

# REFERENCE ARCHITECTURE

## Dell EMC Ready Bundle for VDI

Design, configuration and implementation of a Microsoft RDS environment with Dell EMC PowerEdge Servers and XtremIO Storage

### [Abstract](#)

A Reference Architecture for integrating Dell EMC PowerEdge servers and Microsoft Remote Desktop Services (RDS) and Hyper-V 2016 to create virtual application and virtual desktop environments on 14th generation Dell EMC PowerEdge Servers.

February 2018

## Revisions

Date	Description
May 2017	Initial release
September 2017	Replaced server hardware with latest generation (14G) Dell EMC PowerEdge servers, Skylake CPUs and storage hardware with Dell EMC XtremIO 2
February 2018	Solution name change to fit Ready Solutions for VDI nomenclature

## Acknowledgements

This paper was produced by the following members of the Dell EMC storage engineering team:

Author: Keith Keogh – Lead Architect

Peter Fine – Chief Architect

Support: Cormac Woods – Systems Engineer

Other: Rick Biedler – Engineering Director

David Hulama – Sr. Technical Marketing Advisor

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.

© 2017 – 2018 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

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

Revisions.....	2
Acknowledgements.....	2
Executive summary.....	6
<b>1 Introduction.....</b>	<b>7</b>
1.1 Objective.....	7
1.2 What's new.....	7
<b>2 Solution architecture overview.....</b>	<b>8</b>
2.1 Introduction.....	8
2.2 Physical architecture overview.....	8
2.3 Solution layers.....	9
2.3.1 Networking.....	10
2.3.2 Compute.....	10
2.3.3 Management.....	10
2.3.4 Storage.....	10
2.4 Local Tier 1.....	11
2.5 Shared Tier 1 for rack servers.....	11
2.5.1 Network architecture – ST1 racks.....	13
2.5.2 Rack cabling (HA) – ST1 racks.....	14
2.5.3 Storage scaling guidance – ST1 racks.....	14
2.6 Shared Tier 1 for blade servers.....	14
2.6.1 Network architecture – ST1 blades.....	16
2.6.2 Rack cabling (HA) – ST1 blades.....	17
2.6.3 Storage scaling guidance – ST1 blades.....	17
<b>3 Hardware components.....</b>	<b>19</b>
3.1 Network.....	19
3.1.1 Dell Networking S-Series S3048 (1Gb ToR Switch).....	19
3.1.2 Dell Networking S-Series S4048 (10Gb ToR Switch).....	19
3.1.3 Brocade 6510 (FC ToR Switch).....	21
3.1.4 Brocade M5424 (FC Blade Interconnect).....	22
3.1.5 PowerEdge M I/O Aggregator (10Gb Blade Interconnect).....	23
3.2 Servers.....	24
3.2.1 PowerEdge R640.....	25
3.2.2 PowerEdge M630.....	27
3.3 Storage.....	28
3.3.1 XtremIO X2 X-Brick – Combined Tier 1 and Tier 2.....	28

3.4	Dell Wyse Endpoints .....	31
3.4.1	Wyse 3040 Thin Client (ThinOS, ThinLinux) .....	32
3.4.2	Wyse 5040 AIO Thin Client (ThinOS).....	32
3.4.3	Wyse 5060 Thin Client (ThinOS, ThinLinux, WES7P, WIE10) .....	32
3.4.4	Wyse 7020 Thin Client (WES 7/7P, WIE10, ThinLinux).....	33
3.4.5	Wyse 7040 Thin Client (WES7P, WIE10) .....	33
3.4.6	Latitude 3480 and 5280 Mobile Thin Clients (Win 10 IoT) .....	33
3.4.7	Enhanced Security .....	34
<b>4</b>	<b>Software components.....</b>	<b>35</b>
4.1	Microsoft .....	35
4.1.1	Windows Server 2016 Hyper-V hypervisor .....	35
4.1.2	Remote Desktop Services .....	35
4.1.3	RDSH.....	37
4.1.4	System Center 2016 Virtual Machine Manager.....	39
4.1.5	Licensing.....	40
<b>5</b>	<b>Solution architecture for RDS.....</b>	<b>42</b>
5.1	Management role configuration .....	42
5.1.1	RDSH VM configuration .....	42
5.1.2	SQL databases .....	43
5.1.3	DNS .....	43
5.2	Storage configuration overview .....	43
5.2.1	Local Tier 1 storage .....	44
5.2.2	Shared Tier 1 storage .....	44
5.2.3	Shared Tier 2 storage .....	44
5.2.4	Storage networking – XtremIO Fiber Channel (FC) .....	45
5.3	Virtual networking .....	45
5.3.1	Local Tier 1 .....	45
5.3.2	Shared Tier 1 – FC .....	47
5.4	Scaling guidance .....	48
5.5	Solution high availability .....	50
5.5.1	SQL Server high availability .....	51
5.6	Microsoft RDS communication flow.....	52
<b>6</b>	<b>Solution Performance and Testing.....</b>	<b>53</b>
6.1	Test and performance analysis methodology.....	53
6.1.1	Testing process .....	53
6.1.2	Resource monitoring .....	55

6.1.3 Resource utilization .....	55
6.2 Test configuration details.....	56
6.2.1 Compute VM Configurations .....	56
6.2.2 Platform Configurations .....	57
6.3 Test results and analysis .....	57
6.3.1 R740 Compute.....	58
7 Related resources .....	62

## Executive summary

This document provides the reference architecture for integrating Dell EMC PowerEdge servers, Dell EMC storage, and Microsoft Remote Desktop Services (RDS) software to create virtual shared session and virtual desktop environments. The available server choices include the PowerEdge R640, R740, and M630 servers. Shared storage for the solution is the Dell EMC XtremIO X2 X-Brick.

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.

