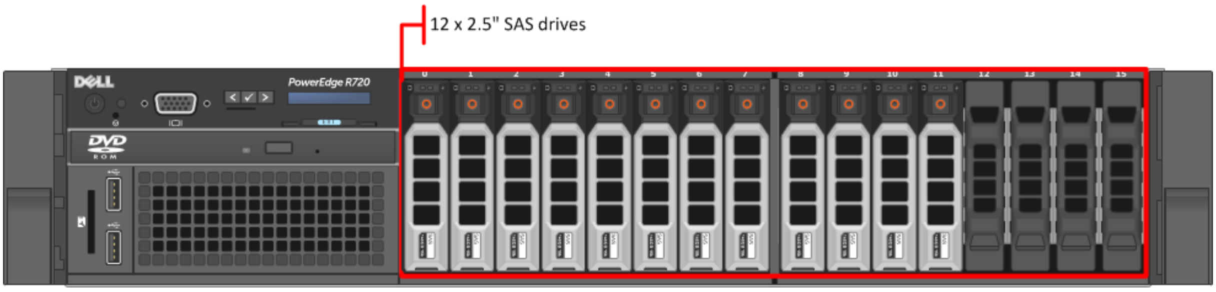

Local Tier 1 Management Host – PowerEdge R720
2 x Intel Xeon E5-2680 Processor (2.7Ghz)
32GB Memory (4 x 8GB DIMMs @ 1600Mhz)
Microsoft Windows Server 2012 Hyper-V
2 x 146GB SAS 6Gbps 15k Disks (OS)
PERC H710 Integrated 1GB RAID Controller – RAID1
Broadcom 5720 1Gb QP NDC (LAN/iSCSI)
Broadcom 5719 1Gb QP NIC (LAN/iSCSI)
iDRAC7 Enterprise w/ vFlash, 8GB SD
2 x 750W PSUs

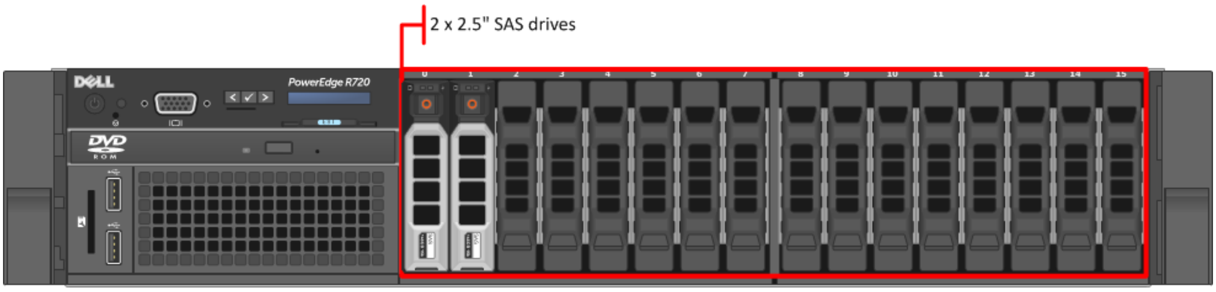

For more information on the Dell PowerEdge R720 server and other server offerings from Dell, please visit: LINK

3.3 Equallogic Storage

Equallogic shared storage is used in the Local Tier 1 enterprise solutions.

3.3.1 PS4100E

Model	Features	Options	Uses
Equallogic PS4100E	12 drive bays (NL-SAS/ 7200k RPM), dual HA controllers, Snaps/Clones, Async replication, SAN HQ, 1Gb	12TB – 12 x 1TB HDs	Tier 2 array for 1000 total users or less in local Tier 1 solution model (1Gb)
		24TB – 12 x 2TB HDs	
		36TB – 12 x 3TB HDs	

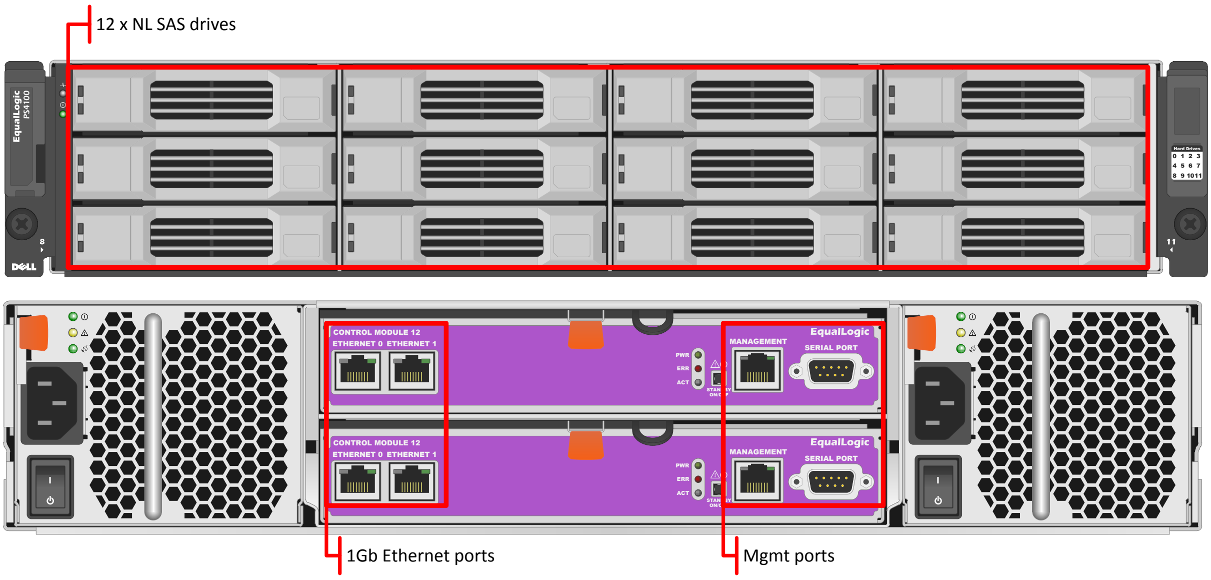

For more information on the Dell Equallogic PS4100E and other networked storage options from Dell, please visit: LINK

3.4 Dell Wyse End Points

3.4.1 Display Choices for Dell Wyse Endpoints

Good	**E2213*** 1680x1050 VGA, DVI	
Better	**P2212H** 1920x1080 VGA, DVI, USB	
Best	**U3011*** 2560x1600 VGA, DVI, DP, USB	

*U3011 Not Available In Mainland China

*E2213 Not Available In South Korea

3.4.2 Dell Wyse R10L

The Dell Wyse R10L thin client combines outstanding power and security with display performance and manageability to help maximize your investment in your virtual desktop infrastructure. With built-in Dell Wyse ThinOS, the R10L offers exceptional speed and security, while high-performance central and graphics processors offer an outstanding user experience. Enjoy scalable enterprise-wide management, including simple deployment, patching and updates, and future-proof connectivity to accommodate a wide variety of peripherals and interfaces. Drawing only 12 to 15 watts in typical usage, the R10L can reduce carbon emissions for an environmentally-conscious and comfortable working environment.

3.4.3 Dell Wyse D90D7

The Dell Wyse D90D7 is a high-performance Windows Embedded Standard 7 thin client for virtual desktop environments. Featuring a dual-core AMD processor and a revolutionary, unified engine that eliminates performance constraints, the D90D7 achieves outstanding speed and power for the most demanding VDI and embedded Windows applications, rich graphics and HD video. Take a unit from box to productivity in minutes. Just select the desired configuration and the D90D7 does the rest automatically—no need to reboot. And with Microsoft Windows Embedded Device Manager connectivity, you can leverage your existing Microsoft System Center Configuration Manager platform. The D90D7 is an ideal thin client for demanding virtual desktop or cloud applications.

3.4.4 Dell Wyse Z90D7

Dell Wyse Z90D7 is a super high-performance Windows Embedded Standard 7 thin client for virtual desktop environments. Featuring a dual-core AMD processor and a revolutionary, unified engine designed to eliminate performance constraints, the Z90D7 offers incredible speed and power for the most demanding embedded Windows applications, rich graphics and HD video. With Microsoft Windows Embedded Device Manager connectivity, one can leverage your existing Microsoft System Center Configuration Manager platform. The Z90D7 is an ideal thin client for the most demanding mix of virtual desktop or cloud applications.

3.4.5 Dell Wyse X90m7

Dell Wyse X90m7 mobile thin clients powered with Windows Embedded Standard 7 offers great performance and a crisp 14" LED backlit display, along with a dual-core high-performance AMD processor with Radeon HD 6310 graphics for exceptional HD multimedia capabilities. A built-in webcam, integrated wireless a/b/g/n, and support for 3G/4G cards and optional smart card reader offers you rich interactivity and flexibility. With no local hard drive to weigh it down, the X90m7 can eliminate the risk of losing or exposing sensitive data. And with Microsoft Windows Embedded Device Manager connectivity, one can manage it with your existing Microsoft System Center Configuration Manager platform. Seamlessly and securely connect to your IT infrastructure and access some of the most advanced web browser and media player capabilities available in a mobile thin client.

4 Solution Architecture for Microsoft Remote Desktop Services

4.1 Overview

This solution architecture follows a distributed model where solution components exist in layers. The Compute layer is where VDI desktop VMs execute, the Management layer being dedicated to the broker management role VMs. Both layers, while inextricably linked, scale independently.

4.1.1 RDS Options

Server 2012 RDS provides a number of VDI options to meet your needs, all within a single, simple, wizard-driven environment that is easy to set up and manage.

- **Sessions**, hosted by the RDSH role (formerly Terminal Services), provide easy access to a densely shared session environment. Each RDP-based session shares the total available server resources with all other sessions logged in concurrently on the server. This is the most cost effective option and a great place to start with Server 2012 RDS. An RDS CAL is required for each user or device accessing this type environment.
- **Pooled VMs** are the non-persistent user desktop VMs traditionally associated with VDI. Each user VM is assigned a dedicated slice of the host server's resources to guarantee the performance of each desktop. The desktop VM is dedicated to a single user while in use then returned to the pool at logoff or reboot and reset to a pristine gold image state for the next user. Applications can be built into gold images or published via RemoteApps. An RDS CAL is required for each user or device accessing this type environment.
- **Personal VMs** are persistent 1-to-1 desktop VMs assigned to a specific entitled user. All changes made by Personal VM users will persist through logoffs and reboots making this a truly personalized computing experience. An RDS CAL is required for each user or device accessing this type environment.

Please contact Dell or Microsoft for more information on licensing requirements for VDI.

4.2 Compute Server Infrastructure

The Compute host configuration for the enterprise solution varies slightly as to whether it will be hosting RDSH or RDVH roles, or both. The RDVH role must be enabled in the Hyper-V parent partition thus providing one RDVH role per Compute host if pooled or personal VMs are required. The RDSH role should be enabled in up to 4 VMs on a single Compute host to support up to 260 session-based users.

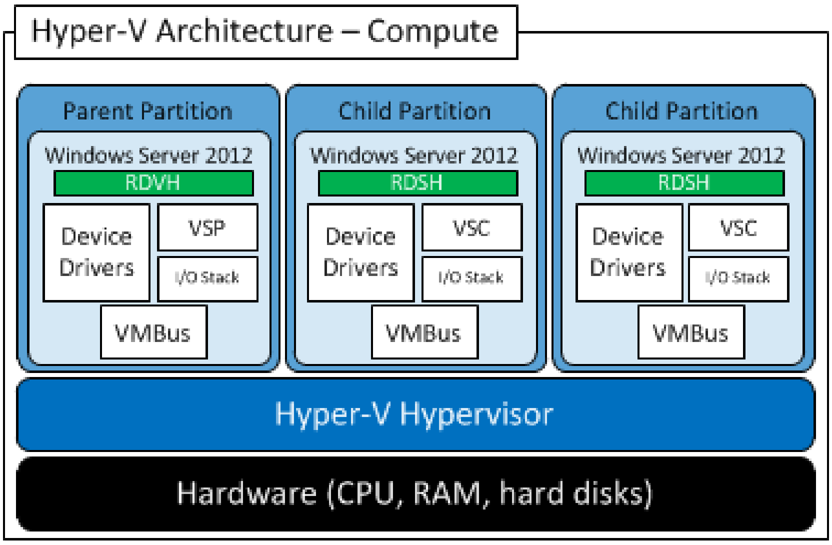

The requirements for RDSH VMs are outlined below. All application and non-OS related files should be installed in the 5GB data disk:

Role	vCPU	Startup RAM (GB)	Dynamic Memory			NIC	OS + Data vDisk (GB)	Tier 2 Volume (GB)
			Min|Max	Buffer	Weight			
RD Session Host	8	16	512MB | 20GB	20%	Med	1	40 + 20	-

4.3 Management Server Infrastructure

The Management host configuration consists of VMs running in Hyper-V child partitions with the pertinent RDS roles enabled. No RDS roles need to be enabled in the root partition for Management hosts.

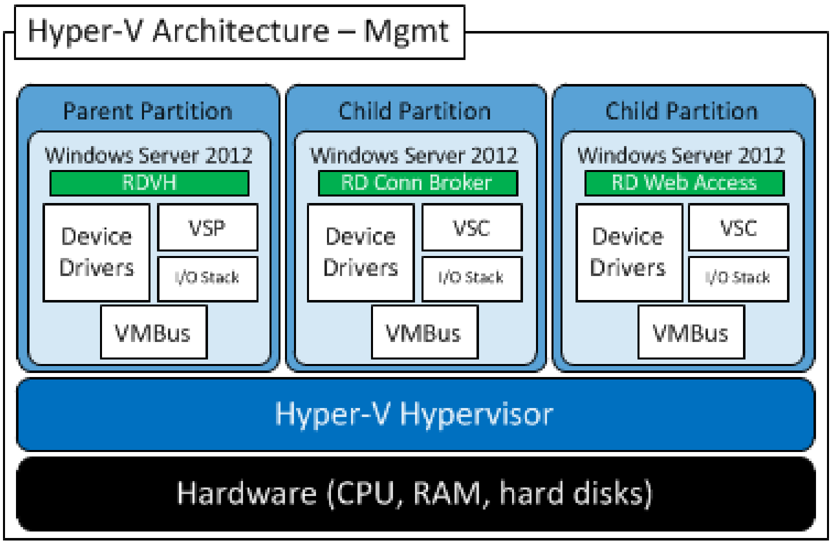

Management role requirements for the base solution are summarized below. Data disks should be used for role-specific application files/ data, logs, IIS web files, etc and should exist in the Management volume on the 4100E array. Please note that the Tier2 volume presented to the file server is designated as a pass-through disk (PTD).

Role	vCPU	Startup RAM (GB)	Dynamic Memory			NIC	OS + Data vDisk (GB)	Tier 2 Volume (GB)
			Min|Max	Buffer	Weight			
RDCB + License Server	1	4	512MB|8GB	20%	Med	1	40 + 10	-
RDWA + RDG	1	4	512MB|8GB	20%	Med	1	40 + 10	-
File Server	1	4	512MB|8GB	20%	Med	1	40 + 10	2048 (PTD)
TOTALS	**3**	**12**				**3**	**120 + 30**	**2048**

4.4 Storage Architecture

4.4.1 Local Tier 1

In the enterprise solution model, Tier 1 storage exists as local hard disks on the Compute hosts. To achieve the required performance level, RAID 10 must be used across all local disks used for VDI. A single volume per local Tier 1 compute host is sufficient to host the provisioned desktop VMs along with their respective write caches. Increased IO performance is provided via the 1GB cache module on the H710 RAID controller. Two of the disks included on each host should be mirrored and used to install the OS.

