

REFERENCE ARCHITECTURE

Dell EMC Ready Bundle for VDI

Design, configuration and implementation of a VMware Horizon environment with Dell EMC PowerEdge Servers and XtremIO Storage

[Abstract](#)

A Reference Architecture for integrating Dell EMC PowerEdge servers and VMware Horizon brokering software on VMware ESXi hypervisor to create virtual application and virtual desktop environments on 14th generation Dell EMC PowerEdge Servers.

February 2018

Revisions

Date	Description
September 2017	Initial release
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
1 Introduction.....	7
1.1 Purpose	7
1.2 Scope.....	7
1.3 What's new	7
2 Solution architecture overview	8
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	11
2.4 Local tier 1	11
2.5 Shared tier 1 for rack servers	11
2.5.1 Shared tier 1 – network architecture.....	13
2.5.2 Shared tier 1 – rack cabling (HA)	14
2.5.3 Shared tier 1 – storage scaling guidance	14
2.6 Shared tier 1 for blade servers	15
2.6.1 Shared tier 1 – network architecture.....	16
2.6.2 Shared tier 1 – cabling.....	17
2.6.3 Shared tier 1 – scaling guidance	18
3 Hardware components	19
3.1 Network.....	19
3.1.1 Dell Networking S3048 (1Gb ToR switch).....	19
3.1.2 Dell Networking S4048 (10Gb ToR switch).....	20
3.1.3 Brocade 6510 (FC ToR switch)	20
3.1.4 Brocade M5424 (FC blade interconnect)	22
3.1.5 PowerEdge M I/O aggregator (10Gb blade interconnect).....	23
3.2 Servers	23
3.2.1 Dell EMC PowerEdge R640	23
3.2.2 PowerEdge R740.....	26
3.2.3 Dell EMC PowerEdge M630.....	27

3.3	Storage	28
3.3.1	XtremIO X2 X-Brick – Combined Tier 1 and Tier 2	28
3.4	GPUs	32
3.4.1	NVIDIA Tesla GPUs	32
3.5	Dell Wyse Thin Clients	34
3.5.1	Wyse 5030 PCoIP Zero Client.....	34
3.5.2	Wyse 5040 AIO Thin Client with PCoIP	34
3.5.3	Wyse 5050 AIO PCoIP Zero Client	35
3.5.4	Wyse 7030 PCoIP Zero Client.....	35
3.5.5	Wyse 5060 Thin Client (ThinOS) with PCoIP.....	35
3.5.6	Wyse 7040 Thin Client with Windows Embedded Standard 7P.....	36
3.5.7	Wyse 7020 Thin Client (WES 7/7P, WIE10, ThinLinux).....	36
3.5.8	Latitude 3480 and 5280 Mobile Thin Clients (Win 10 IoT)	37
4	Software Components.....	38
4.1	VMware.....	38
4.1.1	VMware vSphere 6.x	38
4.1.2	VMware Horizon	39
4.1.3	VMware Horizon Apps.....	41
4.2	NVIDIA GRID vGPU	43
4.2.1	vGPU Profiles	43
5	Solution Architecture for Horizon 7.....	50
5.1	Management Server Infrastructure.....	50
5.1.1	RDSH VM Configuration.....	50
5.1.2	NVIDIA GRID License Server Requirements	51
5.1.3	SQL Databases	51
5.1.4	DNS	52
5.2	Storage Architecture Overview.....	52
5.2.1	Local Tier 1 Storage	52
5.2.2	Shared Tier 1 Storage	52
5.2.3	Shared Tier 2 Storage	53
5.2.4	Storage Networking – XtremIO Fibre Channel (FC).....	53
5.3	Virtual Networking.....	54
5.3.1	Local Tier 1	54
5.3.2	Shared Tier 1	55
5.3.3	VMware NSX	56
5.4	Scaling Guidance.....	58

5.5	Solution High Availability	59
5.5.1	vSphere HA (Shared Tier 1)	61
5.5.2	Management Server High Availability.....	61
5.5.3	Horizon CS High Availability	61
5.5.4	SQL Server High Availability	61
5.6	VMware Horizon communication flow	62
6	Solution Performance and Testing	63
6.1	Test and performance analysis methodology.....	63
6.1.1	Testing process	63
6.1.2	Resource monitoring	65
6.1.3	Resource utilization	66
6.2	Test configuration details.....	66
6.2.1	Compute VM configurations	67
6.2.2	Platform Configuration.....	67
6.3	Test results and analysis	67
6.3.1	B5 Compute.....	69

Executive summary

This document provides the reference architecture for integrating a Dell EMC Ready Bundle for VDI with VMware® Horizon™ software to create virtual application and virtual desktop environments.

The Dell EMC Ready Bundle for VDI is a comprehensive solution that encompasses storage, compute, networking, and virtualization using industry-proven Dell EMC PowerEdge™ server technology.

VMware Horizon provides a complete end-to-end virtualization solution delivering Microsoft® Windows™ virtual desktops or server-based hosted shared sessions to users on a wide variety of endpoint devices

1 Introduction

1.1 Purpose

This document addresses the architecture design, configuration and implementation considerations for the key components of the architecture required to deliver virtual desktops via VMware Horizon on VMware vSphere 6.

1.2 Scope

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

- Define the detailed technical design for the solution.
- Define the hardware requirements to support the design.
- Define the constraints that 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 component selection guidance.

1.3 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 include:

- Pooled Virtual Desktops (Non-persistent)
- RDSH (application/ session virtualization)

2.2 Physical architecture overview

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

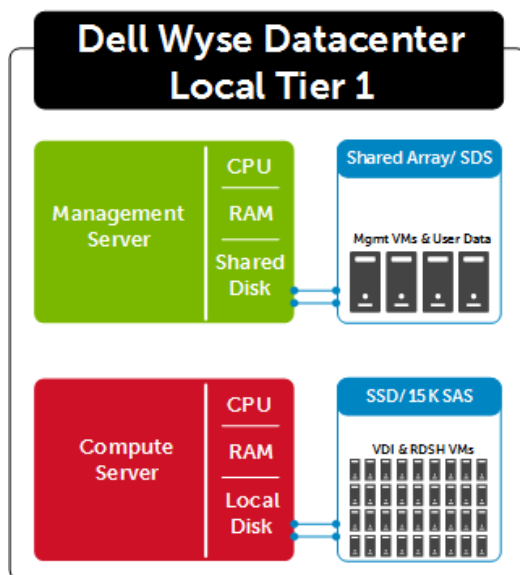


Figure 1 Local Tier 1

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

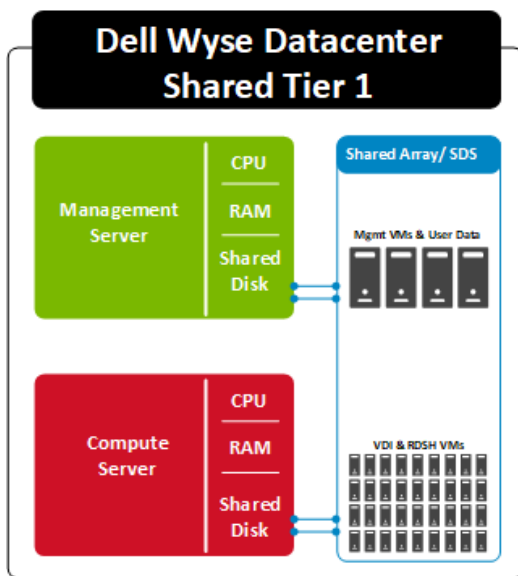


Figure 2 Shared Tier 1

NOTE: At the time of this writing, the 14th 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 section 3.6)

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 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 48-port switch for 10 GB LAN traffic along with a single Brocade fibre channel switch for SAN connectivity. Additional switches are added and stacked as required to provide High Availability for the Network layer.




Figure 3 Networking layer

2.3.2 Compute

The compute layer consists of the server resources responsible for hosting the Horizon user sessions, hosted via the VMware vSphere hypervisor, local or shared tier 1 solution models. Shared Tier 1 rack server pictured below.




Figure 4 Compute layer

2.3.3 Management

VDI management components are dedicated to their own layer so that they do not negatively affect 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 the entire server VMs necessary to support the VDI infrastructure. Shared Tier 1 rack servers pictured below.




Figure 5 Management layer

2.3.4 Storage

The storage layer consists of the Dell EMC XtremIO X2 X-Brick array for combined shared T1 and T2. The configuration shown below is the minimum disk configuration for the X2 array and 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. Additional arrays are added to the solution when scaling beyond 3,500 users. Configurations with and without user data depicted in the image below.




Figure 6 Storage layer

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.

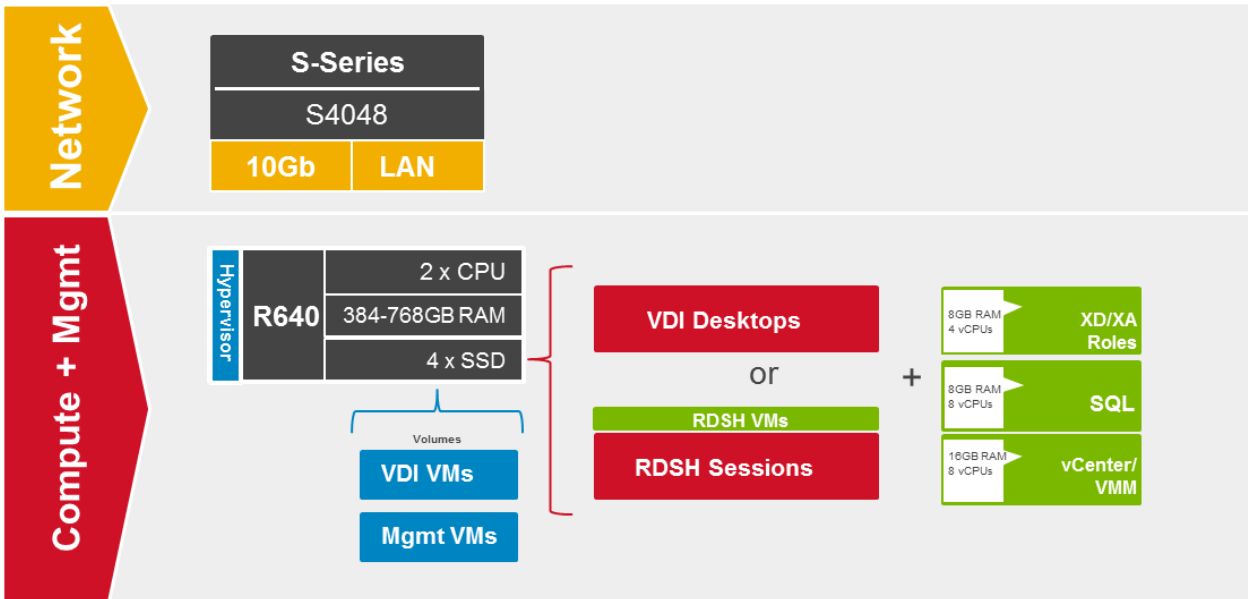


Figure 7 Local Tier 1 storage

Note: 150 user density is based on the Task Worker workload

2.5 Shared tier 1 for rack servers

Shared Tier 1 for rack servers 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 or SD cards for the

ESXi operating system. User data can either be stored on the same array as the VMs or on another storage location.