Dell EMC XtremIO X2 is the next-generation XtremIO all-flash storage array platform that offers consistently high performance with low latency; unmatched storage efficiency with inline, all-the-time data services such as thin provisioning, deduplication, and compression; rich application integrated copy services; and unprecedented management simplicity. The content-aware, in-memory metadata and inline, all-the-time data services have made XtremIO the ultimate shared storage platform for virtual server and desktop environments and workloads that benefit from efficient copy data management.

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. Proposed design choices include rack or blade servers, local disks or shared storage. Guidance contained within this document follows a building block methodology enabling the combination of several different components each with their own scaling capabilities.

## 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.

## 1.2 What's new

- Introducing the latest 14th generation Dell EMC PowerEdge servers with Skylake processors.
- Introducing the Dell EMC XtremIO X2 array.

## 2 Solution architecture overview

### 2.1 Introduction


Dell EMC Ready Bundle for VDI solutions provide a number of deployment options to meet your desktop virtualization requirements. Our solution is able to provide a compelling desktop experience to a range of employees within your organization from task workers to knowledge workers to power users. The deployment options for Dell EMC Ready Bundle for VDI for Microsoft RDS include:

- Desktop-based virtualization using Remote Desktop Virtualization Hosts (RDVH)
  - Pooled Virtual Desktops (Non-persistent)
  - Personal Virtual Desktops (Persistent)
- Session-based virtualization using Remote Desktop Sessions Hosts (RDSH)
  - RemoteApp programs
  - Personal or Shared session desktops


Please refer to the [Remote Desktop Services](#) section for an explanation of acronyms and components.

### 2.2 Physical architecture overview

Dell EMC Ready Bundle for VDI is a 100% virtualized solution architecture. The core Dell EMC Ready Bundle for VDI architecture consists of two models: Local Tier 1 (LT1) for proof of concept (POC) and Shared Tier 1 (ST1) for production use. “Tier 1” in the Dell EMC Ready Bundle for VDI context defines from which disk source the VDI sessions (Compute VMs) execute. LT1 includes rack servers with SSDs while ST1 can include rack or blade servers. Storage utilized for user data and Management VM execution is designated as Tier 2 (T2). Our design places shared T1 and T2 on the same storage array but customers can elect to use separate storage for T2 if so desired. For LT1 configurations (POC), Compute and Management VMs use the same local storage.



In the Shared Tier 1 solution model, all compute and management layer hosts are diskless utilizing the new Boot Optimized Storage Solution (BOSS) device for the hypervisor/operating system.



**NOTE:** At the time of this writing, the 14<sup>th</sup> generation blade servers are not yet available. The boot device options for the existing M630 blade servers include SD cards or local disks.

## 2.3 Solution layers


The Dell EMC Ready Bundle for VDI Solution leverages a core set of hardware and software components consisting of five primary layers:

- Networking Layer
- Compute Server Layer
- Management Server Layer
- Storage Layer
- Thin Client Layer (please refer to the [Dell Wyse Endpoints](#) section)

These components have been integrated and tested to provide the optimal balance of high performance and lowest cost per user. The Dell EMC Ready Bundle for VDI stack is designed to be cost effective allowing IT departments to implement high-performance fully virtualized desktop environments.


### 2.3.1 Networking

Only a single high performance Dell Networking S-Series 48-port switch is required to get started in the network layer for a combined pilot/POC configuration. For all other configurations, you can start with a single Dell Networking S-Series 48-port switch for 10Gb LAN traffic along with a single Brocade fiber channel switch for SAN connectivity. Additional switches are added and stacked as required to provide High Availability for the Network layer.




### 2.3.2 Compute

The compute layer consists of the server resources responsible for hosting the desktop or RDSH VMs on Hyper-V hypervisor with local or shared Tier 1 solution models (shared Tier 1 rack servers pictured below).



### 2.3.3 Management


VDI management components are dedicated to their own layer so as to not negatively impact the user sessions running in the compute layer. This physical separation of resources provides clean, linear, and predictable scaling without the need to reconfigure or move resources within the solution as you grow. The management layer will host all the server VMs necessary to support the VDI infrastructure, shared Tier 1 rack depicted below: RDCB (Remote Desktop Connection Broker), RDWA (Remote Desktop Web Access), and RDG (Remote Desktop Gateway). Refer to the [Remote Desktop Services](#) section for information on these components.



### 2.3.4 Storage


The storage layer consists of the Dell EMC XtremIO X2 X-Brick for combined shared T1 and T2. The configuration shown below is the minimum disk configuration for the X2 array which can support up to 3,500 knowledge worker users. Additional disks and/or larger disk sizes can be used if necessary to provide more

capacity for persistent desktop users or if user data is also stored on the array (via file servers). Additional X-Bricks are added to the solution when scaling beyond 3,500 users.



## 2.4 Local Tier 1


For pilot/POC or small deployments, a single server can be used. This architecture is non-distributed with all VDI, Management, and storage functions on a single host. If additional scaling is desired, you can grow into a larger distributed ST1 architecture seamlessly. Disk size depends on total capacity requirements of all VMs but a minimum of 4 x 960GB SSDs is recommended. SQL server is optional for a pilot/POC deployment, but highly recommended if intending to scale the setup for production use and ensuring the broker is highly available. SQL is required if installing the optional System Center Virtual Machine Manager (SCVMM).



**NOTE:** This configuration can support up to 150 users based on the Task Worker workload.

## 2.5 Shared Tier 1 for rack servers

This solution model provides a high-performance scalable rack-based configuration that incorporates shared T1 and T2 storage for execution of VDI sessions and management VMs. Since all VMs reside on the shared storage array, the servers are diskless and use a BOSS device for the operating system. User data can either be stored on the same array (via a file server) as the VMs or on another storage location. The figure below depicts the shared storage configuration with optional file server VM for user data.




