

Dell EMC Storage Spaces Direct (S2D) Ready Nodes for Citrix XenDesktop – Reference Architecture

Integration of Citrix XenDesktop with Dell EMC S2D Ready Node appliance clusters

Dell Engineering
October 2017

Revisions

Date	Description
October 2017	Initial release

The information in this publication is provided “as is.” Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

Copyright © 2015 – 2017 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA [10/16/2017] [Reference Architecture]

Dell believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

Table of contents

1	Introduction.....	5
1.1	Objective.....	5
2	Solution architecture overview	6
2.1	Introduction.....	6
2.2	Storage Spaces Direct (S2D) overview.....	6
2.3	S2D architecture.....	7
2.4	Dell EMC S2DRN hyper-converged infrastructure.....	10
2.5	Dell EMC S2DRN VDI solution architecture.....	12
3	Hardware components	14
3.1	Network.....	14
3.2	VDI optimized Dell EMC S2D Ready Nodes	16
3.3	Dell Wyse Endpoints	19
4	Software components.....	22
4.1	Citrix.....	22
4.2	Microsoft	34
5	Solution architecture for XenDesktop.....	36
5.1	Management role configuration.....	36
5.2	Storage configuration overview	38
5.3	Virtual networking	39
5.4	Scaling guidance	41
5.5	Solution high availability	42
6	Solution Performance and Testing.....	43
6.1	Test and performance analysis methodology.....	43
6.2	Test configuration details.....	46
6.3	Test results and analysis	48
6.4	Citrix XenDesktop communication flow	67
7	Related resources	68
	Acknowledgements.....	69
	About the authors.....	70

Executive summary

This document provides the reference architecture for integrating Dell EMC Microsoft Storage Spaces Direct Ready Nodes (S2D Ready Nodes) and Citrix XenDesktop and XenApp software suite to create virtual application and virtual desktop environments. The available S2D Ready Node choices include the latest 14th generation PowerEdge R640 and R740xd servers.

The Dell EMC S2D Ready Nodes comprise a hyper-converged solution that combines storage, compute, networking, and virtualization using industry-proven Dell EMC PowerEdge™ servers and Microsoft Storage Spaces Direct technology provided natively in Windows Server 2016 Datacenter edition.

As the foundation for a complete, adaptive IT solution, PowerEdge servers deliver superior agility and reliability, outstanding operational efficiencies and top performance at any scale. With its latest generation of PowerEdge servers, Dell EMC makes server innovations more affordable and accessible, putting more power into the hands of people than ever before.

Citrix XenDesktop provides a complete end-to-end virtualization software solution delivering Microsoft Windows virtual desktops or server-based hosted shared sessions to users on a wide variety of endpoint devices.

1 Introduction

This document addresses the architecture design, configuration and implementation considerations for the key components required to deliver virtual desktops or shared sessions via Citrix XenDesktop and XenApp on Windows Server 2016 Hyper-V hypervisor running on the Dell EMC Microsoft Storage Spaces Direct Ready Nodes (S2D Ready Nodes) hyper-converged infrastructure platform.

NOTE: For step by step deployment instructions, please refer to the Dell EMC Microsoft Storage Spaces Direct Ready Node Deployment Guide located here: [LINK](#)

1.1 Objective

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 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

Dell EMC Microsoft Storage Spaces Direct Ready Nodes (S2D Ready Nodes) combine the Storage Spaces Direct technology included in Microsoft Windows Server 2016 Datacenter with industry-leading Dell EMC PowerEdge servers using local drives to create highly available and scalable software-defined storage at a fraction of the cost of traditional storage arrays. Dell EMC S2D Ready Nodes are pre-configured with certified components, tested and certified by Dell EMC and Microsoft to deliver an outstanding hyper-converged infrastructure for Hyper-V based VDI solutions.

2.2 Storage Spaces Direct (S2D) overview

Storage Spaces Direct (S2D) is the evolution of Storage Spaces, first introduced in Windows Server 2012. It leverages many of the features you know today in Windows Server, such as Failover Clustering, the Cluster Shared Volume (CSV) file system, Server Message Block (SMB) 3, and of course Storage Spaces. It also introduces new technology, most notably the [Software Storage Bus](#) and the [Resilient File System](#) (ReFS).

S2D enables building highly available and scalable storage using servers with local storage. Using servers with local storage decreases complexity, increases flexibility and scalability, as well as enables the use of storage devices with better performance, such as solid state disks and NVMe disk devices. This was previously not possible with clustered Storage Spaces and shared disks.

S2D removes the need for a shared SAS fabric, simplifying deployment and configuration. Instead it uses the network as a storage fabric, leveraging SMB3 and SMB Direct (RDMA) for high-speed, low-latency CPU efficient storage. To scale out, simply add more servers to increase storage capacity and I/O performance.

Key benefits of this solution include:

Simplicity: Go from running Windows Server 2016 to your first Storage Spaces Direct cluster in under 15 minutes.

High Performance: Storage Spaces Direct provides consistent, low latency performance thanks to its hypervisor-embedded architecture, its built-in read/write cache, and support for multiple types of high-speed storage devices.

Fault Tolerance: Built-in resiliency handles drive, server, or component failures with continuous availability. When hardware fails, just swap it out; the software heals itself, with no complicated management steps.

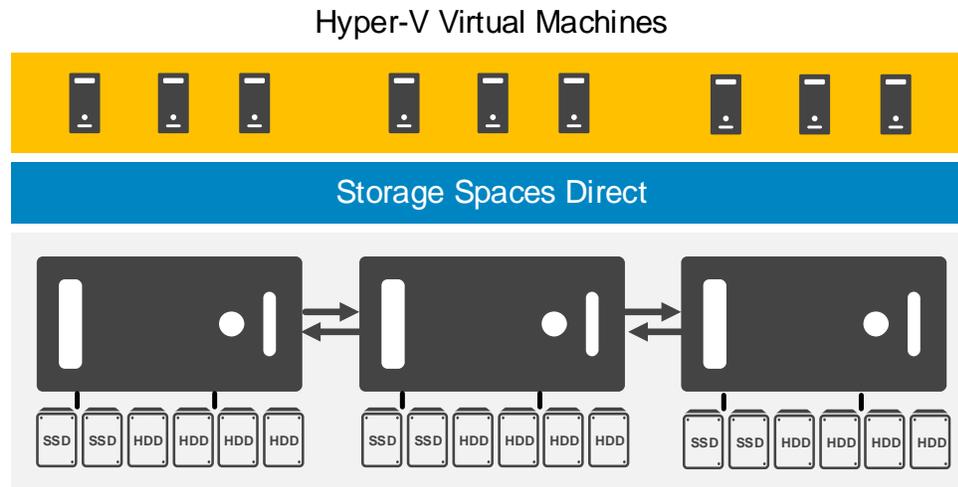
Resource Efficiency: Erasure coding delivers up to 2.4x greater storage efficiency, with unique innovations like Local Reconstruction Codes and ReFS real-time storage tiers to extend these gains to hard disk drives and mixed hot/cold workloads, all while minimizing CPU consumption to give resources back to where they're needed most - the VMs.

Manageability: Use [Storage QoS Controls](#) to keep overly busy VMs in check with minimum and maximum per-VM IOPS limits. The Health Service provides continuous built-in monitoring and alerting, and new APIs make it easy to collect rich, cluster-wide performance and capacity metrics.

Scalability: Go up to 16 servers and over 400 drives, for up to 1 petabyte of storage per cluster. To scale out, simply add drives or add more servers; Storage Spaces Direct will automatically onboard new drives and begin using them. Storage efficiency and performance improve predictably at scale.

S2D can be configured in a converged (also known as disaggregated) deployment option or as a hyper-converged deployment option. The converged option layers in Scale-out File Servers to provide network-

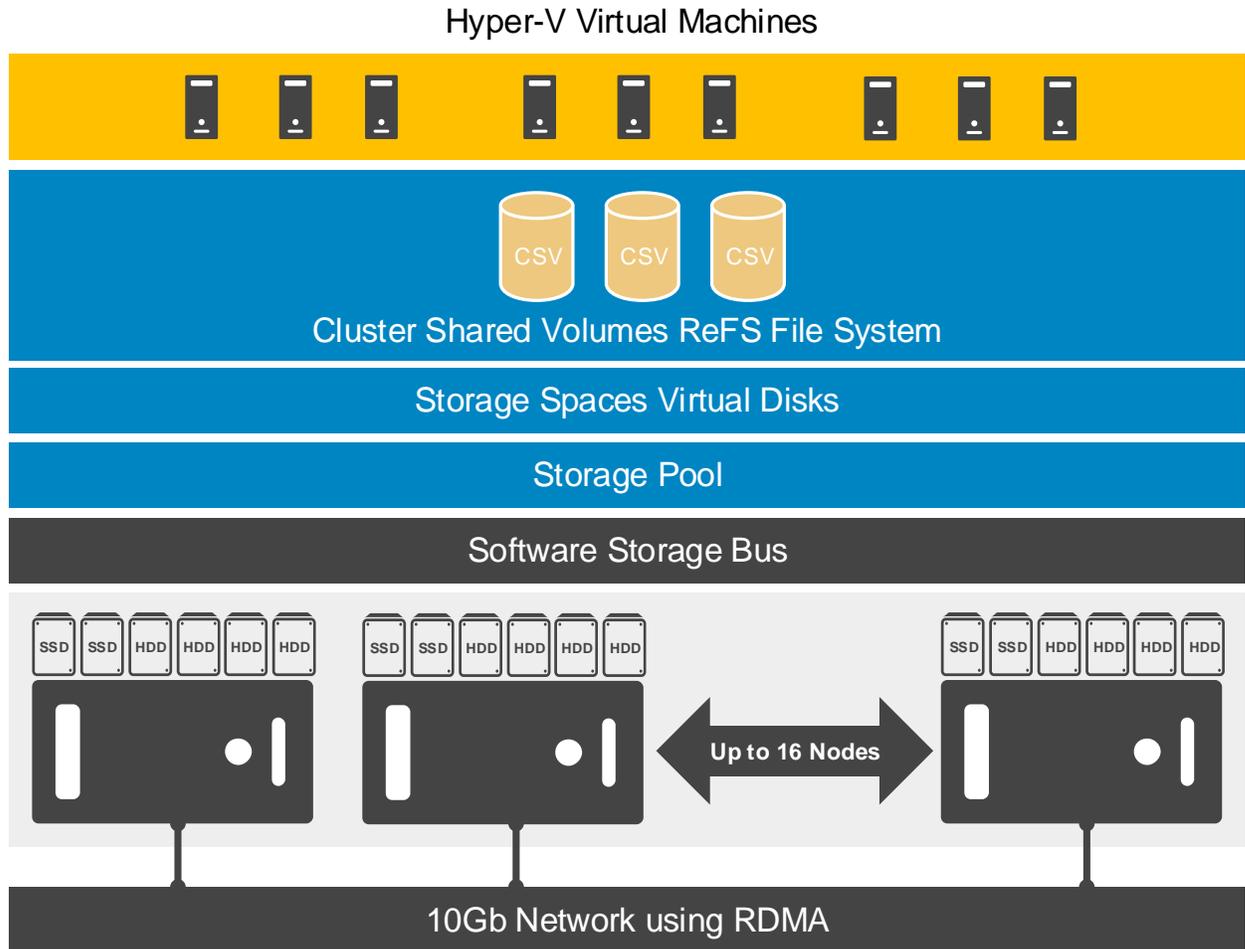
attached storage over SMB3 file shares whereas the hyper-converged option runs Hyper-V virtual machines directly on the servers providing the storage, storing the VM files on the local volumes. This reference architecture only speaks to the hyper-converged deployment option. Hyper-converged deployment with hybrid storage (SSDs + HDDs) shown below. All-flash storage option available as well.



2.3 S2D architecture

S2D uses local disks that are exclusively connected to nodes of a Windows Server 2016 Failover Cluster and allows Storage Spaces to create pools using those disks. Virtual Disks (Spaces) that are configured on the pool will have their redundant data (mirrors or parity) spread across the disks in different nodes of the cluster. Since copies of the data is distributed, this allows access to data even when a node fails or is down for maintenance.

The hyper-converged deployment scenario has the Hyper-V (compute) and Storage Spaces Direct (storage) components on the same cluster. Virtual machine files are stored on local CSVs. This allows for scaling Hyper-V compute and storage within a single contiguous cluster. Once Storage Spaces Direct is configured and the CSVs are available, configuring and provisioning Hyper-V is the same process and uses the same tools that you would use with any other Hyper-V deployment on a failover cluster. Below is an overview of the Storage Spaces Direct stack with Hyper-V:



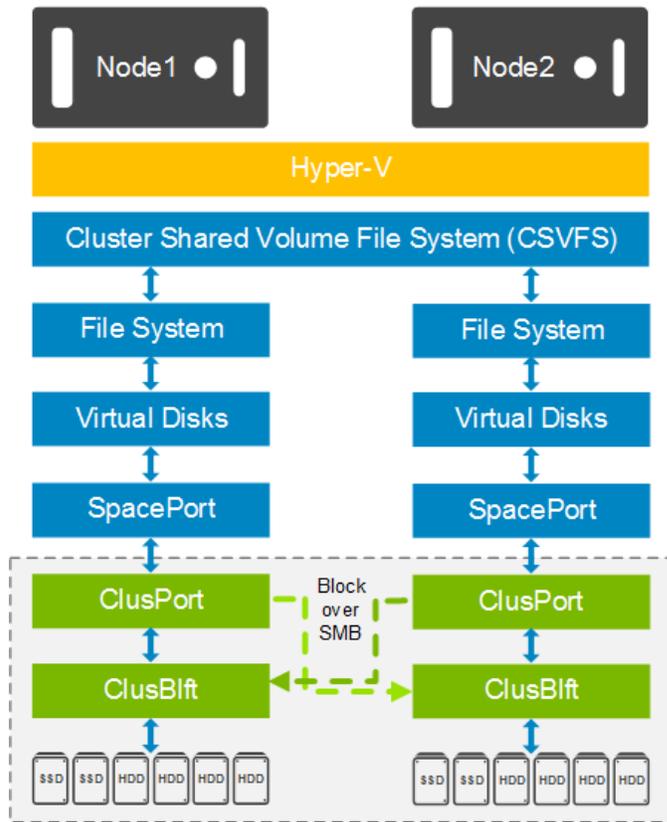
Components of the stack include:

Networking Hardware. Storage Spaces Direct uses SMB3, including SMB Direct and SMB Multichannel, over Ethernet to communicate between servers. We strongly recommend using 10+ GbE with remote-direct memory access (RDMA), either iWARP or RoCE capable NICs. Although RDMA is not required for S2D, it [significantly improves the performance](#) over just TCP/IP.

Storage Hardware. From 2 to 16 servers with local-attached SATA, SAS, or NVMe drives. For hybrid storage configurations, each server must have at least 2 NVMe or SSD drives for caching, and at least 4 additional drives for capacity. For all-flash, each node must have at least 4 SSD drives.

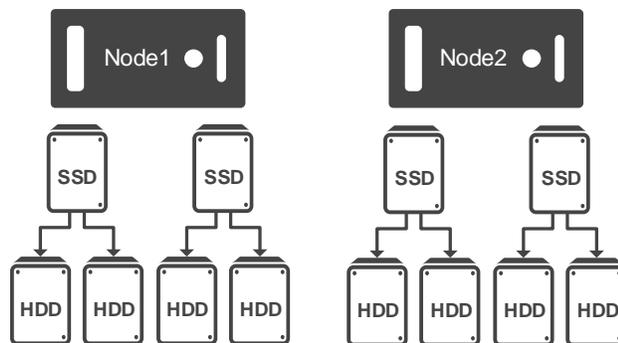
Failover Clustering. The built-in clustering feature of Windows Server is used to connect the servers to provide HA and balance available resources.

Software Storage Bus. The Software Storage Bus is new in Storage Spaces Direct. It spans the cluster and establishes a software-defined storage fabric whereby all the servers can see all of each other's local drives providing full mesh connectivity. Software Storage Bus consists of two components on each server in the cluster: ClusPort and ClusBlft.



- **ClusPort** acts as a virtual host bus adapter (HBA) that allows the node to connect to physical disks in other nodes in the cluster.
- **ClusBlft** virtualizes all physical disks in each node for ClusPort in the other nodes to connect to.
- Uses SMB3 and SMB Direct as transport communication between the nodes.
- Provides IO prioritization for Application IO (VMs) over system IO (rebalance or rebuilds).
- Uses de-randomization of IO going to spinning disk to drive higher performing sequential IO pattern.
- Implements the Storage Bus Cache

Storage Bus Layer Cache. The Software Storage Bus dynamically binds the fastest drives present (e.g. SSD) to slower drives (e.g. HDDs) to provide server-side read/write caching that accelerates IO and boosts throughput.



Key points about Storage Bus Cache:

- Storage Bus Cache (SBC) is scoped to each node and is a per node and per device cache.
- Devices are categorized as caching or capacity devices. Caching devices can be NVMe or SSD. Capacity devices can be SSD or HDD.
- S2D will automatically select the fastest performing device available as the caching device by MediaType (NVMe > SSD)
- When S2D is enabled, SBC automatically consumes all capacity on the caching devices except 32GB which is used for storage pool and virtual disk metadata.
- Caching devices are obfuscated by S2D and will not appear as usable capacity to the administrator.
- Capacity devices are bound to Caching devices in round robin manner. Rebinding will automatically occur if there is a caching device failure or topology change (drive add/ removal).
- Data resiliency is achieved by writing multiple data copies to different nodes in the cluster.
- The SBC capacity is divided into read cache and write cache.

Storage Pool. The collection of drives that form the basis of Storage Spaces is called the storage pool. It is automatically created, and all eligible drives are automatically discovered and added to it. We strongly recommend you use one pool per cluster, with the default settings.

Storage Spaces. Storage Spaces provides fault tolerance to virtual "disks" using mirroring, erasure coding, or both. It is the equivalent of distributed, software-defined RAID using the drives in the pool. In Storage Spaces Direct, these virtual disks typically have resiliency to two simultaneous drive or server failures (e.g. 3-way mirroring, with each data copy in a different server) though chassis and rack fault tolerance is also available.

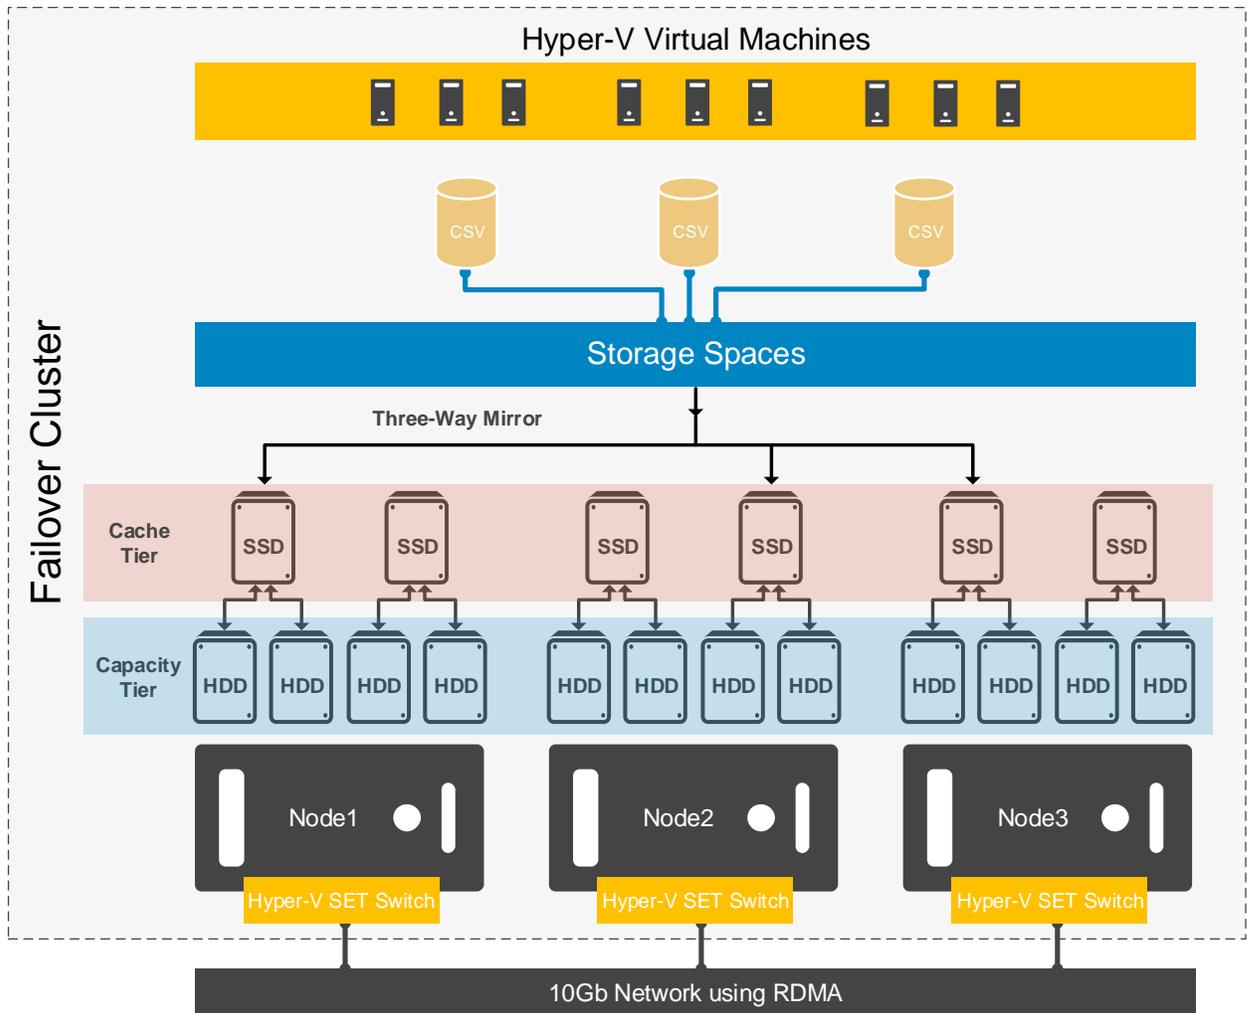
Resilient File System (ReFS). ReFS is the premier Windows filesystem purpose-built for virtualization. It includes dramatic accelerations for VHDX file operations such as creation, expansion, checkpoint merging, and built-in checksums to detect and correct bit errors. It also introduces real-time tiers that rotate data between "hot" and "cold" storage tiers in real-time based on usage. ReFS does not currently have deduplication functionality.

