

Dell EMC Storage Spaces Direct (S2D) Ready Nodes for Microsoft Remote Desktop Services (RDS) – Reference Architecture

Integration of Microsoft RDS with Dell 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 [10/18/2017] [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	11
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	Microsoft	22
5	Solution architecture for RDS	29
5.1	Management role configuration	29
5.2	Storage configuration overview	31
5.3	Virtual networking	33
5.4	Scaling guidance	34
5.5	Solution high availability	36
5.6	Microsoft RDS communication flow	37
6	Solution Performance and Testing	38
6.1	Test and performance analysis methodology	38
6.2	Test configuration details	41
6.3	Test results and analysis	43
7	Related resources	55
	Acknowledgements	56
	About the authors	57

Executive summary

This document provides the reference architecture for integrating Dell EMC Microsoft Storage Spaces Direct Ready Nodes (S2D Ready Nodes) and Microsoft Remote Desktop Services (RDS) software 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.

Microsoft RDS 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 Microsoft Remote Desktop Services (RDS) 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

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 RDSH 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 S-Series 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 S-Series 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 Microsoft RDS 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 5040 AIO Thin Client (ThinOS)



The Dell Wyse 5040 AIO all-in-one (AIO) thin client runs ThinOS (with or without PCoIP), has a 21.5" Full HD display and offers versatile connectivity options for use in a wide range of industries. With four USB 2.0 ports, Gigabit Ethernet and integrated dual band Wi-Fi options, users can link to their peripherals and quickly connect to the network while working with processing-intensive, graphics-rich applications. Built-in speakers, a camera and a microphone make video conferencing and desktop communication simple and easy. It even supports a second attached display for those who need a dual monitor

configuration. A simple one-cord design and out-of-box automatic setup makes deployment effortless while remote management from a simple file server, Wyse Device Manager (WDM), or Wyse Management Suite can help lower your total cost of ownership as you grow from just a few thin clients to tens of thousands. 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 comes with flexible OS options: ThinOS (with or without PCoIP), 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). 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](#)

3.3.4 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 Microsoft

4.1.1 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.

4.1.2 Remote Desktop Services

Microsoft Remote Desktop Services (RDS) accelerates and extends desktop and application deployments to any device, improves remote worker efficiency, and helps secure critical intellectual property while simplifying regulatory compliance. Remote Desktop Services enables virtual desktop infrastructure (VDI), session-based desktops, and applications, allowing users to work anywhere.

The core RDS components include:



Remote Desktop Connection Broker (RDCB)

Remote Desktop Connection Broker (RD Connection Broker) allows users to connect to virtual desktops (RDVH) or RemoteApp programs and session-based desktops (RDSH) while evenly distributing the load to their respective collections.

Remote Desktop Gateway (RDG)

Remote Desktop Gateway (RD Gateway) enables authorized users to connect to virtual desktops, RemoteApp programs, and session-based desktops on an internal corporate network from any Internet-connected device.

Remote Desktop Web Access (RDWA)

Remote Desktop Web Access (RD Web Access) enables users to access RemoteApp and Desktop Connections through the Start menu on a computer that is running Windows 7 through Windows 10, or through a web browser. RemoteApp and Desktop Connection provides a customized view of RemoteApp programs and session-based desktops in a session collection, and RemoteApp programs and virtual desktops in a virtual desktop collection.

Remote Desktop Virtualization Host (RDVH)

Remote Desktop Virtualization Host (RD Virtualization Host) integrates with Hyper-V to deploy pooled or personal virtual desktop collections within your organization.

Remote Desktop Session Host (RDSH)

Remote Desktop Session Host (RDSH) enables a server to host RemoteApp programs or session-based (personal or shared) desktops. Users can connect to RD Session Host servers in a session collection to run programs, save files, and use resources on those servers.

Database

Although RDS can be configured using an internal database, a Microsoft SQL Server database is recommended to store configuration and session information. To implement broker high availability, SQL Server is required. Alternatively, an Azure SQL database can be used. The RDCB must have a persistent connection to the database as it stores data collected and managed by the RDS services.

Remote Desktop Licensing

Remote Desktop Licensing (RD Licensing) manages the licenses required to connect to a Remote Desktop Session Host server or a virtual desktop. You can use RD Licensing to install, issue, and track the availability of licenses.