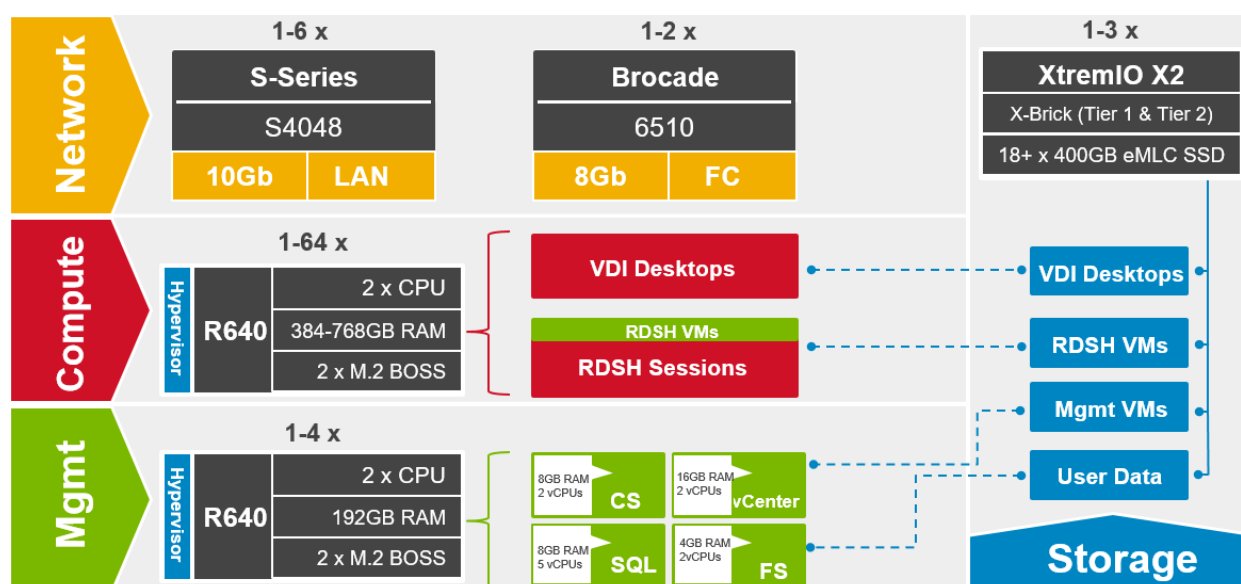


Figure 8 Shared Tier 1 storage - rack

NOTE: If necessary, additional disks can be added to increase capacity.

High-performance graphics capabilities compliment the solution and can easily be added at any time to any new or existing deployment. Simply add the appropriate number of GPU enabled servers to your architecture and provide a superior user experience with NVIDIA GRID vGPU technology. The figure below show the same architectural design without user data stored on the same storage array.

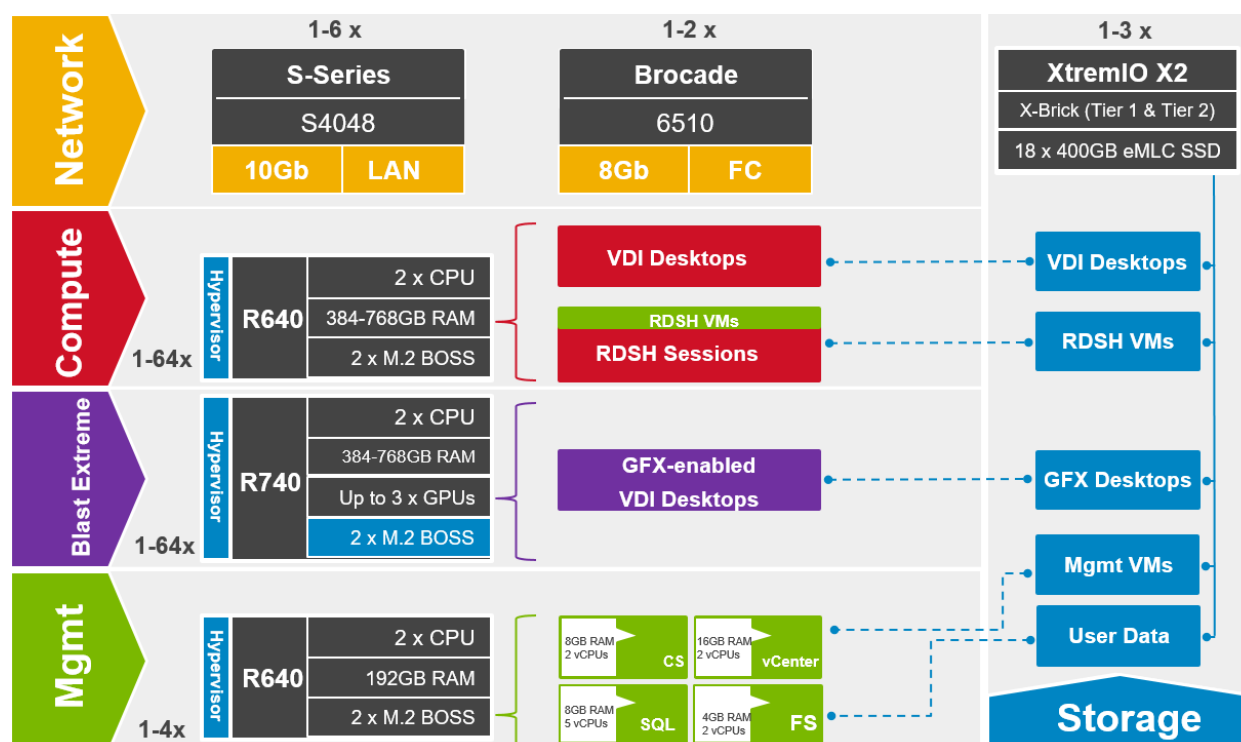


Figure 9 High performance graphics added

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 persistent disk sizes and if user data is also stored on the array.

2.5.1 Shared tier 1 – network architecture

In the Shared Tier 1 architecture for rack servers using FC, a separate switching infrastructure is required for FC. Management and compute servers both connect to shared storage using FC. Both management and compute servers connect to all network VLANs in this model. All ToR traffic has designed to be layer 2/ switched locally, with all layer 3/ routable VLANs routed through a core or distribution switch. The following diagrams illustrate the server NIC to ToR switch connections, vSwitch assignments, as well as logical VLAN flow in relation to the core switch.




Figure 10 Shared Tier 1 network architecture

2.5.2 Shared tier 1 – rack cabling (HA)

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




Figure 11 Shared Tier 1 rack cabling

2.5.3 Shared tier 1 – storage scaling guidance

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

Table 1 Shared tier 1 storage scaling guidance

Shared Tier 1 HW Scaling (Rack - FC)						
User Scale	XtremIO X2	Compute			ToR	ToR
		Servers	CPU Cores	Memory	LAN	8Gb FC
Up to 3500	1 x X-Brick	1 – 20	Up to 800	Up to 15TB	1 x S4048	1 x 6510
3501 – 7000	2 x X-Brick	20 – 39	Up to 1,560	Up to 30TB	1-2 x S4048	1-2 x 6510
7001 – 10000	3 x X-Brick	39 – 56	Up to 2,240	Up to 43TB	2-3 x S4048	2 x 6510

NOTE: For deployments over 10,000 users, create additional pods using sizing guidance contained herein.

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 either mirrored SD cards (ESXi only) for the operating system. User data can either be stored on the same array as the VMs or on another storage location.

NOTE: At the time of this writing, the 14th generation blade servers are not yet available.

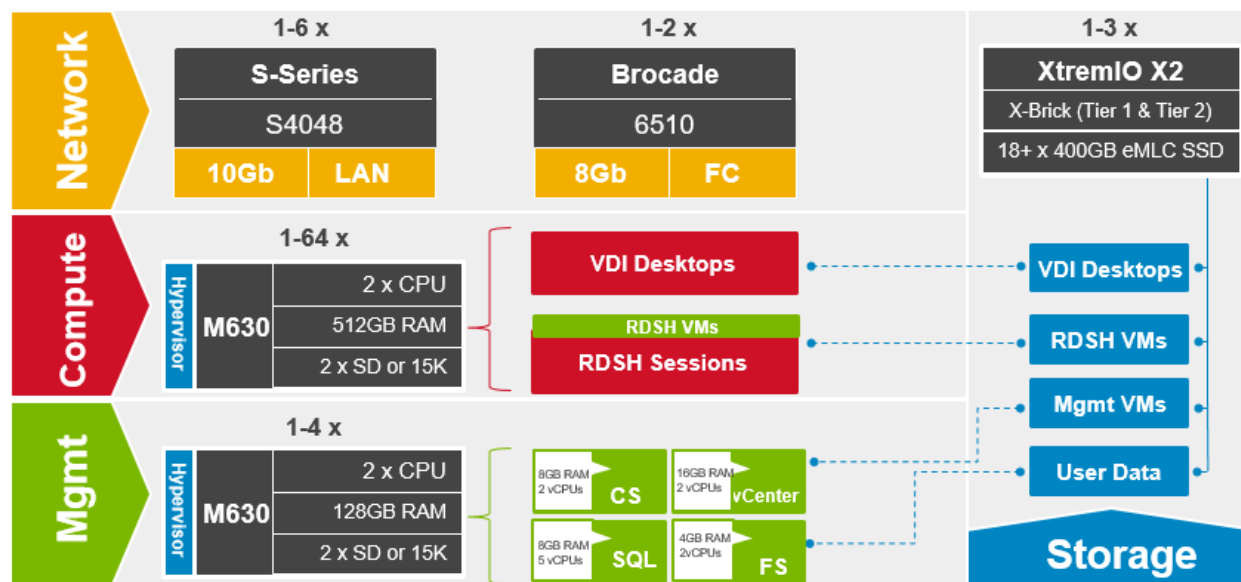


Figure 12 Shared Tier 1 – blade servers

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.