**NOTE:** Minimum disk configuration per XtremIO X2 X-Brick is 18 x 400GB SSDs which is sufficient for up to 3500 VDI users. Additional disks may be required if increased capacity is needed for larger personal disk sizes and if user data is also stored on the array.

**NOTE:** Maximum of 64 nodes supported per cluster, but each cluster limited to a maximum of 8000 running virtual machines.


## 2.5.1 Network architecture – ST1 racks

The architecture for the Shared Tier 1 solution uses a dedicated Fiber Channel (FC) switching infrastructure for the management and compute servers to connect to shared storage. Both management and compute servers connect to all network VLANs in this model. All ToR traffic has been designed to be layer 2 (switched locally), with all layer 3 (routable VLANs) routed through a core or distribution switch. The following diagram illustrates the server NIC to ToR switch connections.



## 2.5.2 Rack cabling (HA) – ST1 racks

The following diagram depicts the cabling for the components in the ST1 rack servers solution.



## 2.5.3 Storage scaling guidance – ST1 racks

**NOTE:** Scaling numbers are based on our density testing for the Knowledge Worker workload. Customer needs may vary.

Shared Tier 1 HW Scaling (Rack - FC)						
User Scale	XtremIO X2	Compute			ToR LAN	ToR 8Gb FC
		Servers	CPU Cores	Memory		
Up to 3500	1 x X-Brick	1 – 13	Up to 520	Up to 10TB	1 x S4048	1 x 6510
3501 – 7000	2 x X-Brick	13 – 26	Up to 1,040	Up to 20TB	1-2 x S4048	1-2 x 6510
7001 – 10000*	3 x X-Brick	26 – 37	Up to 1,480	Up to 29TB	2 x S4048	2 x 6510

**NOTE:** Hyper-V cluster limited to a maximum of 8000 running virtual machines.


**NOTE:** Density is based on empirical data but has not been validated with Windows Server 2016 with Hyper-V.

## 2.6 Shared Tier 1 for blade servers

As is the case in the ST1 for rack servers model, blade servers can be used to provide a high-performance scalable configuration that incorporates shared T1 and T2 storage for execution of VDI sessions and

management VMs. Since all VMs reside on the shared storage array, the blades use mirrored drives for the operating system. User data can either be stored on the same array (via a file server) as the VMs or on another storage location. The figure below depicts the shared storage configuration with optional file server VM for user data.

**NOTE:** At the time of this writing, the 14<sup>th</sup> generation blade servers are not yet available.




**NOTE:** Minimum disk configuration per XtremIO X2 X-Brick is 18 x 400GB SSDs which is sufficient for up to 3500 VDI users. Additional disks may be required if increased capacity is needed for larger personal disk sizes and if user data is also stored on the array.

**Note:** Maximum of 64 nodes supported per cluster, with each cluster limited to a maximum of 8000 running virtual machines.



## 2.6.2 Rack cabling (HA) – ST1 blades

The following diagram depicts the cabling for the components in the **ST1 blade servers** solution.



## 2.6.3 Storage scaling guidance – ST1 blades

**NOTE:** Scaling numbers are based on our density testing for the Knowledge Worker workload. Customer needs may vary.

Shared Tier 1 HW Scaling (Blade - FC)						
User Scale	XtremIO X2	Compute			Blade LAN + FC	ToR 8Gb FC
		Servers	CPU Cores	Memory		
Up to 3500	1 x X-Brick	1 – 13	Up to 520	Up to 7TB	2 x IOA + 2 x M5424	1 x 6510
3501 – 7000	2 x X-Brick	13 – 26	Up to 1,040	Up to 14TB	4 x IOA + 4 x M5424	1 x 6510
7001 – 10000*	3 x X-Brick	26 – 37	Up to 1,480	Up to 19TB	6 x IOA + 6 x M5424	2 x 6510

---

**NOTE:** Hyper-V cluster limited to a maximum of 8000 running virtual machines.

---



---

**NOTE:** Density is based on empirical data but has not been validated with Windows Server 2016 with Hyper-V.

---

## 3 Hardware components


### 3.1 Network

The following sections contain the core network components for the Dell EMC Ready Bundle for VDI 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](#)

### 3.1.3 Brocade 6510 (FC ToR Switch)

The Brocade® 6510 Switch meets the demands of hyper-scale, private cloud storage environments by delivering market-leading speeds up to 16Gb, Fiber Channel (FC) technology and capabilities that support highly virtualized environments. Designed to enable maximum flexibility and investment protection, the Brocade 6510 is configurable in 24, 36, or 48 ports and supports 2, 4, 8, or 16Gb speeds in an efficiently designed 1U package. It also provides a simplified deployment process and a point-and-click user interface making it both powerful and easy to use. The Brocade 6510 offers low-cost access to industry-leading Storage Area Network (SAN) technology while providing “pay-as-you-grow” scalability to meet the needs of an evolving storage environment.

Model	Features	Options	Uses
Brocade 6510	48 x 2/4/8/16Gb Fiber Channel  Additional (optional) FlexIO module  Up to 24 total ports (internal + external)	Ports on demand from 24, 36, and 48 ports	FC ToR switches for all solutions. Optional for blades



For more information on the Brocade 6510 switch, please visit: [LINK](#)

### 3.1.4 Brocade M5424 (FC Blade Interconnect)

The Brocade® M5424 switches and Dell™ PowerEdge™ M1000e Blade enclosures provide robust solutions for FC SAN deployments. Not only does this offering help simplify and reduce the amount of SAN hardware components required for a deployment, but it also maintains the scalability, performance, interoperability and management of traditional SAN environments. The M5424 can easily integrate FC technology into new or existing storage area network (SAN) environments using the PowerEdge™ M1000e Blade enclosure. The Brocade® M5424 is a flexible platform that delivers advanced functionality, performance, manageability, scalability with up to 16 internal Fabric ports and up to 8 2GB/4GB/8GB auto-sensing uplinks and is ideal for larger storage area networks. Integration of SAN switching capabilities with the M5424 also helps to reduce complexity and increase SAN manageability.