Cluster Shared Volumes. The CSV file system unifies all the ReFS volumes into a single namespace accessible through any server so that every volume looks and acts like it's mounted locally.

NOTE: Links to additional information can be found in the [Related Resources](#) section.

2.4 Dell EMC S2DRN hyper-converged infrastructure

Dell EMC S2D Ready Nodes are offered in all-flash or hybrid models (shown below). In a hybrid model, SSDs are used for the caching read and write IO (cache tier) while the HDDs are used for the capacity tier. For all-flash models, the cache tier is only used for writes. The cache tier is also referred to as the performance tier. For either model, usable disk capacity is calculated using only the capacity tier.



Each disk tier is extensible and can be expanded or decreased at any time. Dynamic cache bindings ensure that the proper ratio of cache to capacity disks remain consistent for any configuration regardless of whether cache or capacity disks are added or removed. The same is true in the case of drive failure in which case S2D self-heals to ensure a proper balance. To experience a more even performance, we recommend making the number of capacity drives an even multiple of the number of cache drives.

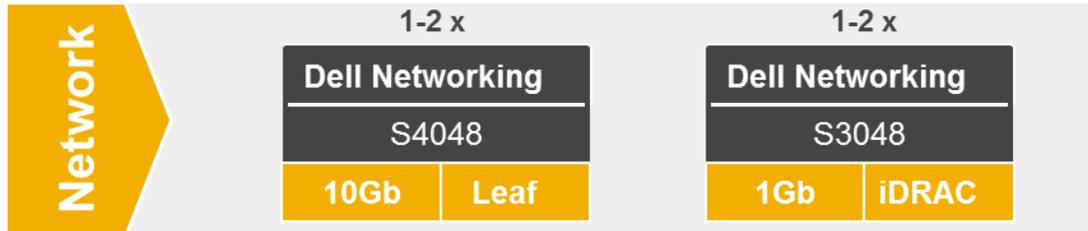
Cluster Shared Volumes are created within the S2D storage pool and shared across all nodes in the cluster. Consumed capacity in the cluster is determined by provisioned volume sizes, not actual storage usage within a given volume.

NOTE: Three-way mirroring recommended but other resiliency settings are available. Refer to the [Fault tolerance and resiliency](#) section for more information.

2.5 Dell EMC S2DRN VDI solution architecture

2.5.1 Networking

The networking layer consists of the 10Gb Dell Networking S4048 utilized to build a leaf/spine architecture with RDMA support. The S3048 provides 1Gb switching for iDRAC/ OOBM connectivity.

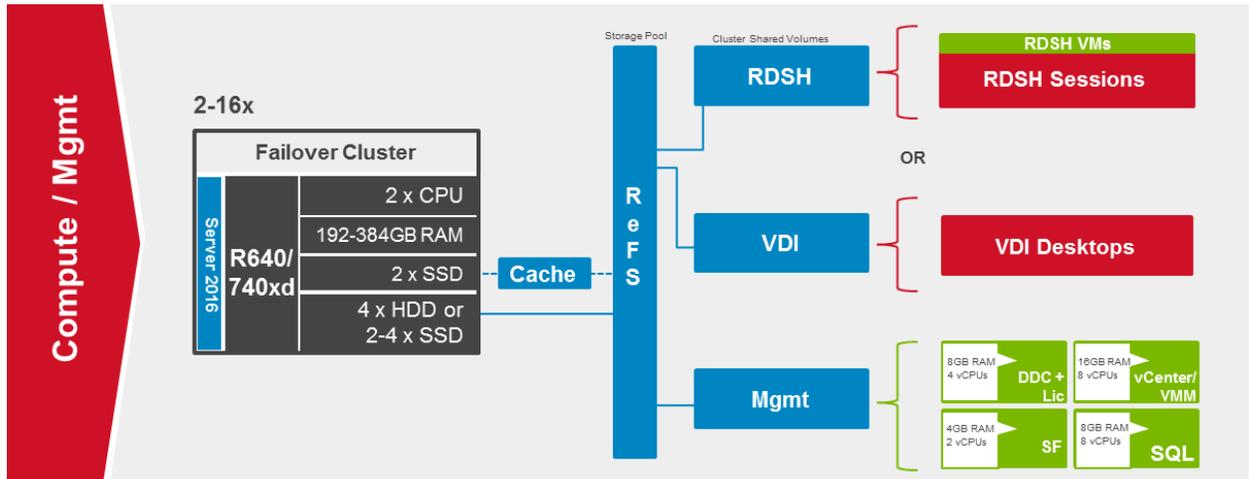


Each Dell EMC S2D Ready Node is equipped with dual port RDMA network adapters for host and storage traffic. Although the solution can function with a single switch, Dell EMC highly recommends redundant top-of-rack S4048 switches. Each RDMA port is connected to a different S4048 switch in this scenario.

NOTE: For cabling details, please refer to the Dell EMC Microsoft Storage Spaces Direct Ready Node Deployment Guide located here: [LINK](#)

2.5.2 S2DRN VDI clusters

The compute, management and storage layers are converged into each Dell EMC S2D Ready Node in the cluster with each cluster supporting up to 16 nodes. For the storage layer, a single S2D storage pool is used for the ReFS cluster shared volumes that house the compute (VDI desktops or XenApp sessions) and management VM disks/files. For this architecture, Dell EMC recommends that the VDI compute and management infrastructure be installed on the same cluster with all nodes available to host both compute and management VMs. Please refer to the [Storage configuration overview](#) section for more details.



NOTE: S2D supports a two-node minimum cluster; however, with two nodes, S2D will automatically create two-way mirror volumes that cannot be upgraded in-place to any other resiliency setting. Therefore, Dell EMC does not support expansion to a larger cluster size from a two-node cluster. Dell EMC S2D Ready Nodes in a two-node configuration are primarily meant for test/Dev or ROBO scenarios. For supported production deployments, please start with at least three Dell EMC S2D Ready Nodes in the same cluster.

3 Hardware components

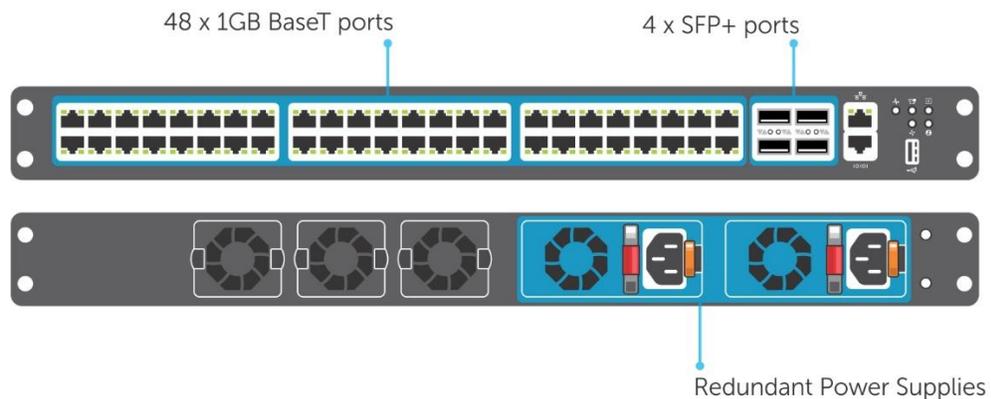
3.1 Network

The following sections contain the core network components for the Dell Wyse Datacenter solutions. General uplink cabling guidance to consider in all cases is that TwinAx is very cost effective for short 10Gb runs and for longer runs use fiber with SFPs.

3.1.1 Dell Networking S3048 (1Gb ToR switch)

For out-of-band management such as iDRAC or in environments where 1Gb networking is sufficient, Dell recommends the S3048 network switch. The S3048 is a low-latency top-of-rack (ToR) switch that features 48 x 1GbE and 4 x 10GbE ports, a dense 1U design, and up to 260Gbps performance. The S3048-ON also supports Open Network Installation Environment (ONIE) for zero-touch installation of alternate network operating systems.

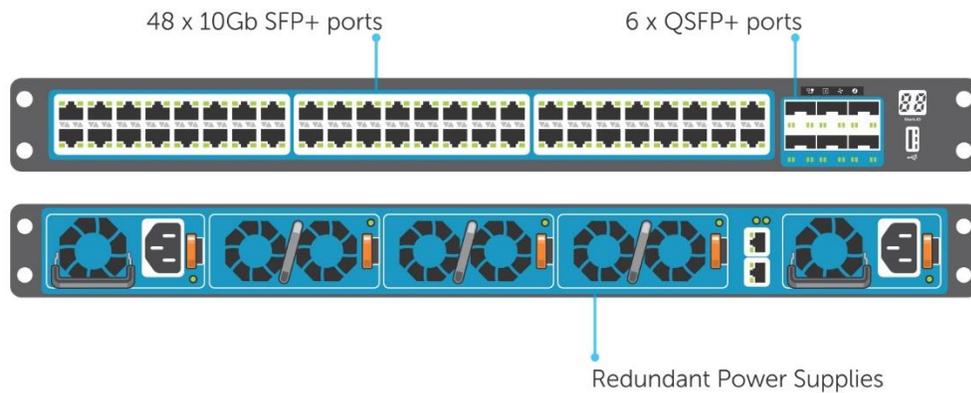
Model	Features	Options	Uses
Dell Networking S3048-ON	48 x 1000BaseT 4 x 10Gb SFP+	Redundant hot-swap PSUs & fans	1Gb connectivity
	Non-blocking, line-rate performance	VRF-lite, Routed VLT, VLT Proxy Gateway	
	260Gbps full-duplex bandwidth	User port stacking (up to 6 switches)	
	131 Mbps forwarding rate	Open Networking Install Environment (ONIE)	



3.1.2 Dell Networking S4048 (10Gb ToR switch)

Optimize your network for virtualization with a high-density, ultra-low-latency ToR switch that features 48 x 10GbE SFP+ and 6 x 40GbE ports (or 72 x 10GbE ports in breakout mode) and up to 720Gbps performance. The S4048-ON also supports ONIE for zero-touch installation of alternate network operating systems. For BaseT connectivity, the S4048T model is available.

Model	Features	Options	Uses
Dell Networking S4048-ON	48 x 10Gb SFP+ 6 x 40Gb QSFP+	Redundant hot-swap PSUs & fans	10Gb connectivity
	Non-blocking, line-rate performance	72 x 10Gb SFP+ ports with breakout cables	
	1.44Tbps bandwidth	User port stacking (up to 6 switches)	
	720 Gbps forwarding rate VXLAN gateway support	Open Networking Install Environment (ONIE)	



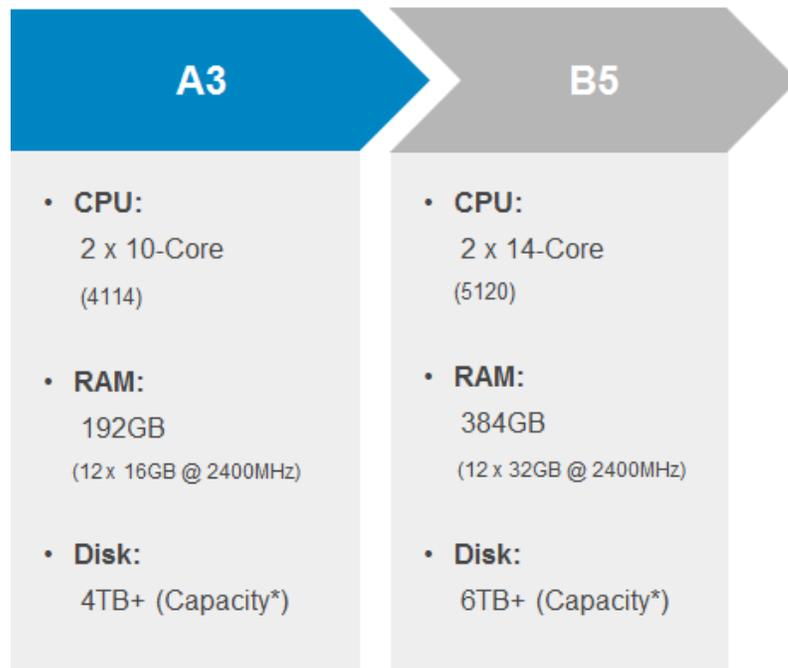
For more information on the S3048, S4048 switches and Dell Networking, please visit: [LINK](#)

3.2 VDI optimized Dell EMC S2D Ready Nodes

Dell EMC S2D Ready Nodes are pre-configured with certified components and validated building blocks that simplify ordering and reduce deployment risks while providing a superb customer support experience. Dell EMC MS2D Ready Nodes are built on the latest 14th generation Dell EMC PowerEdge R640 and R740xd servers. For general purpose virtual desktop and virtual application solutions, Dell recommends the R740xd Hybrid storage configuration.

Optimized for VDI, the Dell EMC S2D Ready Node portfolio has been designed and arranged in two top-level overarching configurations which apply to the available physical platforms showcased below.

- **A3** configuration is perfect for small scale, POC or low density cost-conscious environments.
- **B5** configuration is geared toward larger scale general purpose workloads, balancing performance and cost-effectiveness.

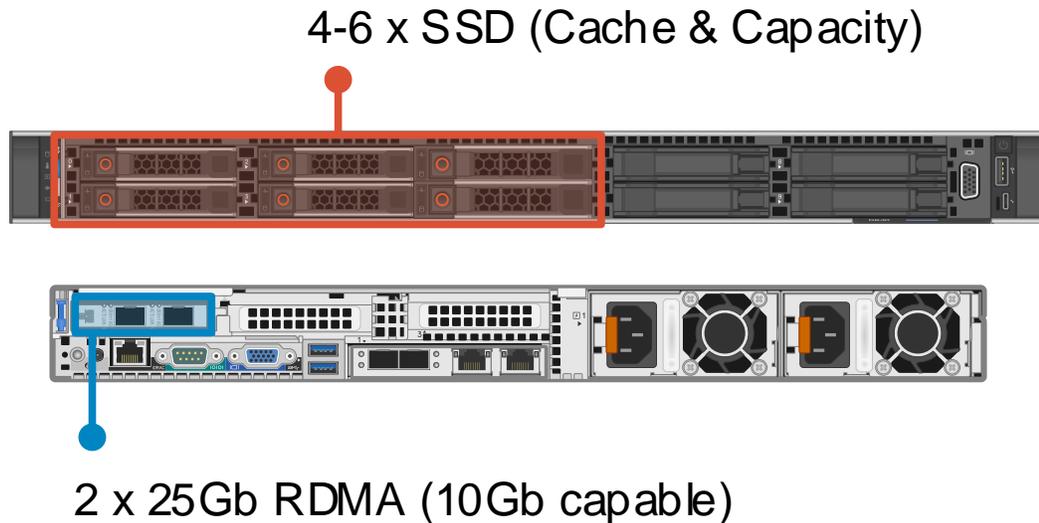


A new Boot Optimized Storage Subsystem (BOSS) card is available with the Dell EMC PowerEdge 14th generation (14G) servers. The Boot Optimized Storage Solution uses one or two M.2 SATA devices instead of 2.5" SSD drives to house the OS, and utilizes a two-port SATA Hardware RAID controller chip to provide Hardware RAID 1 and pass-through capabilities. The M.2 devices offer the same performance as 2.5" SSD drives and frees up an additional drive slot for data needs by consolidating the SSDs and controller chip on a single PCIe adapter card. A BOSS card with dual 240GB M.2 storage configured in RAID1 is the boot device for the Dell EMC S2D Ready Nodes.

The Dell EMC S2D Ready Nodes include a DP 10Gb + DP 1Gb network daughter card (NDC) that is not required for the solution but can be used optionally for further network traffic segmentation.

3.2.1 Dell EMC S2DRN R640 (All-Flash Storage)

The Dell EMC PowerEdge R640 is the ideal dual-socket, 1U platform for dense scale-out data center computing. The R640 combines density, performance and scalability to optimize application performance and data center density. The R640 platform supports the latest Intel Xeon SP processors (up to 28 cores) and up to 24 DDR4 DIMMS for a maximum of 1.5TB of memory. The R640 is the base platform for the Dell EMC S2D Ready Node all-flash configurations. For Storage Spaces Direct, a minimum of four SSD disks is required on this platform.

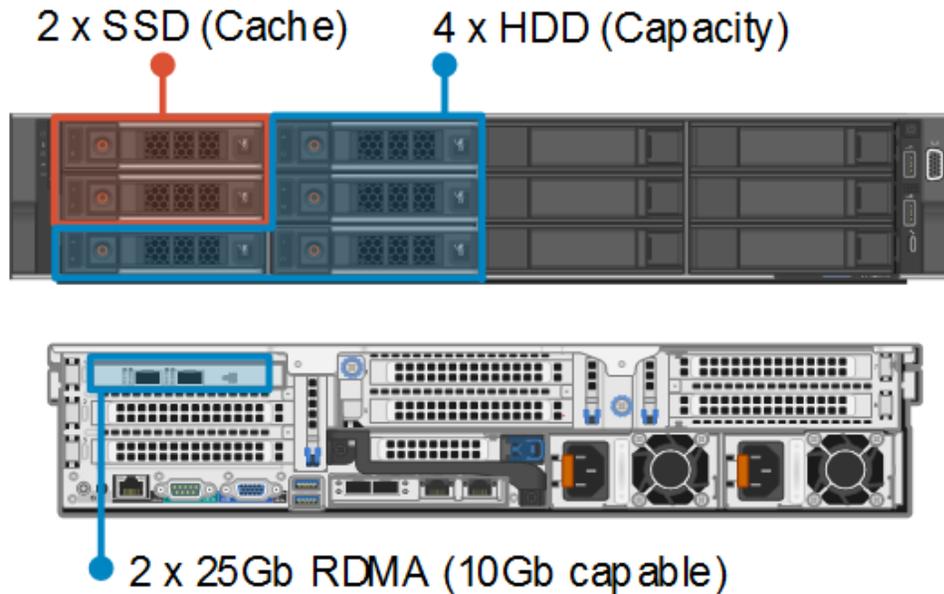


The available VDI configurations are summarized in the table below with primary differences being in the number of disks and amount of memory and processors.