2.6.1 Shared tier 1 – network architecture

In the Shared Tier 1 for blade servers architecture, there is no need to switch LAN ToR since the IOAs in the chassis support LAN to the blades and are uplinked to the core or distribution layers directly. However, a separate switching infrastructure is required for FC. Management and compute servers both connect to shared storage using FC switched via chassis interconnects. Both management and compute servers connect to all network VLANs in this model. For greater redundancy, a ToR switch is used to support iDRAC used outside of the chassis. All ToR traffic has been designed to be layer 2 locally, with all layer 3 VLANs routed through a core or distribution switch. The following diagrams illustrate the server NIC to ToR switch connections, vSwitch assignments, as well as logical VLAN flow in relation to the core switch.




Figure 13 Shared tier 1 network architecture

2.6.2 Shared tier 1 – cabling

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




Figure 14 Shared tier 1 cabling

2.6.3 Shared tier 1 – scaling guidance

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

Table 2 Shared tier 1 scaling guidance

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 – 20	Up to 800	Up to 10TB	4 x IOA + 4 x M5424	1 x 6510
3501 – 7000	2 x X-Brick	20 – 39	Up to 1,560	Up to 20TB	6 x IOA + 6 x M5424	1 x 6510
7001 – 10000	3 x X-Brick	39 – 56	Up to 2,240	Up to 30TB	8 x IOA + 8 x M5424	2 x 6510

Note: For deployments over 10,000 users, create additional pods using sizing guidance contained herein.

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 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.

Table 3 Dell Networking S3048 features

Model	Features	Options	Uses
Dell Networking S3048-ON	<ul style="list-style-type: none">• 48 x 1000BaseT• 4 x 10Gb SFP+• Non-blocking, line-rate performance• 260Gbps full-duplex bandwidth• 131 Mbps forwarding rate	<ul style="list-style-type: none">• Redundant hot-swap PSUs & fans	1Gb connectivity
		<ul style="list-style-type: none">• VRF-lite, Routed VLT, VLT Proxy Gateway	
		<ul style="list-style-type: none">• User port stacking (up to 6 switches)	
		<ul style="list-style-type: none">• Open Networking Install Environment (ONIE)	




Figure 15 Dell Networking S3048

3.1.2 Dell Networking S4048 (10Gb ToR switch)

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

Table 4 Dell Networking S4048 features

Model	Features	Options	Uses
Dell Networking S4048-ON	<ul style="list-style-type: none"> • 48 x 10Gb SFP+ • 6 x 40Gb QSFP+ • Non-blocking, line-rate performance • 1.44Tbps bandwidth • 720 Gbps forwarding rate • VXLAN gateway support 	• Redundant hot-swap PSUs & fans	10Gb connectivity
		• 72 x 10Gb SFP+ ports with breakout cables	
		• User port stacking (up to 6 switches)	
		• Open Networking Install Environment (ONIE)	




Figure 16 Dell Networking S4048

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 Fibre 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.

Table 5 Brocade 6510 features

Model	Features	Options	Uses
Brocade 6510	<ul style="list-style-type: none">• 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




Figure 17 Brocade 6510 FC switch

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.

Table 6 Brocade M5424 features

Model	Features	Options	Uses
Brocade M5424	<ul style="list-style-type: none">• 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



Figure 18 Brocade M5424 FC blade interconnect

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.

Table 7 PowerEdge M I/O aggregator features

Model	Features	Options	Uses
PowerEdge M I/O Aggregator (IOA)	<ul style="list-style-type: none">Up to 32 x 10Gb ports + 4 x external SFP+2 x line rate fixed QSFP+ ports2 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	

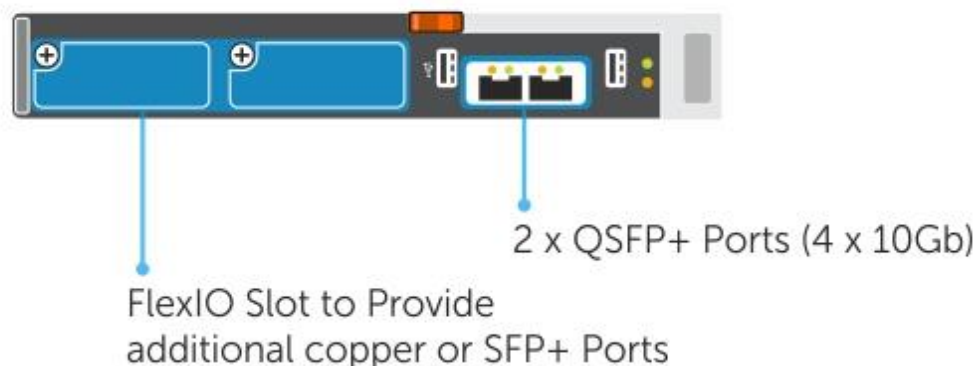


Figure 19 Table 7 PowerEdge M I/O aggregator

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

3.2 Servers

NOTE: Configurations shown below are recommendations for VDI based on stated density determined by our testing and are representative of the C7 classification found in our other reference architectures located here: [LINK](#). They do not represent absolute platform maximums and can be adjusted as needed.

3.2.1 Dell EMC PowerEdge R640

The Dell EMC PowerEdge R640 is the ideal dual-socket, 1U platform for dense scale-out datacenter computing. The R640 combines density, performance and scalability to optimize application performance and datacenter 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 option exists in the form of a Boot Optimized Storage Subsystem (BOSS) card 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).





Figure 20 PowerEdge R640

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 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. The OS can be installed to the BOSS device for vSphere or to dual SD cards. 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 Effective speed: 2667MT/s @ 768GB
Storage Ctrlrs	PERC H730P – RAID10
Storage	2 x 120GB M.2 BOSS in RAID1 (Hypervisor) 4 x 960GB SSD or 10 x 600GB 15K SAS (VMs)
Network	4 x 10Gb SFP+ (BT options available)
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, VDI desktops or RDSH sessions 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. ESXi can be installed to the BOSS device (as shown in the table) or to dual SD cards.

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 Effective speed: 2667MT/s @ 768GB	12 x 16GB 2667MT/s RDIMMs Effective speed: 2400MT/s @ 192GB
Storage Ctrls	PERC H330 – no RAID	PERC H330 – no RAID
Storage	2 x 120GB M.2 BOSS in RAID1 (Hypervisor)	2 x 120GB M.2 BOSS in RAID1 (Hypervisor)
Network	4 x 10Gb SFP+ (BT options available) 2 x QLogic 2562 8Gb DP FC HBA	4 x 10Gb SFP+ (BT options available) 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 R740

The graphics compute host for this solution (ESXi hypervisor only) is the Dell EMC PowerEdge R740 server. The PowerEdge R740 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 PowerEdge R740 can be outfitted with 3 double-wide GPU accelerators. Recommended system components are shown in the table below; however, CPUs and memory can be adjusted to suit customer requirements based on the desired graphics profiles used. Refer to the vGPU profiles section for the graphics profiles and densities possible with GPU-enabled VMs.




Figure 21 PowerEdge R740

R740 Shared T1	
	Compute
CPU	2 x Intel Gold 6138 (20C, 2.0GHz)
Memory	24 x 32GB 2667MT/s RDIMMs Effective speed: 2667MT/s @ 768GB
Storage Ctrls	PERC H330 – no RAID
Storage	2 x 120GB M.2 BOSS in RAID1 (Hypervisor)
GPUs	Up to 3 x FLDW GPU (NVIDIA)
Network	4 x 10Gb SFP+ (BT options available) 2 x QLogic 2562 8Gb DP FC HBA
iDRAC	iDRAC9 Enterprise with OpenManage Essentials
Power	2 x 2000W PSUs

Management server configuration is the same as shown above for Shared Tier 1 Rack. For more information on the R740, please visit: [Link](#)

NOTE: Depending on the number of GPU cards and Graphics profiles used, higher core processors and memory may be necessary.

3.2.3 Dell EMC 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-processor 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 processors(Broadwell) and 24 DIMMs (768GB RAM) of DDR4 memory in an extremely compact half-height blade form factor.




Figure 22 PowerEdge M630

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

3.2.3.1 Shared tier 1 blade (FC)

The Shared Tier 1 blade server varies slightly from the rack server equivalent since the latest 14th 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 Effective speed: 2400MT/s @ 512GB	8 x 16GB 2400MT/s RDIMMs Effective speed: 2400MT/s @ 128GB
	Storage Ctrls	PERC H330 – RAID1	PERC H330 – RAID1
	Storage	2 x SSD or 15K SAS (Hypervisor)	2 x SSD or 15K SAS (Hypervisor)
	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)
- Environment Consolidation (using XtremIO Virtual Copies or VMware's XCOPY)

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.

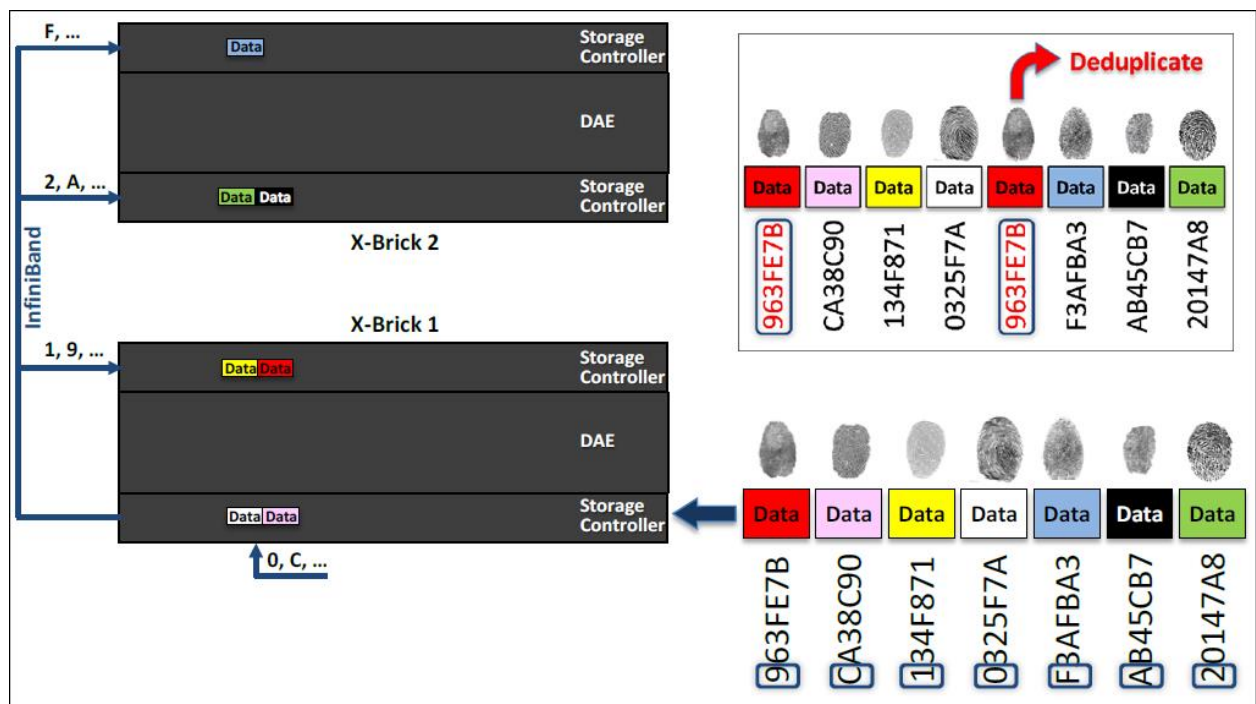


Figure 23 XtremIO streamlined data

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.