Model	Features	Options	Uses
Brocade M5424	16 x internal Fabric ports Up to 8 2/4/8Gb auto-sensing uplinks	Ports on demand from 12 to 24 ports	Blade switch for FC in Shared Tier 1 model




For more information on the Brocade M5424 switch, please visit: [LINK](#)

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

Simplify network management and increase server bandwidth with the PowerEdge™ M I/O Aggregator, enabling easy, plug-and-play data center convergence.

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




For more information on the PowerEdge IOA switch, please visit: [LINK](#)

## 3.2 Servers

This reference architecture is built on the latest 14<sup>th</sup> generation Dell EMC PowerEdge R640 servers for rack installations and the M630 servers for blades. Optimized for VDI, the Dell EMC PowerEdge portfolio has been designed and arranged in three 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.
- **C7** configuration is the premium configuration offering an abundance of high performance to maximize user density.




---

**NOTE:** The configurations below do not represent absolute platform maximums and can be adjusted as needed.

---

### 3.2.1 PowerEdge R640

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. Local drive options include 2.5” or 3.5” disks (3.5” drive chassis shown below). A new Boot Optimized Storage Subsystem (BOSS) card is also available which allows the separation of the operating system from the data drives using M.2 SATA SSDs that can be configured in a hardware RAID mirror (RAID1).



For more information on the R640, please visit: [Link](#)

### 3.2.1.1 Local Tier 1 Rack

For small deployments such as ROBO or POC/pilot setups, the Local Tier 1 model combines Compute and Management on the same server with VDI desktops (or RDSH sessions) and management role VMs executing from local storage. Windows Server 2016 OS is installed to the BOSS device so that the disks are used solely for data. To provide sufficient capacity and IOPS, use at least 4 x SSDs for all-flash. Optionally, 10 x 15K SAS drives can be substituted for the SSDs.

R640 Local T1		Compute + Management
CPU	2 x Intel Gold 6138 (20C, 2.0GHz)	
Memory	24 x 32GB 2667MT/s RDIMMs <b>Effective speed: 2667MT/s @ 768GB</b>	
Storage Ctrls	PERC H730P – RAID10	
Storage	2 x 120GB M.2 BOSS in RAID1 ( <b>Hypervisor</b> ) 4 x 960GB SSD or 10 x 600GB 15K SAS ( <b>VMs</b> )	
Network	4 x 10Gb SFP+/BT	
iDRAC	iDRAC9 Enterprise with OpenManage Essentials	
Power	2 x 1100W PSUs	


### 3.2.1.2 Shared Tier 1 Rack (FC)

In the Shared Tier 1 model, desktop or session-based VMs execute on shared storage so there is no need for local disks on each server to host VMs. Fibre Channel is leveraged as the block storage protocol for Compute and Management hosts with Tier 1 and Tier 2 storage. All configuration options are identical except for CPU and RAM which are reduced on the Management host.

R640 Shared T1		Compute	Management
CPU		2 x Intel Gold 6138 (20C, 2.0GHz)	2 x Intel Xeon Silver 4114 (10C, 2.2GHz)
Memory		24 x 32GB 2667MT/s RDIMMs <b>Effective speed: 2667MT/s @ 768GB</b>	12 x 16GB 2667MT/s RDIMMs <b>Effective speed: 2400MT/s @ 192GB</b>
Storage Ctrls		PERC H330 – no RAID	PERC H330 – no RAID
Storage		2 x 120GB M.2 BOSS in RAID1 <b>(Hypervisor)</b>	2 x 120GB M.2 BOSS in RAID1 <b>(Hypervisor)</b>
Network		4 x 10Gb SFP+/BT 2 x QLogic 2562 8Gb DP FC HBA	4 x 10Gb SFP+/BT 2 x QLogic 2562 8Gb DP FC HBA
iDRAC		iDRAC9 Enterprise with OpenManage Essentials	iDRAC9 Enterprise with OpenManage Essentials
Power		2 x 1100W PSUs	2 x 1100W PSUs

### 3.2.2 PowerEdge M630

The blade server platform recommendation for the Dell EMC Ready Bundle for VDI solution is the PowerEdge M630. This half-height blade server is a feature-rich, dual-CPU platform that offers a blend of density, performance, efficiency and scalability. The M630 offers remarkable computational density, scaling up to 22 cores, 2 socket Intel Xeon CPUs (Broadwell) and 24 DIMMs (768GB RAM) of DDR4 memory in an extremely compact half-height blade form factor.



For more information on the PowerEdge M630, please visit: [Link](#)

#### 3.2.2.1 Shared Tier 1 Blade (FC)

The Shared Tier 1 blade server varies slightly from the rack server equivalent since the latest 14<sup>th</sup> generation blade servers are not yet available at the time of this writing. However, the processor cores are the same as

the rack servers. Two network interconnect Fabrics are configured for the blades: the A-Fabric dedicated to 10Gb LAN traffic and the B-Fabric dedicated to 8Gb FC.