S2DRN R640	A3	B5
CPU	2 x 4114 (10C, 2.2GHz)	2 x 5120 (14C, 2.2GHz)
Memory	12 x 16GB 2667MT/s RDIMMs (192GB) Effective speed: 2400MT/s	12 x 32GB 2667MT/s RDIMMs (384GB) Effective speed: 2400MT/s
Storage Ctrls	PERC HBA330	PERC HBA330
Storage	2 x 240GB BOSS M.2 (Boot) 4 x 960GB SSD 2.5" (Cache + Capacity)	2 x 240GB BOSS M.2 (Boot) 6 x 960GB SSD 2.5" (Cache + Capacity)
Network	2 x 10Gb, 2 x 1Gb SFP+/ BT 2 x 25Gb Mellanox DA/SFP (RDMA)	2 x 10Gb, 2 x 1Gb SFP+/ BT 2 x 25Gb Mellanox DA/SFP (RDMA)
iDRAC	iDRAC9 Enterprise	iDRAC9 Enterprise
Power	2 x 1100W PSUs	2 x 1100W PSUs

3.2.2 Dell EMC S2DRN R740xd (Hybrid Storage)

The PowerEdge R740xd was designed to accelerate application performance leveraging accelerator cards and storage scalability. The 2-socket, 2U platform supports the latest Intel Xeon SP processors (up to 28 cores) and up to 24 DDR4 DIMMs for a maximum of 1.5TB of memory. The R740xd is the platform for the Dell EMC S2D Ready Node hybrid storage offering and requires a minimum of two SSD disks and four HDDs.



The available VDI configurations are summarized in the table below with primary differences being in the amount of memory and processors.

S2DRN R740xd	A3	B5
CPU	2 x 4114 (10C, 2.2GHz)	2 x 5120 (14C, 2.2GHz)
Memory	12 x 16GB 2667MT/s RDIMMs (192GB) Effective speed: 2400MT/s	12 x 32GB 2667MT/s RDIMMs (384GB) Effective speed: 2400MT/s
Storage Ctrls	PERC HBA330	PERC HBA330
Storage	2 x 240GB BOSS M.2 (Boot) 2 x 960GB SSD 3.5" (Cache) 4 x 2TB HDD 3.5" (Capacity)	2 x 240GB BOSS M.2 (Boot) 2 x 960GB SSD 3.5" (Cache) 4 x 2TB HDD 3.5" (Capacity)
Network	2 x 10Gb, 2 x 1Gb SFP+/ BT 2 x 25Gb Mellanox DA/SFP (RDMA)	2 x 10Gb, 2 x 1Gb SFP+/ BT 2 x 25Gb Mellanox DA/SFP (RDMA)
iDRAC	iDRAC9 Enterprise	iDRAC9 Enterprise
Power	2 x 1100W PSUs	2 x 1100W PSUs

3.3 Dell Wyse Endpoints

The following Dell Wyse clients will deliver a superior Citrix user experience and are the recommended choices for this solution.



3.3.1 Wyse 3040 Thin Client (ThinOS, ThinLinux)

The Wyse 3040 is the industry's first entry-level Intel x86 quad-core thin client, powered by a quad-core Intel Atom 1.44GHz processor, delivering robust connectivity options with a choice of Wyse ThinOS or ThinLinux operating systems. The Wyse 3040 is Dell's lightest, smallest and most power-efficient thin client – it consumes 3.3 Watts in idle state – and offers superb performance and manageability for task and basic productivity users. Despite its small size, the 3040 includes all typical interfaces such as four USB ports including USB 3.1, two DisplayPort interfaces and wired and wireless options. It is highly manageable as it can be monitored, maintained, and serviced remotely via Wyse Device Manager (WDM) or Wyse Management Suite. For more information, please visit: [Link](#)



3.3.2 Wyse 5010 Thin Client (ThinOS)



Designed for knowledge workers and power users, the Wyse 5010 is a high performance thin client based on Wyse ThinOS, the virus-resistant firmware base designed for optimal thin client security, performance, and ease-of-use. Highly secure, compact and powerful, it combines a dual-core AMD 1.4 GHz CPU with a revolutionary unified graphics engine for an outstanding user experience. It addresses the performance challenges of processing-intensive applications like computer-aided design, multimedia, HD video and 3D modelling. Scalable 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 Wyse 5010 offers a unique combination of performance, efficiency, and affordability. For more information, please visit: [Link](#)

3.3.3 Wyse 5060 Thin Client (ThinOS, ThinLinux, WES7P, WIE10)

The Wyse 5060 offers high performance and reliability, featuring all the security and management benefits of Dell thin clients. It come with flexible OS options: ThinOS, ThinLinux, Windows Embedded Standard 7P (WES7P) or Windows 10 IoT Enterprise (WIE10). Designed for knowledge workers demanding powerful virtual desktop performance, and support for unified communications solutions like Skype for Business, the Wyse 5060 thin client delivers the flexibility, efficiency and security organizations require for their cloud environments. It is powered by a quad-core AMD 2.4GHz processor, supports dual 4K (3840x2160) monitors and provides multiple connectivity options with six USB ports, two of which are USB 3.0 for high-speed peripherals, as well as two DisplayPort connectors, wired networking or wireless 802.11 a/b/g/n/ac. The Wyse 5060 can be monitored, maintained, and serviced remotely via Wyse Device Manager (WDM), cloud-based Wyse Management Suite or Microsoft SCCM (5060 with Windows versions). For more information, please visit: [Link](#)



3.3.4 Wyse 7020 Thin Client (WES 7/7P, WIE10, ThinLinux)

The versatile Dell Wyse 7020 thin client is a powerful endpoint platform for virtual desktop environments. It is available with Windows Embedded Standard 7/7P (WES), Windows 10 IoT Enterprise (WIE10), Wyse ThinLinux operating systems and it supports a broad range of fast, flexible connectivity options so that users can connect their favorite peripherals while working with processing-intensive, graphics-rich applications. This 64-bit thin client delivers a great user experience and support for local applications while ensuring security.



Designed to provide a superior user experience, ThinLinux features broad broker support including Citrix Receiver, VMware Horizon and Amazon Workspace, and support for unified communication platforms including Skype for Business, Lync 2013 and Lync 2010. For additional security, ThinLinux also supports single sign-on and VPN. With a powerful quad core AMD G Series APU in a compact chassis with dual-HD monitor support, the Wyse 7020 thin client delivers stunning performance and display capabilities across 2D, 3D and HD video applications. Its silent diskless and fan less design helps reduce power usage to just a fraction (it only consumes about 15 watts) of that used in traditional desktops. Wyse Device Manager (WDM)

helps lower the total cost of ownership for large deployments and offers remote enterprise-wide management that scales from just a few to tens of thousands of cloud clients. Customers choosing WIE10 licenses can save about \$50/device/year as WIE10 qualifies under Microsoft Software Insurance, without the need to have more expensive VDA licenses to connect to a Windows virtual desktop. For more information, please visit [Link](#)

3.3.5 Wyse 7040 Thin Client (WES7P, WIE10)

The Wyse 7040 is a high-powered, ultra-secure thin client running Windows Embedded Standard 7P (WES7P) or Windows 10 IoT Enterprise (WIE10) operating systems. Equipped with an Intel i5/i7 processors, it delivers extremely high graphical display performance (up to three displays via display-port daisy-chaining, with 4K resolution available on a single monitor) for seamless access to the most demanding applications. The Wyse 7040 is compatible with both data center hosted and client-side virtual desktop environments and is compliant with all relevant U.S. Federal security certifications including OPAL compliant hard-drive options, VPAT/Section 508, NIST BIOS, Energy-Star and EPEAT. Wyse enhanced WES7P OS provides additional security features such as BitLocker. The Wyse 7040 offers a high level of connectivity including dual NIC, 6 x USB3.0 ports and an optional second network port, with either copper or fiber SFP interface. Wyse 7040 devices are highly manageable through Intel vPRO, Wyse Device Manager (WDM), Microsoft System Center Configuration Manager (SCCM) and Dell Command Configure (DCC). Customers choosing WIE10 licenses can save about \$50/device/year as WIE10 qualifies under Microsoft Software Insurance, without the need to have more expensive VDA licenses to connect to a Windows virtual desktop. For more information, please visit: [Link](#)



3.3.6 Latitude 3480 and 5280 Mobile Thin Clients (Win 10 IoT)

Designed to securely deliver virtual desktops and applications to mobile users who want to connect a broad range of peripherals, the Latitude 3480 and 5280 mobile thin clients run **Windows 10 IoT Enterprise**. They support a wide variety of connection brokers including Citrix XenDesktop/XenApp, Microsoft RDS and VMware Horizon right out of the box, and are an ideal alternative to much less secure Chromebooks.



The Latitude 3480 features an Intel dual core processor with integrated graphics for a rich multimedia experience, and delivers great value with a 14" Full-HD display and robust connectivity with plenty of ports.



The Latitude 5280 delivers excellent performance with 12.5-inch, Full HD display. It offers the ability to support a 4K monitor via an optional docking station, and it supports a broad mix of peripheral attachments and network connections.

They are easily manageable through Wyse Device Manager (WDM), Wyse Management Suite and Microsoft's System Center Configuration Manager (SCCM). For enhanced security, optional advanced threat protection in the form of Dell Threat Defense offers proactive malware protection. For more information, please visit the following pages for: [Latitude 3480](#) , [Latitude 5280](#)

3.3.7 Enhanced Security

Note that all the above thin clients running Windows Embedded Standard 7 or Windows 10 IoT can be protected against viruses, ransomware and zero-day threats by installing [Dell Threat Defense](#), a revolutionary anti-malware software solution using artificial intelligence and mathematical modeling and is not signature-based. Threat Defense prevents 99% of executable malware, far above the average 50% of threats identified by the top anti-virus solutions. It doesn't need a constant internet connection nor frequent updates (only about twice a year), it only uses 1-3% CPU and has only a ~40MB memory footprint, making it an ideal choice to protect thin clients without impacting the end user productivity.

If you also want to protect virtual desktops against such malware and threats with a similar success, Dell recommends using [Dell Endpoint Security Suite Enterprise](#), a full suite featuring advanced threat prevention and data-centric encryption using an on-premise management console. This suite can also be used to protect physical PCs, MAC OS X systems and Windows Server.

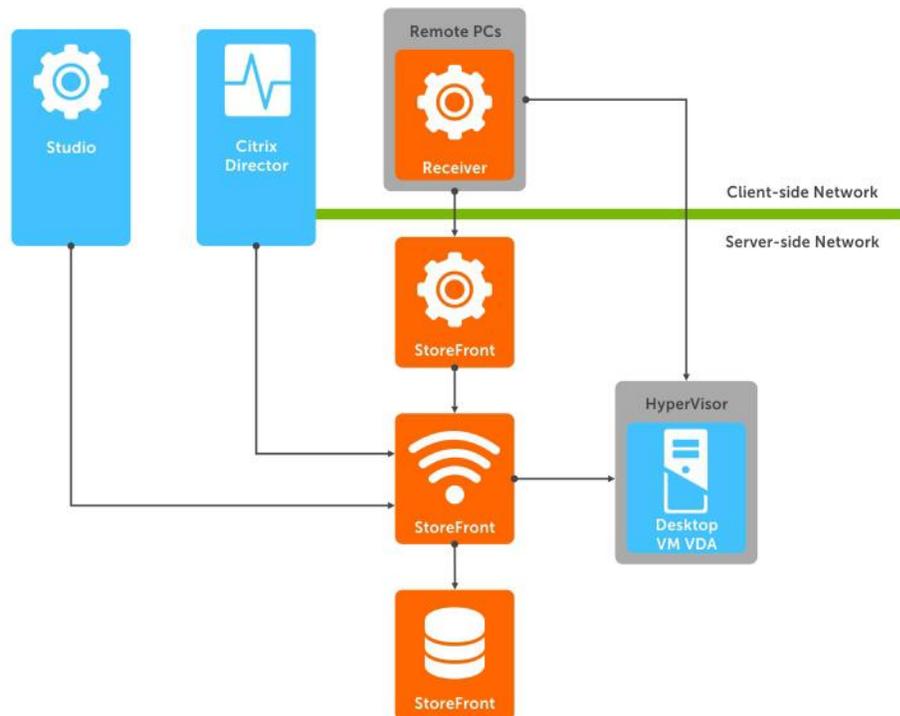
4 Software components

4.1 Citrix

4.1.1 Citrix XenDesktop

The solution is based on Citrix XenDesktop 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.

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. The Controller also manages the state of desktops, starting and stopping them based on demand and administrative configuration.

Database – At least one Microsoft SQL Server database is required for every XenApp or XenDesktop Site to store configuration and session information. The Delivery Controller must have a persistent connection to the database as it stores data collected and managed by the Controller services.

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. Starting in version 7.12, Director now includes detailed descriptions for connection and machine failures, one month historical data (Enterprise edition), custom reporting, and notifications via SNMP traps.

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. For devices that cannot install the Receiver software, Citrix Receiver for HTML5 provides connectivity through a HTML5-compatible web browser.

StoreFront – StoreFront authenticates users to sites hosting resources and manages stores of desktops and applications that user's access. StoreFront version 3.8 (released with XenDesktop 7.12) and above includes ability to create and use multiple IIS websites each having its own domain name.

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.

Provisioning Services (PVS) – The Provisioning Services infrastructure is based on software-streaming technology. This technology allows computers to be provisioned and re-provisioned in real-time from a single shared-disk image.

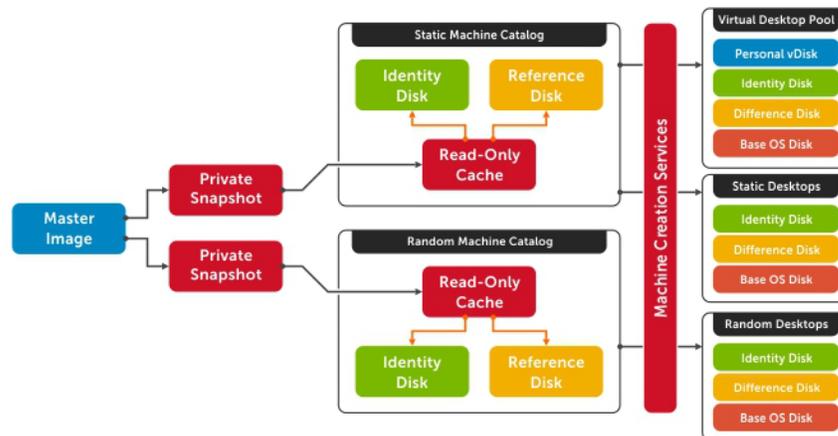
Virtual Delivery Agent (VDA) – The Virtual Desktop Agent is a transparent plugin that is installed on every virtual desktop or XenApp host (RDSH) and enables the direct connection between the virtual desktop and users' endpoint devices. Windows and Linux VDAs are available.

4.1.2 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.

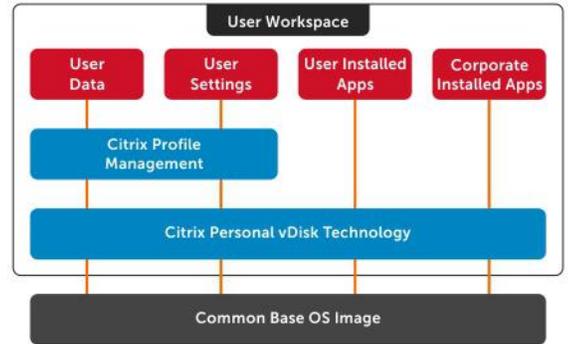


While traditionally used for small to medium sized XenDesktop deployments, MCS can bring along with it some substantial Tier 1 storage cost savings because of the snapshot/identity/difference disk methodology. The Tier 1 disk space requirements of the identity and difference disks when layered on top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

4.1.3 Personal vDisk

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

Used in conjunction with a static desktop experience, Citrix Personal vDisk allows each user to receive personal storage in the form of a layered vDisk (3GB minimum). This personal vDisk enables users to personalize and persist their desktop environment while providing storage for any user or departmental apps.

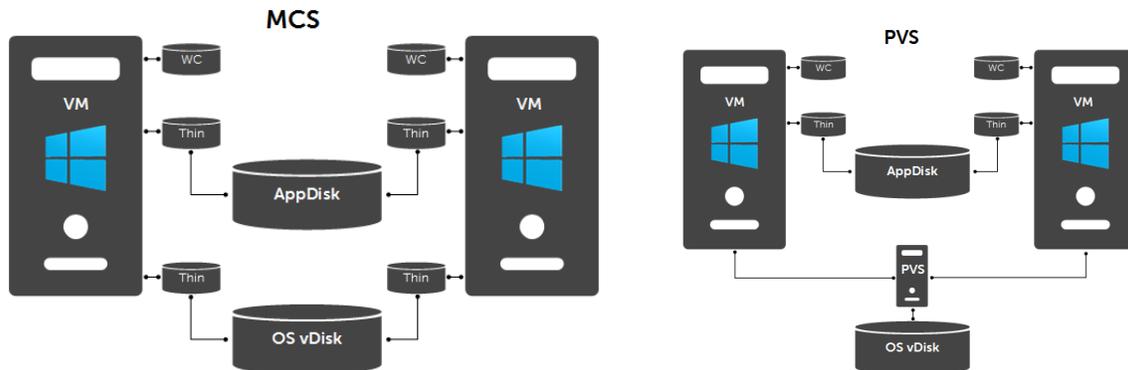


Personal vDisk provides the following benefits to XenDesktop:

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

4.1.4 AppDisks

Citrix AppDisk, included in XenDesktop 7.8 and higher, provides layering technology to manage departmental applications as an independent storage layer. AppDisk eases the management burden of maintaining multiple departmental images by instantly layering applications onto a single golden image that remains separate and pristine. AppDisks can be associated with either published desktops or published applications via XenApp. AppDisk does not replace the functionality provided by Personal vDisk but currently cannot be used within the same golden image either. AppDisks, when integrated with AppDNA, provide the ability to analyze OS and application performance, compatibility, and remediation capabilities.



4.1.5 HDX 3D Pro

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

Use HDX 3D Pro technologies with:

- Computer-aided design, manufacturing, and engineering (CAD/CAM/CAE) applications
- Geographical information system (GIS) software
- Picture Archiving Communication System (PACS) workstations for medical imaging
- Latest OpenGL, DirectX, CUDA and CL versions supported
- Latest NVIDIA Grid cards
- Shared or dedicated GPUs or a mix of both on desktop or server OS VMs

HDX 3D Pro provides the best user experience over any bandwidth using Framehawk integration:

- On wide area network (WAN) connections: Deliver an interactive user experience over WAN connections with bandwidths as low as 1.5 Mbps.
- On local area network (LAN) connections: Deliver a user experience equivalent to that of a local desktop on LAN connections.

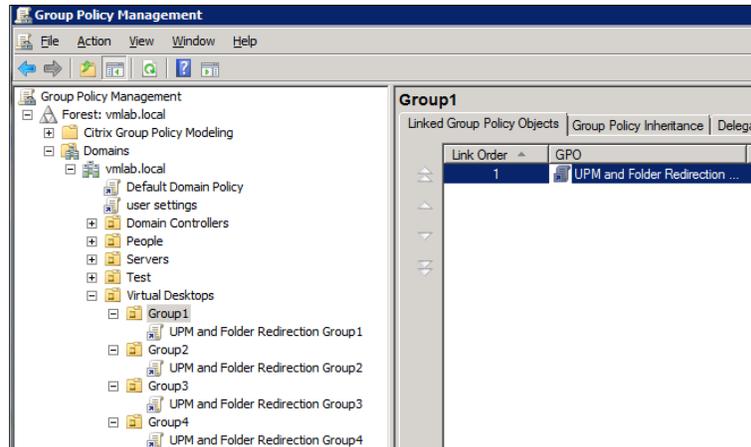
Framehawk is a display remoting technology implemented as an ICA virtual channel that optimizes delivery of virtual desktops and applications to users on broadband wireless connections where high packet loss or congestion occurs.

4.1.6 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, installed on any device where the user profiles is managed, and a Group Policy Administrative Template, which is imported to a group policy.

In order to further optimize, the profile management folders within the user profile is redirected the users' home drive. The folder redirection is managed via group policy objects within Active Directory. The following folders are redirected:

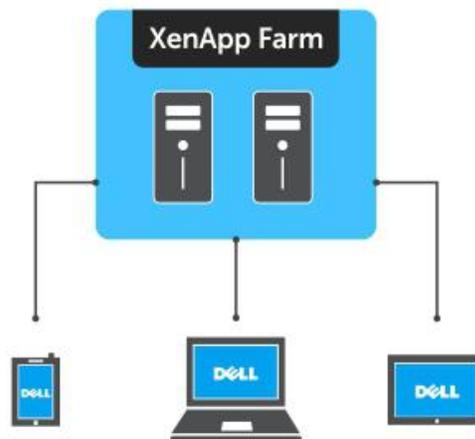
- Contacts
- Downloads
- Favorites
- Links
- My Documents
- Searches
- Start Menu
- Windows
- My Music
- My Pictures
- My Videos
- Desktop



4.1.7 Citrix XenApp

Citrix XenApp version 7.6 and above includes enhancements in the areas of faster access to virtual apps with higher connection resiliency, improved graphics rendering, and new app-usage reporting and monitoring tools.