Figure 24 Capacity savings with in-line deduplication and compression

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)




Figure 25 X-Brick features diagram

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 GPUs

3.4.1 NVIDIA Tesla GPUs

Accelerate your most demanding enterprise data center workloads with NVIDIA® Tesla® GPU accelerators. Scientists can now crunch through petabytes of data up to 10x faster than with CPUs in applications ranging from energy exploration to deep learning. In addition, Tesla accelerators deliver the horsepower needed to run bigger simulations faster than ever before. For enterprises deploying VDI, Tesla accelerators are perfect for accelerating virtual desktops.

3.4.1.1 NVIDIA Tesla M10

The NVIDIA® Tesla® M10 is a dual-slot 10.5 inch PCI Express Gen3 graphics card featuring four mid-range NVIDIA Maxwell™ GPUs and a total of 32GB GDDR5 memory per card (8GB per GPU). The Tesla® M10 doubles the number of H.264 encoders over the NVIDIA® Kepler™ GPUs and improves encoding quality, which enables richer colors, preserves more details after video encoding, and results in a high-quality user experience.



The NVIDIA® Tesla® M10 GPU accelerator works with NVIDIA GRID™ software to deliver the industry's highest user density for virtualized desktops and applications. It supports up to 64 desktops per GPU card (up to 128 desktops per server) and gives businesses the power to deliver great graphics experiences to all of their employees at an affordable cost.

Table 8 Tesla M10 specifications

Specs	Tesla M10
Number of GPUs	4 x NVIDIA Maxwell™ GPUs
Total CUDA cores	2560 (640 per GPU)
GPU Clock	Idle: 405MHz / Base: 1033MHz
Total memory size	32GB GDDR5 (8GB per GPU)
Max power	225W
Form Factors	Dual slot (4.4" x 10.5")
Aux power	8-pin connector
PCIe	x16 (Gen3)
Cooling solution	Passive

3.4.1.2 NVIDIA Tesla M60

The NVIDIA® Tesla® M60 is a dual-slot 10.5 inch PCI Express Gen3 graphics card featuring two high-end NVIDIA Maxwell™ GPUs and a total of 16GB GDDR5 memory per card. This card utilizes NVIDIA GPU Boost™ technology which dynamically adjusts the GPU clock to achieve maximum performance. Additionally, the Tesla M60 doubles the number of H.264 encoders over the NVIDIA® Kepler™ GPUs.



Accelerate your most demanding enterprise data center workloads with NVIDIA® Tesla® GPU accelerators. Scientists can now crunch through petabytes of data up to 10x faster than with CPUs in applications ranging from energy exploration to deep learning. In addition, Tesla accelerators deliver the horsepower needed to run bigger simulations faster than ever before. For enterprises deploying VDI, Tesla accelerators are perfect for accelerating virtual desktops.

Table 9 Tesla M60 specifications

Specs	Tesla M60
Number of GPUs	2 x NVIDIA Maxwell™ GPUs
Total CUDA cores	4096 (2048 per GPU)
Base Clock	899 MHz (Max: 1178 MHz)
Total memory size	16GB GDDR5 (8GB per GPU)
Max power	300W
Form Factors	Dual slot (4.4" x 10.5")
Aux power	8-pin connector
PCIe	x16 (Gen3)
Cooling solution	Passive/ Active

3.5 Dell Wyse Thin Clients

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

3.5.1 Wyse 5030 PCoIP Zero Client



For uncompromising computing with the benefits of secure, centralized management, the Dell Wyse 5030 PCoIP zero client for VMware Horizon is a secure, easily managed zero client that provides outstanding graphics performance for advanced applications such as CAD, 3D solids modeling, video editing and advanced worker-level office productivity applications.

Smaller than a typical notebook, this dedicated zero client is designed specifically for VMware Horizon. It features the latest processor technology from Teradici to process the PCoIP protocol in silicon and includes client-side content caching to deliver the highest level of performance available over 2 HD displays in an extremely compact, energy-efficient form factor. The Dell Wyse 5030 delivers a rich user experience while resolving the challenges of provisioning, managing, maintaining and securing enterprise desktops. For more information, please visit: [Link](#).

3.5.2 Wyse 5040 AIO Thin Client with PCoIP



The Dell Wyse 5040 AIO all-in-one (AIO) thin client runs ThinOS with 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 Thin Client Manager 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](#)

3.5.3 Wyse 5050 AIO PCoIP Zero Client



The Wyse 5050 All-in-One (AIO) PCoIP zero client has a 23.6" Full HD display and combines the security and performance of the Wyse 5030 PCoIP zero client for VMware with the elegant design of Dell's best-selling P24 LED monitor. The Wyse 5050 AIO provides a best-in-class virtual experience with superior manageability – at a better value than purchasing a zero client and high resolution monitor separately. A dedicated hardware PCoIP engine delivers the highest level of display performance available for advanced applications, including CAD, 3D solids modeling, video editing and more. Elegant in appearance and energy efficient, the Wyse 5050 AIO is a fully functional

VMware Horizon endpoint that delivers a true PC-like experience. It offers the full benefits of an efficient and secure centralized computing environment, like rich multimedia, high-resolution 3D graphics, HD media, and full USB peripheral interoperability locally (LAN) or remotely (WAN). For more information, please visit: [Link](#).

3.5.4 Wyse 7030 PCoIP Zero Client



The Wyse 7030 PCoIP zero client from Dell offers an outstanding rich graphics user experience with the benefits of secure, centralized management. It is a secure, easily managed zero client that provides outstanding graphics performance for advanced applications such as CAD, 3D solids modeling, video editing and advanced worker-level office productivity applications. About the size of a notebook, this dedicated zero client designed specifically for VMware Horizon. It features the latest processor technology from Teradici to process the PCoIP protocol in silicon and includes client-side content caching to deliver the highest level of display performance available over 4 HD displays in a

compact, energy-efficient form factor. The Dell Wyse 7030 delivers a rich user experience while resolving the challenges of provisioning, managing, maintaining and securing enterprise desktops. For more information, please visit: [Link](#)

3.5.5 Wyse 5060 Thin Client (ThinOS) with PCoIP

The Wyse 5060 offers high performance, reliability and flexible OS options, featuring all the security and management benefits of Dell thin clients. 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. This quad core thin client 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 Cloud Client Manager (CCM) or Microsoft SCCM (5060 with Windows versions). For more information, please visit: [Link](#).

3.5.6 Wyse 7040 Thin Client with Windows Embedded Standard 7P



The Wyse 7040 is a high-powered, ultra-secure thin client. Equipped with 6th generation 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 Windows Embedded Standard 7P 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](#)

3.5.7 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](#)

3.5.8 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](#)



4 Software Components

4.1 VMware

4.1.1 VMware vSphere 6.x

The vSphere hypervisor also known as ESXi is a bare-metal hypervisor that installs directly on top of your physical server and partitions it into multiple virtual machines. Each virtual machine shares the same physical resources as the other virtual machines and they can all run at the same time. Unlike other hypervisors, all management functionality of vSphere is done through remote management tools. There is no underlying operating system, reducing the install footprint to less than 150MB.

VMware vSphere includes three major layers: Virtualization, Management and Interface. The Virtualization layer includes infrastructure and application services. The Management layer is central for configuring, provisioning and managing virtualized environments. The Interface layer includes the vSphere web client.

Throughout the Dell EMC Ready Bundle for VDI solution, all VMware and Microsoft best practices and prerequisites for core services are adhered to (NTP, DNS, Active Directory, etc.). The vCenter used in the solution is a vCenter Server Appliance (VCSA) residing on a host in the management Tier. Horizon Composer is installed on a standalone Windows Server 2012 R2 VM when using the VCSA.

VMware vSphere® is the next-generation infrastructure for next-generation applications. It provides a powerful, flexible, and secure foundation for business agility that accelerates the digital transformation to cloud computing and promotes success in the digital economy.

Improved Appliance Management

vCenter Server Appliance also exclusively provides improved appliance management capabilities. The vCenter Server Appliance Management interface continues its evolution and exposes additional configuration data. In addition to CPU and memory statistics, it now shows network and database statistics, disk space usage and health data. This reduces reliance on a command-line interface for simple monitoring and operational tasks.

VMware vCenter High Availability

vCenter Server has a new native high availability solution that is available exclusively for vCenter Server Appliance. This solution consists of active, passive, and witness nodes that are cloned from the existing vCenter Server instance. The VMware vCenter® High Availability (vCenter HA) cluster can be enabled, disabled, or destroyed at any time. There is also a maintenance mode that prevents planned maintenance from causing an unwanted failover. vCenter HA uses two types of replication between the active and passive nodes: Native PostgreSQL synchronous replication is used for the vCenter Server database; a separate asynchronous file system replication mechanism is used for key data outside of the database.

Failover can occur when an entire node is lost—host failure, for example—or when certain key services fail. For the initial release of vCenter HA, a recovery time objective (RTO) of about 5 minutes is expected, but this can vary slightly depending on the load, size, and capabilities of the underlying hardware.

Backup and Restore

New in vCenter Server is native backup and restore for the vCenter Server Appliance. This new, out-of-the-box functionality enables users to back up vCenter Server and Platform Services Controller appliances

directly from the VAMI or API. The backup consists of a set of files that is streamed to a storage device of the user's choosing using SCP, HTTP(S), or FTP(S) protocols. This backup fully supports VCSA instances with both embedded and external Platform Services Controller instances.

vSphere HA Support for NVIDIA GRID vGPU Configured VMs

vSphere HA now protects VMs with the NVIDIA GRID vGPU shared pass-through device. In the event of a failure, vSphere HA attempts to restart the VMs on another host that has an identical NVIDIA GRID vGPU profile. If there is no available healthy host that meets this criterion, the VM fails to power on.