Volumes	Size (GB)	RAID	Storage	Purpose	File System
OS	135	1	Tier 1	Host Operating System	NTFS
VDI	680	10	Tier 1	Pooled + Shared VDI	NTFS

4.4.2 Shared Tier 2

Tier 2 is shared iSCSI storage used to host the Management server VMs and user data. The Equallogic 4100 series arrays will be used for smaller scale deployments up to 500 pooled VDI users (Local Tier 1 only). The table below outlines the minimum volume requirements for Tier 2 storage. Larger disk sizes can be chosen to meet the capacity needs of the customer. The user data volume can be presented either via a VHDX or native NTFS pass-through disk to simplify a future upgrade to NAS. All VM disks should be presented as VHDX.

Volumes	Size (GB)	RAID	Storage Array	Purpose	File System
Management	500	50	Tier 2	RDS VMs, File Server	NTFS
User Data	2048	50	Tier 2	File Server	NTFS
User Profiles	20	50	Tier 2	User profiles	NTFS
Templates/ ISO	200	50	Tier 2	ISO/ gold image storage (optional)	NTFS

4.4.3 Virtual Hard Disk Format

The VHDX disk format provides numerous advantages over the older VHD specification and should be used for all virtual disks in the solution. Larger vDisk support, up to 64TB, corruption protection during power failures, and larger sector disk alignment are a few of the new features.

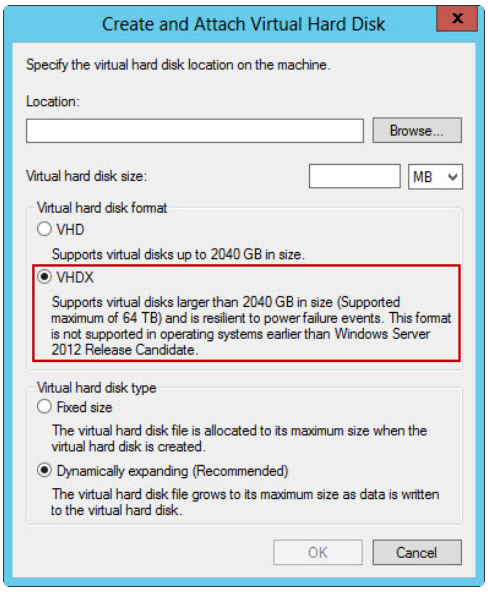

4.4.4 DNS

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

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

4.4.5 File Services

The File Services role will be provided via a dedicated VM. In the interest of portability and providing a clean path to an optional HA upgrade, the volumes can be presented to the file server VM in the form of a Pass-Through Disk. This will ensure a cleaner transition for customers who upgrade to HA and add a NAS head to their environments by keeping the data on the storage array and not inside a VHDX that will need to be copied out.

4.4.6 User Profile Disks

User Profile Disks is a component of the Server 2012 RDS solution which is used to manage user profiles. Profile Disks provide a cohesive method to manage user documents and profile data in a VDI environment. Profile disks can be enabled in either session or virtual desktop collections and provide options to customize based on the need of the implementation. The solution file server will be used to host user profile disks and home drive data via SMB shares, which can be separated using an optional user data volume, if desired. Otherwise, all user data should be stored within a profile disk.

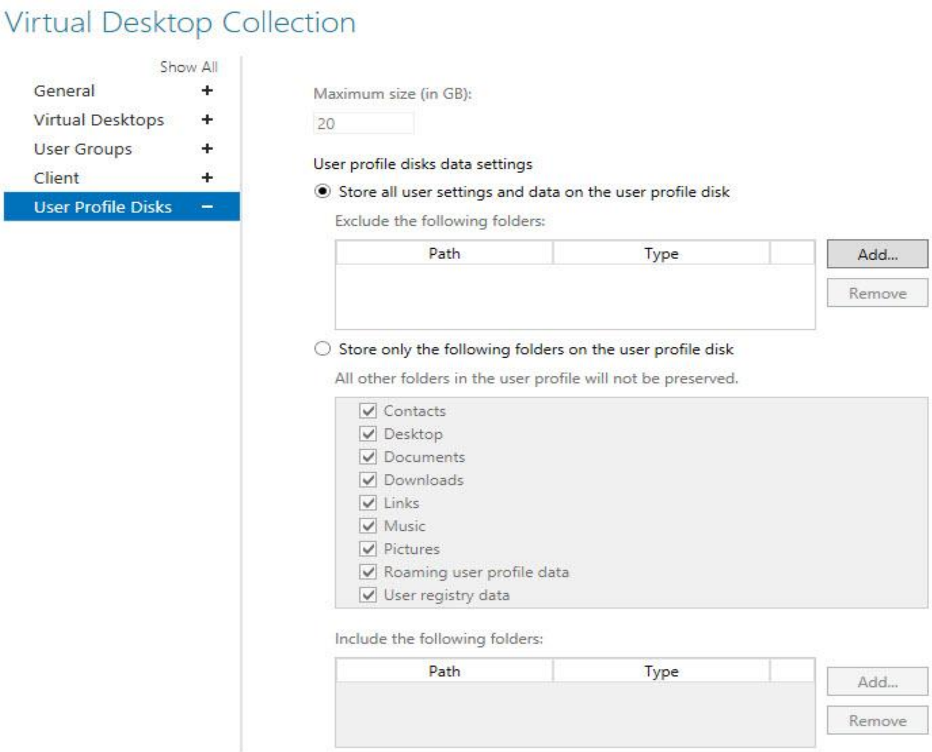

4.5 Hyper-V Configuration

The Local Tier 1 solution will be built upon the Server 2012 Hyper-V hypervisor. All Microsoft best practices and prerequisites should be adhered to (NTP, DNS, Active Directory, etc).

Solution Architecture Components	
Hypervisor	Microsoft Windows Server 2012 Hyper-V
VDI Broker	Remote Desktop Services 2012

Server OS	Microsoft Windows Server 2012 Standard edition for Management and Compute hosts.
Desktop OS	Microsoft Windows 7 Enterprise (x86)

Virtual Machine Configuration	
Hard Drive	SCSI Controller
Disk Type	VHDX – Dynamically Expanding
Smart Paging File	Store with VM
Virtual CPU	1 per VDI VM
Dynamic Memory (VDI)	512MB – Minimum, 2GB – Maximum

4.5.1 Core Components

Each Compute and Management host will run the full GUI version of Server 2012 in this solution. All RDS component roles, except for RDVH, will exist as VMs yielding 100% virtualized architecture in both Compute and Management server layers. RD Session Hosts will be enabled in dedicated VMs on the Compute hosts, while the RDS infrastructure components will be enabled in dedicated VMs on the Management hosts. The RDVH role will be enabled in the parent partition of selected Compute hosts as discussed in section 4.2.

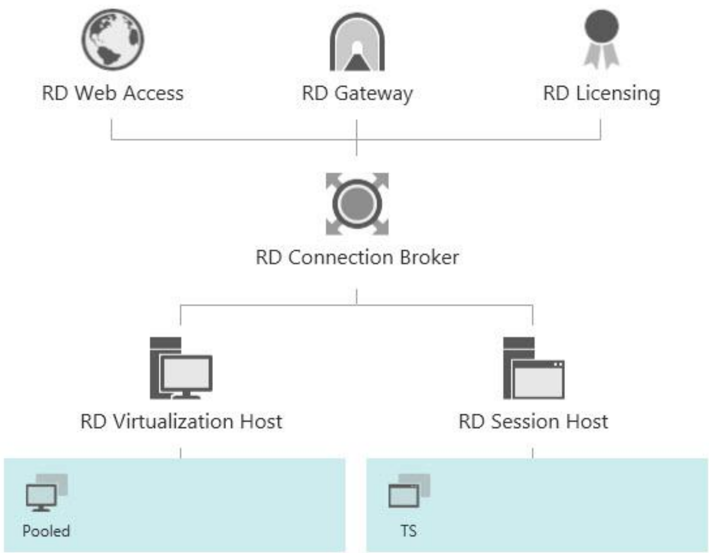

4.5.2 Hyper-V Networking (Local Tier 1)

The network configuration in this model will vary slightly between the Compute and Management hosts. The Compute hosts will not need access to iSCSI storage since they are hosting the VDI sessions on local disk. 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 Hyper-V 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 Hyper-V Management traffic – L3 routed via core switch
    - iSCSI VLAN: Configured for iSCSI traffic – L2 switched only via ToR switch
    - VDI Management VLAN: Configured for VDI infrastructure traffic – L3 routed via core switch
- An optional DRAC VLAN can be configured for all hardware management traffic, which should be L3 routed via core switch

In this solution architecture, LAN and iSCSI traffic will be segmented in dedicated VLANs but combined within a single switch to minimize the initial network investment. Following best practices and in solutions that may desire larger scales, this traffic should be separated into discrete switches. Each Local Tier 1 Compute host will have a quad port NDC as well as an add-on 1Gb dual port PCIe NIC. The LAN traffic from the server to the ToR switch should be configured as a LAG to maximize bandwidth. The Compute hosts will require 2 vSwitches, one for VDI LAN traffic, and another for the Hyper-V Management.

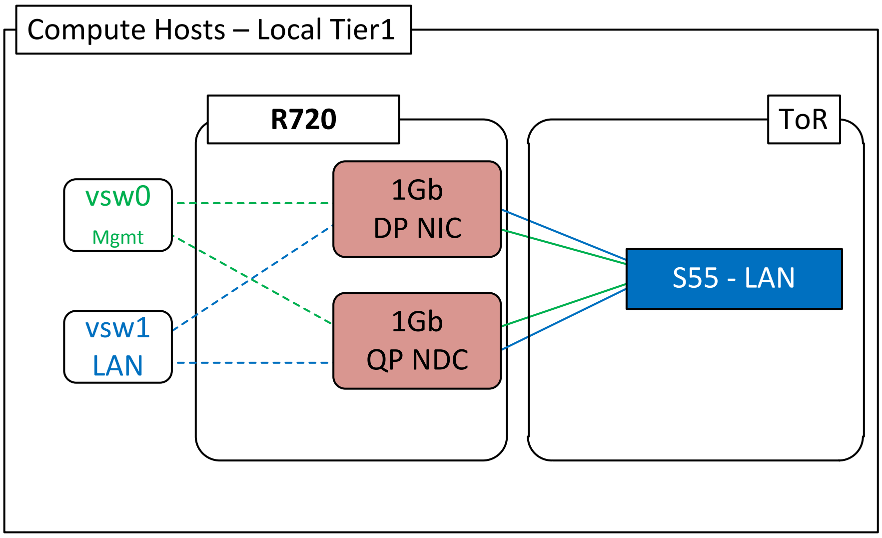

The Management hosts have a slightly different configuration since they will additionally access iSCSI storage. The add-on NIC for the Management hosts will be a 1Gb quad port NIC. 3 ports of both the NDC and add-on NIC will be used to form 3 virtual switches. iSCSI should be isolated onto its own vSwitch with teamed NICs and connections from all 3 vSwitches should pass through both the NDC and add-on NIC per the diagram below. Care should be taken to ensure that all vSwitches are assigned redundant NICs that are NOT from the same PCIe device. The LAN traffic from the server to the ToR switch should be configured as a LAG. VLAN IDs should be specified in all virtual switches used within the Compute layer Hyper-V host.

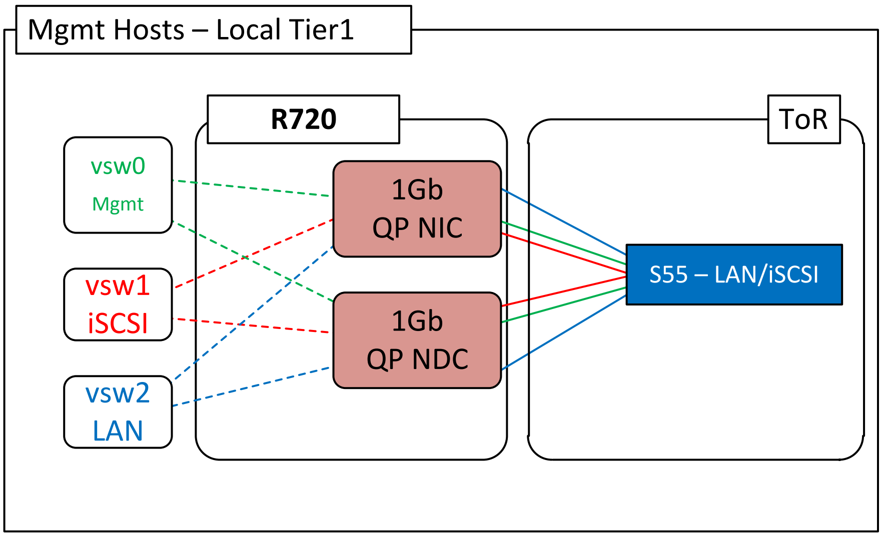

NIC teaming should be configured in the Hyper-V host using Dell drivers to ensure that the proper NICs from differing PCIe devices are bonded. The resulting teamed virtual NIC should then be assigned to the appropriate virtual switch within Hyper-V. VLAN IDs should be specified in all virtual switches used within the Management layer Hyper-V host. All NICs and switch ports should be set to auto negotiate.

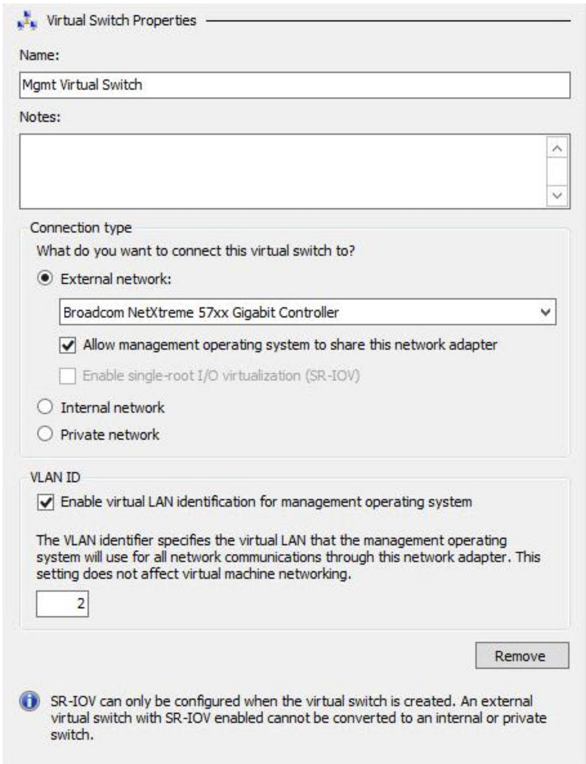

4.6 Solution High Availability

Each layer in the solution architecture can be individually protected to prevent an extended service outage. The Network layer only requires an additional switch configured in a stack with the first. Please refer to section 3.1.1.1 that covers Force10 switch stacking.