Citrix XenApp delivers Windows apps as secure mobile services. With XenApp, IT can mobilize the business - increasing user productivity, while reducing costs by centralizing control and security of intellectual property. XenApp delivers high-performance apps to any PC, Mac, laptop, tablet or smartphone that enable the delivery of a native experience that is optimized for the type of device, as well as the network. XenApp is built on a 3rd generation FlexCast Management Architecture (FMA) and is the only hybrid cloud-ready platform that separates the management plane from the workload to enable IT to securely deliver published apps on-premises, and manage workers and mobile workspaces either on-premises or in the cloud.



Benefits of hosted desktop sessions and applications:

- Management of applications (single instance)
- Management of simple desktop images (no applications installed)
- PVS to stream XenApp servers as well as user desktops
- Scalability of XenDesktop compute hosts: CPU and IOPS reduction via application offload
- Shared storage scalability: less IOPS = more room to grow

Citrix XenDesktop with XenApp integration can effectively deliver a desktop/application hybrid solution as well. Specifically where a single or small number of shared VDI desktop images are deployed via XenDesktop, each with common shared applications installed within the golden image. A user-specific application set is then deployed and made accessible via the hosted application compute infrastructure, accessible from within the virtual desktop.

User Environment	XenDesktop	XenApp
User-Specific Applications		✓
Profile and User Data	✓	✓
Shared Applications	✓	
Shared Virtual Desktop Image	✓	

Alternatively, XenApp provides a platform for delivering Windows server-based sessions to users who may not need a full desktop VM. Hosted desktops increase infrastructure resource utilization while reducing complexity as all applications and sessions are centrally managed.

User Environment	XenDesktop	XenApp
User-Specific Applications		✓
Profile and User Data		✓
Dedicated Virtual Desktop Image		✓

4.1.7.1 XenApp integration

The XenApp servers can exist as physical or virtualized instances of Windows 2016. A minimum of one (1), up to a maximum of ten (10) virtual servers are installed per physical compute host. Since XenApp instances are easily added to an existing XenDesktop stack, the only additional components required are one or more Windows Server OS instances running the Citrix VDA added to the XenDesktop site.

The total number of required virtual XenApp servers is dependent on application type, quantity and user load. Deploying XenApp virtually and in a multi-server farm configuration increases overall farm performance, application load balancing as well as farm redundancy and resiliency.

4.1.7.2 XenDesktop with XenApp and Personal vDisk integration

In a XenDesktop implementation that leverages hosted applications, these execute from a centralized Windows Server and are then accessed via the Citrix Receiver. There are some instances, however, where certain departmental or custom applications cannot run using XenApp. At the same time, for organizational policy or certain storage considerations, delivering these applications as a part of a base image is not possible either. In this case, Citrix Personal vDisk technology is the appropriate solution.

With Citrix Personal vDisk, each user of that single shared virtual desktop image also receives a personal layered vDisk, which enables the user to personalize their desktop and receive native application execution within a Windows client OS and not from a server. When leveraging the integration of XenApp within XenDesktop, all profile and user data is seamlessly accessed within both environments.

User Environment	XenDesktop	XenApp
User-Specific Applications		✓
Profile and User Data	✓ PvDisk	✓
Departmental Applications	✓	
Shared Applications	✓	

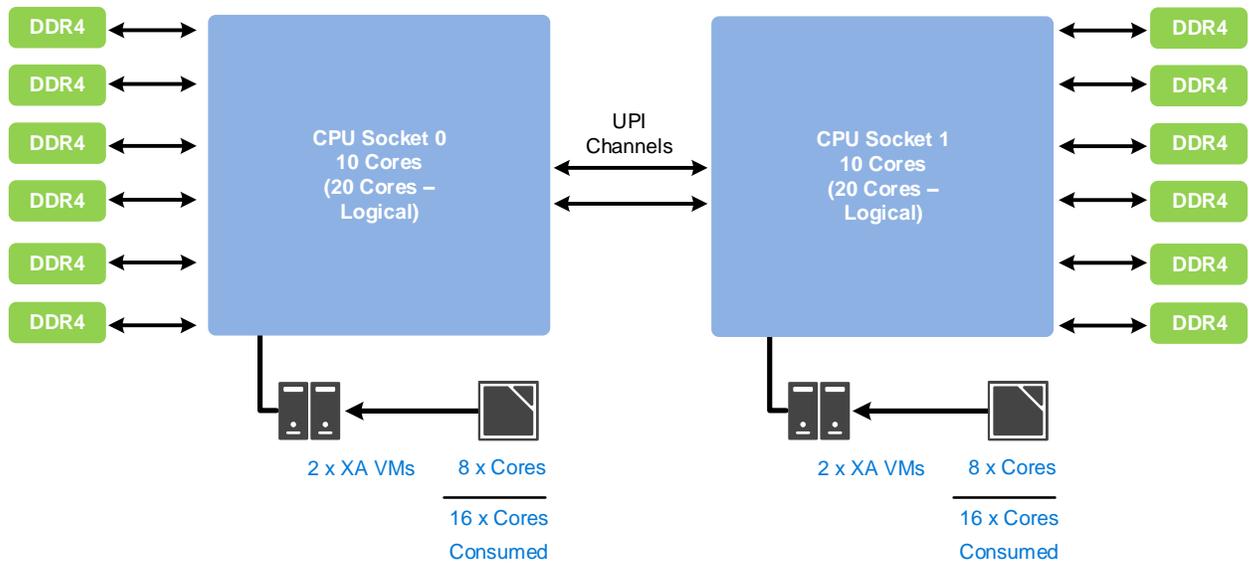
4.1.7.3 NUMA architecture considerations

Best practices and testing has showed that aligning RDSH design to the physical Non-Uniform Memory Access (NUMA) architecture of the server CPUs results in increased and optimal performance. NUMA alignment ensures that a CPU can access its own directly-connected RAM banks faster than those banks of the adjacent processor which are accessed via the Ultra Path Interconnect (UPI) which was formally the Quick Path Interconnect (QPI) with older processors. The same is true of VMs with large vCPU assignments, best performance will be achieved if your VMs receive their vCPU allotment from a single physical NUMA node. Ensuring that your virtual RDSH servers do not span physical NUMA nodes will ensure the greatest possible performance benefit.

The general guidance for RDSH NUMA-alignment on the Dell EMC XC appliance is as follows:

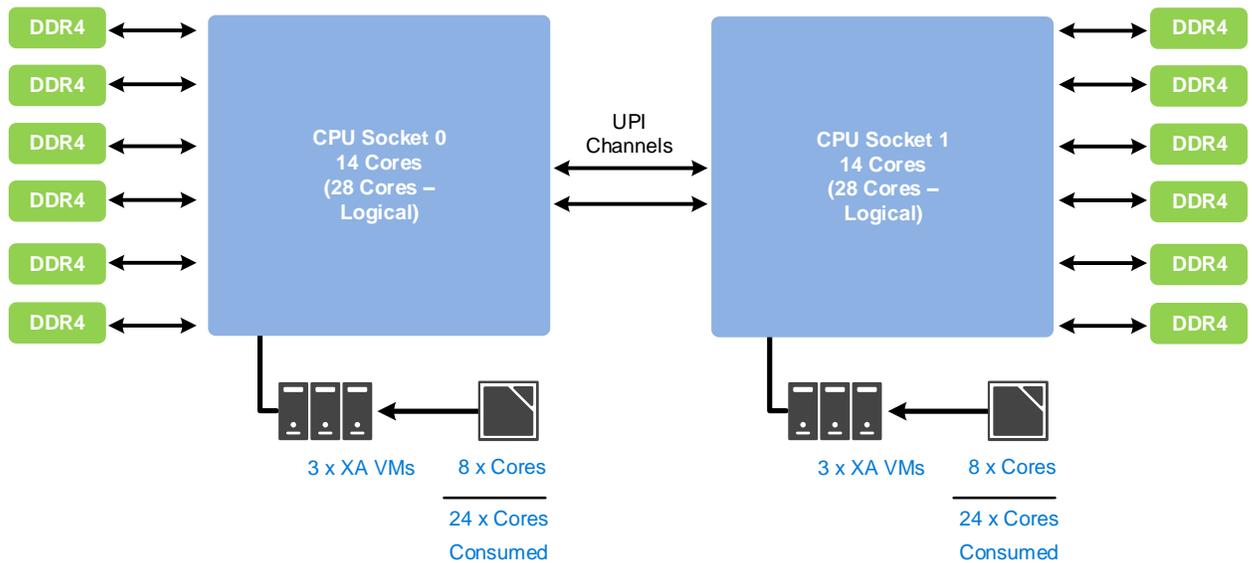
4.1.7.4 A3 NUMA alignment

10 physical cores per CPU given the Intel Xeon Silver 4114 part, 20 logical cores with Hyper-Threading active, gives us a total of 40 consumable cores per node. Configuring the XenApp VMs as shown below (XA VMs) will ensure that no physical NUMA node spanning occurs which could lower performance for an effected VM.



4.1.7.5 B5 NUMA alignment

14 physical cores per CPU in the B5 configuration, 28 logical with Hyper-threading active, gives a total of 56 consumable cores per node. Configuring the XenApp VMs as shown below (XA VMs) will ensure that no physical NUMA node spanning occurs which could lower performance for an effected VM.



NOTE: To ensure no NUMA spanning, remove the checkmark for “Allow virtual machines to span physical NUMA nodes” in the Hyper-V settings of the compute host(s). Also ensure “Maximum NUMA nodes allowed on a socket” is set to a value of 1. This setting is found under Processor > NUMA for each RDSH VM.

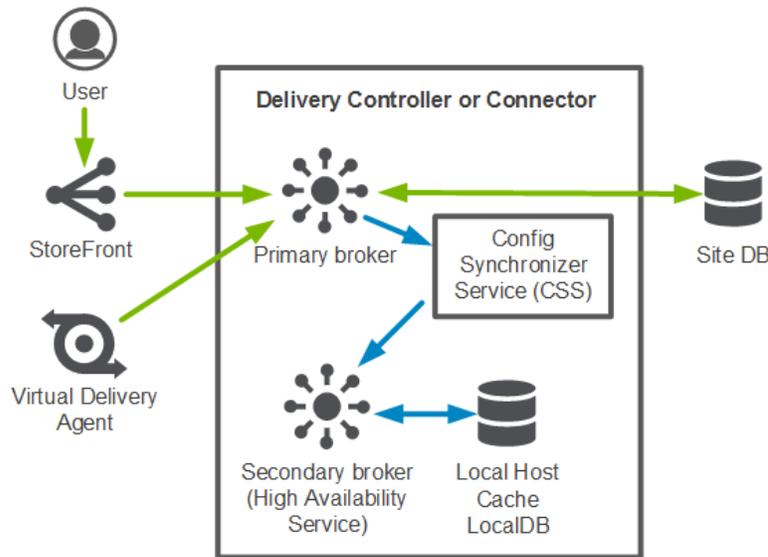
4.1.8 Local Host Cache

In XenApp and XenDesktop version 7.12 and above, the Local Host Cache (LHC) feature allows connection brokering operations to continue when connectivity to the Site database has been interrupted. This includes both failures between the Delivery Controller and Site database in on-premises deployments and when the WAN link between the Site and Citrix control plane fails in a Citrix Cloud environment.

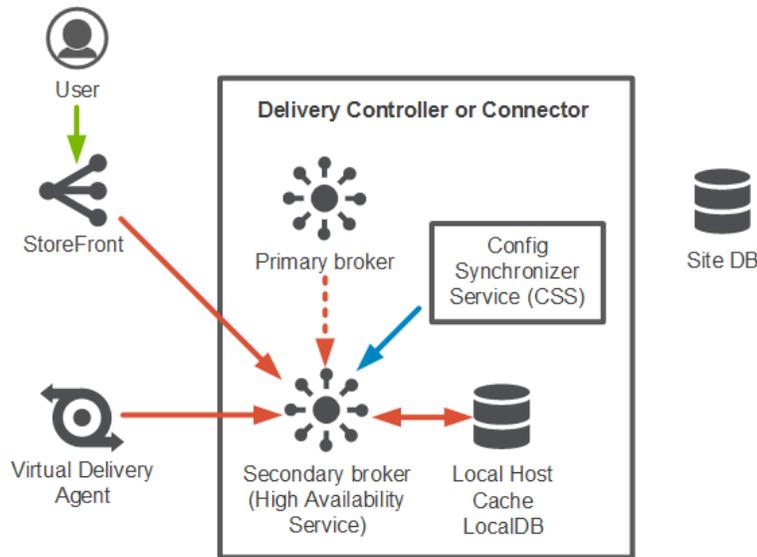
LHC replaces the connection leasing feature as the recommended XenApp and XenDesktop high availability solution. During an outage, LHC will support new users and existing users launching new resources, as well as users accessing pooled resources (shared desktops). Earlier versions of XenApp had a feature named Local Host Cache but this is an entirely different implementation that is more robust and immune to corruption.

The principal broker on a delivery controller accepts requests and communicates with the Site database to connect users. A check is made every two minutes to determine if changes have been made to the principal broker's configuration and if so, the information is synchronized with the secondary broker. All configuration data is copied to ensure the LocalDB database matches the site database.

The following diagram shows the communication paths during normal operations.



The diagram below illustrates changes in communication when the principal broker is unable to connect to the Site database.



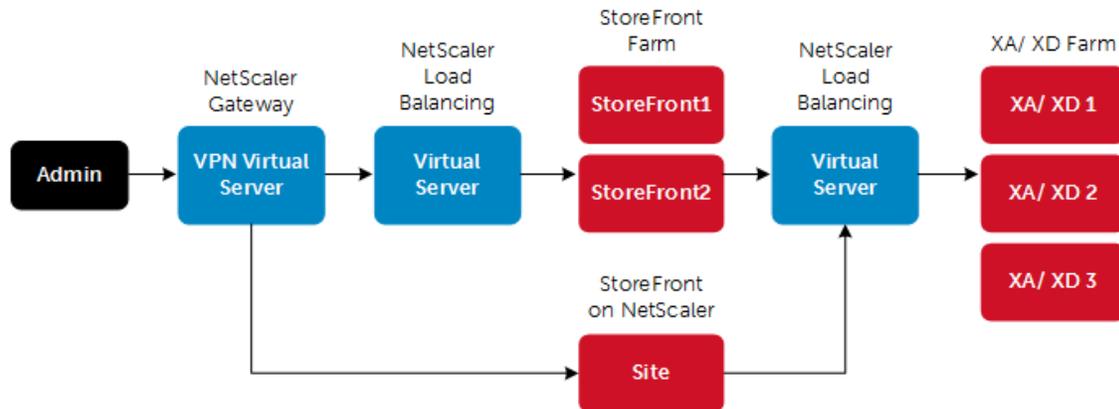
The principal broker stops listening for requests and instructs the secondary broker to begin listening and processing requests. When a VDA communicates with the secondary broker, a re-registration process is triggered during which current session information is delivered. During this time, the principal broker continually monitors the connection to the Site database. Once restored, the principal broker resumes brokering operations and instructs the secondary broker to stop listening for connection information.

4.1.9 Citrix NetScaler

Citrix NetScaler is an all-in-one web [application delivery controller](#) that makes applications run better, reduces web application ownership costs, optimizes the user experience, and makes sure that applications are always available by using:

- Proven application acceleration such as [compression](#) and [caching](#)
- High application availability through advanced L4-7 [load balancer](#)
- Application security with an integrated application firewall
- Server offloading to significantly reduce costs and consolidate servers

A NetScaler appliance resides between the clients and the servers, so that client requests and server responses pass through it. In a typical installation, virtual servers (vservers) configured on the NetScaler provide connection points that clients use to access the applications behind the NetScaler. In this case, the NetScaler owns public IP addresses that are associated with its vservers, while the real servers are isolated in a private network. It is also possible to operate the NetScaler in a transparent mode as an L2 bridge or L3 router, or even to combine aspects of these and other modes. NetScaler can also be used to host the StoreFront function eliminating complexity from the environment.



Global Server Load Balancing

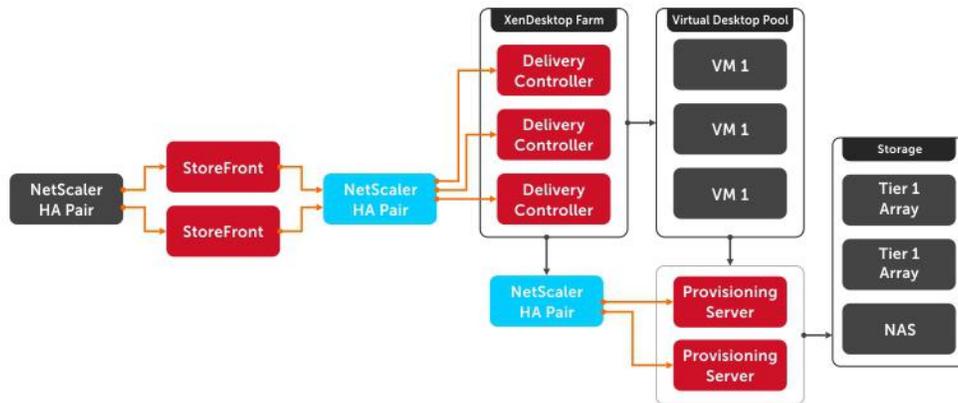
GSLB is an industry standard function. It is in widespread use to provide automatic distribution of user requests to an instance of an application hosted in the appropriate data center where multiple processing facilities exist. The intent is to seamlessly redistribute load on an as required basis, transparent to the user community. These distributions are used on a localized or worldwide basis. Many companies use GSLB in its simplest form.

They use the technology to automatically redirect traffic to Disaster Recovery (DR) sites on an exception basis. That is, GSLB is configured to simply route user load to the DR site on a temporary basis only in the event of a catastrophic failure or only during extended planned data center maintenance. GSLB is also used to distribute load across data centers on a continuous load balancing basis as part of normal processing.

NetScaler and XenDesktop Deployment Guide: [Link](#)

Several of the management components of the XenDesktop stack are made highly-available using NetScaler to load balance traffic. The following management components require the use of a load balancer to function in a high availability mode:

- StoreFront Servers
- Licensing Server
- XenDesktop XML Service
- XenDesktop Desktop Director
- Provisioning Services TFTP Service
- Framework UDP virtual channel (supported on NetScaler Gateway 11.0.62.10 or later and NetScaler Unified Gateway 11.0.64.34 or later)



4.2 Microsoft

4.2.1 Windows Server 2016 Hyper-V hypervisor

Hyper-V is Microsoft's hardware virtualization product built in to Windows Server. As with other hypervisors, Hyper-V allows you to create and run virtual machines. Hyper-V works with virtualization-aware hardware to tightly control the resources available to each virtual machine running in its own isolated space. Dell EMC servers include virtualization-aware CPUs and network adapters. Windows Server 2016 is the required server version for this architecture.

Hyper-V includes the following features:

Computing environment – A Hyper-V virtual machine includes the same basic parts as a physical computer, such as memory, processor, storage, and networking. All these parts have features and options that you can configure in different ways to meet different needs.

Disaster recovery and backup – For disaster recovery, Hyper-V Replica creates copies of virtual machines, intended to be stored in another physical location, so you can restore the virtual machine from the copy. For backup, Hyper-V offers two types. One uses saved states and the other uses Volume Shadow Copy Service (VSS) so you can make application-consistent backups for programs that support VSS.

Optimization – Each supported guest operating system has a customized set of services and drivers, called integration services that make it easier to use the operating system in a Hyper-V virtual machine.

Portability – Features such as live migration, storage migration, and import/export make it easier to move or distribute a virtual machine.

Remote connectivity – Hyper-V includes Virtual Machine Connection, a remote connection tool for use with both Windows and Linux. Unlike Remote Desktop, this tool gives you console access, so you can see what's happening in the guest even when the operating system isn't booted yet.

Security – Secure boot and shielded virtual machines help protect against malware and other unauthorized access to a virtual machine and its data.

Starting with Windows Server 2016, all physical CPU cores in a host must be licensed. Once all cores are licensed, customers are entitled to use 2 virtual machines with Standard Edition or unlimited virtual machines with Datacenter Edition. For Standard Edition, 2 additional VMs can be entitled by licensing all CPU cores again. For example, on a server with a total of 40 cores, 10 virtual machines can be used with Standard Edition by purchasing 200 core licenses (10 divided by 2 = 5, 5 x 40 cores = 200 core licenses). For VDI use where large numbers of VMs per host are possible, Datacenter Edition will typically be the most economical choice.