For more information on VMware vSphere and what's new in this release, visit [link](#).

4.1.2 VMware Horizon

The solution is based on VMware Horizon, which provides a complete end-to-end solution delivering Microsoft Windows virtual desktops to users on a wide variety of endpoint devices. Virtual desktops are dynamically assembled on demand, providing users with pristine, yet personalized, desktops each time they log on.

VMware Horizon provides a complete virtual desktop delivery system by integrating several distributed components with advanced configuration tools that simplify the creation and real-time management of the virtual desktop infrastructure. For the complete set of details, please see the Horizon resources page at <http://www.vmware.com/products/horizon-view/resources.html>.

The Horizon License matrix can be found [here](#). The Horizon Enterprise license will cover Just in time desktops and App Volumes whereas these new features are not covered under the Standard and Advanced Horizon licenses.

The core Horizon components include:

Horizon Connection Server (HCS) – Installed on servers in the data center and brokers client connections, The VCS authenticates users, entitles users by mapping them to desktops and/or pools, establishes secure connections from clients to desktops, support single sign-on, sets and applies policies, acts as a DMZ security server for outside corporate firewall connections and more.

Horizon Client – Installed on endpoints. Is software for creating connections to Horizon desktops that can be run from tablets, Windows, Linux, or Mac PCs or laptops, thin clients and other devices.

Horizon Portal – A web portal to access links for downloading full Horizon clients. With HTML Access Feature enabled enablement for running a Horizon desktop inside a supported browser is enabled.

Horizon Agent – Installed on all VMs, physical machines and Terminal Service servers that are used as a source for Horizon desktops. On VMs the agent is used to communicate with the Horizon client to provide services such as USB redirection, printer support and more.

Horizon Administrator – A web portal that provides admin functions such as deploy and management of Horizon desktops and pools, set and control user authentication and more.

vCenter Server – This server provides centralized management and configuration to entire virtual desktop and host infrastructure. It facilitates configuration, provision, management services. It is installed on a Windows Server 2008 host (can be a VM).

Horizon Transfer Server – Manages data transfers between the data center and the Horizon desktops that are checked out on the end users' desktops in offline mode. This Server is required to support desktops that

run the Horizon client with Local Mode options. Replications and synchronizing are the functions it will perform with offline images.

4.1.2.1 What's new in this release of Horizon

This new release of VMware Horizon delivers following important new features and enhancements:

Just in time delivery with Instant Clone Technology

Reduce infrastructure requirements while enhancing security with Instant Clone technology and App Volumes. Instantly deliver brand new personalized desktop and application services to end users every time they log in. Just in Time Delivery with Instant Clone Technology is turning the traditional VDI provisioning model on its head.

The booted-up parent VM can be “hot-cloned” to produce derivative desktop VMs rapidly, leveraging the same disk and memory of the parent, with the clone starting in an already “booted-up” state. This process bypasses the cycle time incurred with traditional cloning where several power cycle and reconfiguration calls are usually made.

When Instant Clone technology is used in conjunction with VMware App Volumes and User Environment Manager, administrators can use Instant Clone Technology to rapidly spin up desktops for users that retain user customization and persona from session to session, even though the desktop itself is destroyed when the user logs out. Virtual desktops benefit from the latest O/S and application patches automatically applied between user logins, without any disruptive recompose.

Transformational user experience with Blast Extreme

A new VMware controlled protocol for a richer app & desktop experience Protocol optimized for mobile and overall lower client TCO. All existing Horizon remote experience features work with Blast Extreme and updated Horizon clients. Deliver rich multimedia experience in lower bandwidth Rapid client proliferation from strong Horizon Client ecosystem.

Blast Extreme is network-friendly, leverages both TCP and UDP transports, powered by H.264 to get the best performance across more devices, and reduces CPU consumption resulting in less device power consumed for longer battery life.

Modernize application lifecycle management with App Volumes

Transform application management from a slow, cumbersome process into a highly scalable, nimble delivery mechanism that provides faster application delivery and application management while reducing IT costs by up to 70%.

VMware App Volumes is a transformative solution that delivers applications to Horizon virtual desktops. Applications installed on multi-user AppStacks or user-specific writable volumes attach instantly to a desktop at user login. The App Volumes user experience closely resembles that of applications natively installed on the desktop with App Volumes, applications become VM-independent objects that can be moved easily across data centers or to the cloud and shared with thousands of virtual machines.

Smart policies with streamlined access

Improve end user satisfaction by simplifying authentication across all desktop and app services while improving security with smarter, contextual, role-based policies tied to a user, device or location.

- **Policy-Managed Client Features**, which enables IT to use policy to define which specific security-impacting features, are accessible upon login. These include clipboard redirection, USB, printing, and client-drives. All of these can be enforced contextually, based on role, evaluated at logon/logoff, disconnect/reconnect and at pre-determined refresh intervals for consistent application of policy across the entirety of the user experience. For example, a user logging in from a network location consider unsecured, can be denied access to USB and printing. Additionally, PCoIP bandwidth profile settings allow IT to customize the user experience based on user context and location.
- **True SSO** streamlines secure access to a Horizon desktop when users authenticate via VMware Identity Manager. A short lived VMware Horizon virtual certificate is generated, enabling a password-free Windows login, bypassing the usual secondary login prompt users would encounter before getting to their desktop.

4.1.3 VMware Horizon Apps

Horizon Apps (Deliver Virtual Applications to Any Device, Anywhere!!)

The ability to support published application with Horizon has been available with VMware Horizon 7 Enterprise but now we have the standalone options of VMware Horizon Apps Standard & Advance. Horizon provides a platform to deliver an enterprise-class application publishing solution as well as virtual desktops. VMware Horizon leverages Microsoft Remote Desktop Session Host (RDSH) to deliver published applications as well as published desktops running on the Microsoft RDSH Server.

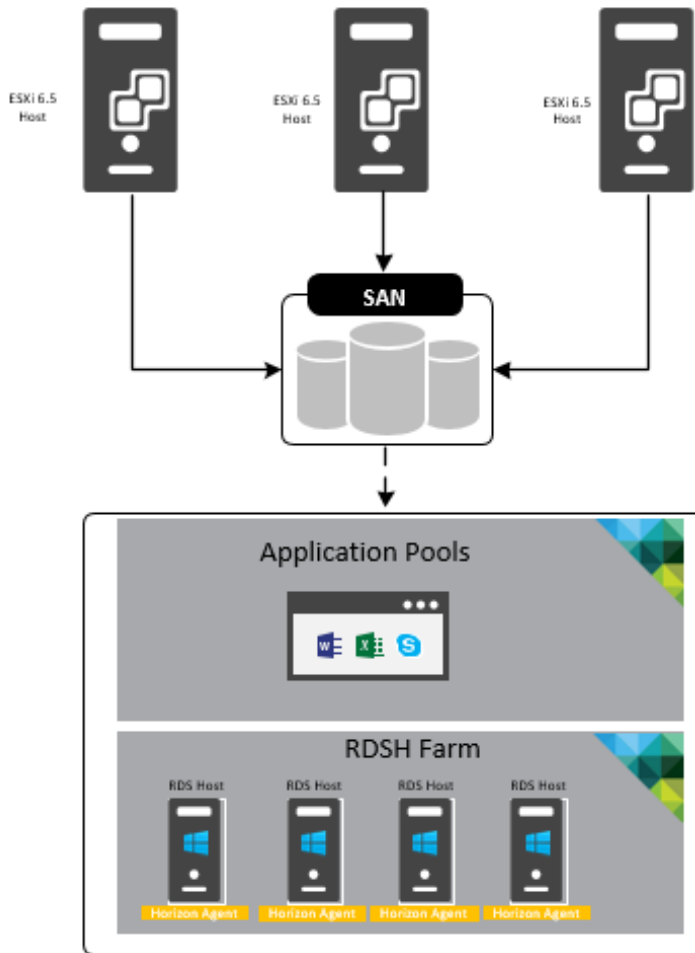


Figure 26 Horizon Apps

VMware Horizon features and components such as the Blast Extreme display protocol, instant-clone provisioning, App Volumes application delivery, and User Environment Manager are heavily integrated into RDSH to provide a seamless user experience and an easy-to-manage, scalable solution.

Next Generation Application and Delivery Platform via Just-in-Time Management Platform (JMP) are available via Horizon Apps Advanced Edition or Horizon 7 Enterprise. MP apps offer simple image management, quick scaling and zero downtime while providing simple and powerful user and group policy controls at the push of a button.

Horizon JMP Apps are composed of:

VMware Instant Clones Technology: RDSH farms can be provisioned rapidly via instant cloned Microsoft RDSH Servers.

VMware App Volumes: real time application delivery via Appstacks mapped to the RDSH Servers. App Volumes allows you to separate the Windows OS image from the application images. Groups of applications can be installed into virtual disks called AppStacks. The appropriate AppStack can then be assigned to the RDSH farm to personalize the applications delivered.

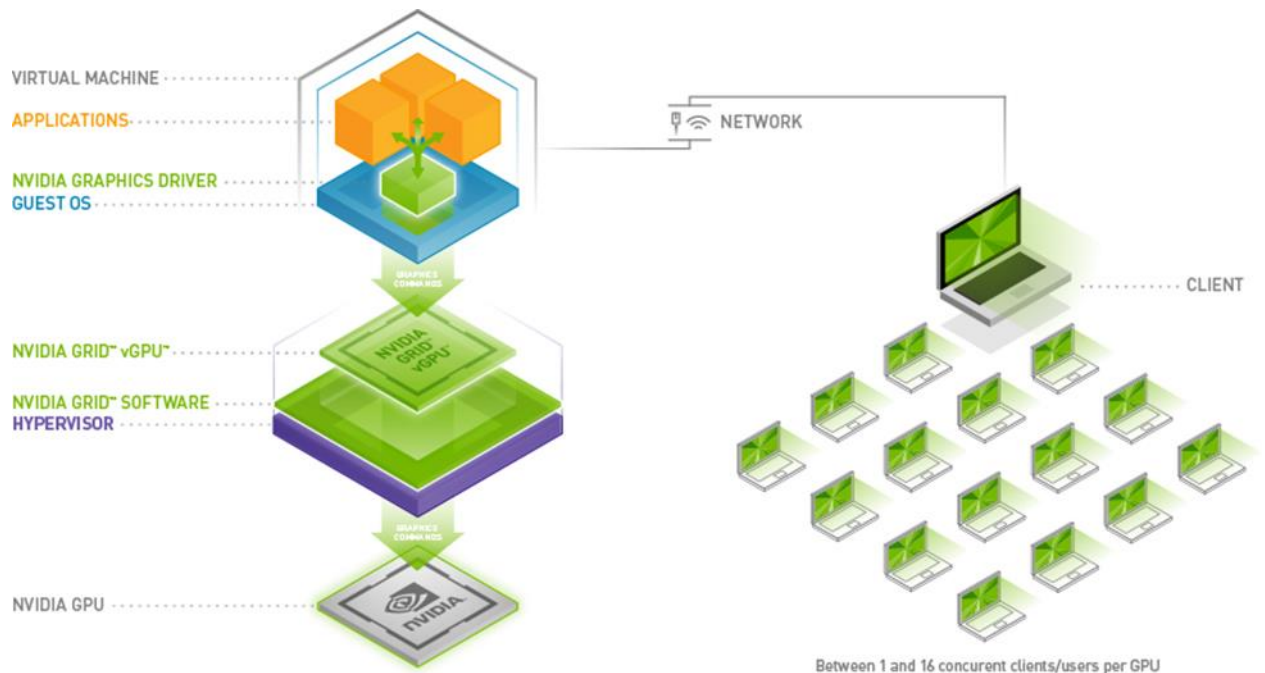
VMware User Environment Manager (UEM): VMware UEM simplifies end-user profile management by offering personalization and dynamic policy configuration to the RDSH Server you can configure fine-grained policies for folder redirection, mapping the user's home drive, configuring location-based printers, and application blocking—all based on user accounts. You can use the Horizon 7 Smart Policies feature to enable or disable client features based on user device, location, and other defined conditions.