RDS deployment terms and options include:

Collections

As the name suggests, a collection is a group or pool of virtual machines that the broker will connect users to. A collection can be **managed** meaning the RDCB creates and maintains the collection using a template VM including recreating VMs as needed or a collection can be **unmanaged** meaning the RDCB can broker connections to the pool but there is no template VM. The RDCB does not create or manage the VMs in an unmanaged collection.

Pooled Virtual Desktops

Pooled virtual desktops are based off of a template VM and are reverted back to a pristine state when users log off making them non-persistent. They are thin provisioned, checkpointed, and only consume a fraction of the storage used by the original template VM. User Profile Disks (UPD) can be used to store user profile data so that it's available the next time a user logs on to a pooled desktop.

Personal Virtual Desktops

Personal virtual desktops are based off of a template VM but changes made by the user are stored to the VM virtual disk and available after their session ends. Users are reconnected to the same VM each time creating a persistent experience. Refer to the [System Center 2016 Virtual Machine Manager](#) section for integration details with personal virtual desktops.

Session-based virtualization

Users connect to RDSH servers and run their applications or desktops in Windows Server 2016 sessions (multiple users connecting to the same server). Session-based desktops can be configured as **shared** where no session info is retained after logging off or as **personal** where users are reconnected to the same RDSH server allowing them to make changes to and save their session info.

For additional information about RDS, please visit: [LINK](#)

4.1.3 RDSH

Compute hosts with the RDSH role provide easy access to a densely shared session environment. Each RDP-based session shares the total available server resources with all other sessions logged in concurrently on the server. An RDS CAL and Server CAL is required for each user/device accessing this type of environment but neither SA nor VDA licenses are required since the underlying OS for RDSH VMs are Server based instead of Client based.

Benefits of hosted desktop sessions and applications:

- Management of applications (single instance)
- Management of simple desktop images (no applications installed)
- Scalability of compute hosts: CPU and IOPS reduction via application offload
- Shared storage scalability: less IOPS = more room to grow

RDSH 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. Starting with Windows Server 2016, RDSH now provides the ability to use personal session desktops where users always connect to the same RDSH server giving them a persistent desktop experience.

4.1.3.1 RDSH Integration

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


4.1.3.2 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 CPU 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 Wyse Datacenter solution is as follows:


4.1.3.3 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 RDSH VMs as shown below will ensure that no physical NUMA node spanning occurs which could lower performance for an effected VM.



4.1.3.4 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 RDSH VMs as shown below 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.

4.1.4 System Center 2016 Virtual Machine Manager

An optional 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.

It's important to note that SCVMM doesn't integrate as seamlessly with RDS as one might expect. Service templates can be used with SCVMM to build out, maintain, and scale the RDS infrastructure components such as the connection broker (RDCB) and SQL. When it comes to compute VM collections, SCVMM can be used as a provisioning & management tool for unmanaged collections but doesn't integrate directly with the RDS farm or RDCB. Therefore, SCVMM can't be used to provision any VM intended to exist within a managed pool owned by the RDCB. As stated in the [Remote Desktop Services](#) section, a managed collection is a pool created and maintained by the RDCB by using a template VM in RDS which allows the RDCB to recreate and provision more VMs based on the template as needed.

SCVMM can be used with unmanaged collections in RDS but it will require some manual steps which can be handled via PowerShell. RDCB can broker connections to VMs in an unmanaged collection but it can't be used to create/maintain them. However, SCVMM can manage the VMs in the collection and make use of features like Intelligent Placement as well as handling virtual MAC assignments which can become an issue on a server hosting a large number of VMs. In order to do so, create the unmanaged collection in RDS first. Next, create/clone the desired number of VMs using SCVMM. Cloning multiple copies of VMs is possible in SCVMM via PowerShell. Once the VMs are created, they must be manually added to the unmanaged collection via PowerShell or Server Manager. Alternatively, SCVMM service templates can be used to also provision the compute VMs. There are many online articles describing how to configure these templates for an RDSH deployment.

4.1.5 Licensing

Several licensing aspects must be taken in to account when deploying a Microsoft RDS environment including:

- Windows Server 2016 Hyper-V
- Windows Server CALs
- RDS CALs
- Software Assurance (SA) / Virtual Desktop Access (VDA) client licensing
- SCVMM (optional)

Windows Server 2016 Hyper-V

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](#)

Windows Server and RDS CALs

Windows 2016 licensing follows a core + CAL licensing model. Any user/device accessing a Datacenter or Standard edition of Windows Server requires a Windows Server CAL. The Server CALs are considered the “base” CAL while RDS CALs are additive. Therefore, for a Microsoft RDS deployment, all users/devices accessing RDS require a Server CAL + a RDS CAL. Some licensing offerings, such as the Microsoft VDI Suite, include the RDS CALs.