Protecting the Compute layer for RDSH and RDVH is provided by adding an additional host to a collection, thus effectively increasing the hosting capacity of a given collection. Session requests will be fulfilled by all hosts in the collection and as a result, each will have reserve capacity to insure against a host failure. **Care needs to be taken to ensure that user provisioning does not exceed the overflow capacity provided by the additional node.** A simple fail-safe measure would be to ensure that the appropriate number of users entitled to connect to the environment be tightly controlled via Active Directory. In a failure scenario users working on a failed host would simply reconnect to a fresh session on a surviving Compute host.

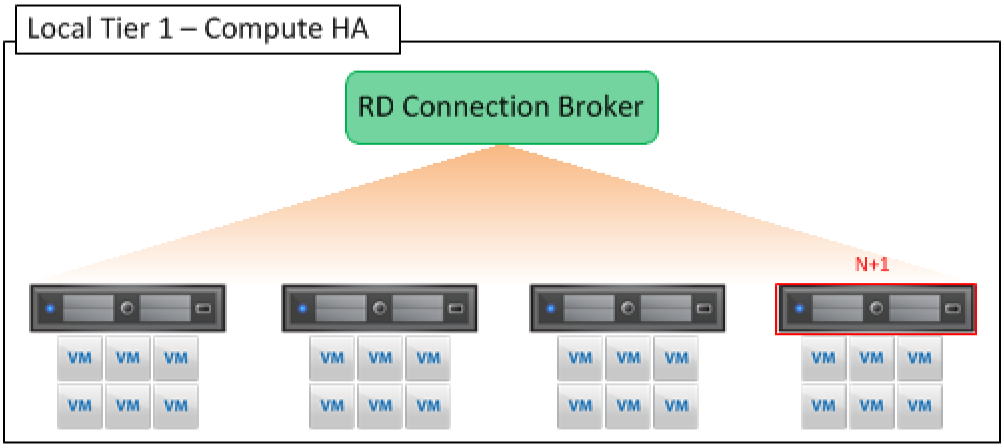

To implement HA for the Management layer, we will also add an additional host but will add a few more layers of redundancy. The following will protect each of the critical infrastructure components in the solution:

- The Management hosts will be configured in a Hyper-V cluster (Node and Disk Majority).
- The storage volume that hosts the Management VMs will be upgraded to a Cluster Shared Volume (CSV).
- SQL Server will be added to the environment to support RD Connection Broker HA.
    - Optionally SQL mirroring can be configured to further protect SQL.
- The RD Connection Broker will be configured for HA with a second RDCB introduced.

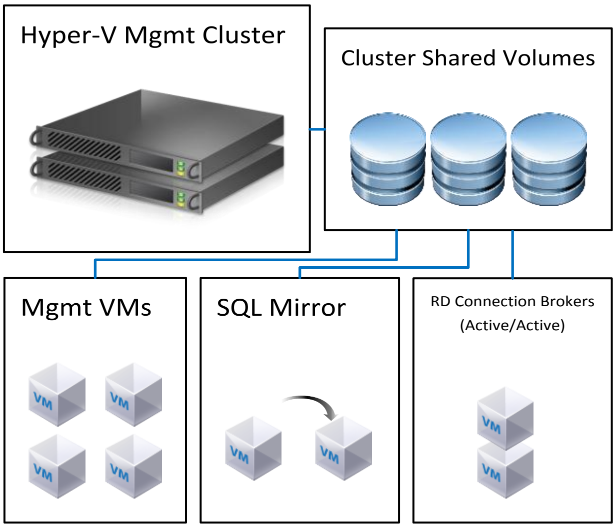

The following storage volumes are applicable in a 2-node Management layer HA scenario:

Volumes	Host	Size (GB)	RAID	Storage Array	Purpose	File System	CSV
Management	1	500	50	Tier 2	RDS VMs, File Server	NTFS	Yes
Management	2	500	50	Tier 2	RDS VMs, File Server	NTFS	Yes
SQL Data	2	100	50	Tier 2	SQL Data Disk	NTFS	Yes
SQL Logs	2	100	50	Tier 2	SQL Logs Disk	NTFS	Yes
SQL TempDB Data	2	5	50	Tier 2	SQL TempDB Data Disk	NTFS	Yes
SQL TempDB Logs	2	5	50	Tier 2	SQL TempDB Logs Disk	NTFS	Yes
SQL Witness	1	1	50	Tier 2	SQL Witness Disk	NTFS	Yes
Quorum 1	-	500MB	50	Tier 2	Hyper-V Cluster Quorum	NTFS	Yes
User Data	-	2048	50	Tier 2	File Server	NTFS	No
User Profiles	-	20	50	Tier 2	User profiles	NTFS	No
Templates/ ISO	-	200	50	Tier 2	ISO/ gold image storage (optional)	NTFS	Yes

For more information on building redundancy for the RD Connection Broker please visit: LINK

4.6.1 SQL Database HA

HA for SQL will be provided via an optional 3-server synchronous mirror configuration that includes a witness (High safety with automatic failover). This configuration will protect all critical data stored within the database from physical server as well as virtual server problems. The principal VM that will host the primary copy of the data should exist on the first Mgmt host. The mirror and witness VMs should exist on the second or later Mgmt hosts. All critical databases should be mirrored to provide HA protection.

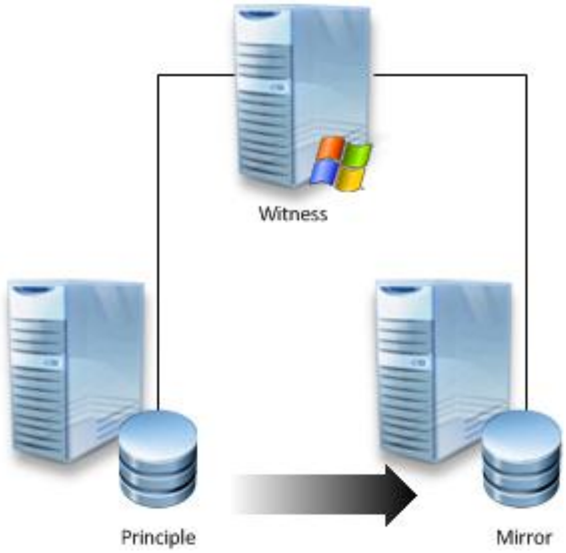

There are a number of steps required to successfully set up SQL mirroring per Microsoft best practices.

The following article details the step-by-step mirror configuration: LINK

Additional resources can be found in TechNet: LINK1 and LINK2

4.7 Application Virtualization

Microsoft Application Virtualization (App-V) provides multiple methods to deliver virtualized applications to RDS environments, virtual desktops, physical desktops, connected as well as disconnected clients. App-V can help reduce the costs and time associated with managing gold master VM and PC images with integrated applications. App-V also removes the problems of application conflicts since virtualized applications are never installed on an end point. Once an application has been packaged using the Microsoft Application Virtualization Sequencer, it can be saved to removable media, streamed to desktop clients or presented to session-based users on a RDSH host. App-V provides a scalable framework that can be managed by System Center Configuration Manager for a complete management solution.

To learn more about application virtualization and how it integrates into a RDS environment please visit: LINK

4.8 Disaster Recovery and Business Continuity

DR and BC can be achieved natively via Hyper-V Replicas. This technology can be used to replicate VMs from a primary site to a DR or BC site over the WAN asynchronously. Hyper-V Replicas are unbiased as to underlying hardware platform and can be replicated to any server, network, or storage provider. Once the initial replica is delivered from the primary site to the replica site, incremental VM write changes are replicated using log file updates. Multiple recovery points can be stored and maintained, using snapshots, to restore a VM to a specific point in time.

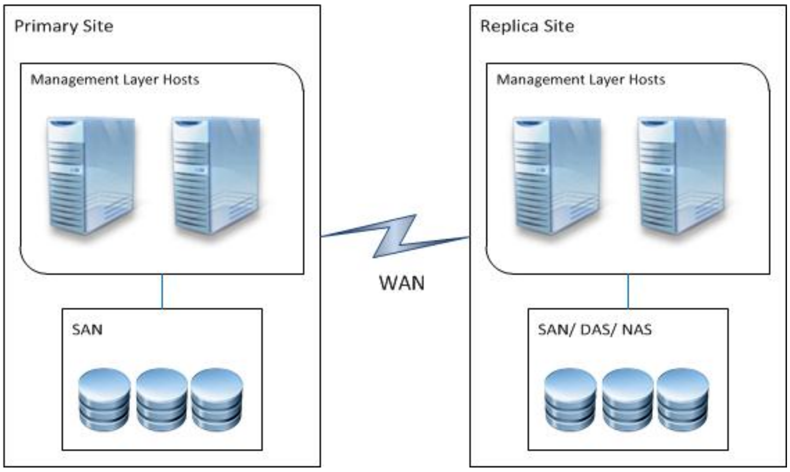
Primary Site
Management Layer Hosts
SAN
WAN
Replica Site
Management Layer Hosts
SAN/ DAS/ NAS

4.9 RDS 2012 Data Flow

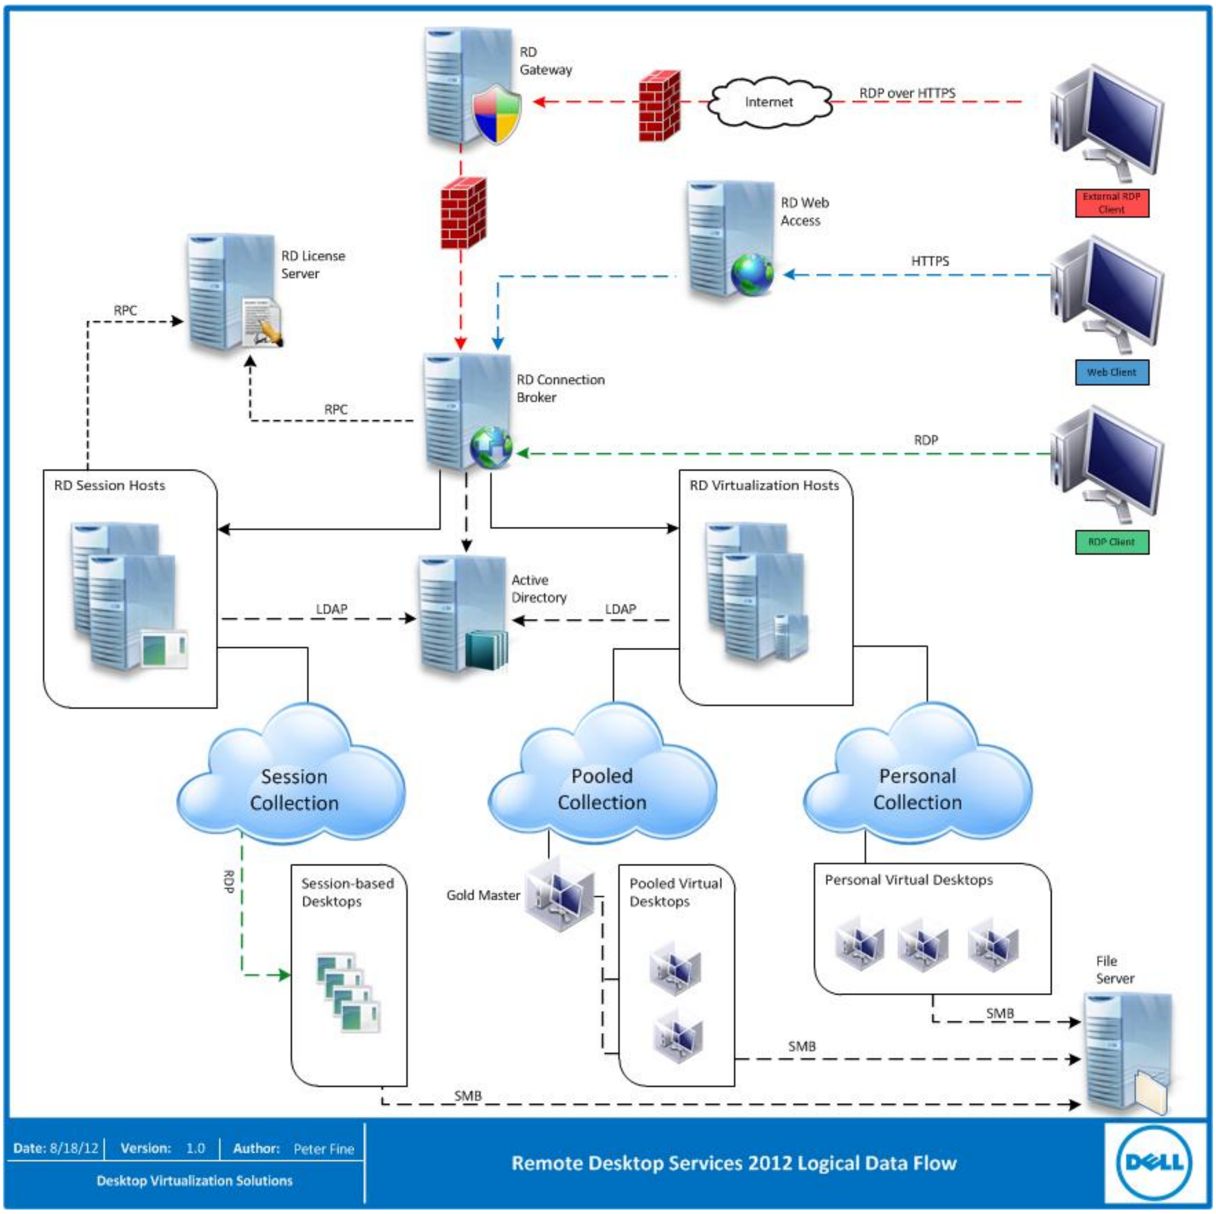

4.10 Summary

The Dell Windows Server 2012 RDS solution provides a robust and scalable VDI platform for pooled, personal and Session host deployments. Using VDI-optimized hardware in a configuration that has been validated and proven by Dell DVS Engineering, you can deploy Microsoft based VDI that is both cost effective and high performing. Our layered architecture provides flexibility to maximize your infrastructure investment with the capability to expand and contract where necessary.

Appendix A - Dell DVS 10-Seat Trial Kit

The 10 User POC bundle was purpose-built to provide high performance VDI using a modicum of infrastructure. Only 11 1Gb Ethernet ports are required (1 x server + 10 x end points) which can be provided using existing customer network infrastructure. If suitable network capacity is not in place, Dell recommends using a Force10 or PowerConnect model switch.