For additional details, Windows Server 2016 Licensing Datasheet found here: [LINK](#)

4.2.2 System Center 2016 Virtual Machine Manager

The management platform for a Hyper-V based virtualization environment is Microsoft Systems Center Virtual Machine Manager (SCVMM). SCVMM provides centralized and powerful management, monitoring, and self-service provisioning for virtual machines of an unmanaged personal collection only. SCVMM host groups are a way to apply policies and to check for problems across several VMs at once. Groups are organized by owner, operating system, or by custom names such as “Development” or “Production”. The interface also incorporates Remote Desktop Protocol (RDP); double-click a VM to bring up the console for that VM—live and accessible from the management console. SCVMM is required when Hyper-V is the hypervisor used in a Citrix XenDesktop and XenApp environment.

As with Windows Server 2016, System Center 2016 has transitioned from processors-based to core-based licensing, providing a consistent licensing metric between on-premises and cloud environments.

With System Center 2016:

- All physical cores in the server are required to be licensed.
- Each physical processor is required to be licensed with a minimum of 8 physical cores.
- Each physical server is required to be licensed with a minimum of 16 physical cores.
- Core licenses are sold in packs of two (i.e. 2-pack core licenses).

Datacenter edition is the appropriate choice for highly virtualized environments. For additional details, refer to the [Pricing and licensing overview for System Center 2016](#).

5 Solution architecture for XenDesktop

5.1 Management role configuration

The Management role recommendations 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.

Role	vCPU	Startup RAM (GB)	Dynamic Memory			NIC	OS + Data vDisk	
			Min Max	Buffer	Weight		Min Size (GB)	SDS Location
DDC + Lic	4	8	4GB 10GB	20%	Med	1	60	Mgmt
Storefront	2	4	2GB 6GB	20%	Med	1	60	Mgmt
Primary SQL*	8	8	4GB 16GB	20%	Med	1	60 + 300 (VHDX)	Mgmt
SCVMM	8	8	4GB 16GB	20%	Med	1	60 + 50 (VHDX)	Mgmt
PVS	4	24	12GB 24GB	20%	Med	1	60 + 200 (VHDX)	Mgmt
Total	26 vCPUs	52GB	26GB 72GB	-	-	5 vNICs	850GB	-

***NOTE:** Set memory to static 10GB if using SQL Standard version

5.1.1 XenApp VM configuration

The recommended number of XenApp VMs and their configurations on Hyper-V are summarized below and take into account proper NUMA balancing based on the applicable hardware configuration. For more information on NUMA please refer to the [NUMA Architecture Considerations](#) section.

HW Config	VMs per Host	vCPUs per VM	Startup RAM (GB)	Dynamic Memory			NIC	OS + Data vDisk	
				Min Max	Buffer	Weight		Min Size (GB)	SDS Location
A3	4	8	16	8GB 32GB	20%	Med	1	60	Mgmt
B5	6	8	32	16GB 48GB	20%	Med	1	60	Mgmt

5.1.2 SQL databases

The Citrix and Microsoft are hosted by a single dedicated SQL 2016 (or higher) 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
- vCenter or SCVMM

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.

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).

5.1.3 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.

5.1.3.1 DNS for SQL

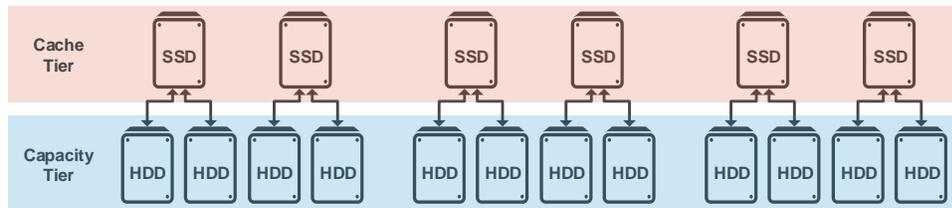
To access the SQL data sources, either directly or via ODBC, a connection to the server name\ instance name must be used. To simplify this process, as well as protect for future scaling (HA), instead of connecting to server names directly, alias these connections in the form of DNS CNAMEs. So instead 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.

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

5.2 Storage configuration overview

All Dell EMC S2D Ready Nodes come with two tiers of storage by default, SSD for performance and HDD for capacity. Additionally, all-flash configurations are available utilizing only SSD disks. Usable disk capacity is calculated using only the disks that participate in the capacity tier. Each disk tier is extensible and can be increased/decreased at any time; however, hardware minimums must be adhered to and all nodes in the cluster must have the exact same disk configurations. When the number of disks are modified, the Storage Bus Layer Cache dynamically adjusts the disk bindings to ensure a proper ratio of cache to capacity disks is maintained. This occurs when drives are intentionally modified as well as during self-healing when there is a drive failure. Example below is showing hybrid storage. Caching in all-flash is only used for writes.



When Storage Spaces Direct is enabled, a single new storage pool is created for the entire cluster. Dell EMC strongly recommends using only one pool per cluster. All eligible drives in the cluster nodes are automatically discovered and added to the pool. When you scale out the cluster, the drives in the new nodes are also automatically added to the pool and if any drives fail, they are retired and removed from the pool. Unlike the prior version of Storage Spaces, you typically do not need to manually perform tasks to manage the pool.

Once the pool has been created, virtual disks and volumes are manually provisioned either via the GUI or PowerShell cmdlet. When creating a volume, either NTFS or ReFS can be chosen as the file system. The preferred and recommended file system for S2D is ReFS as it provides advanced data corruption detection and recovery. The volumes created will become Cluster Shared Volumes in the failover cluster. Multiple volumes can be created to suit customer needs. Dell EMC recommends creating at least two volumes to separate compute and management VMs. The Microsoft recommended limitation is 32 volumes per cluster.

For additional details, see the Planning volumes in Storage Spaces Direct article here: [LINK](#)

5.2.1 Fault tolerance and resiliency

Fault tolerance in Storage Spaces is typically referred to as resiliency. Depending on the number of nodes in the cluster, different resiliency options will be available (see table below). Mirroring, as the name implies, creates 2 copies (two-way) or 3 copies (three-way) of data on discrete hosts. Parity uses bitwise algorithms to provide fault tolerance and is commonly referred to as Erasure Coding. If there are exactly three nodes in the cluster, S2D will automatically use three-way mirroring when a new volume is created. If there are four or more nodes, you add a parameter to choose a resiliency type when creating a volume. For best performance, Dell recommends using three-way mirroring for VDI even when there are three or more nodes in the cluster.

Resiliency	Failure Tolerance	Storage Efficiency	Minimum Hosts
Two-way mirror	1	50%	2
Three-way mirror	2	33.3%	3
Dual parity	2	50-80%	4
Mixed	2	33-80%	4

For more details, refer to the Fault tolerance and storage efficiency article here: [LINK](#)

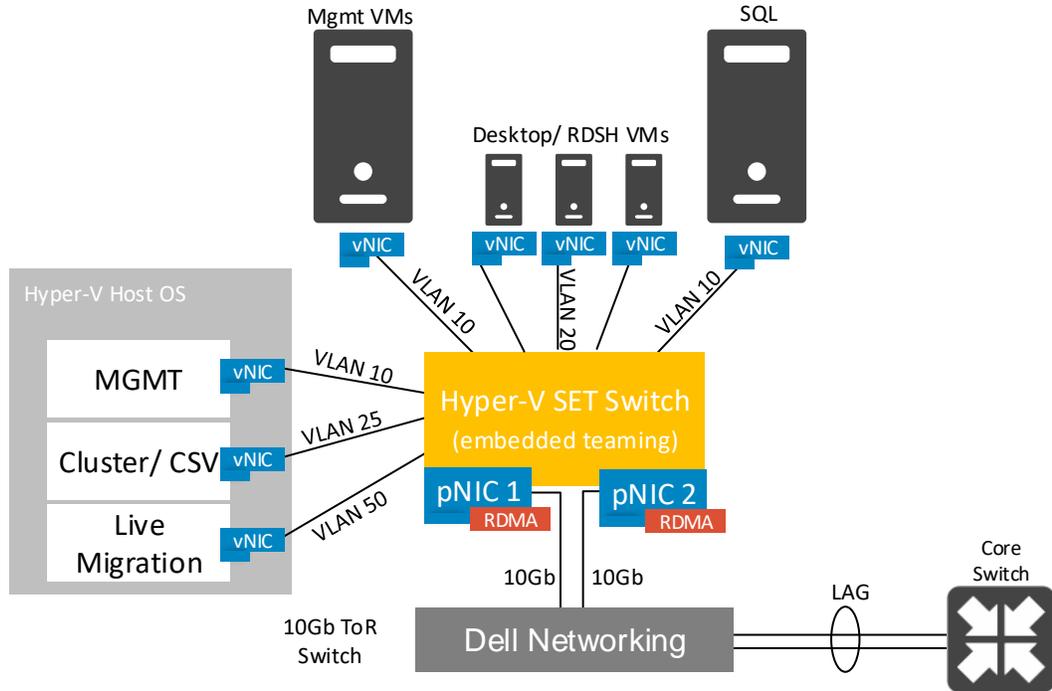
5.3 Virtual networking

The network configuration for the Dell EMC S2D Ready Nodes utilizes 2 x 25Gb RDMA NICs with a new virtual switch option introduced in Windows Server 2016 called Switch Embedded Teaming (SET). As the name suggest, SET has network teaming built in and allows the same two physical NIC ports to be used for VMs as well as the parent partition of the server for RDMA connectivity.

At a minimum, two vNICs for the management OS must be created. The vNICs can then be tagged with different or the same VLAN IDs (if using the same, ensure unique IP addresses). To further segment traffic, create a vNIC per VLAN designation. The diagram below shows the following suggested VLANs included.

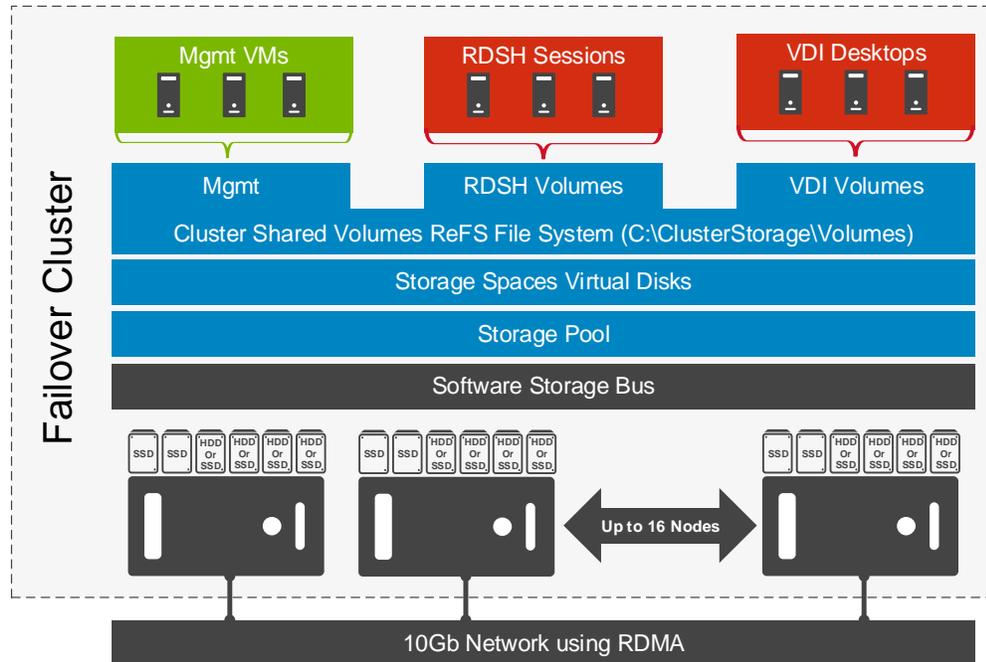
- Management VLAN: Configured for hypervisor infrastructure traffic – L3 routed via spine layer
- VDI VLAN: Configured for VDI session traffic – L3 routed via spine layer

- Cluster/CSV: Configured for failover cluster operations/traffic – L2 switched via leaf layer
- Live Migration VLAN: Configured for Live Migration traffic – L2 switched via leaf layer



5.4 Scaling guidance

Each component of the solution architecture scales independently according to the desired number of supported users. Additional nodes can be added at any time to expand the cluster and storage pool in a modular fashion.



- The components are scaled either horizontally (by adding additional physical and virtual servers to the server pools) or vertically (by adding virtual resources to the infrastructure)
 - Eliminate bandwidth and performance bottlenecks as much as possible
- Allow future horizontal and vertical scaling with the objective of reducing the future cost of ownership of the infrastructure.

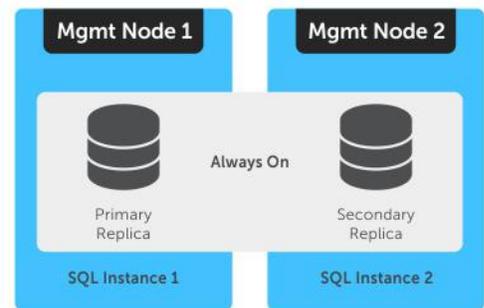
Component	Metric	Horizontal scalability	Vertical scalability
Ready Nodes	VMs per physical host	Add additional nodes	Additional RAM or CPU compute power
Desktop Delivery Servers	Desktops per instance (dependent on SQL performance as well)	Additional servers added to the XenDesktop Site	Additional virtual machine resources (RAM and CPU)
Provisioning Servers	Desktops per instance	Additional servers added to the Provisioning Server farm	Additional network and I/O capacity added to the servers

XenApp Servers	Desktops per instance	Additional virtual servers added to the XenDesktop Site	Additional nodes to host virtual XenApp servers.
Storefront Servers	Logons/ minute	Additional servers added to the Storefront environment	Additional virtual machine resources (RAM and CPU)
Database Services	Concurrent connections, responsiveness of reads/writes	Migrate databases to a dedicated SQL server and increase the number of management nodes	Additional RAM and CPU for the management nodes

5.5 Solution high availability

High availability (HA) is offered to protect each architecture solution layer, individually if desired. Following the N+1 model, additional ToR switches are added to the Network layer and stacked to provide redundancy as required, additional nodes are added, SQL is configured for AlwaysOn or clustered and NetScaler is leveraged for load balancing.

The HA options provide redundancy for all critical components in the stack while improving the performance and efficiency of the solution as a whole.



- Additional switches added to the existing thereby equally spreading each host's network connections across multiple switches.
 - Additional nodes added to provide N+1 protection (up to cluster maximum)
- Applicable Citrix infrastructure server roles are duplicated and spread amongst nodes where connections to each are load balanced via the addition of virtual NetScaler appliances.
- Each PVS Server instance is responsible for the Write Cache of each desktop that it is hosting. This Write Cache is readable by the resilient Provisioning Server and as it is held on the target device. In the event of a Provisioning Server failure, all desktops that were hosted will transfer to an alternate provisioning server for that site and users will be unaware of the failure.
- SQL Server databases also are protected through the addition and configuration of an "AlwaysOn" Failover Cluster Instance or Availability Group.

Please refer to these links for more information: [LINK1](#) and [LINK2](#)

6 Solution Performance and Testing

At the time of publication, here are the available user density recommendations per compute server. Please refer to the Platform Configurations section for hardware specifications.

User density summary

Host Config	Hypervisor	Deployment / Provisioning	Workload	Template	User Density
S2DRN B5	Hyper-V 2016	XenDesktop 7.15 LTSR	Task Worker	Windows 10 & Office 2016	175
S2DRN B5	Hyper-V 2016	XenDesktop 7.15 LTSR	Knowledge Worker	Windows 10 & Office 2016	140
S2DRN B5	Hyper-V 2016	XenDesktop 7.15 LTSR	Power Worker	Windows 10 & Office 2016	135

The detailed validation results and analysis of these reference designs are in the following sections.

6.1 Test and performance analysis methodology

6.1.1 Testing process

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.

Login VSI is currently the load-generation tool used during PAAC of Dell Wyse Datacenter solutions. Each user load is tested against multiple runs. First, a pilot run to validate that the infrastructure is functioning and valid data can be captured, and then, subsequent runs allowing correlation of data.

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

6.1.1.1 Load generation

Login VSI by Login Consultants is the de-facto industry standard tool for testing VDI environments and server-based computing (RDSH 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). The launchers and Login VSI environment are configured and managed by a centralized management console.

Additionally, the following login and boot paradigm is used:

- Users are logged in within a login timeframe of 1.5 hours. Exception to this login timeframe occurs when testing low density solutions such as GPU/graphics based configurations. With those configurations, users are logged on every 10-15 seconds.
 - All desktops are pre-booted in advance of logins commencing.
- All desktops run an industry-standard anti-virus solution. Windows Defender is used for Windows 10 due to issues implementing McAfee.

6.1.1.2 Profiles and workloads

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 with two targeted at graphics-intensive use cases (graphics currently not tested with this solution). We will present more detailed information in relation to these workloads and profiles below but first it is useful to define the terms "profile" and "workload" 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.)

Load-testing on each profile is carried out using an appropriate workload that is representative of the relevant use case and summarized in the table below:

Profile to workload mapping:

Profile Name	Workload
Task Worker	Login VSI Task worker
Knowledge Worker	Login VSI Knowledge worker
Power Worker	Login VSI Power worker

Login VSI workloads are summarized in the sections below. Further information for each workload can be found on Login VSI's [website](#).

Login VSI Task Worker Workload

The Task Worker workload runs fewer applications than the other workloads (mainly Excel and Internet Explorer with some minimal Word activity, Outlook, Adobe, copy and zip actions) and starts/stops the applications less frequently. This results in lower CPU, memory and disk IO usage.

Login VSI Knowledge Worker Workload

The Knowledge Worker workload is designed for virtual machines with 2vCPUs. This workload and contains the following activities:

- 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.
- Various copy and zip actions.

Login VSI Power Worker Workload

The Power Worker workload is the most intensive of the standard workloads. The following activities are performed with this workload:

- Begins by opening four instances of Internet Explorer which remain open throughout the workload.
- Begins by opening two instances of Adobe Reader which remain open throughout the workload.
- There are more PDF printer actions in the workload as compared to the other workloads.
- Instead of 480p videos a 720p and a 1080p video are watched.
- The idle time is reduced to two minutes.
- Various copy and zip actions.

6.1.2 Resource monitoring

The following sections explain respective component monitoring used across all Dell Wyse Datacenter solutions where applicable.

6.1.2.1 Microsoft Performance Monitor

Microsoft Performance Monitor is used for Hyper-V based solutions to gather key data (CPU, Memory, Disk and Network usage) from each of the compute hosts during each test run. This data is exported to .csv files for single hosts and then consolidated to show data from all hosts (when multiple are tested). While the report does not include specific performance metrics for the Management host servers, these servers are monitored during testing to ensure they are performing at an expected performance level with no bottlenecks.

6.1.3 Resource utilization

Poor end-user experience is one of the main risk factors when implementing desktop virtualization but a 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 does not happen, PAAC on Dell Wyse Datacenter solutions monitors the relevant resource utilization parameters and applies relatively conservative thresholds as shown in the table below. 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. Utilization within these thresholds is used to determine the number of virtual applications or 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 (Hyper-V)	85%
Physical Host Memory Utilization	90%
Network Throughput	85%
Storage IO Latency	20ms
LVSI Failed/Unresponsive Sessions	5%

6.2 Test configuration details

The following components were used to complete the validation testing for the solution:

Hardware and software test components

Component	Description/Version
Hardware platform(s)	Dell EMC Storage Spaces Direct Ready Node (B5)
Hypervisor(s)	Microsoft Hyper-V 2016
Broker technology	Citrix XenDesktop 7.15 LTSR
Management VM OS	Microsoft Windows Server 2016
Hypervisor Management	Microsoft System Center VMM 2016
Virtual desktop OS	Microsoft Windows 10 64-bit
Office application suite	Microsoft Office 2016 Professional Plus
Login VSI test suite	4.1.25

6.2.1 Compute VM Configurations

The following table summarizes the compute VM configurations for the various profiles/workloads tested.

Desktop VM specifications:

User Profile	vCPUs	Hyper-V Startup Memory	Hyper-V Min Max Dynamic	Screen Resolution	Operating System
Task Worker	2	1GB	1GB 2GB	1280 X 720	Windows 10 Enterprise 64-bit
Knowledge Worker	2	1.5GB	1.5GB 3GB	1920 X 1080	Windows 10 Enterprise 64-bit
Power Worker	2	2GB	1GB 4GB	1920 X 1080	Windows 10 Enterprise 64-bit

RDSH VM specifications:

Platform Config	vCPUs	Hyper-V Startup Memory	Hyper-V Min Max Dynamic	Operating System
XenApp VMs	8	16GB	8GB 48GB	Windows Server 2016

6.2.2 Platform Configurations

The hardware configurations that were tested are summarized in the table(s) below.