4.2 NVIDIA GRID vGPU

NVIDIA GRID™ vGPU™ brings the full benefit of NVIDIA hardware-accelerated graphics to virtualized solutions. This technology provides exceptional graphics performance for virtual desktops equivalent to local PCs when sharing a GPU among multiple users.

GRID vGPU is the industry's most advanced technology for sharing true GPU hardware acceleration between multiple virtual desktops—without compromising the graphics experience. Application features and compatibility are the same as they would be at the user's desk.

With GRID vGPU technology, the graphics commands of each virtual machine are passed directly to the GPU, without translation by the hypervisor. This allows the GPU hardware to be time-sliced to deliver the ultimate in shared virtualized graphics performance.



(Image provided by NVIDIA Corporation. Copyright NVIDIA Corporation)

4.2.1 vGPU Profiles

Virtual Graphics Processing Unit, or GRID™ vGPU™, is technology developed by NVIDIA® that enables hardware sharing of graphics processing for virtual desktops. This solution provides a hybrid shared mode allowing the GPU to be virtualized while the virtual machines run the native NVIDIA video drivers for better performance. Thanks to OpenGL support, VMs have access to more graphics applications. When utilizing vGPU, the graphics commands from virtual machines are passed directly to the GPU without any hypervisor translation. All this is done without sacrificing server performance and so is truly cutting edge.

The combination of Dell servers, NVIDIA GRID vGPU™ technology and NVIDIA GRID™ cards enable high-end graphics users to experience high fidelity graphics quality and performance, for their favorite applications at a reasonable cost.

For more information about NVIDIA GRID vGPU, please visit: [LINK](#)

The number of users per server is determined by the number of GPU cards in the system (max 2), vGPU profiles used for each GPU in a card (2 GPUs per card), and GRID license type. The same profile must be used on a single GPU but profiles can differ across GPUs in a single card.

Table 10 NVIDIA® Tesla® M10 GRID vGPU Profiles:

Card	vGPU Profile	Graphics Memory (Frame Buffer)	Virtual Display Heads	Maximum Resolution	Maximum Graphics-Enabled VMs		
					Per GPU	Per Card	Per Server (3 cards)
Tesla M10	M10-8Q	8GB	4	4096x2160	1	4	12
	M10-4Q	4GB	4	4096x2160	2	8	24
	M10-2Q	2GB	4	4096x2160	4	16	48
	M10-1Q	1GB	2	4096x2160	8	32	96
	M10-0Q	512MB	2	2560x1600	16	64	192
	M10-1B	1GB	4	2560x1600	8	32	96
	M10-0B	512MB	2	2560x1600	16	64	192
	M10-8A	8GB	1	1280x1024	1	4	12
	M10-4A	4GB			2	8	24
	M10-2A	2GB			4	16	48
	M10-1A	1GB			8	32	96

Card	vGPU Profile	Guest VM OS Supported*		GRID License Required
		Win	64bit Linux	
Tesla M10	M10-8Q	●	●	GRID Virtual Workstation
	M10-4Q	●	●	
	M10-2Q	●	●	
	M10-1Q	●	●	
	M10-0Q	●	●	
	M10-1B	●		GRID Virtual PC
	M10-0B	●		
	M10-8A	●		GRID Virtual Application
	M10-4A	●		
	M10-2A	●		
	M10-1A	●		

Supported Guest VM Operating Systems*	
Windows	Linux
Windows 7 (32/64-bit)	RHEL 6.6 & 7
Windows 8.x (32/64-bit)	CentOS 6.6 & 7
Windows 10 (32/64-bit)	Ubuntu 12.04 & 14.04 LTS
Windows Server 2008 R2	
Windows Server 2012 R2	
Windows Server 2016	

***NOTE:** Supported guest operating systems listed as of the time of this writing. Please refer to NVIDIA's documentation for latest supported operating systems.

Table 11 NVIDIA® Tesla® M60 GRID vGPU Profiles:

Card	vGPU Profile	Graphics Memory (Frame Buffer)	Virtual Display Heads	Maximum Resolution	Maximum Graphics-Enabled VMs		
					Per GPU	Per Card	Per Server (3 cards)
Tesla M60	M60-8Q	8GB	4	4096x2160	1	2	6
	M60-4Q	4GB	4	4096x2160	2	4	12
	M60-2Q	2GB	4	4096x2160	4	8	24
	M60-1Q	1GB	2	4096x2160	8	16	48
	M60-0Q	512MB	2	2560x1600	16	32	96
	M60-1B	1GB	4	2560x1600	8	16	48
	M60-0B	512MB	2	2560x1600	16	32	96
	M60-8A	8GB	1	1280x1024	1	2	6
	M60-4A	4GB			2	4	12
	M60-2A	2GB			4	8	32
	M60-1A	1GB			8	16	48

Card	vGPU Profile	Guest VM OS Supported*		GRID License Required
		Win	64bit Linux	
Tesla M60	M60-8Q	●	●	GRID Virtual Workstation
	M60-4Q	●	●	
	M60-2Q	●	●	
	M60-1Q	●	●	
	M60-0Q	●	●	
	M60-1B	●		GRID Virtual PC
	M60-0B	●		
	M60-8A	●		GRID Virtual Application
	M60-4A	●		
	M60-2A	●		
	M60-1A	●		

Supported Guest VM Operating Systems*	
Windows	Linux
Windows 7 (32/64-bit)	RHEL 6.6 & 7
Windows 8.x (32/64-bit)	CentOS 6.6 & 7
Windows 10 (32/64-bit)	Ubuntu 12.04 & 14.04 LTS
Windows Server 2008 R2	
Windows Server 2012 R2	
Windows Server 2016	

***NOTE:** Supported guest operating systems listed as of the time of this writing. Please refer to NVIDIA's documentation for latest supported operating systems.

4.2.1.1 GRID vGPU Licensing and Architecture

NVIDIA GRID vGPU™ is offered as a licensable feature on Tesla GPUs. vGPU can be licensed and entitled using one of the three following software editions.

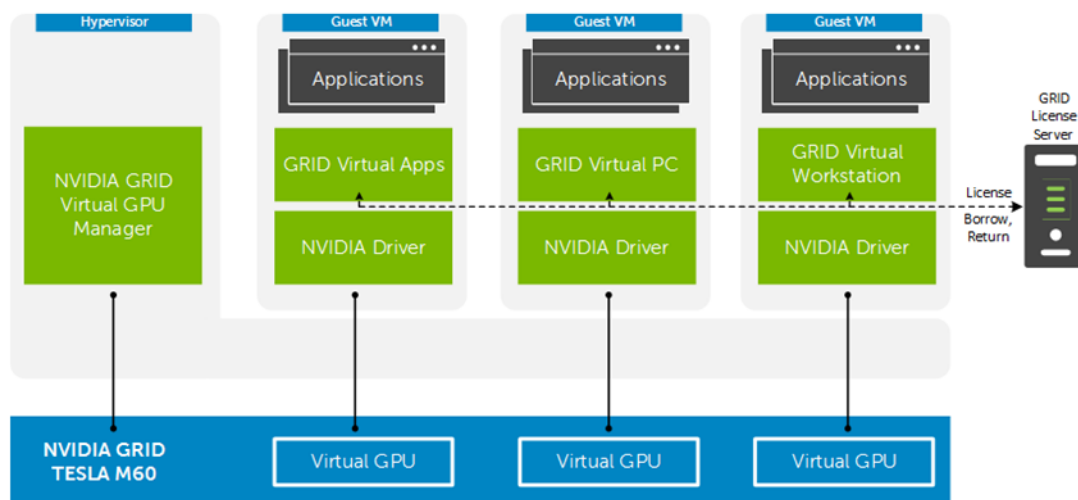


NVIDIA GRID Virtual Applications	NVIDIA GRID Virtual PC	NVIDIA GRID Virtual Workstation
For organizations deploying XenApp or other RDSH solutions. Designed to deliver Windows applications at full performance.	For users who need a virtual desktop, but also need a great user experience leveraging PC applications, browsers, and high-definition video.	For users who need to use professional graphics applications with full performance on any device, anywhere.
Up to 2 displays @ 1280x1024 resolution supporting virtualized Windows applications	Up to 4 displays @ 2560x1600 resolution supporting Windows desktops, and NVIDIA Quadro features	Up to 4 displays @ 4096x2160* resolution supporting Windows or Linux desktops, NVIDIA Quadro, CUDA**, OpenCL** & GPU pass-through

**0Q profiles only support up to 2560x1600 resolution

**CUDA and OpenCL only supported with M10-8Q, M10-8A, M60-8Q, or M60-8A profiles

The GRID vGPU Manager, running on the hypervisor installed via the VIB, controls the vGPUs that can be assigned to guest VMs. A properly configured VM obtains a license from the GRID license server during the boot operation for a specified license level. The NVIDIA graphics driver running on the guest VM provides direct access to the assigned GPU. When the VM is shut down, it releases the license back to the server. . If a vGPU enabled VM is unable to obtain a license, it will run at full capability without the license but users will be warned each time it tries and fails to obtain a license.