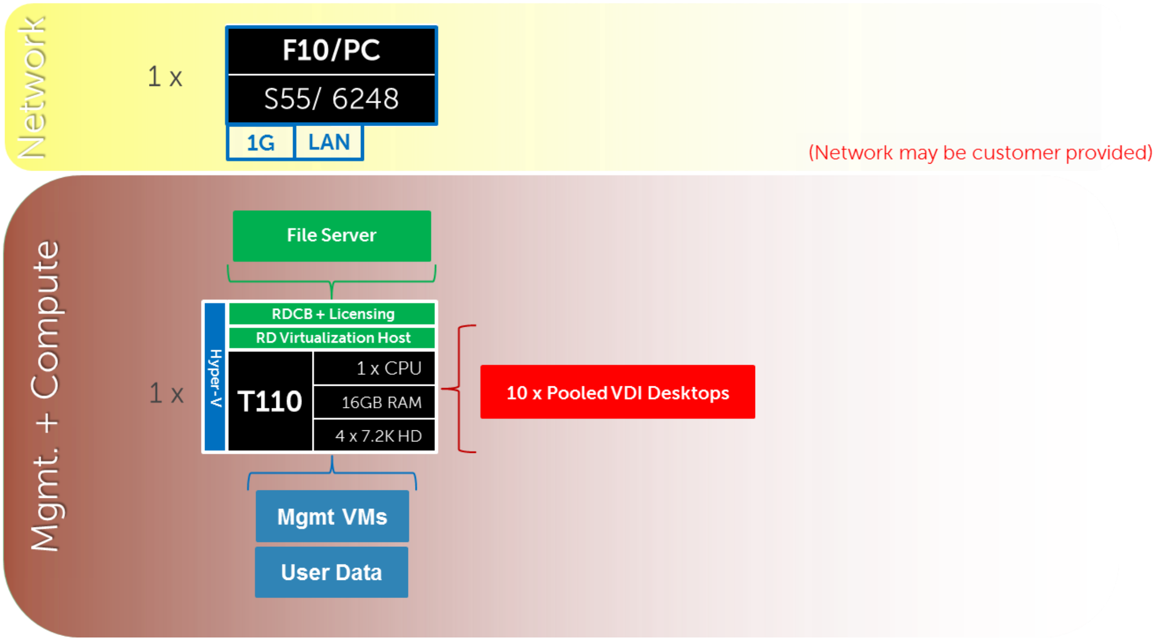

A.1 Server Configuration

The PowerEdge T110 II is the server platform of choice for this offering, providing high performance at an extremely low price of entry. Supporting the Intel Xeon E3-1200 series of CPUs and up to 32GB RAM, the T110 provides a solid server platform to get started with VDI.

All VDI server roles and desktop sessions are hosted on a single server in this model so there is no need for external storage. Higher scale and HA options are not offered with this bundle.

10 User Compute Host – PowerEdge T110 II
1 x Intel Xeon E3-1220 V2 (3.1Ghz)
16GB Memory (4 x 4GB DIMMs @ 1333Mhz) (**RDVH**)
Microsoft Windows Server 2012 Hyper-V
4 x 500GB SATA 7.2k Disks **RAID 10** (OS + VDI)
PERC H200 Integrated RAID Controller
Broadcom 5722 1Gb NIC (LAN)
305W PSU

Based on the server hardware configuration, 10 users will experience excellent performance with additional resource headroom available in reserve. The consumption numbers below are based on average performance:

Task Worker Users	CPU (%)	RAM (GB Consumed)	Disk (IOPS)	Network (Kbps)
10	40	7.5	50	262

A.2 Management and Compute Infrastructure

The solution architecture for the 10 user POC bundle combines the Compute, Management, and Storage layers onto a single server-based platform. To maximize server resources, the connection broker and license server roles are enabled within the Hyper-V parent partition, while the File server and VDI sessions exist as VMs within child partitions.

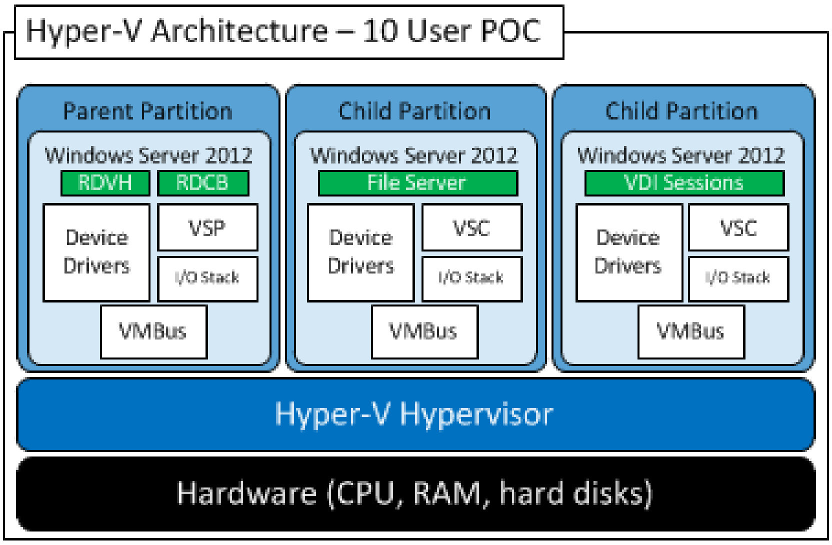

Since the RDCB and Licensing roles will be enabled within the Hyper-V parent partition, only the file server VM requires that specific physical resources be assigned.

Role	vCPU	Startup RAM (GB)	Dynamic Memory			NIC	OS + Data vDisk (GB)	Tier 2 Volume (GB)
			Min|Max	Buffer	Weight			
File Server	1	1	512MB|2GB	20%	Med	1	40 + 10	50
Pooled VDI VMs	1	512MB	512MB|2GB	20%	Med	1	20	-

A.3 Storage Configuration

The 10 User POC solution includes 4 total hard drives configured in RAID10 to host the Windows Server OS as well as VDI sessions. This configuration will maximize available performance and data protection.

Volumes	Size	RAID	Storage	Purpose	File System
OS + VDI	1TB	10	Tier 1	Host OS/ Mgmt roles + VDI Sessions	NTFS

Appendix B – Pooled VDI Performance Analysis Results

Performance analysis of the above architecture was carried out using Login VSI software. Login VSI is a widely used tool to generate workloads that are representative of typical corporate IT users of centralized desktop environments such as Server Based Computing (SBC) and Virtual Desktop Infrastructure (VDI). The workload produced by Login VSI for the current performance analysis effort was representative of a typical set of activities performed by (i) a task worker (the basic workload) and (ii) a knowledge worker (the standard workload). Resource utilisation on the compute node was monitored using Microsoft best practices for measuring performance on Hyper-V as detailed at

http://technet.microsoft.com/en-us/library/cc768535.aspx

In addition to the above, end-user experience was monitored using the Liquidware Labs Stratusphere UX tool. This tool provides comprehensive information (including reports and charts) for IT personnel in relation to end-user experience in a centralized desktop environment; among these charts is a "golden quadrant" type chart, which aggregates parameters that contribute to the end-user experience seen by a centralized desktop environment user into a single chart; this is the chart used during the current performance analysis activity.

For the task worker scenario, the performance analysis scenario used was to pre-boot all virtual desktops and then login 150 task worker (basic) workloads using a login interval of 30 seconds. Once all users have logged in, all 150 users run workload activities at steady-state for 15 minutes and then logoffs commence. For the knowledge worker scenario, a similar test methodology was used with the standard workload and 120 users.

It should be noted that in order to replicate a real corporate user environment, an enterprise-level anti-virus infrastructure was deployed, with McAfee VirusScan Enterprise 8.7 installed on all virtual desktops and McAfee ePolicy Orchestrator 4.5 used for management and deployment purposes.

B.1 Configuration Summary – Workload Independent Settings

The table shown below summarizes the important configuration information in relation to the pooled VDI environment used for performance analysis of both the task and knowledge worker scenarios.

	Hyper-V Compute Host
CPU Resource	32 logical cores (16 physical cores)
Memory (GB)	192
Page File	25.6GB (Default Setting)

B.2 Task Worker (Basic Workload)

The basic workload runs a small number of applications that are representative of applications used by task workers (e.g. call center). The applications are closed immediately after use, resulting in relatively low memory and CPU consumption when compared to the standard workload. The applications used are Internet Explorer, Microsoft Word and Microsoft Outlook, with only 2 of these applications used simultaneously. The user idle time is approximately 17% of total run time.

B.2.1 Configuration Summary – Task Worker

The table shown below gives the significant configuration parameters for the virtual desktop VM used for a task worker.

Configuration Parameter	Setting
Virtual Desktop Density	150 Virtual Desktops Per Host
Number of vCPUs per VM	1
Configured Memory on VM	1GB
Pagefile Size on VM	1.5GB

B.2.2 CPU Resource Utilization Performance Analysis Results

The CPU graphs shown below show logical and virtual processor utilization for the compute host during performance analysis. Hyper-V provides hypervisor performance objects to monitor the performance of both logical and virtual processors. A logical processor correlates directly to the number of processors or cores that are installed on the physical computer. For example, 2 quad core processors installed on the physical computer would correlate to 8 logical processors. Virtual processors are what the virtual machines (VMs) actually use, and all execution in the root and child partitions occurs in virtual processors. The results shown below show sustained logical processor % runtime peaking at approximately 80%, while sustained virtual processor % runtime peaks at approximately 17%. Logical processor % runtime is the key parameter for performance analysis of guest operating systems and a peak of 80% is ideal for an environment that is operating at the optimal combination of maximising density while providing sufficient headroom to ensure that end-user experience is not diminished. A high logical processor % runtime combined with a low virtual processor % runtime is typical of an environment where there are more processors allocated to VMs than are physically available on the compute host, which is the case for this VDI environment.

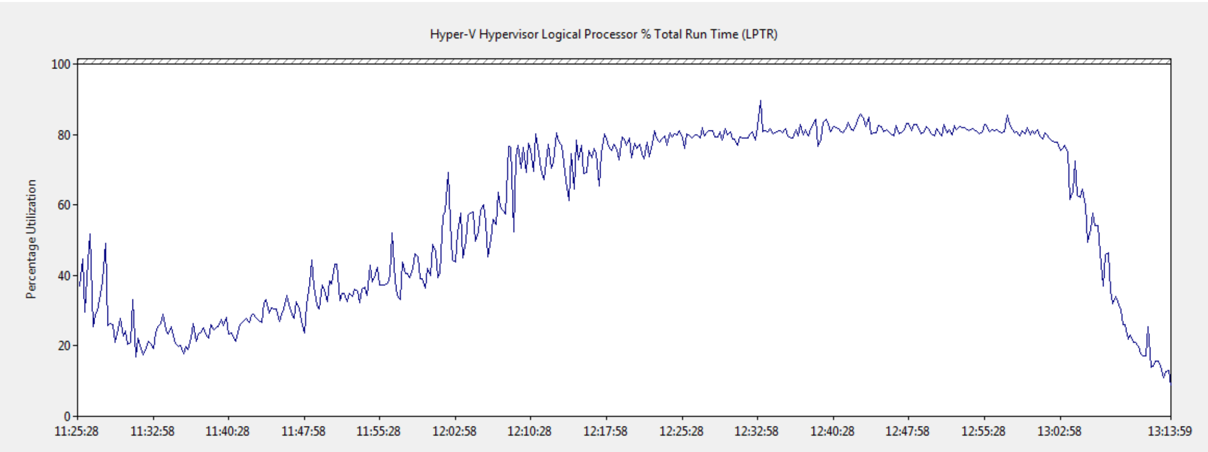

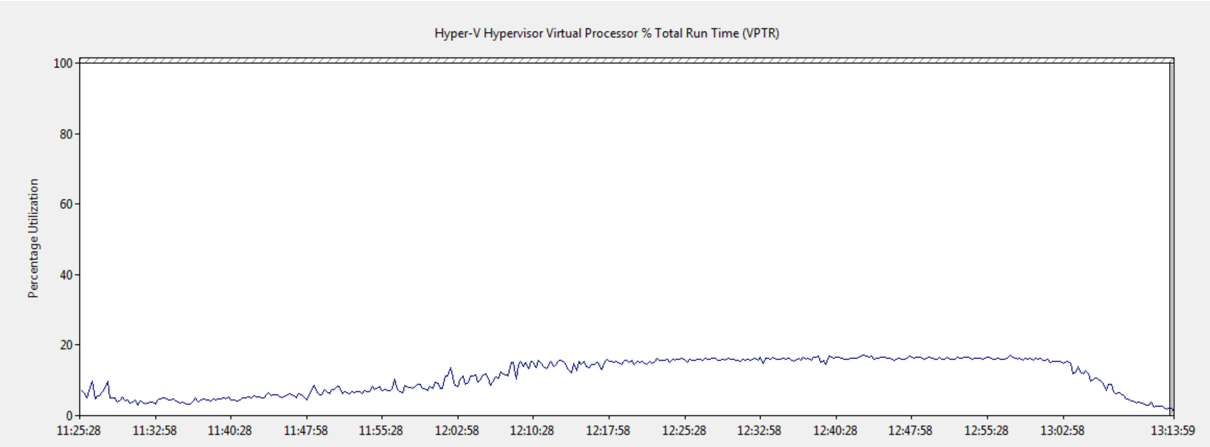

B.2.3 Memory Resource Utilization Performance Analysis Results

The memory graph shown below show available memory during performance analysis. It can be seen that available memory never drops below 75GB, which illustrates that significant memory headroom exists in the environment: maximum memory utilisation is 61%, since the total memory on the compute host is 192GB. The total memory configured on the VMs is approximately 150GB and the fact that memory utilisation never exceeds 117GB (192GB-75GB) demonstrates the efficiencies provided by Microsoft's dynamic memory technology.

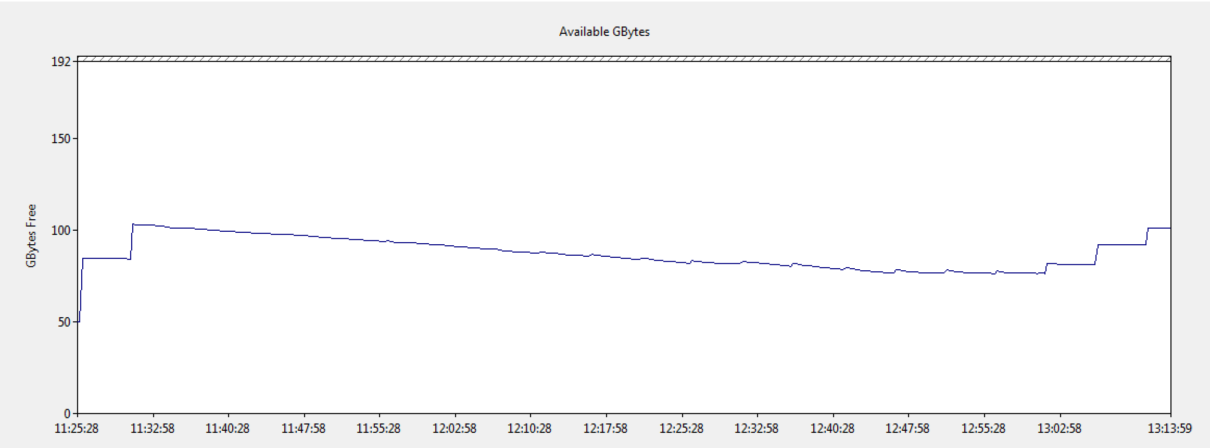

B.2.4 Disk IO Resource Utilization Performance Analysis Results

The Disk IO and latency graphs shown below are illustrative of a VDI environment that is performing to the expected levels for a task worker. The maximum sustained IOPS level reached is 1,500, which represents disk IO activity of 10 per user. Sustained disk IO latency never exceeds 5ms during validation, which is well within the healthy range of 1ms to 15ms specified by Microsoft in the TechNet best practices article described above. It should be noted that these disk IO figures are for the compute host D: drive, which is where the virtual desktops reside: disk activity on the C: drive is minimal.