S2DRN hardware configuration

Enterprise Platform	Platform Config	CPU	Memory	RAID Ctr	HD Config	Network
S2DRN	B5	5120 Gold (14 Core 2.2 GHz)	384GB @2400 MT/s	Dell HBA 330 Mini	2x 960 GB SSD 4x 2TB HDD	2x Mellanox ConnectX-4 Lx

6.3 Test results and analysis

The following table summarizes the test results for the compute hosts using the various workloads and configurations. Refer to the prior section for platform configuration details.

Test result summary:

Platform Config	Hypervisor	Broker & Provisioning	Login VSI Workload	Density Per Host	Avg CPU	Avg Mem Consumed	Avg IOPS / User	Avg Net Mbps / User
S2DRN	Hyper-V 2016	XenDesktop MCS	Task Worker	175	57	348	3.57	5.15
S2DRN	Hyper-V 2016	XenDesktop MCS	Knowledge Worker	140	59	345	5.27	10.20
S2DRN	Hyper-V 2016	XenDesktop MCS	Power Worker	135	77	344	6.35	15.02

Density per Host: Density reflects number of users per compute host that successfully completed the workload test within the acceptable resource limits for the host. For clusters, this reflects the average of the density achieved for all compute hosts in the cluster.

Avg CPU: This is the average CPU usage over the steady state period. For clusters, this represents the combined average CPU usage of all compute hosts.

Avg Consumed Memory: The amount of physical memory used by a host during the steady state phase. For clusters, this is the average consumed memory across all compute hosts over the steady state period.

Avg IOPS/User: IOPS calculated from the average Disk IOPS figure over the steady state period divided by the number of users.

Avg Net Mbps/User: Amount of network usage over the steady state period divided by the number of users. For clusters, this is the combined average of all compute hosts over the steady state period divided by the number of users on a host.

Provisioning data:

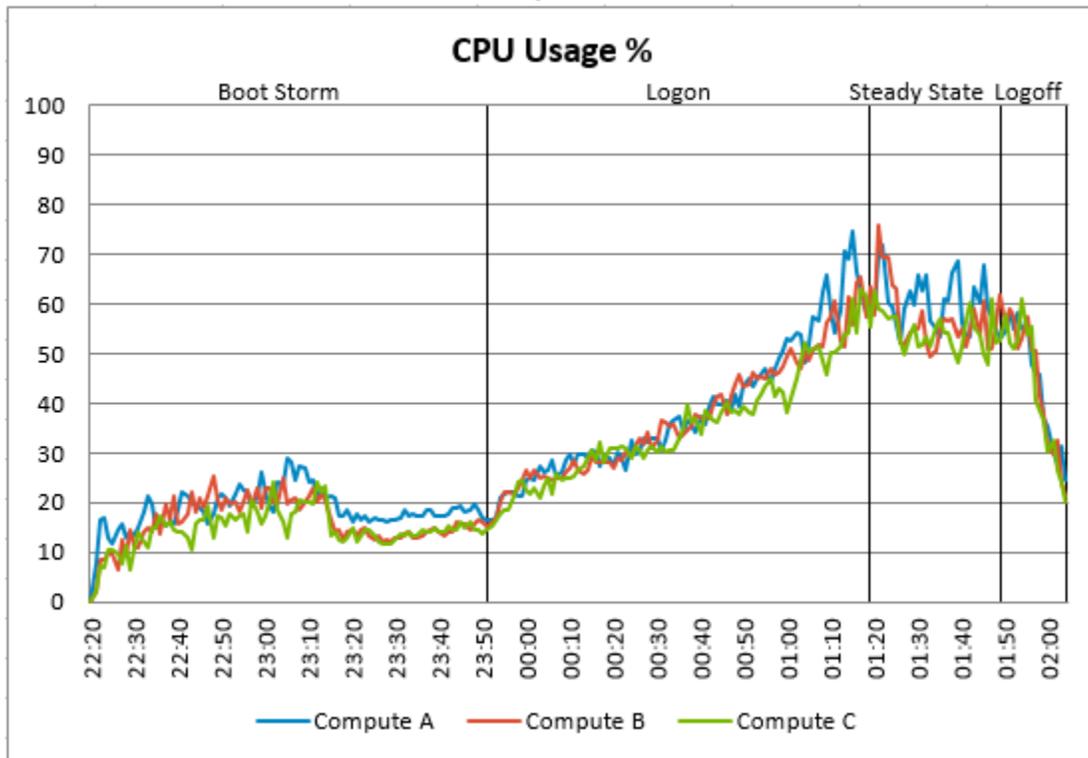
Platform Config	Broker & Provisioning	VM Profile	Pool Size	Provisioning Time	Time per VM	VM Concurrency
S2DRN-B5	XenDesktop MCS	Task Worker	525	1:42	:10	1
S2DRN-B5	XenDesktop MCS	Knowledge Worker	450	1:21	:10	1
S2DRN-B5	XenDesktop MCS	Power Worker	405	1:19	:11	1

6.3.1 Dell EMC S2DRN B5 Configuration with XenDesktop

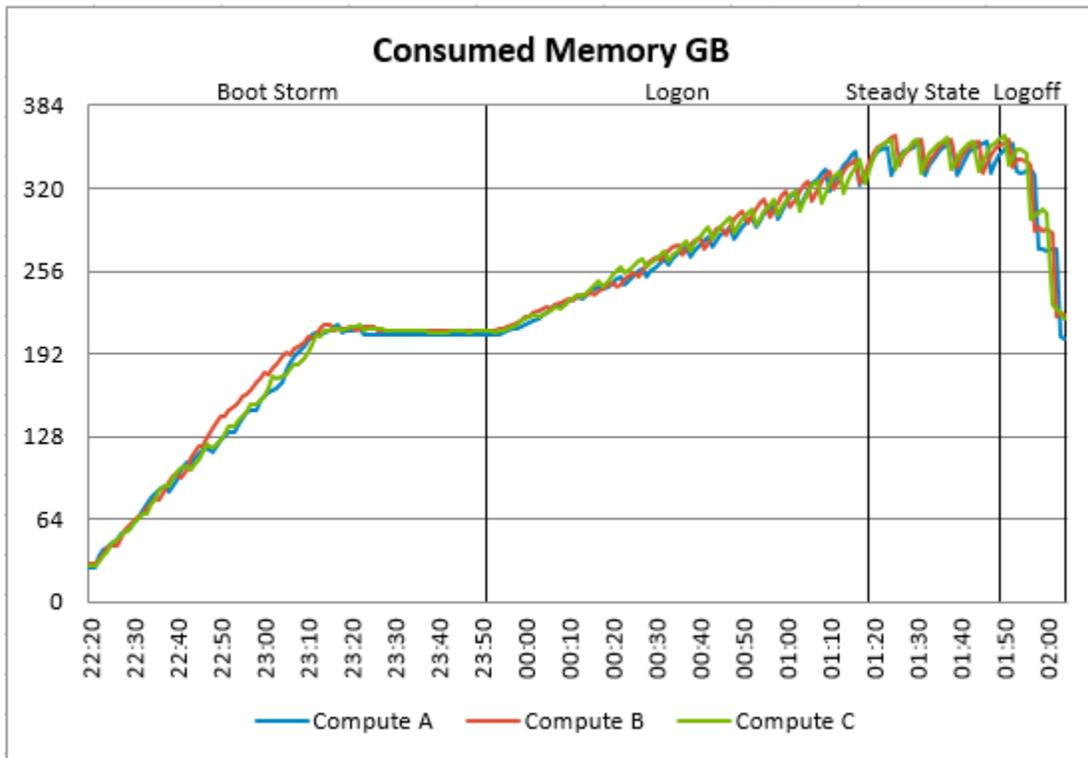
Refer to the Platform Configurations section for hardware configuration details.

6.3.1.1 Task Worker, 175 Users, Hyper-V 2016, XD + MCS

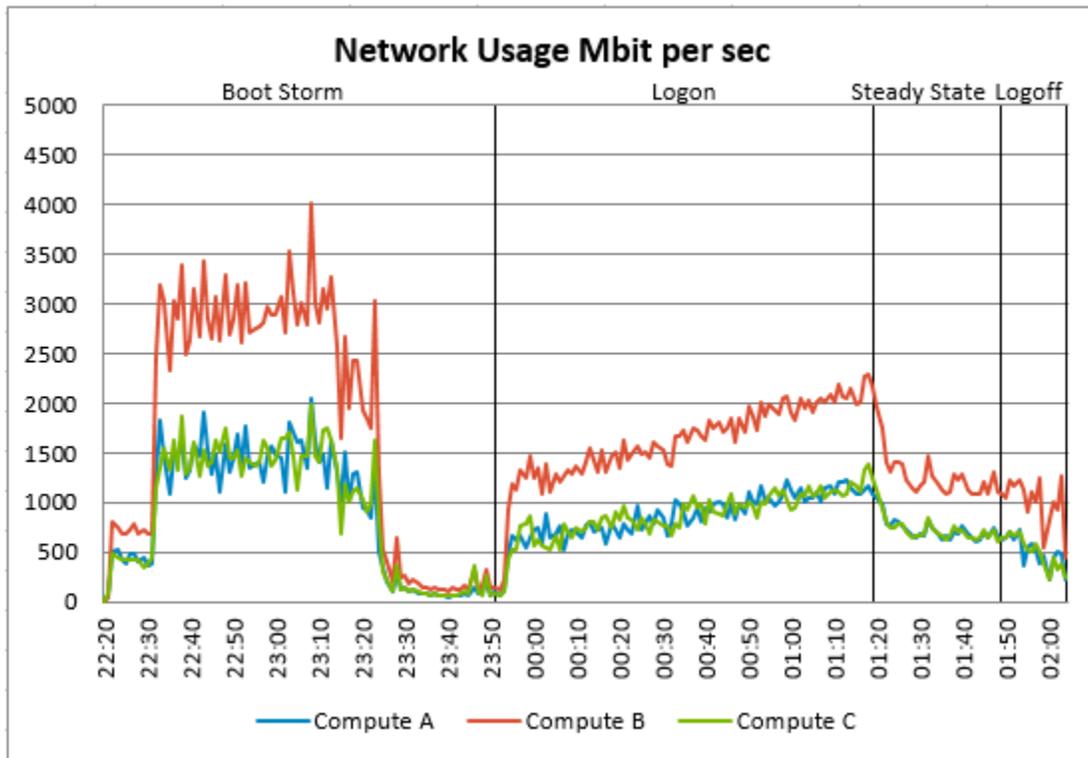
The Task Worker test run provisioned 175 VMs on each host (525 in the cluster) along with the XenDesktop management VM roles. The Peak CPU usage was 76% on one node at the end of the Logon phase, while the steady state average CPU usage was 57%.



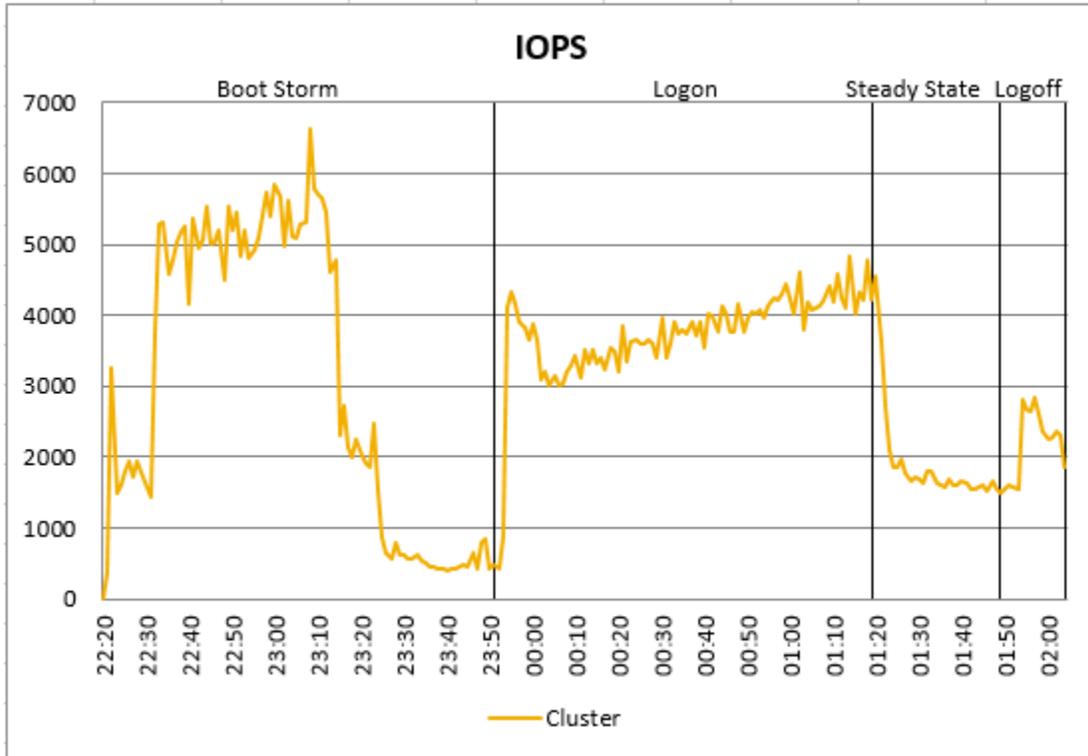
The memory used during the test run averaged 348 GB, and the peak memory usage was 361 GB on one node during Steady State. The average usage was 90.6% of the total memory, slightly over the threshold but acceptable. We estimate that each user session used 1.94 GB during steady state. The saw-tooth pattern reflects the hypervisor managing the memory demand of the individual VMs as they cycled through the workload loops.



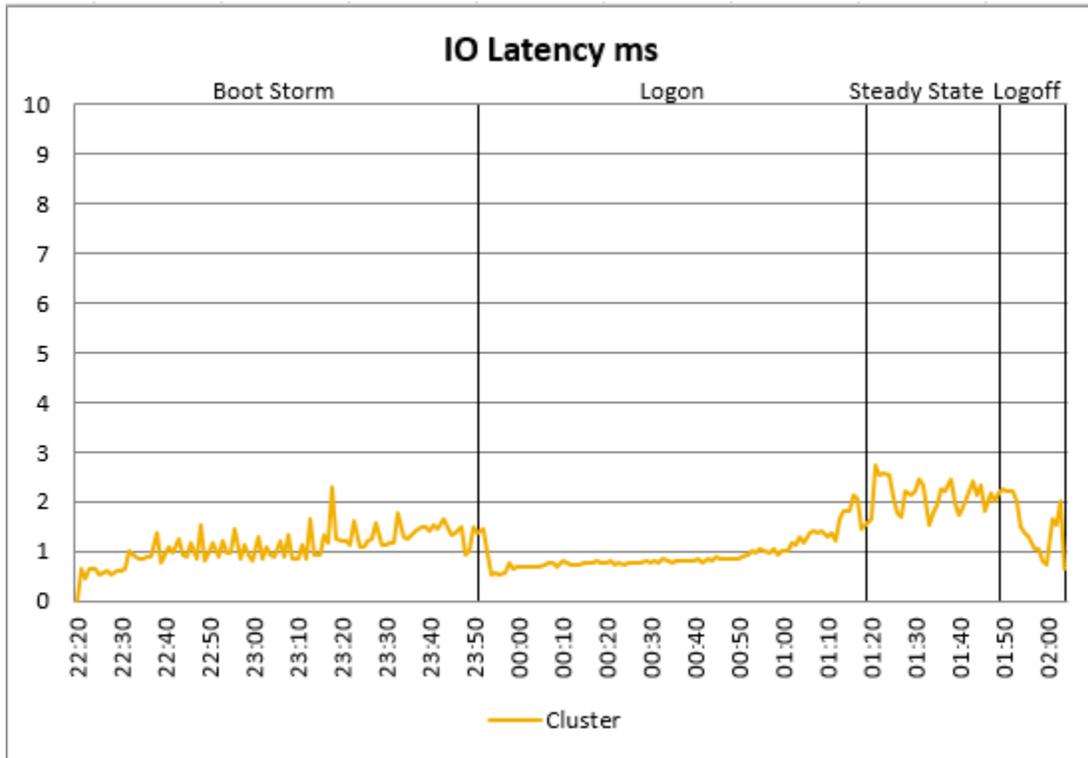
The Steady State average network usage was 900Mbps. The peak was 4012Mbps during the Boot Storm on node Compute B. In this test run Compute B was the node owner of the Cluster Shared Volume hosting all of the desktop VMs, therefore it saw more network traffic than the other nodes. With ReFS in use as the file system, IO is redirected to the owner (“coordinating node”) of the volume. On average each desktop consumed 5.15Mbps during steady state.



The peak Cluster IOPS for the test run was 6619 IOPS during Logon phase, while the average in Steady State was 1877 IOPS. This data was captured from the sum of the “PhysicalDisk(x)\Disk Transfers/sec” metrics for the mounted virtual disks only. Based on these numbers each user session generated 3.57 IOPS.

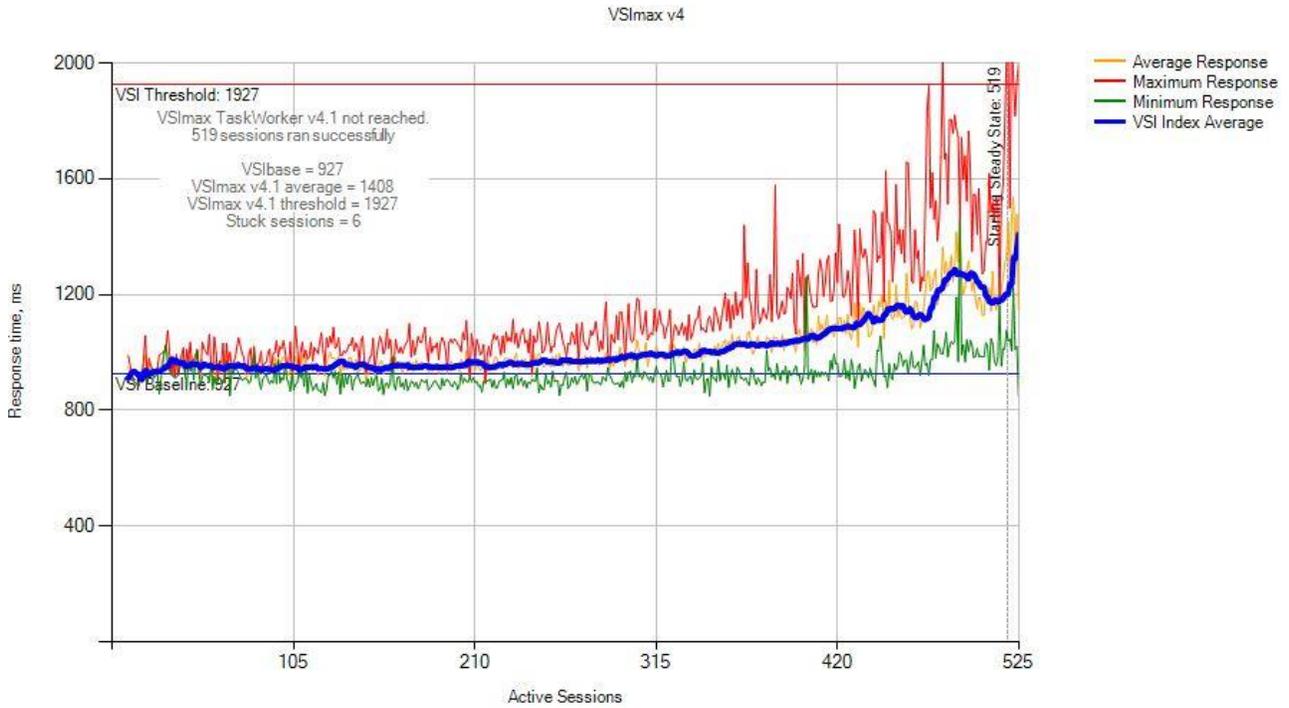


The IO Latency was captured using the “PhysicalDisk(x)\Avg. sec/Transfer” metrics for the mounted virtual disks. The peak IO Latency was 2.7 ms during the Boot Storm. The average IO latency during steady state was 2.1 ms. The chart clearly shows a very steady and very low level of IO Latency throughout the test run.