(Image provided courtesy of NVIDIA Corporation, Copyright NVIDIA Corporation)

5 Solution Architecture for Horizon 7

5.1 Management Server Infrastructure

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

Role	vCPU	RAM (GB)	NIC	OS + Data vDisk (GB)	Tier 2 Volume (GB)
VMware vCenter	2	16	1	290	-
Horizon Connection Server	2	8	1	60	-
SQL Server	5	8	1	60	210 (VMDK)
File Server	2	4	1	60	2048 (RDM)
Total	11	36	4	470	2358

5.1.1 RDSH VM Configuration

The recommended number of RDSH VMs and their configurations on ESXi are summarized below and take into account proper NUMA balancing assuming the CPU. The amount of RDSH VMs per Server depend on the CPU configuration and for more information on NUMA please refer to the [NUMA Architecture Considerations](#) section.

RDSH VM configuration on ESXi

Role	vCPU	RAM (GB)	NIC	OS vDisk (GB)	Tier 2 Volume (GB)
RDSH VM	8	32	1	80	-

5.1.2 NVIDIA GRID License Server Requirements

When using NVIDIA Tesla M60 cards, graphics enabled VMs must obtain a license from a GRID License server on your network to be entitled for vGPU. To configure, a virtual machine with the following specifications must be added to a management host in addition to the management role VMs.

Role	vCP U	RAM (GB)	NIC	OS + Data vDisk (GB)	Tier 2 Volume (GB)
NVIDIA GRID License Srv	2	4	1	40 + 5	-

GRID License server software can be installed on a system running the following operating systems:

- Windows 7 (x32/x64)
- Windows 8.x (x32/x64)
- Windows 10 x64
- Windows Server 2008 R2
- Windows Server 2012 R2
- Red Hat Enterprise 7.1 x64
- CentOS 7.1 x64

Additional license server requirements:

- A fixed (unchanging) IP address. The IP address may be assigned dynamically via DHCP or statically configured, but must be constant.
- At least one unchanging Ethernet MAC address, to be used as a unique identifier when registering the server and generating licenses in NVIDIA's licensing portal.
- The date/time must be set accurately (all hosts on the same network should be time synchronized).

5.1.3 SQL Databases

The VMware databases will be hosted by a single dedicated SQL 2016 (or higher) Server VM (check DB compatibility at [Link](#)). Use caution during database setup to ensure that SQL data, logs and TempDB are properly separated onto their respective volumes. Create all Databases that will be required for:

- Horizon Connection Server
- vCenter
- Horizon Composer

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 VMware are to be adhered to, to ensure optimal database performance.

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

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

5.1.4 DNS




DNS plays a crucial role in the environment not only as the basis for Active Directory but will be used to control access to the various VMware 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, VMware 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.

5.1.4.1 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 Architecture 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 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 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 compute host servers use 10 locally installed hard drives to house the user desktop VMs. In this model, Tier 1 storage exists as local hard disks or SSDs on the Compute hosts themselves. To achieve the required performance level, RAID 10 is recommended for use across all local disks. A single volume per local Tier 1 Compute host is sufficient to host the provisioned desktop VMs along with their respective write caches.

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 VMDK format.

Volumes	Size (GB)	Storage Array	Purpose	File System
Management	350	Tier 2	vCenter, Horizon Connection Server, File and SQL	VMFS
User Data	2048	Tier 2	File Server/ NAS	RDM/NTFS
User Profiles	20	Tier 2	User profiles	VMFS
SQL DATA	100	Tier 2	SQL	VMFS
SQL LOGS	100	Tier 2	SQL	VMFS
TempDB Data	5	Tier 2	SQL	VMFS
TempDB Logs	5	Tier 2	SQL	VMFS
SQL Witness	1	Tier 2	SQL (optional)	VMFS
Templates/ ISO	200	Tier 2	ISO storage (optional)	VMFS

5.2.4 Storage Networking – XtremIO Fibre 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:

- Set the “connection options” field to 1 for point to point only
- Set the “login retry count” field to 60 attempts
- Set the “port down retry” count field to 60 attempts
- Set the “link down timeout” field to 30 seconds
- LUN queue depth set to 256
- HBA queue depth set to 65535

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



For ESXi hosts, we recommend using the latest [EMC Virtual Storage Integrator \(VSI\) plug-in for VMware vCenter](#). The VSI plug-in interacts with XtremIO to create volumes of the required size, map them to the appropriate Initiator Groups, and create a VMFS datastore ready for use. The VSI plug-in is also used to enforce the following best practice settings:

- Enable VAAI
- Set Queue depth on FC HBA to 256
- Set multi-pathing policy to "round robin" on each of the XtremIO SCSI Disks
- Set I/O path switching parameter to 1
- Set outstanding number of IO request limit to 256
- Set the "SchedQuantum" parameter to 64
- Set the maximum limit on disk I/O size to 4096

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 as shown in the diagram below. Round Robin or Fixed Paths are supported. You can leverage Dell EMC ViPR software to ease storage management in a heterogeneous environment.

5.3 Virtual Networking

5.3.1 Local Tier 1

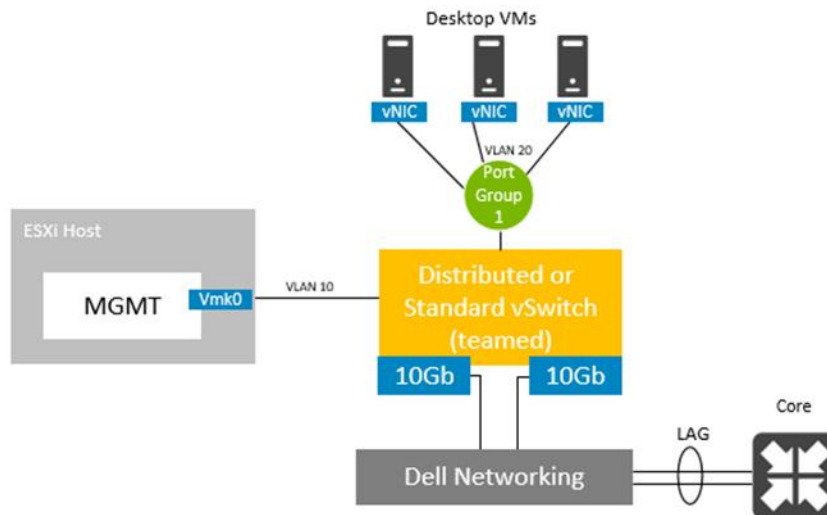
In this model, the servers do not need access to FC storage since they are hosting VDI VMs on local disk. Since this solution model is only recommended for single server deployments, VLAN requirements are reduced:

- 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

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.

vDSwitches should be used as desired for VM traffic especially in larger deployments to ease the management burden across numerous hosts. In the Local Tier 1 rack model the MGMT hosts connect to shared storage and require additional VMK ports. Network share values should be configured equally among the VMkernel port groups that share a physical set of network adapters.

The benefit of using a VMware Distributed Switch (vDS) is that it brings a consistent configuration across all hosts. The vDS is configured at the vCenter level and provides central management and monitoring to all hosts configured on the vDS.



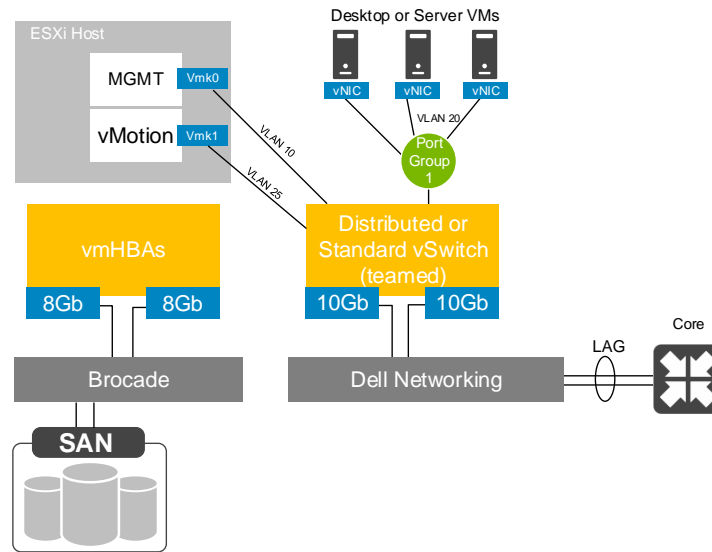
5.3.2 Shared Tier 1

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 (Hyper-V only)
 - 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 (Hyper-V only)
 - 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.

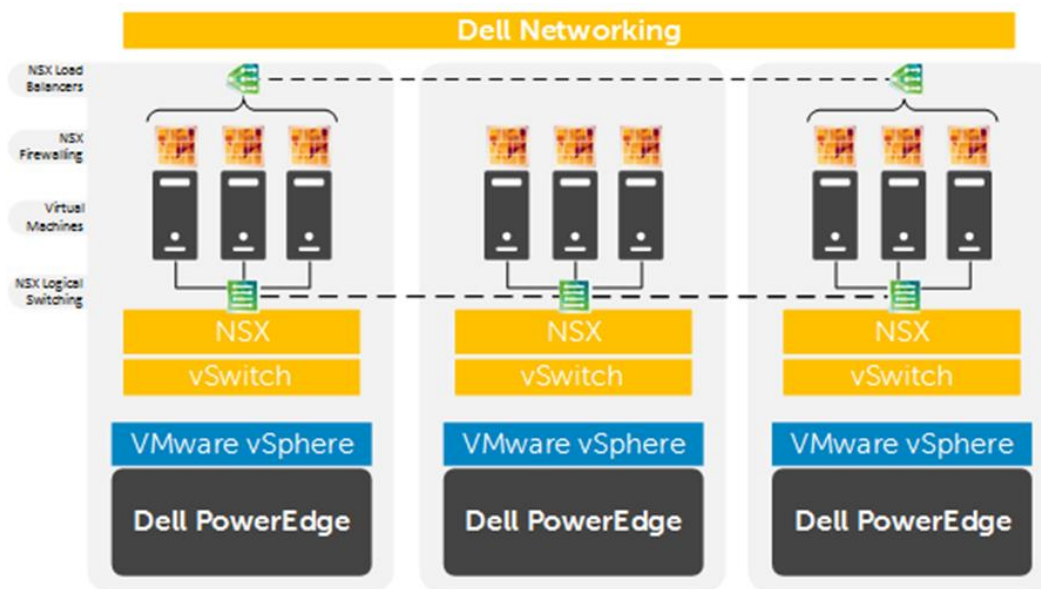
The same basic configuration applies to rack or blade servers although the physical NIC and switching components differ. Network share values should be configured equally among the VMkernel port groups that share a physical set of network adapters.