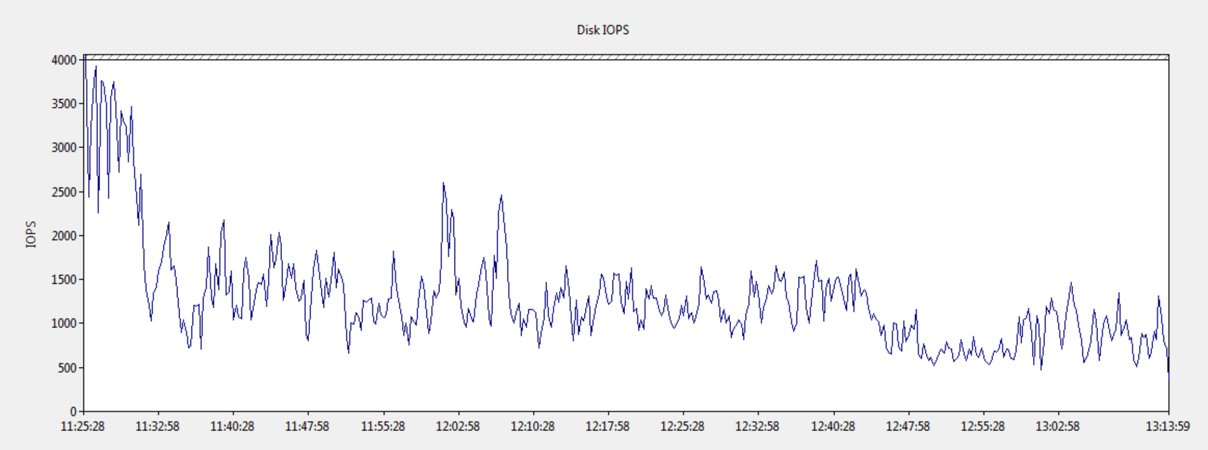

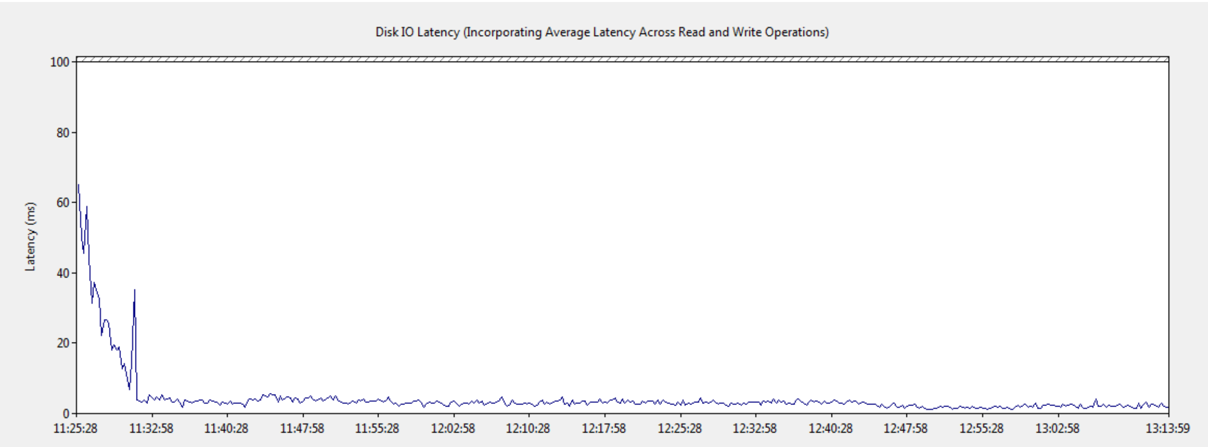

B.2.5 Network Resource Utilization Performance Analysis Results

Network resource utilization in the RDVH environment for a task user scenario was seen to be low. At the Hyper-V hypervisor level, network resource utilisation reaches a maximum of approximately 18 MB/s, representing network bandwidth utilization of approximately 14%, based on a 1Gb/s switching infrastructure. The hypervisor shows zero output queue length throughout the duration of the performance analysis activity. These results illustrate that network resource utilization is very low and does not approach a level at which it could become a bottleneck in the environment.

B.2.6 End-User Experience Performance Analysis Results

The Stratusphere UX scatter plot shown below shows all users in the "golden quadrant" that represents good user experience for VDI users; the positioning of users in this graph is determined by an aggregation of information such as user login time. It should be noted that in addition to the use of Stratusphere UX for analysis of End-User Experience, a real user logged in at the peak of resource utilisation during performance analysis and the user experience perceived by that user was good.

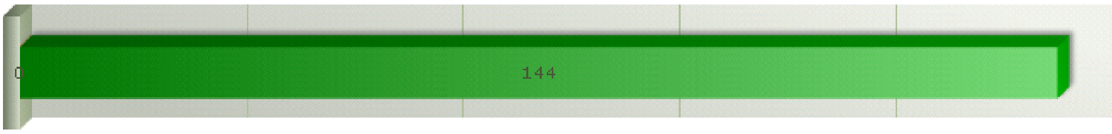

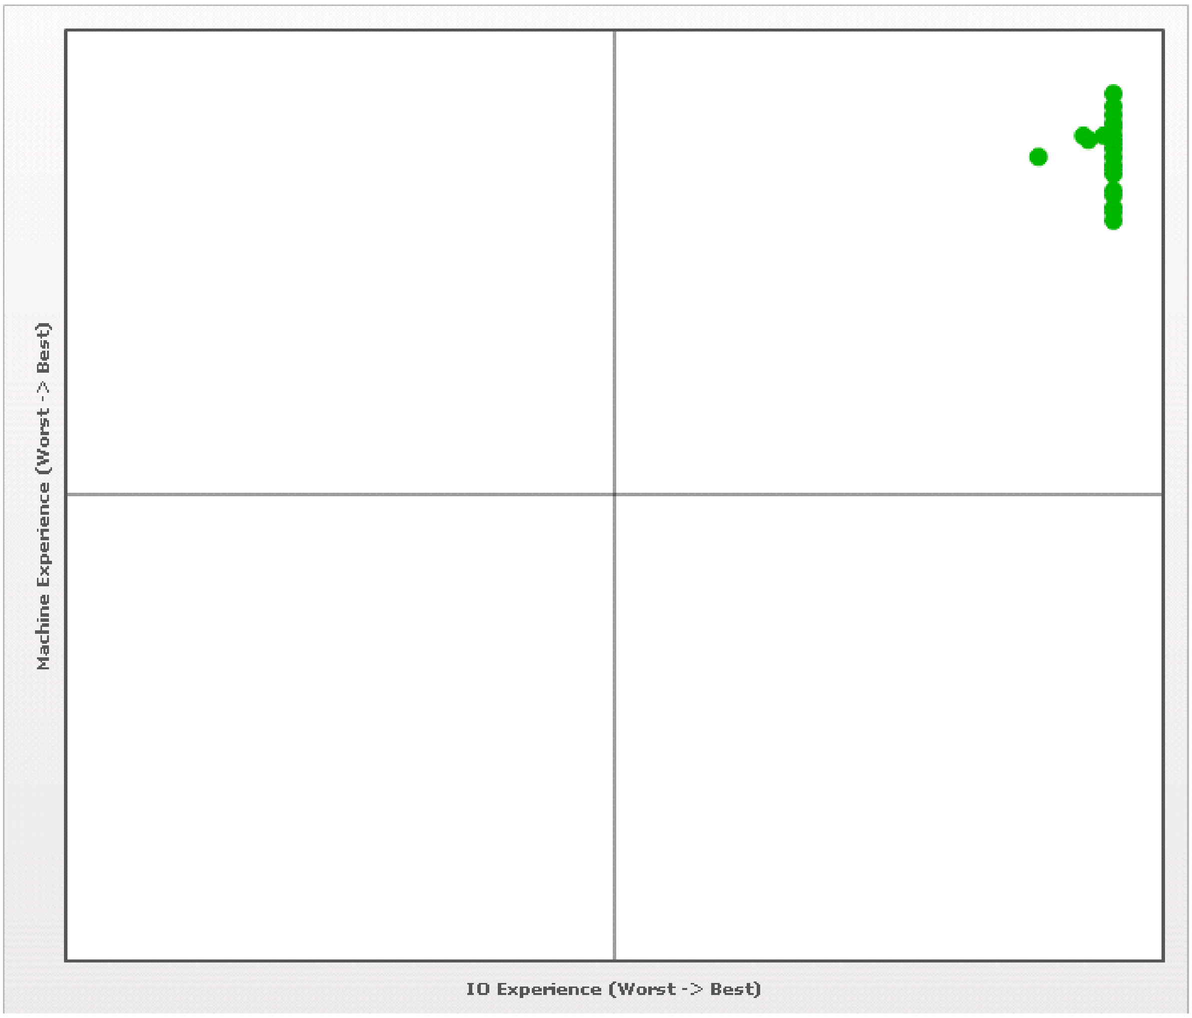

B.3 Knowledge Worker (Standard Workload)

The standard workload runs a number of applications that are representative of applications used by knowledge workers (e.g. accountants). The applications used are Internet Explorer, a number of Microsoft Office applications (Excel, Outlook, Powerpoint and Word), Adobe Acrobat Reader, Bullzip PDF printer and 7-zip file compression software. Relative to the task worker workload discussed above, idle time is slightly lower as a percentage of overall runtime and a maximum of 5 applications are open simultaneously (compared to 2 for the task worker).

B.3.1 Configuration Summary – Knowledge Worker

The table shown below gives the significant configuration parameters for the virtual desktop VM used for a knowledge worker.

Configuration Parameter	Setting

Virtual Desktop Density	120 Virtual Desktops Per Host
Number of vCPUs per VM	1
Configured Memory on VM	1.5GB
Pagefile Size on VM	2.3GB

B.3.2 CPU Resource Utilization Performance Analysis Results

The CPU graphs shown below show logical and virtual processor utilization during performance analysis. The results shown below show sustained logical processor % runtime peaking at approximately 83%, while sustained virtual processor % runtime peaks at approximately 22%. Sustained peak logical processor % runtime of 83% is ideal for an environment that is operating at the optimal combination of maximising density while providing sufficient headroom to ensure that end-user experience is not diminished. As discussed above, a high logical processor % runtime combined with a low virtual processor % runtime is typical of an environment where there are more processors allocated to VMs than are physically available on the compute host, which is the case for this VDI environment.

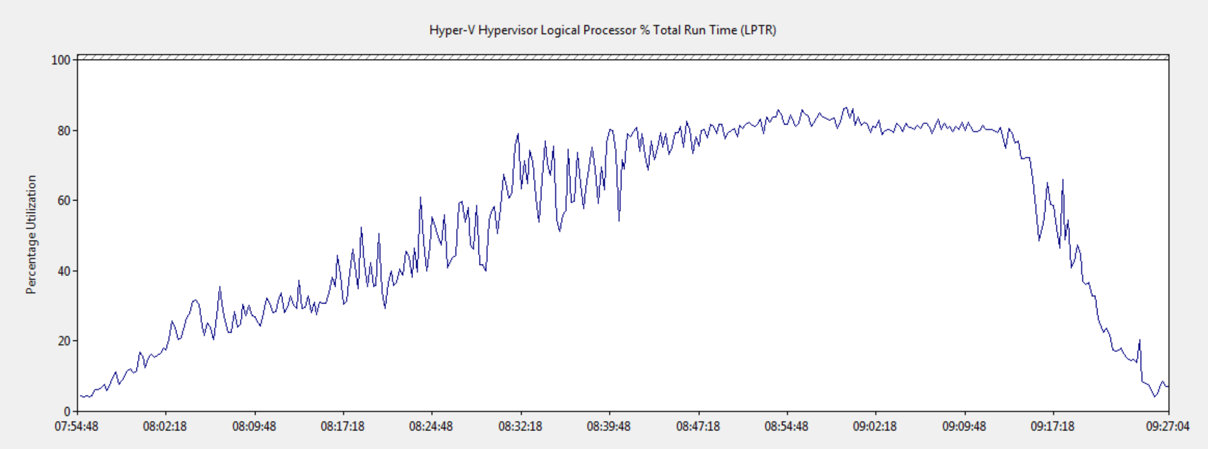

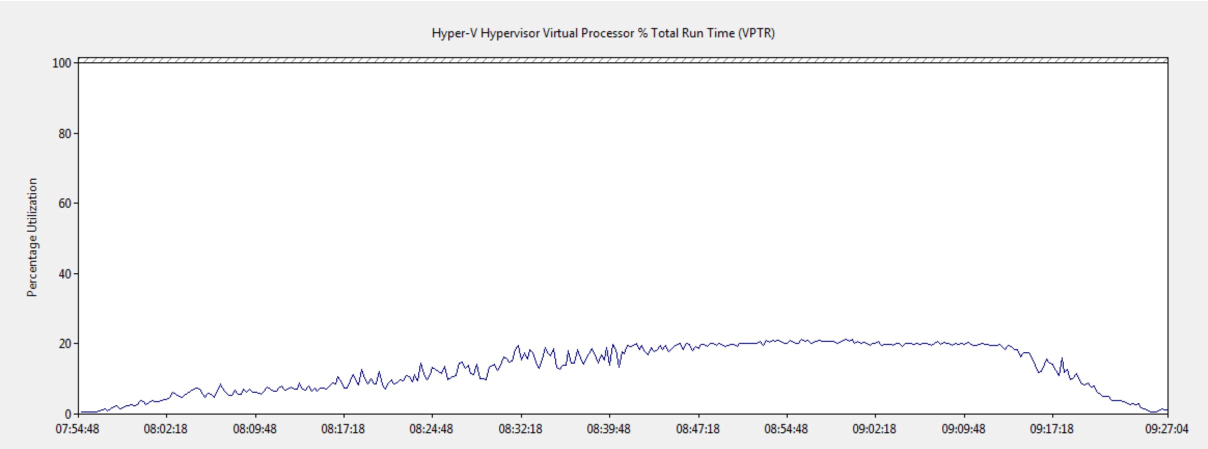

B.3.3 Memory Resource Utilization Performance Analysis Results

The memory graph shown below show available memory during performance analysis. It can be seen that available memory never drops below 85GB, which illustrates that significant memory headroom exists in the environment: maximum memory utilisation is 56%, since the total memory on the compute host is 192GB. The total memory configured on VMs is approximately 180GB and the fact that memory utilisation never exceeds 107GB (192GB-85GB) demonstrates the efficiencies provided by Microsoft's dynamic memory technology.