M630 Shared T1		
	Compute	Management
CPU	2 x E5-2698v4 (20C, 2.2GHz)	2 x E5-2660v4 (14C, 2.0GHz)
Memory	16 x 32GB 2400MT/s RDIMMs <b>Effective speed: 2400MT/s @ 512GB</b>	8 x 16GB 2400MT/s RDIMMs <b>Effective speed: 2400MT/s @ 128GB</b>
Storage Ctrls	PERC H330 – RAID1	PERC H330 – RAID1
Storage	2 x SSD or 15K SAS <b>(Hypervisor)</b>	2 x SSD or 15K SAS <b>(Hypervisor)</b>
Network	QLogic 57810S-k 10Gb DP KR NDC 1 x QLogic QME2572 8Gb FC mezz	QLogic 57810S-k 10Gb DP KR NDC 1 x QLogic QME2572 8Gb FC mezz
iDRAC	iDRAC8 Enterprise	iDRAC8 Enterprise

### 3.3 Storage

#### 3.3.1 XtremIO X2 X-Brick – Combined Tier 1 and Tier 2

Dell EMC's XtremIO is an enterprise-class scalable all-flash storage array that provides rich data services with high performance. It is designed from the ground up to unlock flash technology's full performance potential by uniquely leveraging the characteristics of SSDs and uses advanced inline data reduction methods to reduce the physical data that must be stored on the disks.

XtremIO's storage system uses industry-standard components and proprietary intelligent software to deliver unparalleled levels of performance, achieving consistent low latency for up to millions of IOPS. It comes with a simple, easy-to-use interface for storage administrators and fits a wide variety of use cases for customers in need of a fast and efficient storage system for their datacenters, requiring very little planning to set-up before provisioning.



XtremIO leverages flash to deliver value across multiple dimensions:


- Performance (consistent low-latency and up to millions of IOPS)
- Scalability (using a scale-out and scale-up architecture)
- Storage efficiency (using data reduction techniques such as deduplication, compression and thin-provisioning)
- Data Protection (with a proprietary flash-optimized algorithm named XDP)

XtremIO X2 is the new generation of the Dell EMC's All-Flash Array storage system. It adds enhancements and flexibility in several aspects to the already proficient and high-performant storage array's former generation. Features such as scale-up for a more flexible system, write boost for a more sensible and high-


performing storage array, NVRAM for improved data availability and a new web-based UI for managing the storage array and monitoring its alerts and performance stats, add the extra value and advancements required in the evolving world of computer infrastructure.

The XtremIO X2 Storage Array uses building blocks called X-Bricks. Each X-Brick has its own compute, bandwidth and storage resources, and can be clustered together with additional X-Bricks to grow in both performance and capacity (scale-out). Each X-Brick can also grow individually in terms of capacity, with an option to add to up to 72 SSDs in each brick.

XtremIO architecture is based on a metadata-centric content-aware system, which helps streamlining data operations efficiently without requiring any movement of data post-write for any maintenance reason (data protection, data reduction, etc. – all done inline). The system lays out the data uniformly across all SSDs in all X-Bricks in the system using unique fingerprints of the incoming data and controls access using metadata tables. This contributes to an extremely balanced system across all X-Bricks in terms of compute power, storage bandwidth and capacity. The diagram below shows an incoming data stream with duplicate blocks and fingerprints as well as illustrates the stream after duplicates are removed as it is being written to the array.



Using the same unique fingerprints, XtremIO is equipped with exceptional always-on in-line data deduplication abilities, which highly benefits virtualized environments. Together with its data compression and thin provisioning capabilities (both also in-line and always-on), it achieves incomparable data reduction rates. The figure below demonstrates capacity savings with in-line deduplication and compression prior to the data being written.



System operation is controlled by storage administrators via a stand-alone dedicated Linux-based server called the XtremIO Management Server (XMS). An intuitive user interface is used to manage and monitor the storage cluster and its performance. The XMS can be either a physical or a virtual server and can manage multiple XtremIO clusters.

With its intelligent architecture, XtremIO provides a storage system that is easy to set-up, needs zero tuning by the client and does not require complex capacity or data protection planning, as the system handles it on its own.

### 3.3.1.1 Architecture


An XtremIO X2 Storage System is comprised of a set of X-Bricks that form together a cluster. This is the basic building block of an XtremIO array. There are two types of X2 X-Bricks available: X2-S and X2-R. X2-S is for environments whose storage needs are more IO intensive than capacity intensive, as they use smaller SSDs and less RAM. An effective use of the X2-S is for environments that have high data reduction ratios (high compression ratio or a lot of duplicated data) which lower the capacity footprint of the data significantly. X2-R X-Bricks clusters are made for the capacity intensive environments, with bigger disks, more RAM and a bigger expansion potential in future releases. The two X-Brick types cannot be mixed together in a single system, so the decision as to which type is suitable for your environment must be made in advance. The X2-S is the recommended X-Brick for Dell EMC Ready Bundle for VDI solutions.

Each X-Brick is comprised of:

- Two 1U Storage Controllers (SCs) with:
  - Two dual socket Haswell CPUs
  - 346GB RAM (for X2-S) or 1TB RAM (for X2-R)
  - Two 1/10GbE iSCSI ports
  - Two user interface interchangeable ports (either 4/8/16Gb FC or 1/10GbE iSCSI)
  - Two 56Gb/s InfiniBand ports
  - One 100/1000/10000 Mb/s management port
  - One 1Gb/s IPMI port
  - Two redundant power supply units (PSUs)
- One 2U Disk Array Enclosure (DAE) containing:

- Up to 72 SSDs of sizes 400GB (for X2-S) or 1.92TB (for X2-R)
- Two redundant SAS interconnect modules
- Two redundant power supply units (PSUs)

- 
- 



The Storage Controllers on each X-Brick are connected to their DAE via redundant SAS interconnects. An XtremIO storage array can have one or multiple X-Bricks. Multiple X-Bricks are clustered together into an XtremIO array, using an InfiniBand switch and the Storage Controllers' InfiniBand ports for back-end connectivity between Storage Controllers and DAEs across all X-Bricks in the cluster. The system uses the Remote Direct Memory Access (RDMA) protocol for this back-end connectivity, ensuring a highly-available ultra-low latency network for communication between all components of the cluster. The InfiniBand switches are the same size (1U) for both X2-S and X2-R cluster types, but include 12 ports for X2-S and 36 ports for X2-R. By leveraging RDMA, an XtremIO system is essentially a single shared-memory space spanning all of its Storage Controllers.