SA/VDA Client Licenses

Remotely accessing a VM running a Windows desktop OS requires that the user/device have a Software Assurance (SA) or Virtual Desktop Access (VDA) license to access the virtualized Windows desktop OS. Endpoints with an OEM licensed Windows desktop operating system are generally not licensed to access a VDI environment and require an appropriate SA/VDA license to do so. Certain endpoints, such as those running a non-Windows OS, will require a VDA license. Some endpoints, such as Windows embedded Thin Clients, may already include a VDA license with their purchase. Neither license type is required when accessing session-based desktops as the underlying OS for RDSH session-based desktops is Windows Server 2016 and not an actual desktop OS.

SCVMM (optional)

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](#).

NOTE: For exact licensing details including pricing, please contact a Microsoft Software Licensing specialist at Dell.

5 Solution architecture for RDS

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
RD Broker & Licensing	4	8	4GB 10GB	20%	Med	1	60	Mgmt
RD Gateway & Web Access	4	4	2GB 10GB	20%	Med	1	60	Mgmt
Primary SQL	8	8	4GB 16GB	20%	Med	1	60 + 300 (VHDX)	Mgmt
SCVMM (optional)	8	8	4GB 16GB	20%	Med	1	60 + 50 (VHDX)	Mgmt
Total	24 vCPUs	28GB	14GB 52GB	-	-	5 vNICs	590GB	-

***NOTE:** Set memory to static 10GB if using SQL Standard version. SCVMM is optional but a recommended configuration is provided here in case it will be used.

5.1.1 RDSH VM configuration

The recommended number of RDSH 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 Microsoft databases 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 databases for:

- Microsoft RDS (SQL is not required by highly recommended for high availability)
- SCVMM (if in use)

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 Microsoft 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 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, RDS 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, three vNICs for the management OS must be created, one for management functions and two for storage. The vNICs can then be tagged with the same or different VLAN IDs to further segment traffic. Separate VLANs are recommend for desktop and RDS role VMs. The diagram below shows the recommended vNIC configuration.

- Management: Configured for hypervisor infrastructure traffic – L3 routed via spine layer
- VDI: Configured for VDI session traffic – L3 routed via spine layer
- RDMA1, RDMA2: Configured for converged IO/ SMB traffic. CSV, Live Migration and heartbeating operations will also share these interfaces – 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 cluster) 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
Broker Servers	Desktops per instance (dependent on SQL performance as well)	Additional servers added to the configuration	Additional virtual machine resources (RAM and CPU)
RDSH Servers	Apps/Desktops per instance	Additional virtual servers added to the collection	Additional physical servers to host virtual RDSH servers
RD Gateway/Web Access Servers	Logons/ minute	Additional servers added to the configuration	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

Additional scalability details specific to Windows Server 2016 Hyper-V can be found here: [LINK](#)

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 RDS infrastructure server roles are duplicated and spread amongst management host instances where connections to each are load balanced.
- 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](#)

5.6 Microsoft RDS communication flow



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	RDVH – Pooled Desktops	Task User	Windows 10 x64 & Office 2016	220
S2DRN B5	Hyper-V 2016	RDVH – Pooled Desktops	Knowledge Worker	Windows 10 x64 & Office 2016	175
S2DRN B5	Hyper-V 2016	RDVH – Pooled Desktops	Power Worker	Windows 10 x64 & Office 2016	155

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 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	Microsoft Remote Desktop Services 2016
Management VM OS	Microsoft Windows Server 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
RDSH	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 Ctlr	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	RDS 2016	Task Worker	220	54	343	3.06	4.41
S2DRN	Hyper-V 2016	RDS 2016	Knowledge Worker	175	63	331	3.90	7.26
S2DRN	Hyper-V 2016	RDS 2016	Power Worker	155	63	339	3.91	9.90

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	RDS 2016	Task Worker	660	50:43	-	30
S2DRN-B5	RDS 2016	Power Worker	450	46:00	-	30

6.3.1 S2DRN B5 Configuration


Refer to the [Platform Configurations](#) section for hardware configuration details.