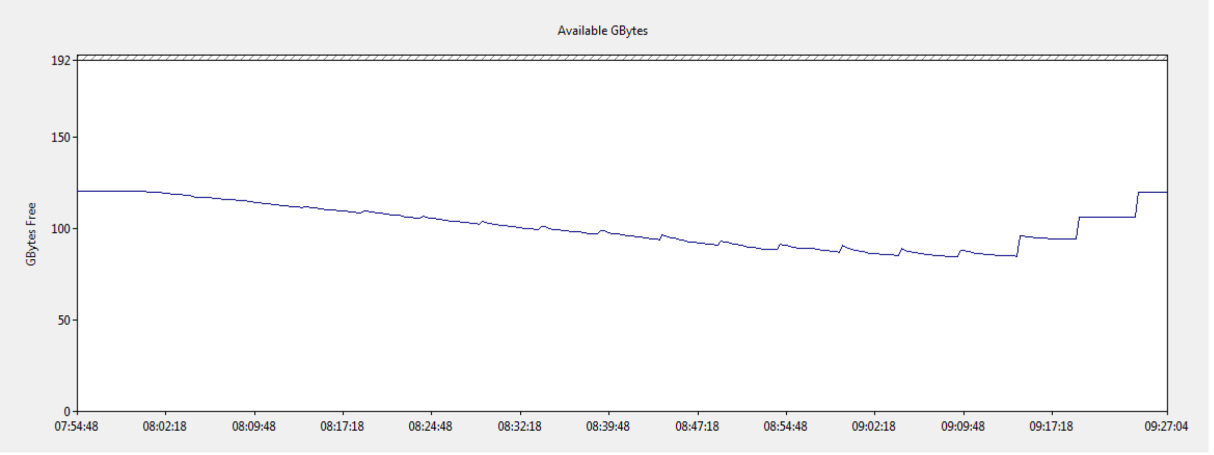

B.3.4 Disk IO Resource Utilization Performance Analysis Results

The Disk IO and latency graphs shown below are illustrative of a VDI environment that is performing to the expected levels for a task worker. The maximum sustained IOPS level reached is 1550, which represents disk IO activity of 13 per user. Sustained disk IO latency never exceeds 5ms during validation, which is well within the healthy range of 1ms to 15ms specified by Microsoft in the TechNet best practices article described above. It should be noted that these disk IO figures are for the compute host D: drive, which is where the virtual desktops reside: disk activity on the C: drive is minimal.

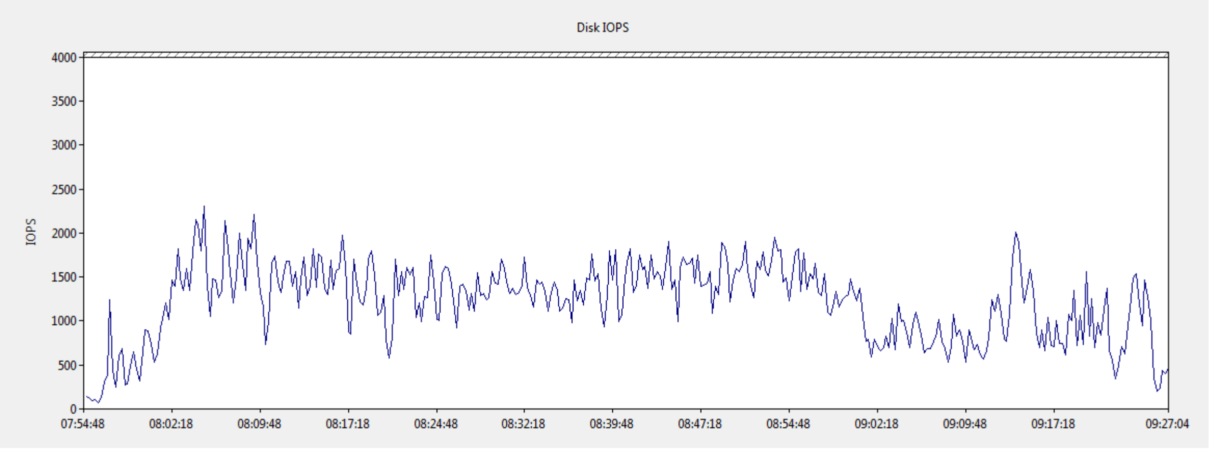

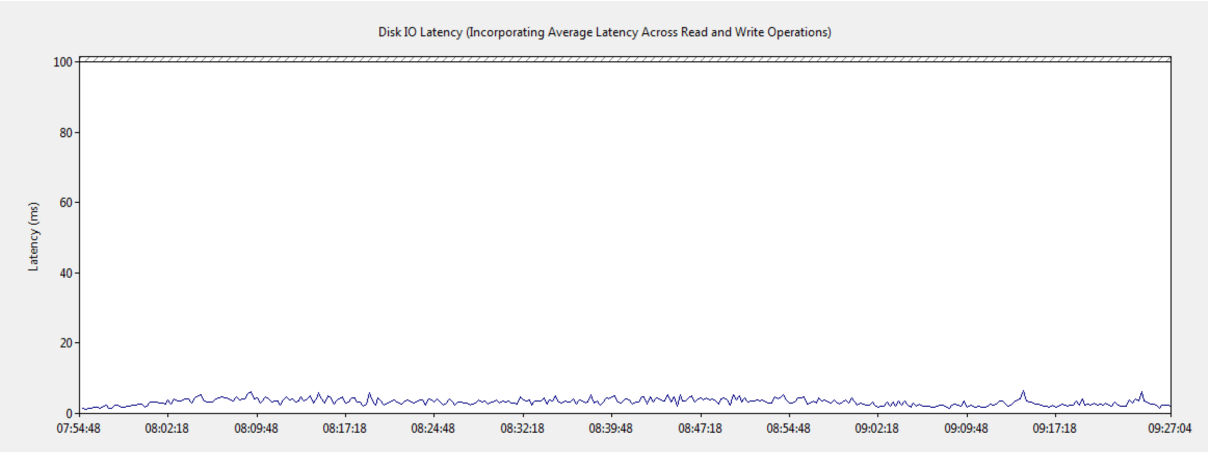

B.3.5 Network Resource Utilization Performance Analysis Results

Network resource utilization in the RDVH environment for a task user scenario was seen to be low. At the Hyper-V hypervisor level, network resource utilisation reaches a maximum of approximately 44 MB/s, representing network bandwidth utilization of approximately 35%, based on a 1Gb/s switching infrastructure. The hypervisor shows zero output queue length throughout the duration of the performance analysis activity. This level of network utilization results illustrate that network resource utilization is very low and does not approach a level at which it could become a bottleneck in the environment.

B.3.6 End-User Experience Performance Analysis Results

The Stratusphere UX scatter plot shown below shows all users in the "golden quadrant" that represents good user experience for VDI users; the positioning of users in this graph is determined by an aggregation of information such as user login time. It should be noted that in addition to the use of Stratusphere UX for analysis of End-User Experience, a real user logged in at the peak of resource utilisation during performance analysis and the user experience perceived by that user was good.

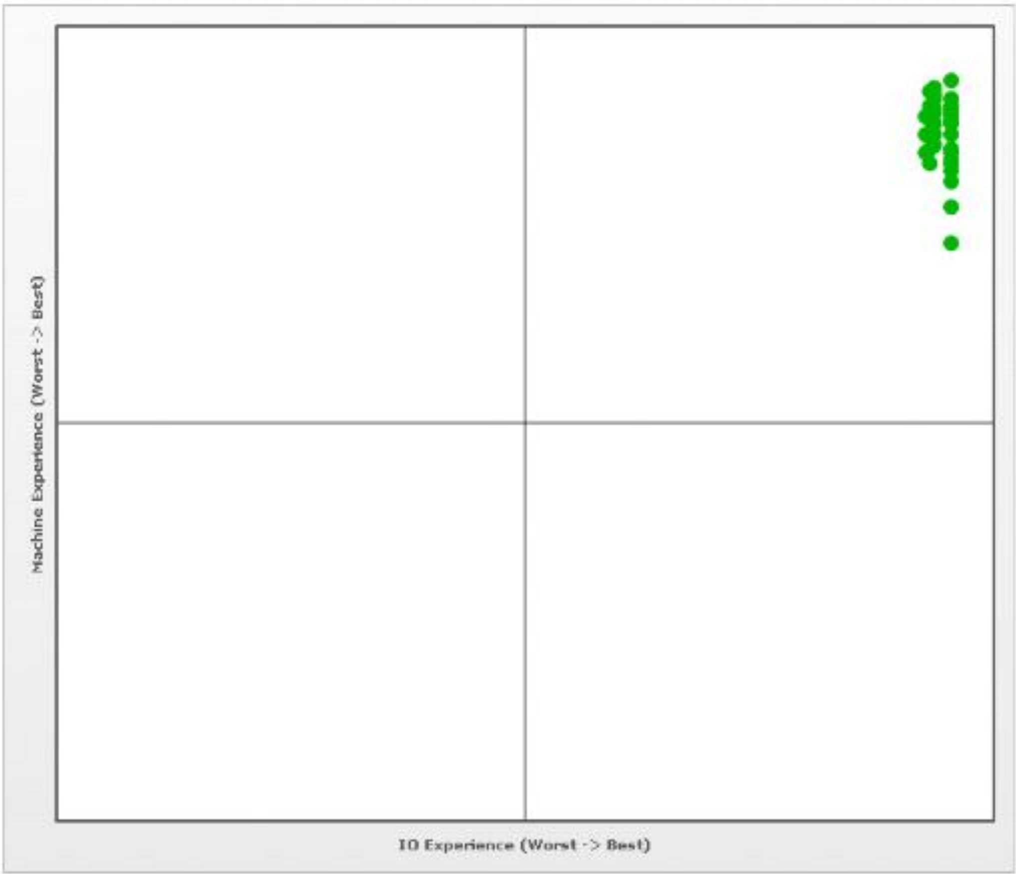

Machine Experience (Worst -> Best)
IO Experience (Worst -> Best)

Appendix C – Remote Desktop Session Performance Analysis Results

The performance of the Remote Desktop Session Host (RDSH) environment was analysed using an identical Login VSI test infrastructure to that described for pooled VDI in Appendix 1 above. Performance analysis was carried out using the perfmon tool and appropriate RDSH related perfmon parameters and the significant results arising from this performance analysis are shown below, illustrating the capacity of the solution to support 260 task workers or 200 knowledge workers. In order to replicate a real corporate user environment, an enterprise-level anti-virus infrastructure was deployed, with McAfee VirusScan Enterprise 8.7 installed on all virtual desktops and McAfee ePolicy Orchestrator 4.5 used for management and deployment purposes. For both basic and standard workloads, it should be noted that in addition to results shown below, a real user logged in at the peak of resource utilisation during performance analysis and the user experience perceived by that user was good.

For each set of performance analysis results shown below (i.e. CPU, Memory, Disk IO and Network), resource utilization graphs are presented for both the hypervisor and a representative VM with the RDSH role. As discussed above, when maximally configured (i.e. with 260 task workers or 200 knowledge workers), each RDS collection uses 4 VMs with the RDSH role; the use of 4 RDSH VMs allows scaling up of the RDSH environment from a small number of users with a single VM to maximal load with 4 VMs. The use of 4 RDSH VMs also allows equal distribution of the 16 physical cores (32 Hyper-Threaded logical cores) between RDSH VMs.

C.1 Impact of Hyper-Threading

Hyper-Threading technology (which is activated in Dell servers using the "Logical Processor" functionality in the BIOS) allows 2 logical cores to be addressed for each physical core that is actually present, thereby facilitating a doubling of the amount of processes that can be scheduled by an OS. It is instructive to assess the benefits of using Hyper-Threading in any infrastructure deployment in order to ensure that the environment is optimally configured.

For the purposes of performance analysis of the current RDSH environment, a particular configuration (i.e. a particular user density) was tested with Hyper-Threading enabled and Hyper-Threading disabled. It should be noted that the user density used for this comparison activity was not the optimal user density described in subsequent sections; hence CPU resource utilization shown below is larger than the normally acceptable resource utilization thresholds.

C.1.1 CPU Resource Utilization Performance Analysis Results

CPU utilization for a configuration without Hyper-Threading enabled is shown below; the graph shown shows the CPU utilization parameter "Hyper-V Hypervisor Logical Processor % Total Run Time." It can be seen that utilization plateaus at 100% for approximately 45 minutes.

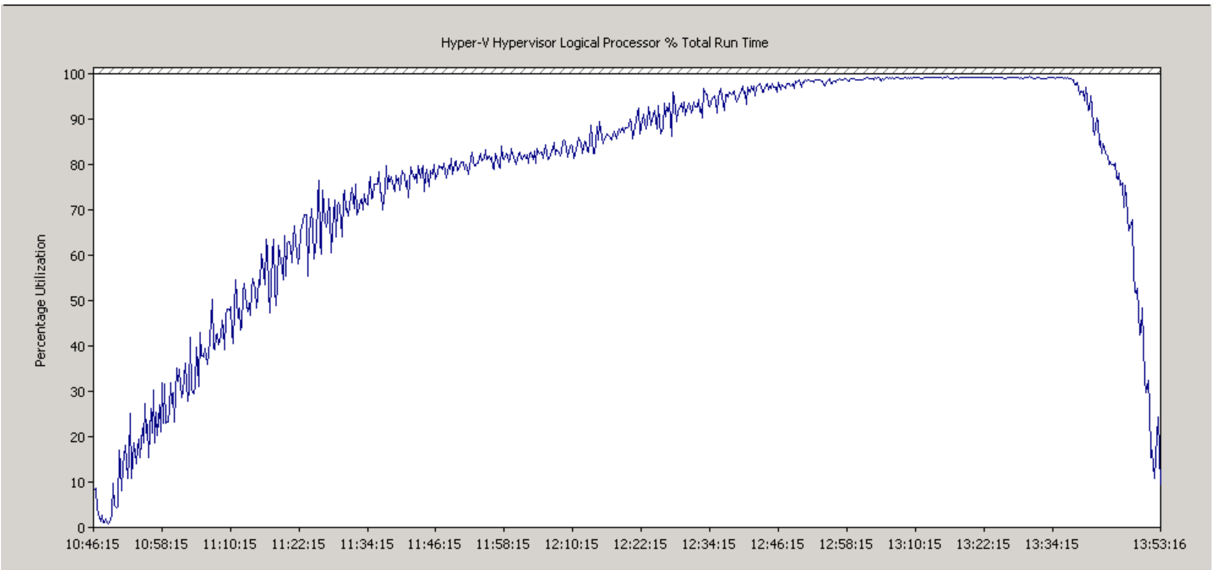

A plot of the "Hyper-V Hypervisor Logical Processor % Total Run Time" CPU utilization parameter for a configuration with Hyper-Threading enabled is shown below. It can be seen that utilization plateaus at 50% for approximately 45 minutes.