The 1GB port for management is configured with an IPv4 address. The XMS, which is the cluster's management software, communicates with the Storage Controllers via the management interface. Through this interface, the XMS communicates with the Storage Controllers and sends storage management requests such as creating an XtremIO Volume, mapping a Volume to an Initiator Group, etc. The second 1GB/s port for IPMI interconnects the X-Brick's two Storage Controllers. IPMI connectivity is strictly within the bounds of an X-Brick, and will never be connected to an IPMI port of a Storage Controller in another X-Brick in the cluster.

### 3.4 Dell Wyse Endpoints

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

### 3.4.1 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](#)



### 3.4.2 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](#)

### 3.4.3 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 (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). For more information, please visit: [Link](#)



### 3.4.4 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. For more information, please visit [Link](#)

### 3.4.5 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). For more information, please visit: [Link](#)



### 3.4.6 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](#)

### 3.4.7 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 that's built in to Windows Server. As with other hypervisors, Hyper-V allows you to create and run virtual machines (software version of computers). 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 server version used with 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:**



**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 is Server based instead of Desktop 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 into Dell EMC Ready Bundle for VDI Architecture

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 EMC Ready Bundle for VDI solution is as follows:


#### 4.1.3.3 NUMA alignment – Skylake processor (R640)

20 physical cores per CPU given the Gold 6138 part, 40 logical cores with Hyper-Threading active, gives us a total of 80 consumable cores per compute node and falls in line with a 2x oversubscription rate. 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 NUMA Alignment – Broadwell processor (M630)

20 physical cores per CPU given the 2698v4 part and 40 logical cores with Hyper-Threading active, gives us a total of 80 consumable cores per compute node and falls in line with a 2x oversubscription rate. Configuring the RDSH VMs as shown below will ensure that no physical NUMA node spanning occurs which could lower performance for an effected VM.



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 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		Size (GB)	Tier2 Vol
RD Broker & Licensing	4	8	4GB   10GB	20%	Med	1	60	-
RD Gateway & Web Access	4	4	2GB   10GB	20%	Med	2	60	-
Primary SQL	8	8	4GB   16GB	20%	Med	1	60	300 (VHDX)
SCVMM (optional)	8	8	4GB   16GB	20%	Med	1	60	50 (VHDX)
File Server	2	4	2GB   6GB	20%	Med	1	60	2048 (PTD)
<b>Total</b>	<b>30</b>	<b>32GB</b>	<b>16GB   58GB</b>	-	-	<b>7</b>	<b>300GB</b>	<b>2398GB</b>

**NOTE:** SCVMM is optional for unmanaged personal collections.

#### 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 assuming the CPU in use is the Intel Xeon Gold 6138 or E5-2698v4. For more information on NUMA please refer to the [NUMA Architecture Considerations](#) section.

Role	VMs per host	vCPUs per VM	Startup RAM (GB)	Dynamic Memory			NIC	OS vDisk	
				Min Max	Buffer	Weight		Size (GB)	Location
RDSH VM	10	8	32	16GB   48GB	20%	Med	1	80	Tier 1

## 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
- 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

The Dell EMC Ready Bundle for VDI solution has a wide variety of Tier 1 and Tier 2 storage options to provide maximum flexibility to suit any use case. Customers have the choice to leverage best-of-breed Dell EMC storage solutions using Fibre Channel (FC) or iSCSI while being assured the storage Tiers of the Dell EMC Ready Bundle for VDI solution will consistently meet or outperform user needs and expectations. This solution architecture is using the Dell EMC XtremIO X2 X-Brick with FC as the storage array for combined Tier 1 and Tier 2.

## 5.2.1 Local Tier 1 storage

Selecting the local Tier 1 storage model means that the host servers use four to ten locally installed drives (hard disks or SSDs) to house the compute and management VMs. To achieve the required performance level, RAID 10 is recommended for use across all local disks. Density is reduced since compute and management VMs are combined on the same single server. This model is suitable for ROBO or proof-of-concept/pilot deployments.

## 5.2.2 Shared Tier 1 storage

Selecting the Shared Tier 1 model means that the virtualization compute hosts are deployed without Tier 1 local storage and leverage shared storage hosted on a high performance array. In this model, shared storage is leveraged for Tier 1 and used for VDI execution. Considering that the maximum density per compute server is well under 500 VMs, we recommend a single LUN per server. The size of each volume will depend on the workload density (number of VMs per server) and the type of desktops (pooled or personal).

## 5.2.3 Shared Tier 2 storage

Tier 2 is shared storage used to host the Management server VMs and user file data. In this solution architecture, shared Tier 2 storage is using the same X-Brick array as the Tier 1 storage. The X-Brick has no built-in filer capability so user file data is stored via file server VMs residing on the storage. A single management volume is sufficient for the VMs hosted by the management servers or optionally, multiple volumes designating logical separation. The example table below provides guidance for 500 users with ~4GB of data/user presented via a file server VM. Volume sizes should be adjusted accordingly if user data is not stored on the array or if larger per user data sizes are required. The solution as designed presents all SQL disks using VHDX format.

Volumes	Size (GB)	Storage Array	Purpose	File System
Management	350	Tier 2	RDS roles, SCVMM (optional), File & SQL (both optional)	NTFS
User Data	2048	Tier 2	File Server	PTD
User Profiles	200	Tier 2	User profiles	NTFS
SQL DATA	100	Tier 2	SQL	NTFS
SQL LOGS	100	Tier 2	SQL	NTFS
TempDB Data	5	Tier 2	SQL	NTFS
TempDB Logs	5	Tier 2	SQL	NTFS
Templates/ISO	200	Tier 2	ISO storage (optional)	NTFS

---

**NOTE:** Table above is for example purposes. Number and size of volumes can vary depending on customer requirements.

---

## 5.2.4 Storage networking – XtremIO Fiber Channel (FC)

The XtremIO all-flash array provides built-in intelligence and automation to dynamically manage enterprise data throughout its lifecycle. Together, block-level intelligence, storage virtualization, integrated software and modular, platform-independent hardware enable exceptional efficiency, simplicity and security.