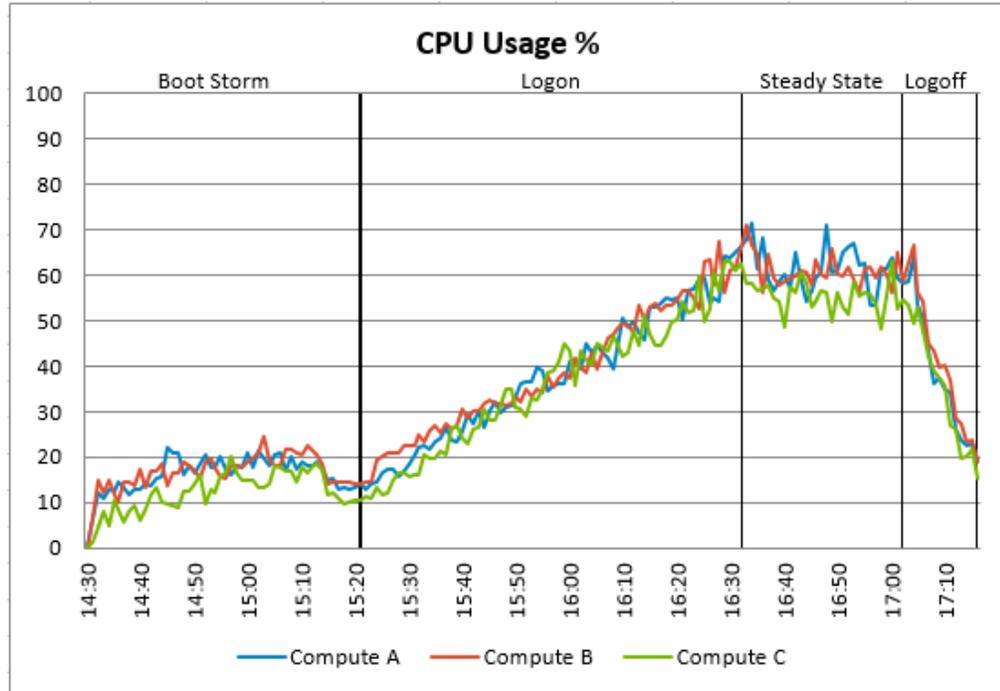
The baseline performance of 927 indicated that the user experience for this test run was Good. The Index average reached 1408, well below the threshold of 1927. However, 6 sessions were stuck during the test.

Login VSI Baseline	VSI Index Average	VSI Max Reached	VSI Threshold
927	1408	No	1927

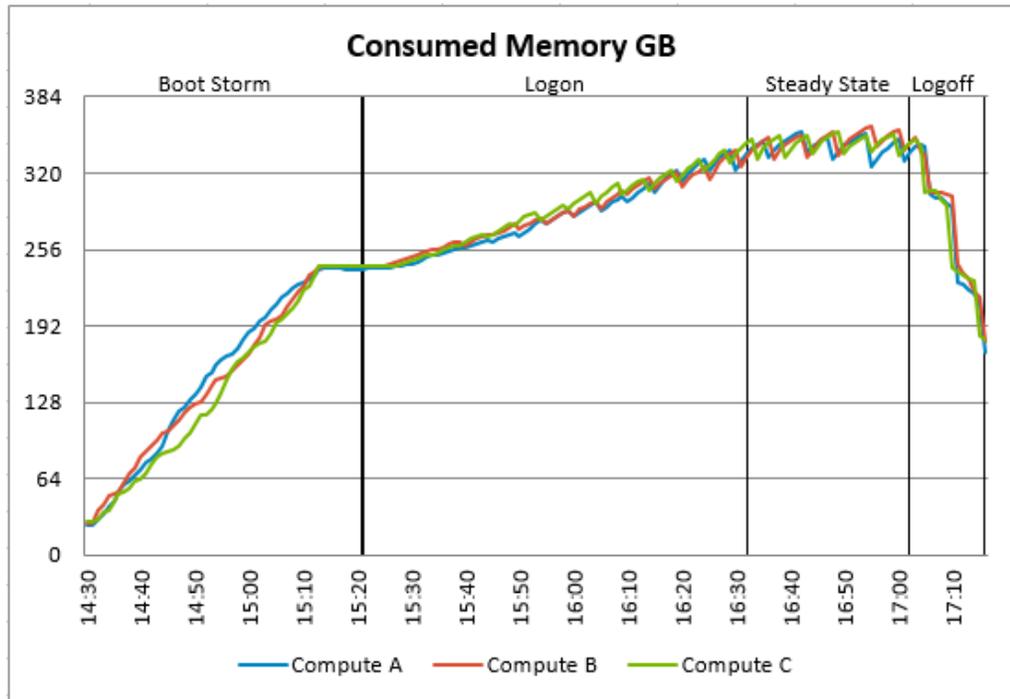


6.3.1.2 Knowledge Worker, 140 Users, Hyper-V 2016, XD + MCS

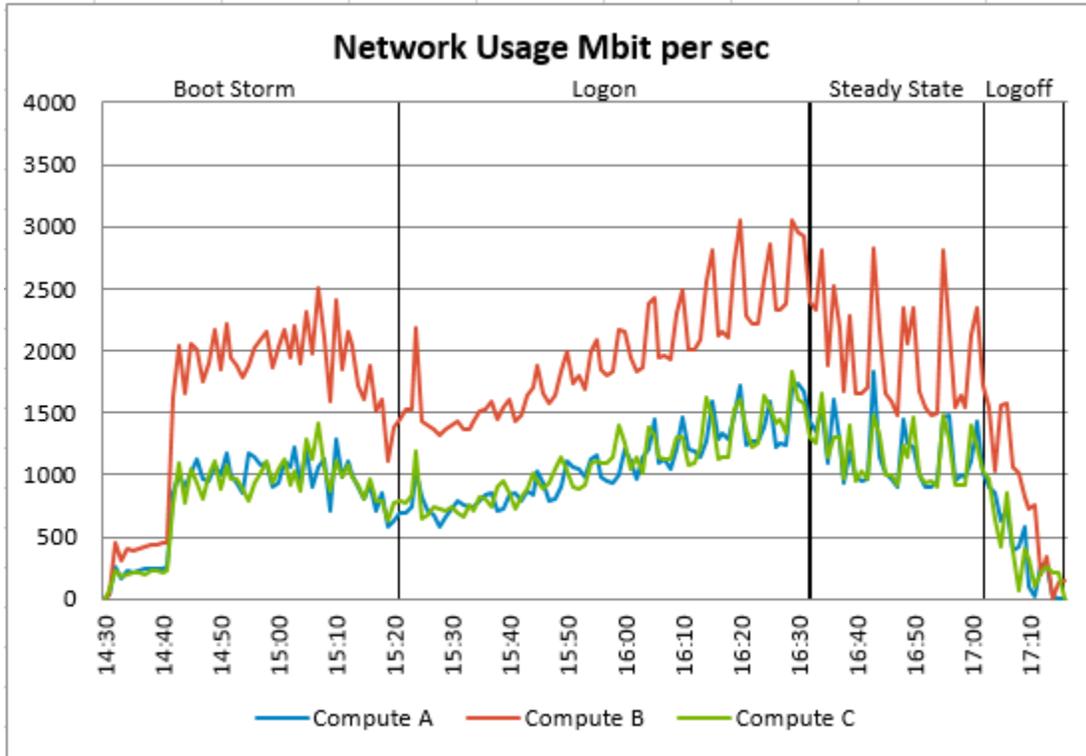
In this workload test run the hosts each had 140 user sessions (420 in the cluster) in addition to the various management VMs. The peak CPU Usage was 72% on one host during logon phase, while the Steady State average was 59% across all hosts.



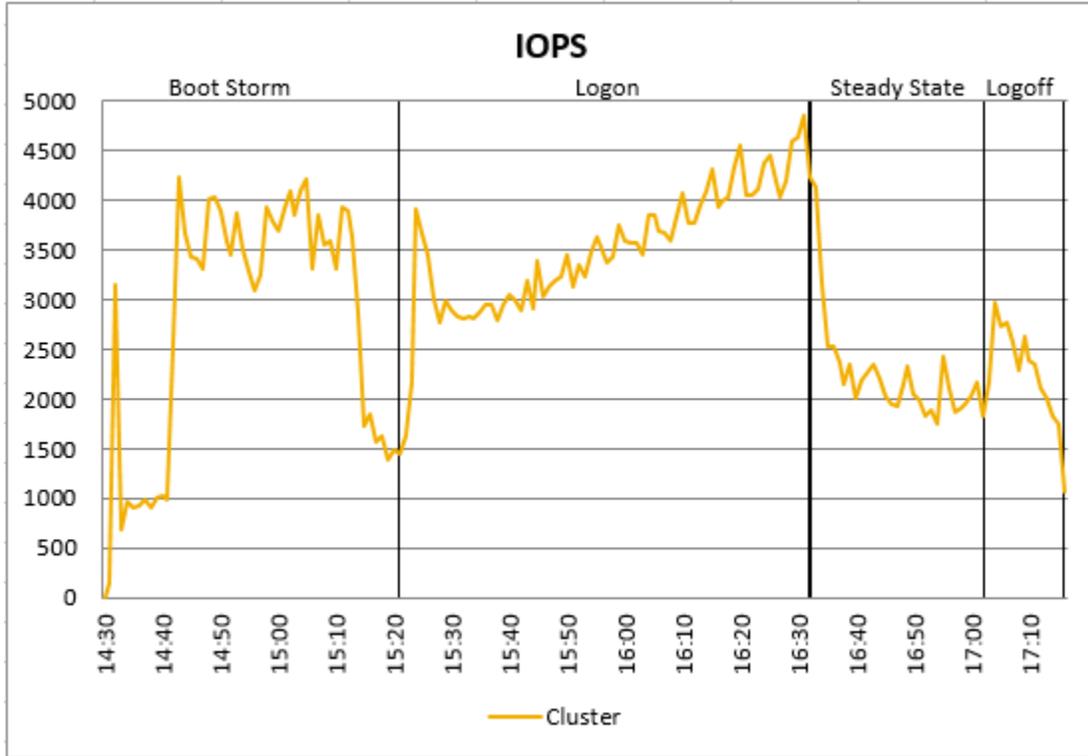
The memory consumption averaged 345 GB in steady state, and the peak usage on any host was 360 GB during the Steady State phase. The peak usage was 93.6%, while the Steady State average usage was 89.8%, just under the 90% threshold. We estimate that each session used 2.41 GB during Steady State.



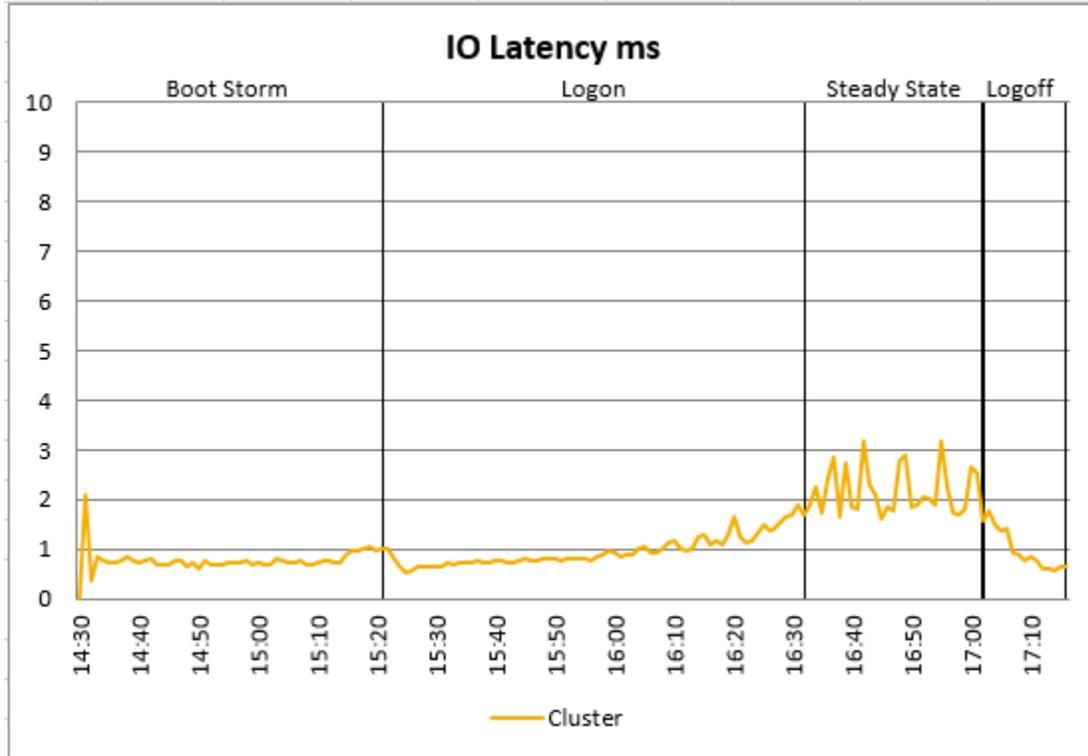
Network usage peaked at 3047 Mbps during Steady State on one host, and the average network usage for all hosts was 1428 Mbps during Steady State. The Compute Cluster Shared Volume was owned by host Compute B during this test run, therefore it received higher network traffic than the other nodes.



The peak Cluster IOPS for the test run was 4859 IOPS at the end of Logon phase, while the average in Steady State was 2214 IOPS. This data was captured from the sum of the “PhysicalDisk(x)\Disk Transfers/sec” metrics for the mounted virtual disks only. Based on these numbers each user session generated 5.27 IOPS.

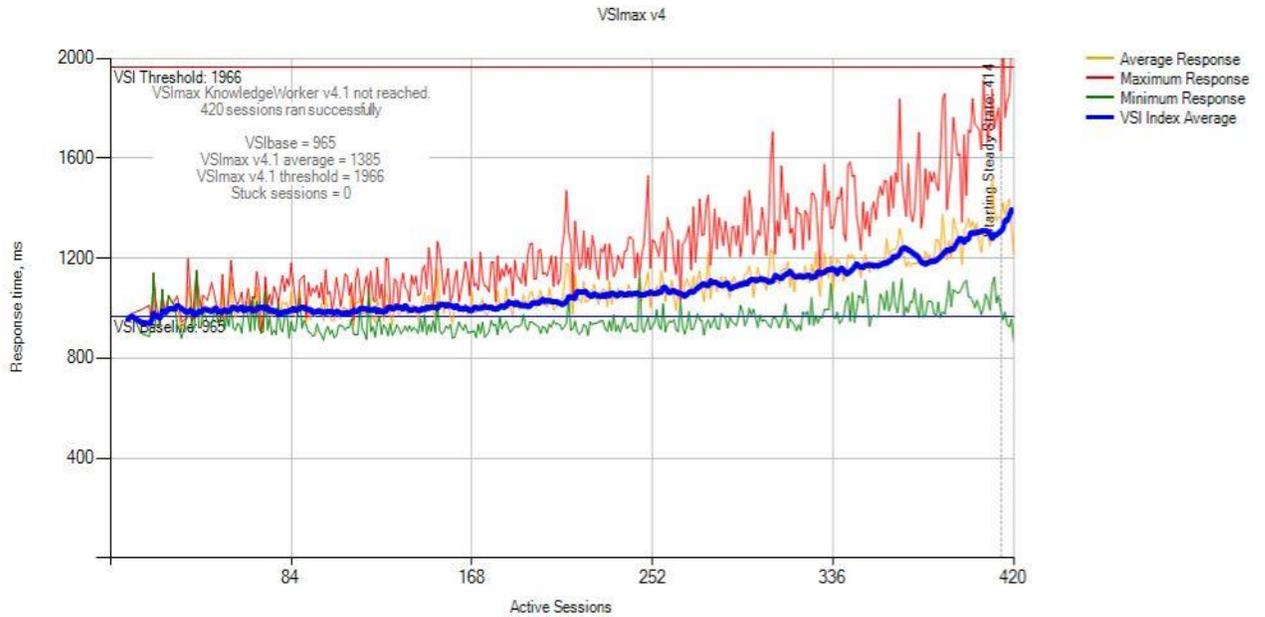


The IO Latency was captured using the “PhysicalDisk(x)\Avg. sec/Transfer” metrics for the mounted virtual disks. The peak IO Latency was 3.2 ms during Steady State. The average IO latency during steady state was 2.2 ms.



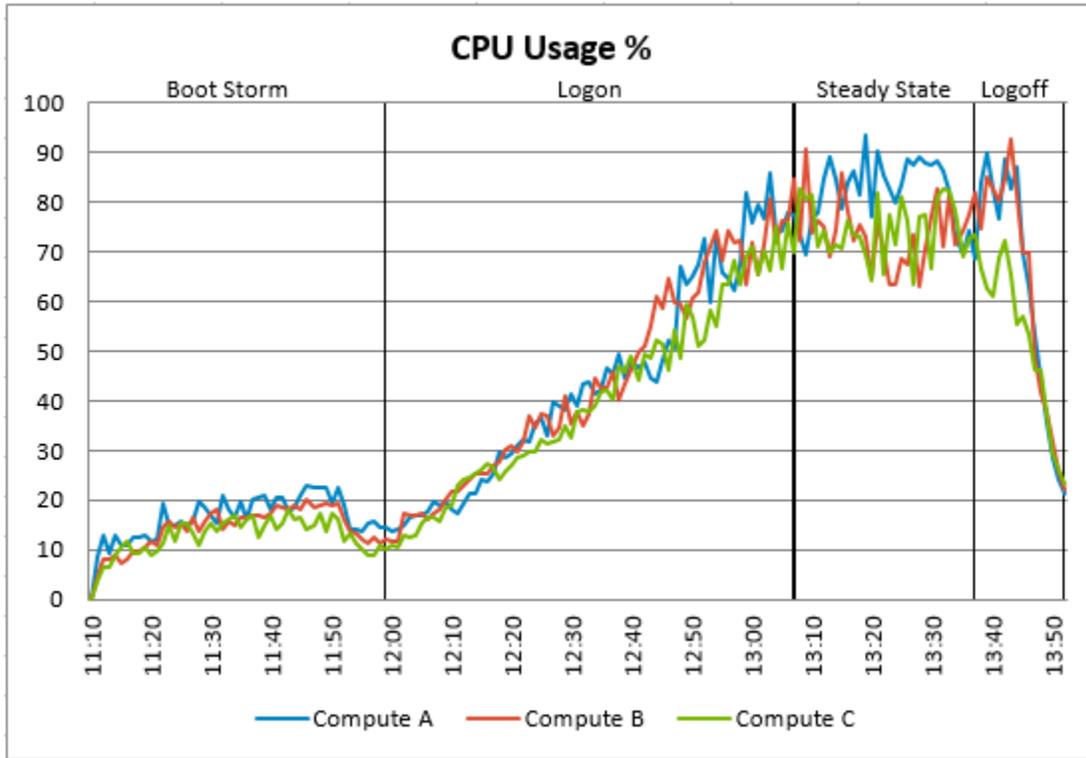
The baseline performance of 965 indicated that the user experience for this test run was Good. The Index average reached 1385, well below the threshold of 1966.