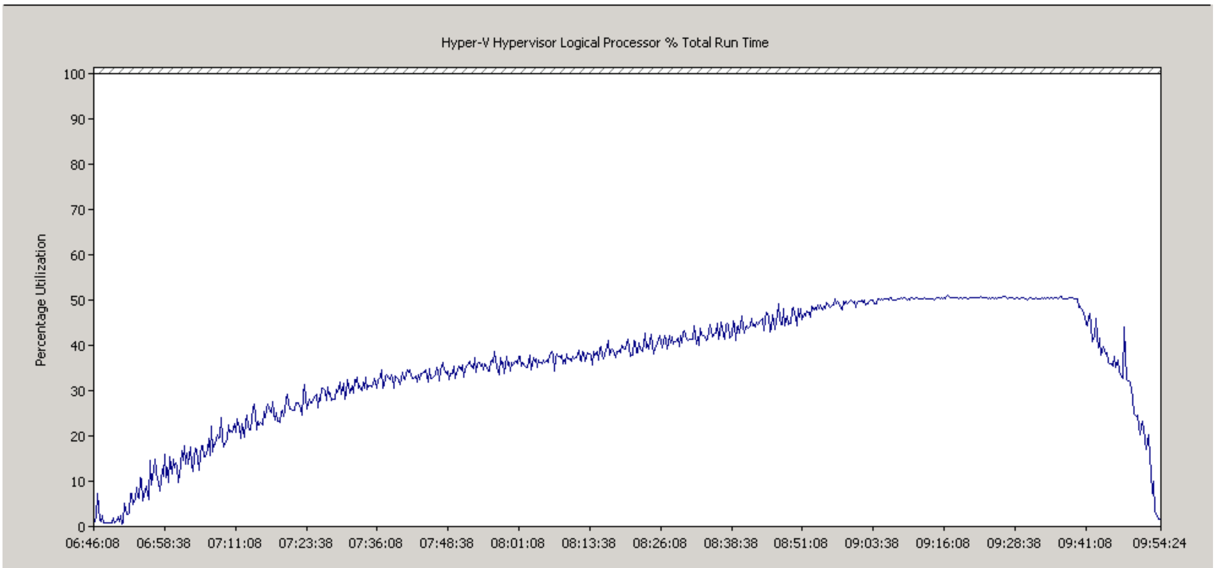

Comparing the "without Hyper-Threading" and the "with Hyper-Threading" scenarios above, the 50% level reached in the with Hyper-Threading scenario suggests possible spare capacity, with the fact that the possibility of spare capacity is combined with plateauing at 50% suggesting that the under-utilization is resulting from only half of the logical processors being used (since the 4 * RDSH VMs were configured with a total of 16 virtual processors versus a total of 32 available logical processors). Consequently, it was decided to use 8 virtual processors per RDSH VM, rather than 4 virtual processors per RDSH VM in order to more fully utilize the available logical processors. This 8 CPU configuration resulted in reduced CPU utilization at the RDSH VM level and was used for all RDSH VMs in the final design used in this reference architecture.

C.2 Task Worker (Basic Workload)

C.2.1 Configuration Summary

The table shown below summarizes the important configuration information in relation to the task worker RDSH environment.

	Hyper-V Compute Host	RDSH VM1	RDSH VM2	RDSH VM3	RDSH VM4
CPU Resource	32 logical cores (16 physical cores)	8 vCPUs	8 vCPUs	8 vCPUs	8 vCPUs
Memory (GB)	96	16	16	16	16
RDSH Session Capacity	260	65	65	65	65

C.2.2 CPU Resource Utilization Performance Analysis Results

The Hyper-V CPU utilization graph shown below shows sustained CPU utilization peaking at approximately 73%, illustrating the capacity of the solution to host 260 users with headroom for additional spikes in user activity. It should be noted that only the "Hyper-V Hypervisor Logical Processor % Total Run Time" is shown but the "Hyper-V Hypervisor Virtual Processor % Total Run Time" graph is very similar, peaking at 72%; a close match between these parameters is indicative of a well-tuned environment.

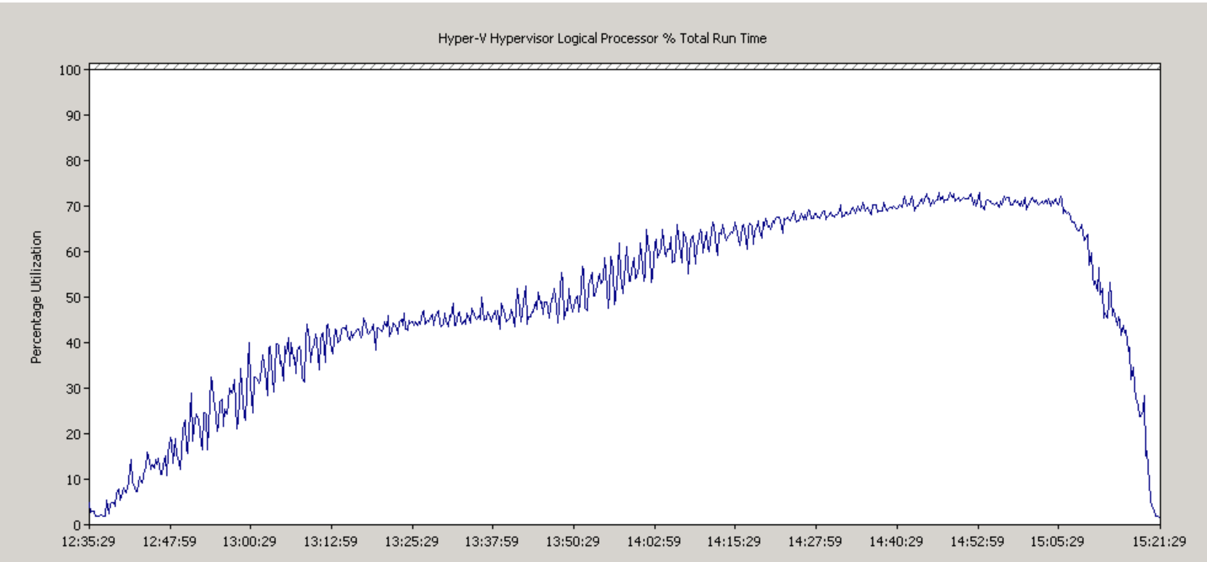

The table shown below shows sustained CPU utilization level at peak load for each of the RDSH VMs. It can be seen that sustained CPU utilization level lies between 82 and 86% for all RDSH VMs; this is a level that represents an ideal combination of maximizing resource utilization while also allowing headroom for spikes in user activity.

	Sustained CPU at Peak Load (%)
RDSH VM1	84
RDSH VM2	86
RDSH VM3	82
RDSH VM4	85

C.2.3 Memory Resource Utilization Performance Analysis Results

The memory graph shown below shows available memory on the Hyper-V host during performance analysis. It can be seen that available memory never drops below approximately 12GB at the hypervisor level, which illustrates that sufficient memory headroom exists in the environment: maximum memory utilization is approximately 88%, since the total memory on the compute host is 96GB.

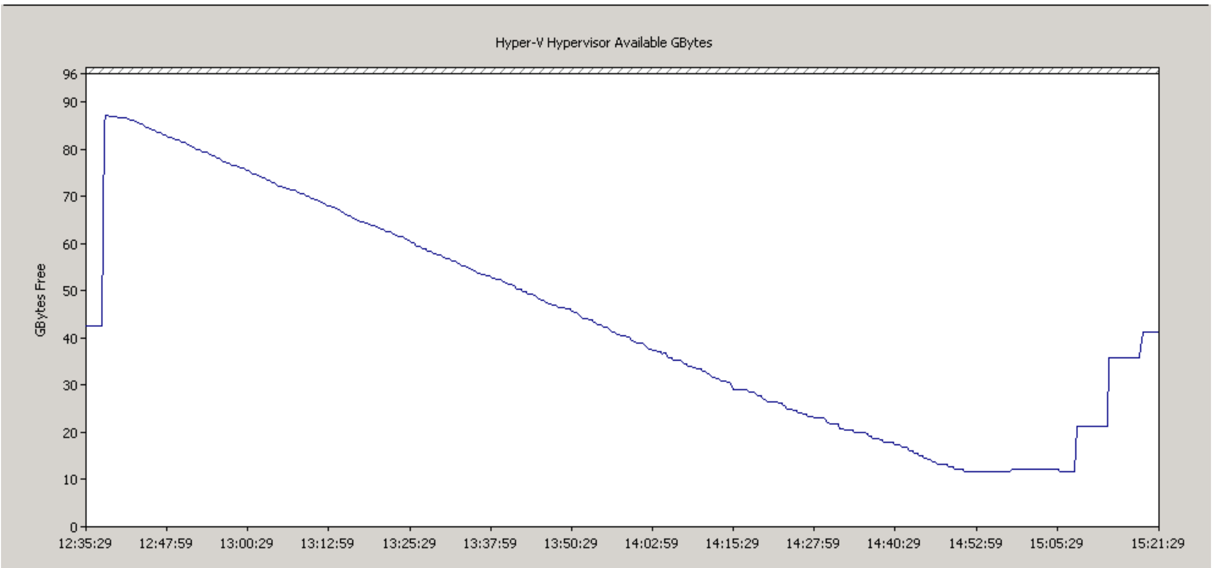

The table shown below shows available memory on the RDSH VMs during performance analysis. It can be seen that available memory at peak load across the 4 VMs doesn't drop below 3.3GB, which illustrates that sufficient memory headroom exists in the environment: maximum memory utilization is approximately 79%, since the total memory on the each VM is 16GB.

	Available Memory at Peak Load (%)
RDSH VM1	21
RDSH VM2	24
RDSH VM3	26

RDSH VM4	25

C.2.4 Disk IO Resource Utilization Performance Analysis Results

The graph shown below shows disk IO on the D: partition of the Windows Server 2012 system that hosts the Hyper-V hypervisor (this is the partition where the RDSH VMs reside). Sustained IOPS at peak load on the hypervisor partition reached a level of less than 100.

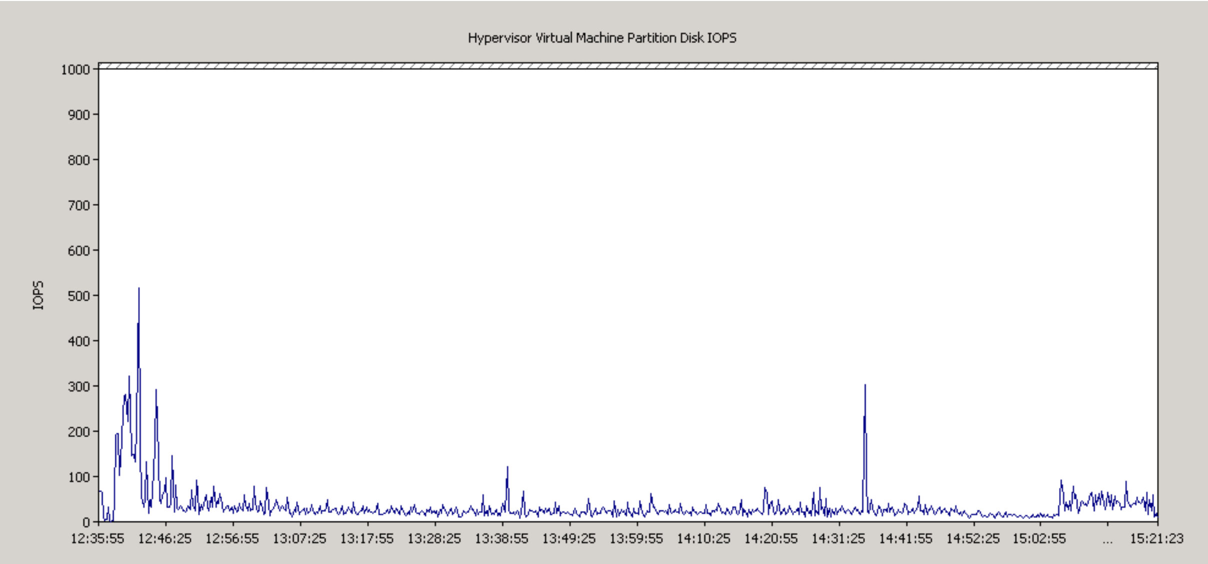

The table shown below shows sustained IOPS at peak load on the RDSH VMs during performance analysis. Peak utilization of 140 IOPS (on RDSH VM1) corresponds with approximately 2.2 IOPS per user.

	Sustained IOPS at Peak Load
RDSH VM1	140
RDSH VM2	137
RDSH VM3	135
RDSH VM4	130

The relevant Hyper-V partition latency information is shown below, it can be seen that latency remains significantly below the 15ms threshold specified by Microsoft in the TechNet best practices article described above for the entire duration of the test.

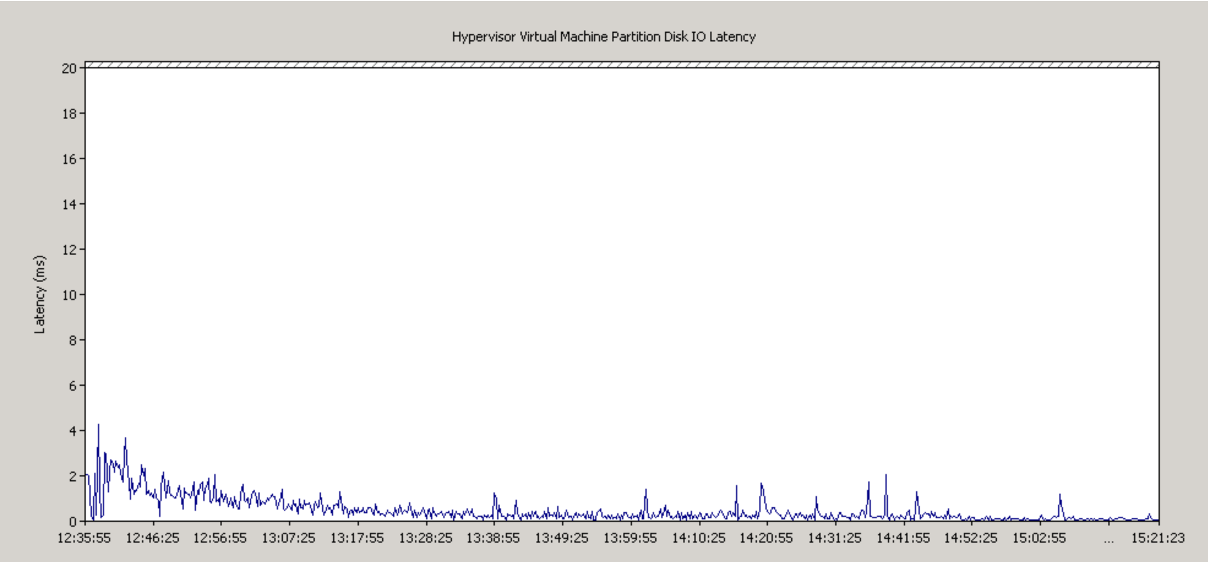

The table shown below shows sustained latency during performance analysis for each of the RDSH VMs; it can be seen that this doesn't exceed 4ms at any stage during performance analysis, indicating that disk IO latency is not a performance bottleneck.

	Maximum Sustained Latency During Performance Analysis (ms)
RDSH VM1	4
RDSH VM2	3
RDSH VM3	4
RDSH VM4	4

C.2.5 Network Resource Utilization Performance Analysis Results

Network resource utilization in the RDSH environment was seen to be low. At the Hyper-V hypervisor level, network resource utilisation reaches a maximum of approximately 26 MB/s, representing network bandwidth utilization of approximately 21%, based on a 1Gb/s switching infrastructure. The hypervisor shows zero output queue length throughout the duration of the performance analysis activity. Individual RDSH VM network utilization is consistent with this overall hypervisor network utilisation (i.e. total VM network utilization equals hypervisor network utilization) and all of the VMs also show zero output queue length throughout the duration of performance analysis. These hypervisor and RDSH VM network utilization results illustrate that network resource utilization is very low and does not approach a level at which it could become a bottleneck in the environment.