XtremIO actively manages data at a block level using real-time intelligence, providing fully virtualized storage at the disk level. Resources are pooled across the entire storage array. All virtual volumes are thin-provisioned. With inline data compression and dedupe, physical storage requirements can be vastly reduced.


Although a single Fabric can be configured to begin with to reduce costs, as a best practice recommendation, the environment is configured with two Fabrics to provide multi-pathing and end-to-end redundancy.

The following QLogic HBA settings are used:

- LUN queue depth set to 256
- HBA queue depth set to 65535

Configure Microsoft Multipath I/O (MPIO) with the least queue depth policy.

Refer to the [EMC XtremIO Storage Array – Host Configuration Guide](#) for setting details.



### 5.2.4.1 FC zoning

Zone at least one port from each server HBA to communicate with each XtremIO controller. The result of this is two distinct FC Fabrics and four redundant paths per server. Round Robin or Fixed Paths are supported. Use a single-Target-per-single-Initiator (1:1) zoning scheme.


## 5.3 Virtual networking

### 5.3.1 Local Tier 1

In this model, the servers do not need access to FC storage since they are hosting VDI VMs on local disk. Since this solution model is only recommended for single server deployments, VLAN requirements are reduced:

- Management VLAN: Configured for hypervisor infrastructure traffic – L3 routed via core switch
- VDI VLAN: Configured for VDI session traffic – L3 routed via core switch
- VDI Management VLAN: Configured for VDI infrastructure traffic (optional) – L3 routed via core switch
- A VLAN for iDRAC is configured for all hardware management traffic – L3 routed via core switch

As shown in the diagram below, native Windows Server NIC Teaming is utilized to load balance and provide resiliency for network connections. A single LBFO NIC team is configured to connect to a Hyper-V switch for external traffic. All vNICs associated with the Management OS connect directly to the external Hyper-V switch. This traffic is combined within a single switch; however, VLANs are required for each traffic type to enable traffic separation. Configure the LAN traffic from the server to the ToR switch as a LAG.




## 5.3.2 Shared Tier 1 – FC

Using Fiber Channel based storage requires additional storage fabrics to be built out in the network stack. The network configuration in this model is identical between the Compute and Management hosts. The benefits of shared storage are available to all hosts such as Live Migration and HA. The following outlines the VLAN requirements for the Compute and Management hosts in this solution model:

- Compute hosts (Shared Tier 1)
  - Management VLAN: Configured for hypervisor Management traffic – L3 routed via core switch
  - Live Migration VLAN: Configured for Live Migration traffic – L2 switched only, trunked from Core
  - Failover Cluster VLAN: Configured for Cluster and Cluster Shared Volume traffic – L2 switched only, trunked from core
  - VDI VLAN: Configured for VDI session traffic – L3 routed via core switch
- Management hosts (Shared Tier 1)
  - Management VLAN: Configured for hypervisor Management traffic – L3 routed via core switch
  - Live Migration VLAN: Configured for Live Migration traffic – L2 switched only, trunked from Core
  - Failover Cluster VLAN: Configured for Cluster and Cluster Shared Volume traffic – L2 switched only, trunked from core
  - VDI Management VLAN: Configured for VDI infrastructure traffic – L3 routed via core switch
- A VLAN for iDRAC is configured for all hardware management traffic – L3 routed via core switch

FC and LAN traffic are physically separated into discrete switching Fabrics. Each Shared Tier 1 Compute and Management host have a quad port NDC (4 x 10Gb) as well as 2 x 8Gb dual port FC HBAs. LAN traffic from the server to the ToR switch is configured as a LAG.

As shown in the diagram below, native Windows Server NIC Teaming is utilized to load balance and provide resiliency for network connections. For the compute or management hosts in this scenario, a single LBFO NIC team is configured to connect to a Hyper-V switch for external traffic. All vNICs associated with the Management OS connect directly to the external Hyper-V switch with MPIO used to connect to shared storage.



## 5.4 Scaling guidance

- 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
Compute Servers	Desktop VMs per physical host based on available CPU	Additional hosts and clusters added as necessary	Additional RAM or CPU compute power
Mgmt. Servers	Number of server VMs per host	Add additional hosts	Add RAM or network adapters
Broker Servers	Desktops per instance (dependent on SQL performance as well)	Additional servers added to the farm	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 farm	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

File Services	Concurrent connections, responsiveness of reads/writes	Split user profiles and home directories between multiple file servers in the cluster. File services can also be migrated to an optional NAS device to provide high availability.	Additional RAM and CPU for the management nodes
---------------	--	---	---

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

### **XtremIO X2 Scaling**

With X2, an XtremIO cluster has both scale-out and scale-up capabilities. Scale-out is implemented by adding X-Bricks to an existing cluster (up to four). The addition of an X-Brick to an existing cluster increases its compute power, bandwidth and capacity linearly. Each X-Brick that is added to the cluster brings with it two Storage Controllers, each with its CPU power, RAM and FC/iSCSI ports to service the clients of the environment, together with a DAE with SSDs to increase the capacity provided by the cluster. Adding an X-Brick to scale-out an XtremIO cluster is intended for environments that grow both in capacity and performance needs.


Scale-up of an XtremIO cluster is implemented by adding SSDs to existing DAEs in the cluster. Adding SSDs to existing DAEs to scale-up an XtremIO cluster is intended for environments that currently grow in capacity needs without the need for extra performance. This occurs, for example, when the same number of users have an increasing amount of data to save, or when an environment grows in both capacity and performance needs but has only reached its capacity limits with still room to grow in performance with its current infrastructure.