6.3.1.1 Task Works, 220 Users, Hyper-V 2016, RDVH Desktops


The Task Worker test run provisioned 220 VMs on each host, 660 in the cluster, along with the RDS management VM roles. The Peak CPU usage was 77% on one node at the end of the Logon phase, while the steady state average CPU usage was 54%. The offsets between the CPU curves show that the RDS Broker prefers to direct sessions to one node at a time until it reaches capacity.




The memory used during the test run averaged 343 GB, and the peak memory usage was 366 GB on one node 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 970Mbps. The peak was 2080Mbps during Logon phase on node Compute C. In this test run Compute C owned the Cluster Virtual Disk hosting all of the desktop VMs, therefore it saw more network traffic than the other nodes.



The peak Cluster IOPS for the test run was 4711 IOPS during Logon phase, while the average in Steady State was 2020 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 3.06 IOPS.




The IO Latency was captured using the “PhysicalDisk(x)\Avg. sec/Transfer” metrics for the mounted virtual disks. The peak IO Latency was 2.8 ms during the Boot Storm. The average IO latency during steady state was 1.6 ms. The chart clearly shows a very steady and very low level of IO Latency throughout the test run.




The baseline performance of 961 indicated that the user experience for this test run was good. The Index average reached 1401, well below the threshold of 1961. However several sessions were stuck during the test and 7 more failed to login.

Login VSI Baseline	VSI Index Average	VSI Max Reached	VSI Threshold
961	1220	NO	1961




6.3.1.2 Knowledge Worker, 175 Users, Hyper-V 2016, RDVH Desktops


In this workload test run, the hosts each had 175 user sessions, 525 in the cluster, in addition to the various management VMs. The peak CPU Usage was 78% on one host during logon phase, while the Steady State average was 63% across all hosts. The CPU usage curves show a preference by the Connection Broker for certain hosts until host saturation has been reached.




The memory consumption averaged 335GB in steady state, and the peak usage on any host was 346GB during the boot storm phase. The peak usage was 90.2%, while the Steady State average usage was 87.2%.




Network usage peaked at 2113Mbps during Steady State on one host, and the average network usage for all hosts was 1271Mbps 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 3792 IOPS at the end of Logon phase, while the average in Steady State was 2046 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 7.22 IOPS.




The IO Latency was captured using the “PhysicalDisk(x)\Avg. sec/Transfer” metrics for the mounted virtual disks. The peak IO Latency was 2.2 ms at the end of the Boot Storm. The average IO latency during steady state was 1.90 ms. The chart clearly shows a very steady and very low level of IO Latency throughout the test run.



The baseline performance of 965 indicated that the user experience for this test run was good. The Index Average reached 1488, well below the threshold of 1965.


Login VSI Baseline	VSI Index Average	VSI Max Reached	VSI Threshold
965	1488	NO	1965

VSI max v4




6.3.1.3 Power Worker, 155 Users, Hyper-V 2016, RDVH Desktops


In this workload test run, the hosts each had 155 user sessions, 465 in the cluster, in addition to the various management VMs. The peak CPU Usage was 88% on one host during logon phase, while the Steady State average was 63% across all hosts. Again, the CPU usage curves show a preference by the Connection Broker for certain hosts until host saturation has been reached.




The memory consumption averaged 339 GB in steady state, and the peak usage on any host was 342 GB during the steady state phase. The peak usage was 89%, just below the 90% threshold, while the Steady State average usage was 88%.




Network usage peaked at 2748Mbps during logon phase on host Compute B, and the average network usage for all hosts during Steady State was 1534Mbps. 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 3152 IOPS during Logon, while the average in Steady State was 1816 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 3.91 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 2.2 ms at the end of the Boot Storm. The average IO latency during steady state was 1.90 ms. The chart clearly shows a very steady and very low level of IO Latency throughout the test run.



The baseline performance of 965 indicated that the user experience for this test run was good. The Index Average reached 1564, well below the threshold of 1965.

Login VSI Baseline	VSI Index Average	VSI Max Reached	VSI Threshold
965	1564	NO	1965



7 Related resources

See the following referenced or recommended resources:

- [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 engineering and testing work for this solution. Geoff Dillon is a Lead Solutions Engineer in the Cloud Client Solutions 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 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.