C.3 Knowledge Worker (Standard Workload)

C.3.1 Configuration Summary

The table shown below summarizes the important configuration information in relation to the knowledge worker RDSH environment.

	Hyper-V Host	RDSH VM1	RDSH VM2	RDSH VM3	RDSH VM4
CPU Resource	32 logical cores (16 physical cores)	8 vCPUs	8 vCPUs	8 vCPUs	8 vCPUs
Memory (GB)	96	16	16	16	16
RDSH Session Capacity	200	50	50	50	50

C.3.2 CPU Resource Utilization Performance Analysis Results

The Hyper-V CPU utilization graph shown below shows sustained CPU utilization peaking at approximately 72%, illustrating the capacity of the solution to host 200 users with headroom for additional spikes in user activity. It should be noted that only the "Hyper-V Hypervisor Logical Processor % Total Run Time" is shown but the "Hyper-V Hypervisor Virtual Processor % Total Run Time" graph is very similar, peaking at 71%; a close match between these parameters is indicative of a well-tuned environment.

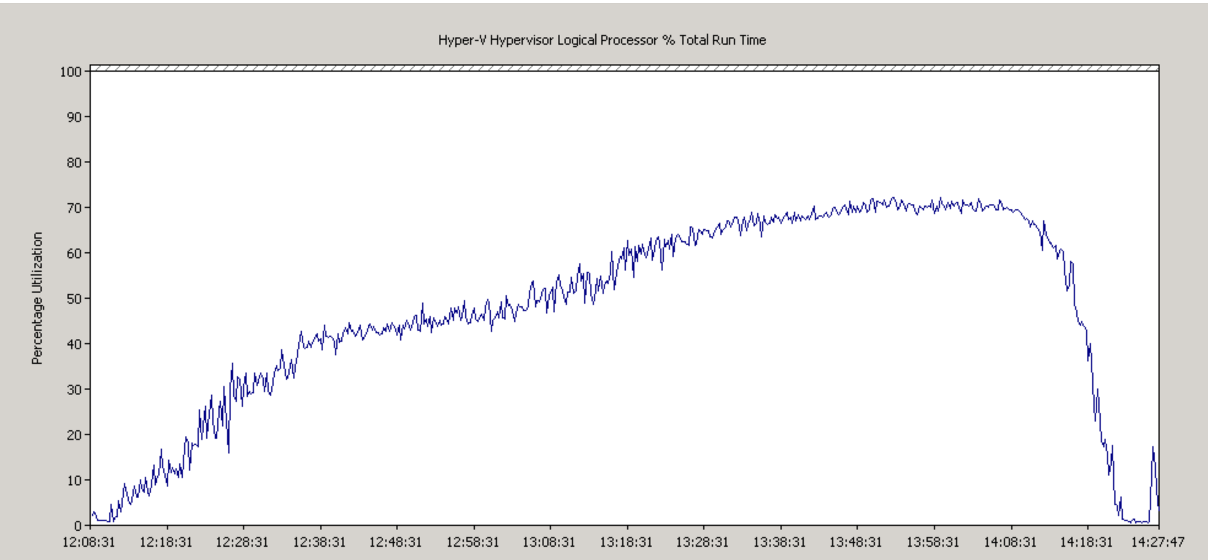

The table shown below shows sustained CPU utilization level at peak load for each of the RDSH VMs. It can be seen that sustained CPU utilization level lies between 85 and 88% for all RDSH VMs; this is a level that represents an ideal combination of maximizing resource utilization while also allowing headroom for spikes in user activity.

	Sustained CPU at Peak Load (%)

RDSH VM1	88
RDSH VM2	88
RDSH VM3	86
RDSH VM4	85

C.3.3 Memory Resource Utilization Performance Analysis Results

The memory graph shown below shows available memory during performance analysis at the Hyper-V hypervisor level. It can be seen that available memory never drops below approximately 16GB at the hypervisor level, which illustrates that sufficient memory headroom exists in the environment: maximum memory utilisation is approximately 83%, since the total memory on the compute host is 96GB.

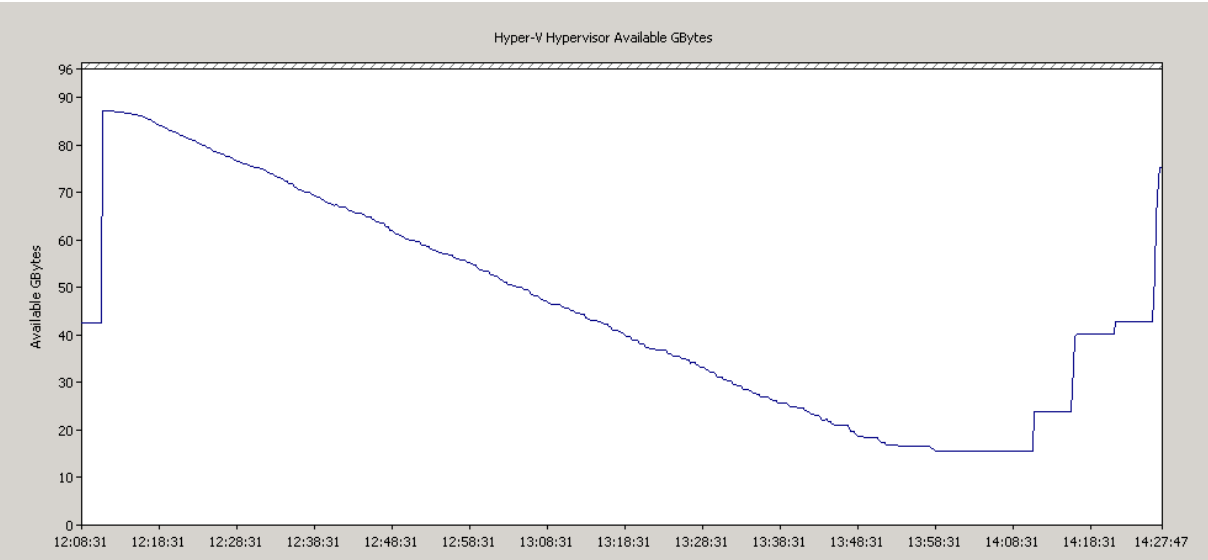

The table shown below shows available memory on the RDSH VMs during performance analysis. It can be seen that available memory at peak load across the 4 VMs doesn't drop below 3.4GB, which illustrates that sufficient memory headroom exists in the environment: maximum memory utilization is approximately 78%, since the total memory on each VM is 16GB.

	Available Memory at Peak Load GB)
RDSH VM1	41
RDSH VM2	23
RDSH VM3	21
RDSH VM4	28

C.3.4 Disk IO Resource Utilization Performance Analysis Results

The graph shown below shows disk IO on the D: partition of the Windows Server 2012 system that hosts the Hyper-V hypervisor (this is the partition where the RDSH VMs reside). Sustained IOPS at peak load on the hypervisor partition reached a level of less than 100.

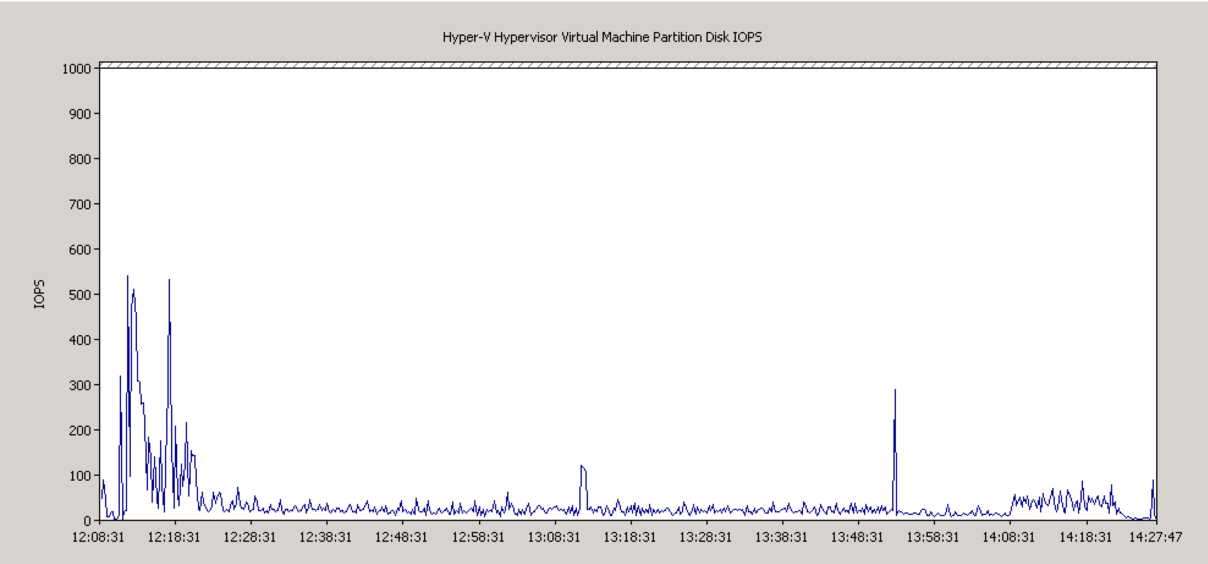

The table shown below shows sustained IOPS at peak load on the RDSH VMs during performance analysis. Peak utilization of 130 IOPS (on RDSH VM1) corresponds with approximately 2.6 IOPS per user.

	Sustained IOPS at Peak Load
RDSH VM1	130
RDSH VM2	115
RDSH VM3	125
RDSH VM4	100

The relevant Hyper-V partition latency information is shown below, it can be seen that latency remains significantly below the 15ms threshold specified by Microsoft in the TechNet best practices article described above for the entire duration of the test.

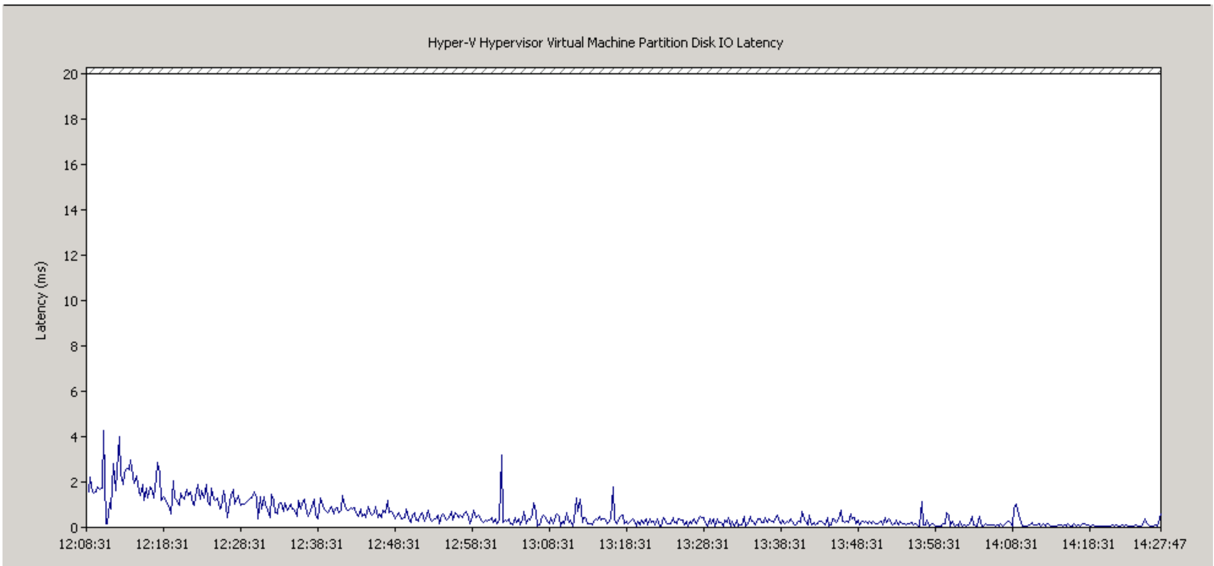

The table shown below shows sustained latency during performance analysis for each of the RDSH VMs; it can be seen that this doesn't exceed 5ms at any stage during performance analysis, indicating that disk IO latency is not a performance bottleneck.

	Sustained Latency at Peak Load (ms)
RDSH VM1	4
RDSH VM2	4
RDSH VM3	5
RDSH VM4	4

C.3.5 Network Resource Utilization Performance Analysis Results

Network resource utilization in the RDSH environment was seen to be low. At the hypervisor level, network resource utilisation reaches a maximum of approximately 31 MB/s, representing network bandwidth utilization of approximately 25%, based on a 1Gb/s switching infrastructure. The hypervisor shows zero output queue length throughout the duration of the performance analysis activity. Individual RDSH VM network utilization is consistent with this overall hypervisor network utilisation (i.e. total VM network utilization equals hypervisor network utilization) and all of the VMs also show zero output queue length throughout the duration of performance analysis. These hypervisor and RDSH VM network utilization results illustrate that network resource utilization is very low and does not approach a level at which it could become a bottleneck in the environment.

Appendix D – Combined RDVH and RDSH on Single Host

Performance analysis was also carried for a basic workload in a mixed RDSH / RDVH environment. The configuration used and a summary of the results obtained are shown below.

D.1 Configuration Summary

	Hyper-V Compute Host
CPU Resource	32 logical cores (16 physical cores)
Memory (GB)	192
User Density	205
RDVH User Density	75
RDSH User Density	130 (consisting of 130 users split between 2 VMs, each with the RDSH role).

D.2 Results Summary

The table shown below shows the sustained resource utilization at peak load for the various resource utilization parameters of interest. All of the resource utilization parameters are at expected levels based on standalone RDVH and standalone RDSH testing and none of the parameters reach a level of resource utilization that would suggest a bottleneck due to that resource. It should be noted that in addition to results shown below, a real user logged in at the peak of resource utilisation during performance analysis and the user experience perceived by that user was good.

Resource Utilization Parameter	**Peak Level**
Hyper-V Hypervisor Logical Processor % Total Run Time (LPTR)	81%
Memory Used	101GB (91GB Free)
D: Partition Disk IOPS	250
D: Partition Disk IO Latency	4ms
Network Usage	22MB/s

5 About the Authors

Peter Fine is the Principal Engineering Architect in the Desktop Virtualization Solutions Group at Dell. Peter has extensive experience and expertise on the broader Microsoft solutions software stack as well as in enterprise virtualization, storage, networking and enterprise datacenter design.

John Kelly is a Solutions Architecture Consultant in the Dell Desktop Virtualization Solutions Group at Dell. He has overall responsibility for performance analysis and characterization of DDVS solutions and has successfully implemented enterprise-level Microsoft, VMware and Citrix server and desktop virtualization solutions.

Keith Keogh is a Systems Engineer in the Desktop Virtualization Solutions Group at Dell with extensive experience validating VDI solutions on Hyper-V, VMware View and Citrix XenDesktop.

Senthil Baladhandayutham is the Solutions Development Manager in the Desktop Virtualization Solutions Group at Dell, managing the development and delivery of Enterprise class Desktop virtualization solutions based on Dell Datacenter components and core virtualization platforms.

Ara Bernardi, is the Principal Program Manager in the Remote Desktop Virtualization team at Microsoft, leading the VDI performance and VDI Reference Architectures work on Windows Server platforms.