Each DAE can hold up to 72 SSDs, and is divided to up to 2 groups of SSDs called Data Protection Groups (DPGs). Each DPG can hold a minimum of 18 SSDs and can grow by increments of 6 SSDs up to the maximum of 36 SSDs. In other words, 18, 24, 30 or 36 is the possible number of SSDs per DPG, when up to 2 DPGs can occupy a DAE.


## 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 compute and management hosts are added to their respective layers, Hyper-V clustering is introduced in both the management and compute layers, SQL is configured with AlwaysOn and native load balancing.

### Rack Servers



## Blade Servers




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 Hyper-V hosts added in the compute or management layers to provide N+1 protection.
- 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.


### 5.5.1 SQL Server high availability

HA for SQL is provided via AlwaysOn using either Failover Cluster Instances or Availability Groups. This configuration protects all critical data stored within the database from physical server as well as virtual server problems. DNS is used to control access to the primary SQL instance. Place the principal VM that will host the primary copy of the data on the first Management host. Additional replicas of the primary database are placed on subsequent Management hosts.



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 [Servers](#) and [Platform Configurations](#) sections for hardware specifications.

User density summary

Host Config	Hypervisor	Deployment / Provisioning	Workload	Template	User Density
B5	Hyper-V 2016	RDVH – Pooled Desktops	Knowledge Worker	Windows 10 x64 & Office 2016	180

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 EMC Ready Bundle for VDI 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 EMC Ready Bundle for VDI 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 EMC Ready Bundle for VDI 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 EMC Ready Bundle for VDI 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 EMC Ready Bundle for VDI 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 EMC Ready Bundle for VDI 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%

---

**NOTE:** The CPU utilization threshold is set to 85% since the additional headroom provided by the processor Turbo Boost feature is correctly reported by Performance Monitor.

---

## 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)	PowerEdge R740
Hypervisor(s)	Windows Server 2016 Hyper-V
Broker technology	RDS
Broker database	Microsoft SQL 2016
Management VM OS	Windows Server 2016
Virtual desktop OS	Windows 10 Enterprise 64-bit
Office application suite	Office Professional 2016
Login VSI test suite	Version 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 Start-up Memory	Hyper-V Min   Max Dynamic	Screen Resolution	Operating System
Task Worker	1	1GB	1GB   2GB	1280 X 720	Windows 10 Enterprise 64-bit
Knowledge Worker	2	1.5GB	1GB   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 Start-up 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.

### Compute hardware configuration

Platform Config	CPU	Memory	RAID Ctlr	HD Config	Network & HBA
R740 (B5)	Intel(R) Xeon(R) Gold 5120 CPU @ 2.20GHz	384GB @2400 MT/s	Dell Perc H740P	Shared Tier1 XtremIO	Intel(R) 10GbE 4P X710 rNDC
					QLogic QLE2562 8Gb Dual Port

### Shared storage hardware configuration

Storage Platform	Controller Software Version	SSD Configuration
XtremIO X2 X-Brick	6.0.0.52	18 x 400GB

## 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	Deployment / Provisioning	Login VSI Workload	Density per Host	Avg CPU	Avg Mem Consumed	Avg IOPS / User	Avg Net Mbps / User
B5	Hyper-V 2016	RDVH – Pooled Desktops	Knowledge Worker	180	70%	335GB	8	1.13

**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.

## 6.3.1 R740 Compute


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

### 6.3.1.1 Knowledge Worker, 180 Users, Hyper-V 2016, RDVH Desktops


The below graph shows the CPU performance data. The CPU reached 80% usage during the logon period and settles to 65% usage during the steady state with 180 users logged in and performing the Knowledge workload.



335GB of memory was consumed by the VMs when all the sessions were logged in. The memory consumption at 87% is just at the acceptable threshold for a VDI solution, leaving a small surplus to allow for extra resource requirements in the solution.





Network bandwidth is not an issue with a peak during logon of 240Mbps and a steady state peak of approximately 205 Mbps. Network usage is affected by the Login VSI content and profile and home folder redirection. The output queues on the network adapters remained at zero (0) throughout the testing.




The host reached a maximum of 22000 IOPs during the boot phase, 5000 during the login phase and dropped to 1300 during the steady state resulting in about 8 IOPS per user in the later state of the test. In all cases the latency remained below 1 ms during the logon and steady state of the test. However latency during

the boot stage did reach 12 ms as measured from the Hyper-V host. The X-Brick Performance Chart is shown following the Hyper-V chart. This shows that the latency was observed as zero (0) during the test with one brief spike to 4.5 ms during the boot stage.



The Login VSI Max user experience score shown below for this test was not reached and there was still plenty of room before the VSI index would reach the threshold.



## 7 Related resources

See the following referenced or recommended resources:

- The [Dell EMC Cloud-Client Computing Datacenter Tech Center](#) page which includes this RA and other CCC reference architectures
- [Hyper-V Technology Overview](#)
- [What's new in Hyper-V on Windows Server 2016](#) and [Hyper-V on Windows Server 2016](#)
- [EMC XtremIO Storage Array Host Configuration Guide](#) – includes HBA and MPIO configuration settings for Windows
- [Dell EMC XtremIO All-Flash Storage With Microsoft Hyper-V 2016](#) – deployment considerations, integration options and best practices when using Dell EMC XtremIO in a Microsoft virtualized environment
- [Windows Server 2016 Licensing Datasheet](#)
- [System Center 2016 pricing and licensing datasheet](#)
- [Windows Server 2016 and System Center 2016 licensing FAQ](#)
- [Microsoft Software Assurance for Thin Clients with Windows 10 IoT Enterprise](#)