Login VSI Baseline	VSI Index Average	VSI Max Reached	VSI Threshold
965	1385	No	1966

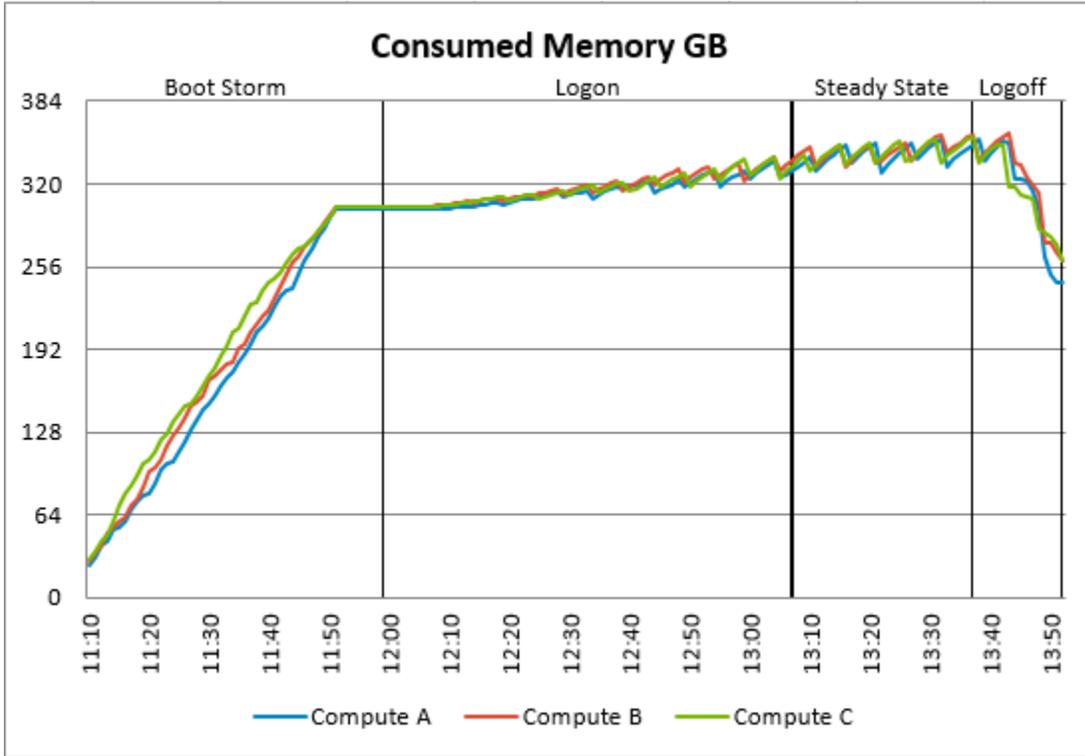


6.3.1.3 Power Worker, 135 Users, Hyper-V 2016, XD + MCS

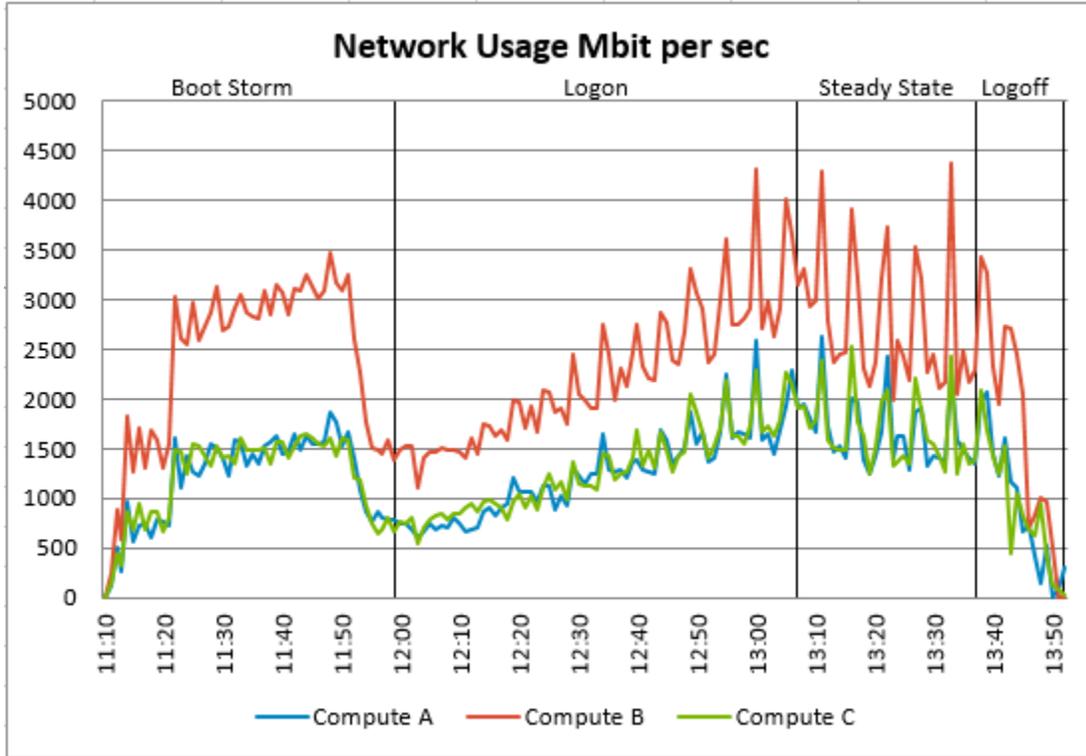
In this workload test run the hosts each had 135 user sessions (405 in the cluster) in addition to the various management VMs. The peak CPU Usage was 93% on one host during Steady State, while the Steady State average was 77% across all hosts. Since the Management roles for XenDesktop were hosted on the cluster, some additional CPU usage occurred besides the desktops themselves.



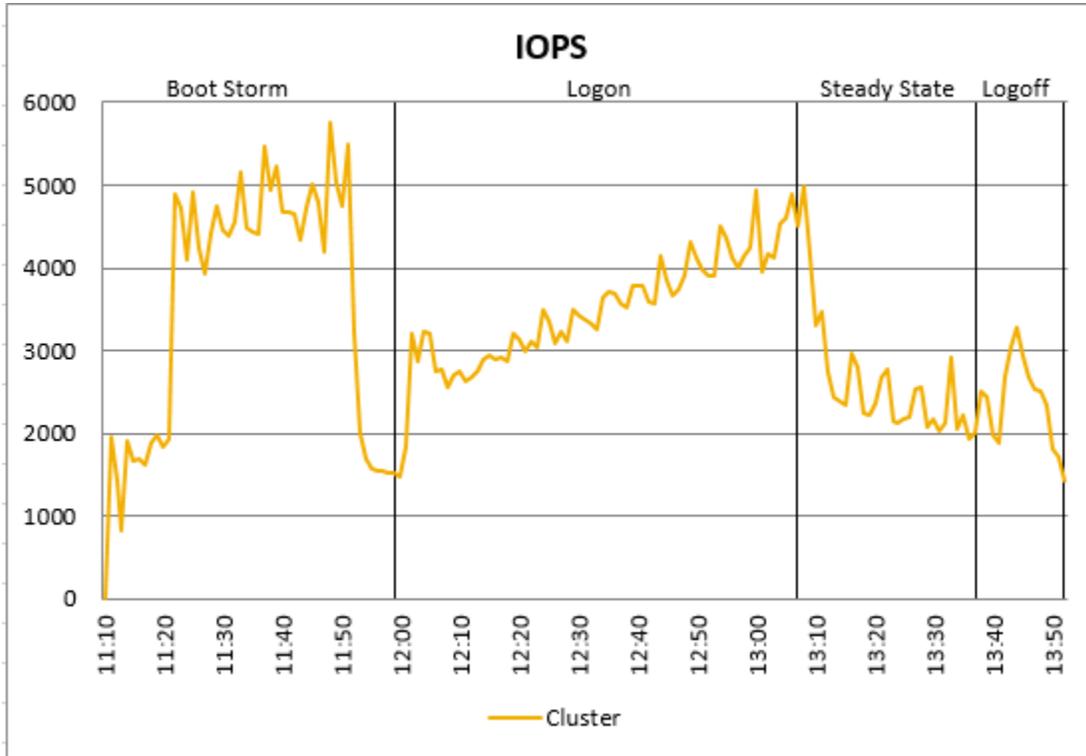
The memory consumption averaged 344 GB in steady state, and the peak usage on any host was 360 GB during the logoff phase. The peak usage percentage was 93.7%, while the Steady State average percentage was 89.6%, just under the 90% threshold (346 GB). The Management Roles used approximately 8GB of memory per host. Accounting for this additional usage we estimate that the desktops averaged 2.49 GB per session during Steady State.



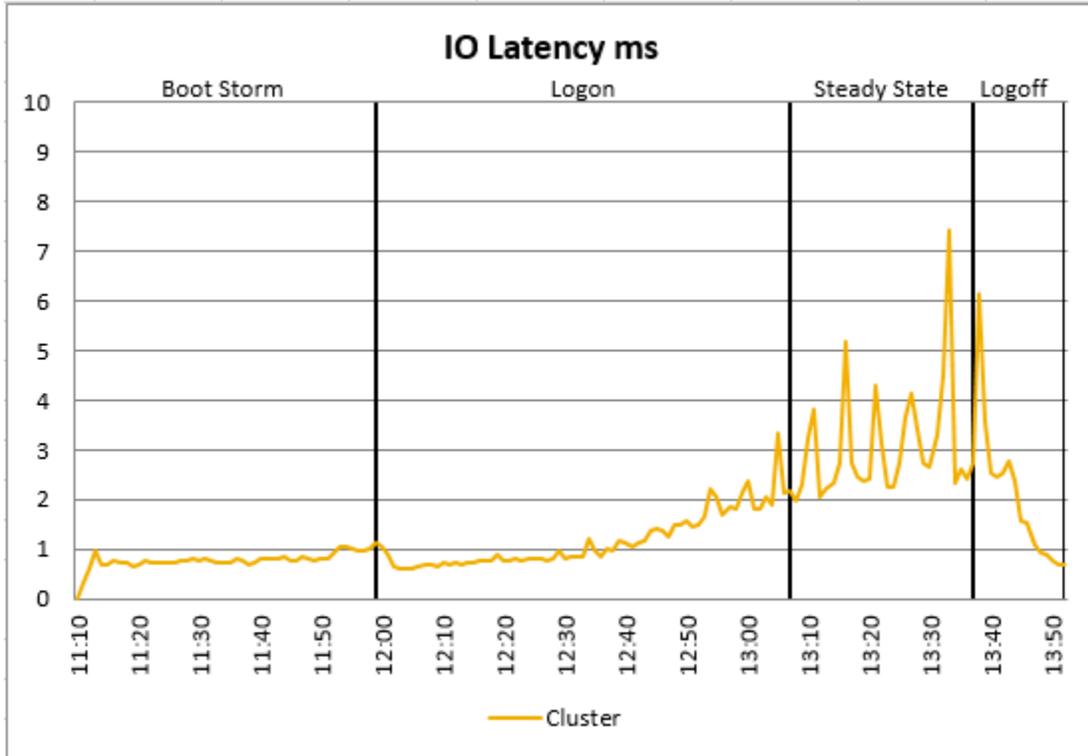
Network usage peaked at 4379 Mbps during Steady State on host Compute B, and the average network usage for all hosts during Steady State was 2028 Mbps. The Compute Cluster Shared Volume was owned by host Compute B during this test run, therefore it received higher network traffic than the other nodes due to Cluster Redirected IO.



The peak Cluster IOPS for the test run was 5747 IOPS during the Boot Storm, while the average in Steady State was 2571 IOPS. This data was captured from the sum of the “PhysicalDisk(x)\Disk Transfers/sec” metrics for the mounted virtual disks only. Based on this chart each user session generated 6.35 IOPS during Steady State.

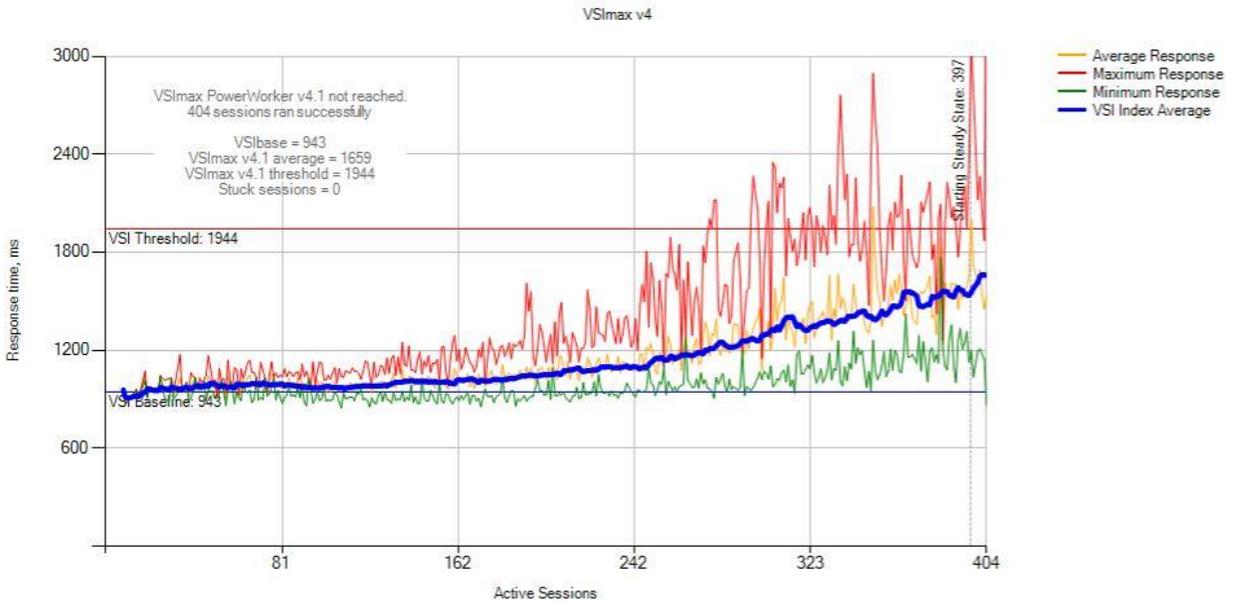


The IO Latency was captured using the “PhysicalDisk(x)\Avg. sec/Transfer” metrics for the mounted virtual disks. The peak IO Latency was 7.4 ms during Steady State. The average IO latency during steady state was 3.1 ms.

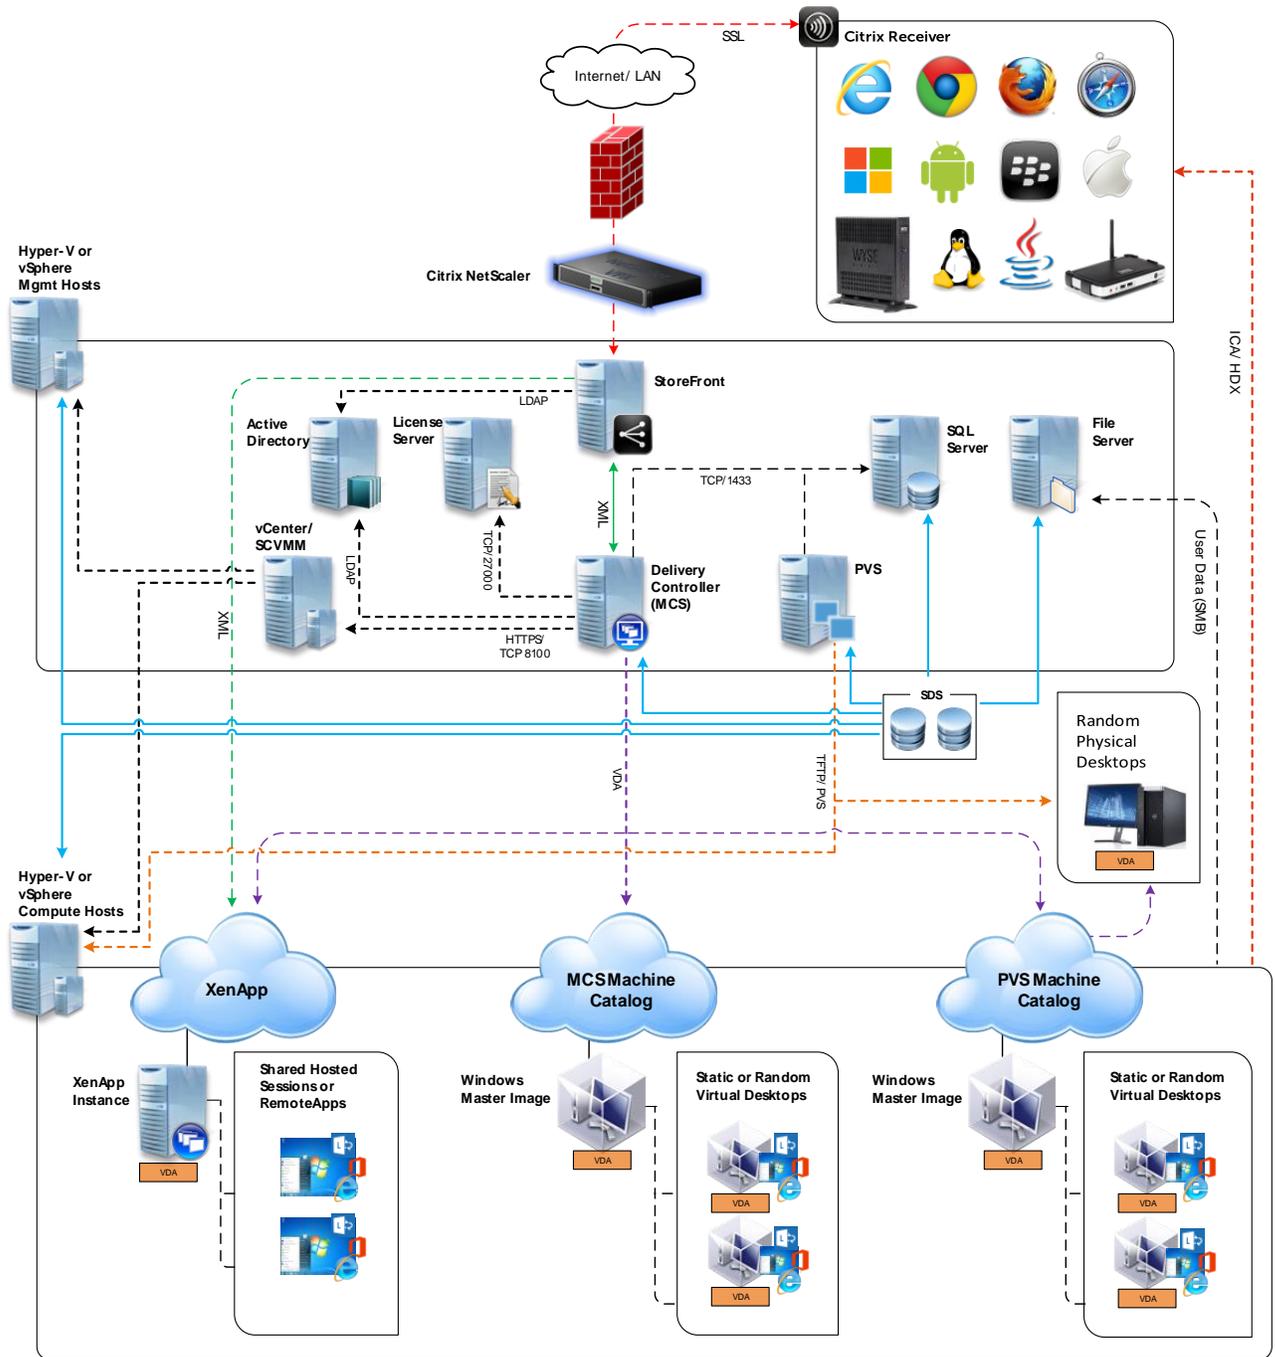


The baseline performance of 943 indicated that the user experience for this test run was Good. The Index average reached 1659, well below the threshold of 1944.

Login VSI Baseline	VSI Index Average	VSI Max Reached	VSI Threshold
943	1659	No	1944



6.4 Citrix XenDesktop communication flow



7 Related resources

See the following referenced or recommended resources:

- The [Dell EMC Cloud-Client Computing Solutions for Citrix Tech Center page](#) which includes this RA and other Citrix XenDesktop based RAs.
- [Storage Spaces Direct in Windows Server 2016](#)
- [Software Storage Bus Overview](#)
- [Planning volumes in Storage Spaces Direct](#)
- [Fault domain awareness in Windows Server 2016](#)
- [Deep Dive: The Storage Pool in Storage Spaces Direct](#)
- [Fault tolerance and storage efficiency in Storage Spaces Direct](#)
- [Understanding the cache in Storage Spaces Direct](#)
- [Resilient File System \(ReFS\) overview](#)

Acknowledgements

Thanks to Geoff Dillon for his engineering and testing work of this solution. Geoff Dillon is a Lead Solutions Engineer in the Datacenter EUC Enterprise Engineering group at Dell. Geoff is experienced in enterprise system deployment and administration and is working on making Dell's VDI solutions stand out in the market.

Thanks to David Hulama of the Dell Wyse Technical Marketing team for his support and assistance with datacenter EUC programs at Dell. David is a Senior Technical Marketing Advisor for VMware VDI solutions at Dell. David has a broad technical background in a variety of technical areas and expertise in enterprise-class virtualization solutions.

Thanks to Nicolas Cuendet of Dell Wyse Product Marketing for his support and assistance with Wyse Datacenter Appliances at Dell and for his contribution to the Dell Wyse Endpoints section of this document.

About the authors

Peter Fine is the Chief Architect and Field CTO of Datacenter EUC Enterprise Engineering at Dell. Peter leads the strategy, architecture and engineering of the datacenter EUC product and solutions portfolio. Follow Peter @ExitTheFastLane or www.ExitTheFastLane.com.

Jerry Van Blaricom is a Lead Architect in the Cloud Client Solutions Engineering Group at Dell. Jerry has extensive experience with the design and implementation of a broad range of enterprise systems and is focused on making Dell's virtualization offerings consistently best in class.