5.3.3 VMware NSX

Dell and VMware's Software Defined Datacenter (SDDC) architecture goes beyond simply virtualizing servers and storage but also extends into the network. VMware NSX is a network virtualization platform deployable on any IP network that is integrated with vSphere Virtual Distributed Switching and provides the same features and benefits to networking as the ESXi hypervisor does to virtual machines. NSX provides a complete set of logical networking elements and services—including logical switching, routing, firewalling, load balancing, VPN, quality of service (QoS), and monitoring. These services are provisioned in virtual networks through any cloud management platform leveraging the NSX APIs. Through Dell's open networking, companies are best able to take advantage of this disaggregation of a virtual network overlay and an open physical underlay. Building a zero-trust security model is easy with NSX as each virtualized workload can be protected with a stateful firewall engine providing extreme policy granularity. Any VM in the datacenter can be rigorously secured or isolated if compromised, especially useful for virtual desktops to prevent malicious code from attacking and spreading through the network.

VMware NSX is implemented via a layered architecture consisting of data, control and management planes. The NSX vSwitch exists within and requires the vSphere Distributed Switch to abstract the physical network while proving access-level switching in the hypervisor. NSX enables the use of virtual load balancers, firewalls, logical switches and routers that can be implemented and scaled seamlessly to suit any deployed architecture. VMware NSX compliments Dell Networking components deployed ToR, leaf/spine or at the core.



Key Features of Dell Open Networking and VMware NSX

Power of Choice	Choose from best-of-breed open networking platforms, operating systems and applications.
Accelerated Innovation	Take advantage of open networking with open source standards-based tools and expertise to help accelerate innovation.
Open Networking Platform	All Dell Networking data center switches support the Open Network Install Environment (ONIE), allowing customers to choose between multiple operating systems and meet their unique needs.
Hardware VTEP Gateway	Layer 2 gateway through VXLAN Tunnel End Points (VTEP) bridges virtual and physical infrastructures.
Virtual Switching	VXLAN based network overlays enable logical layer 2 overlay extensions across a routed (L3) fabric within and across data center boundaries.
Virtual Routing	Dynamic routing between virtual networks performed in a distributed manner in the hypervisor kernel, and scale-out routing with active-active failover with physical routers.
Distributed Firewalling	Distributed stateful firewalling, embedded in the hypervisor kernel for up to 20 Gbps of firewall capacity per hypervisor host.
Load Balancing	L4-L7 load balancer with SSL offload and pass through, server health checks, and App Rules for programmability and traffic manipulation.

For more information on VMware NSX and integrated offers from Dell Networking please see the Dell Networking [Solution Brief](#) and the [Reference architecture](#).

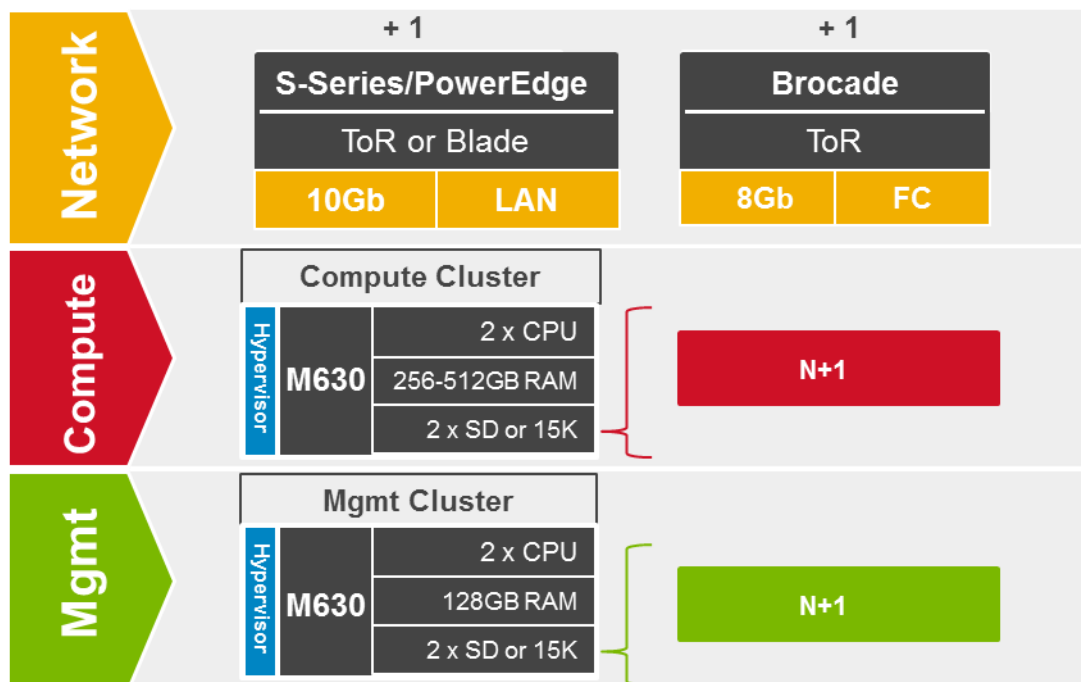
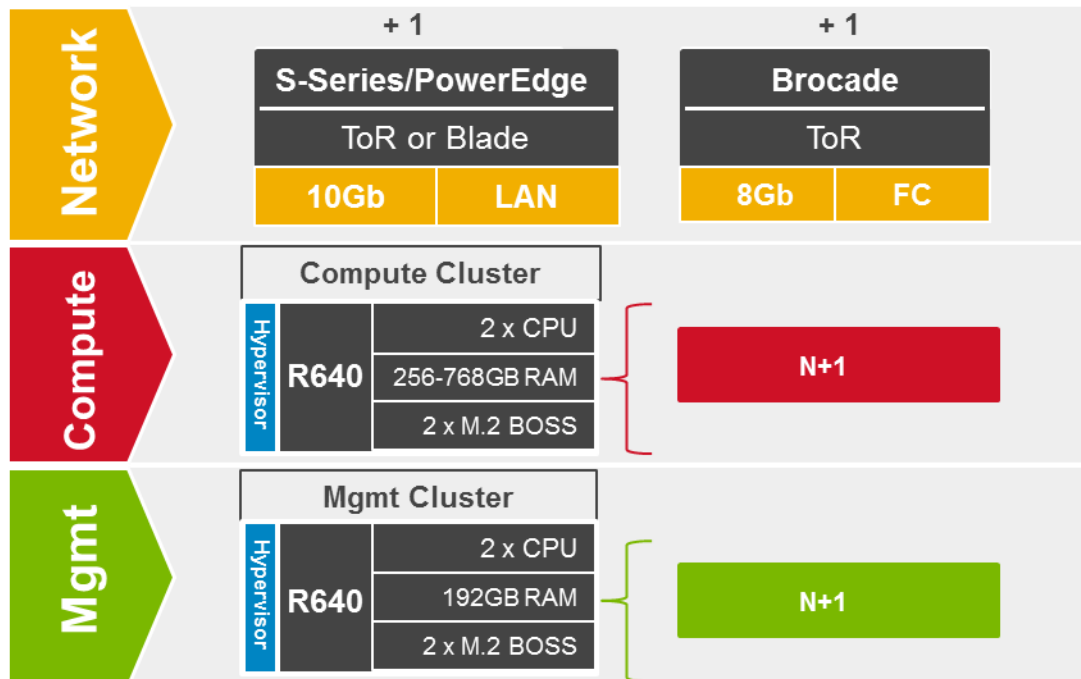
5.4 Scaling Guidance

- 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
Virtual Desktop Host/Compute Servers	VMs per physical host	Additional hosts and clusters added as necessary	Additional RAM or CPU compute power
Horizon Composer	Desktops per instance	Additional physical servers added to the Management cluster to deal with additional management VMs.	Additional RAM or CPU compute power
Horizon Connection Servers	Desktops per instance	Additional physical servers added to the Management cluster to deal with additional management VMs.	Additional HCS Management VMs.
VMware vCenter	VMs per physical host and/or ESX hosts per vCenter instance	Deploy additional servers and use linked mode to optimize management	Additional vCenter Management VMs.
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 the optional NAS device to provide high availability.	Additional RAM and CPU for the management nodes

5.5 Solution High Availability

High availability (HA) is offered to protect each architecture solution layer, individually if desired. Following the N+1 model, additional ToR switches are added to the Network layer and stacked to provide redundancy as required, additional compute and management hosts are added to their respective layers, vSphere clustering is introduced in both the management and compute layers, SQL is configured with AlwaysOn .



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 ESXi hosts added in the compute or mgmt layers to provide N+1 protection.
- A number of enhancements occur at the Management tier, the first of which is the addition of another host. The Management hosts will then be configured in an HA cluster. All applicable Horizon server roles can then be duplicated on the new host where connections to each will be load balanced via the addition of a F5 Load Balancer. SQL will also receive greater protection through the addition and configuration of a SQL mirror with a witness.

5.5.1 vSphere HA (Shared Tier 1)

Both compute and management hosts are identically configured, within their respective tiers and leverage shared storage so can make full use of vSphere HA. The Compute hosts can be configured in an HA cluster following the boundaries of vCenter with respect to limits imposed by VMware (2000 VMs per vCenter). This will result in multiple HA clusters managed by multiple vCenter servers.

A single HA cluster will be sufficient to support the Management layer up to 10K users. An additional host can be used as a hot standby or to thin the load across all hosts in the cluster.

5.5.2 Management Server High Availability

The applicable core Horizon roles will be load balanced via DNS by default. In environments requiring HA, F5 can be introduced to manage load-balancing efforts. Horizon, HCS and vCenter configurations (optionally vCenter Update Manager) are stored in SQL, which will be protected via the SQL mirror.

If the customer desires, some Role VMs can be optionally protected further via the form of a cold stand-by VM residing on an opposing management host. A vSphere scheduled task can be used, for example, to clone the VM to keep the stand-by VM current. Note – In the HA option, there is no file server VM, its duties have been replaced by introducing a NAS head.

The following will protect each of the critical infrastructure components in the solution:

- The Management hosts will be configured in a vSphere cluster.
- SQL Server mirroring is configured with a witness to further protect SQL.

5.5.3 Horizon CS High Availability


The HCS role as a VM and running in a VMware HA Cluster, the HCS server can be guarded against a physical server failure.

For further protection in an HA configuration, deploy multiple replicated Horizon Connection Server instances in a group to support load balancing and HA. Replicated instances must exist on within a LAN connection environment it is not recommended VMware best practice to create a group across a WAN or similar connection.

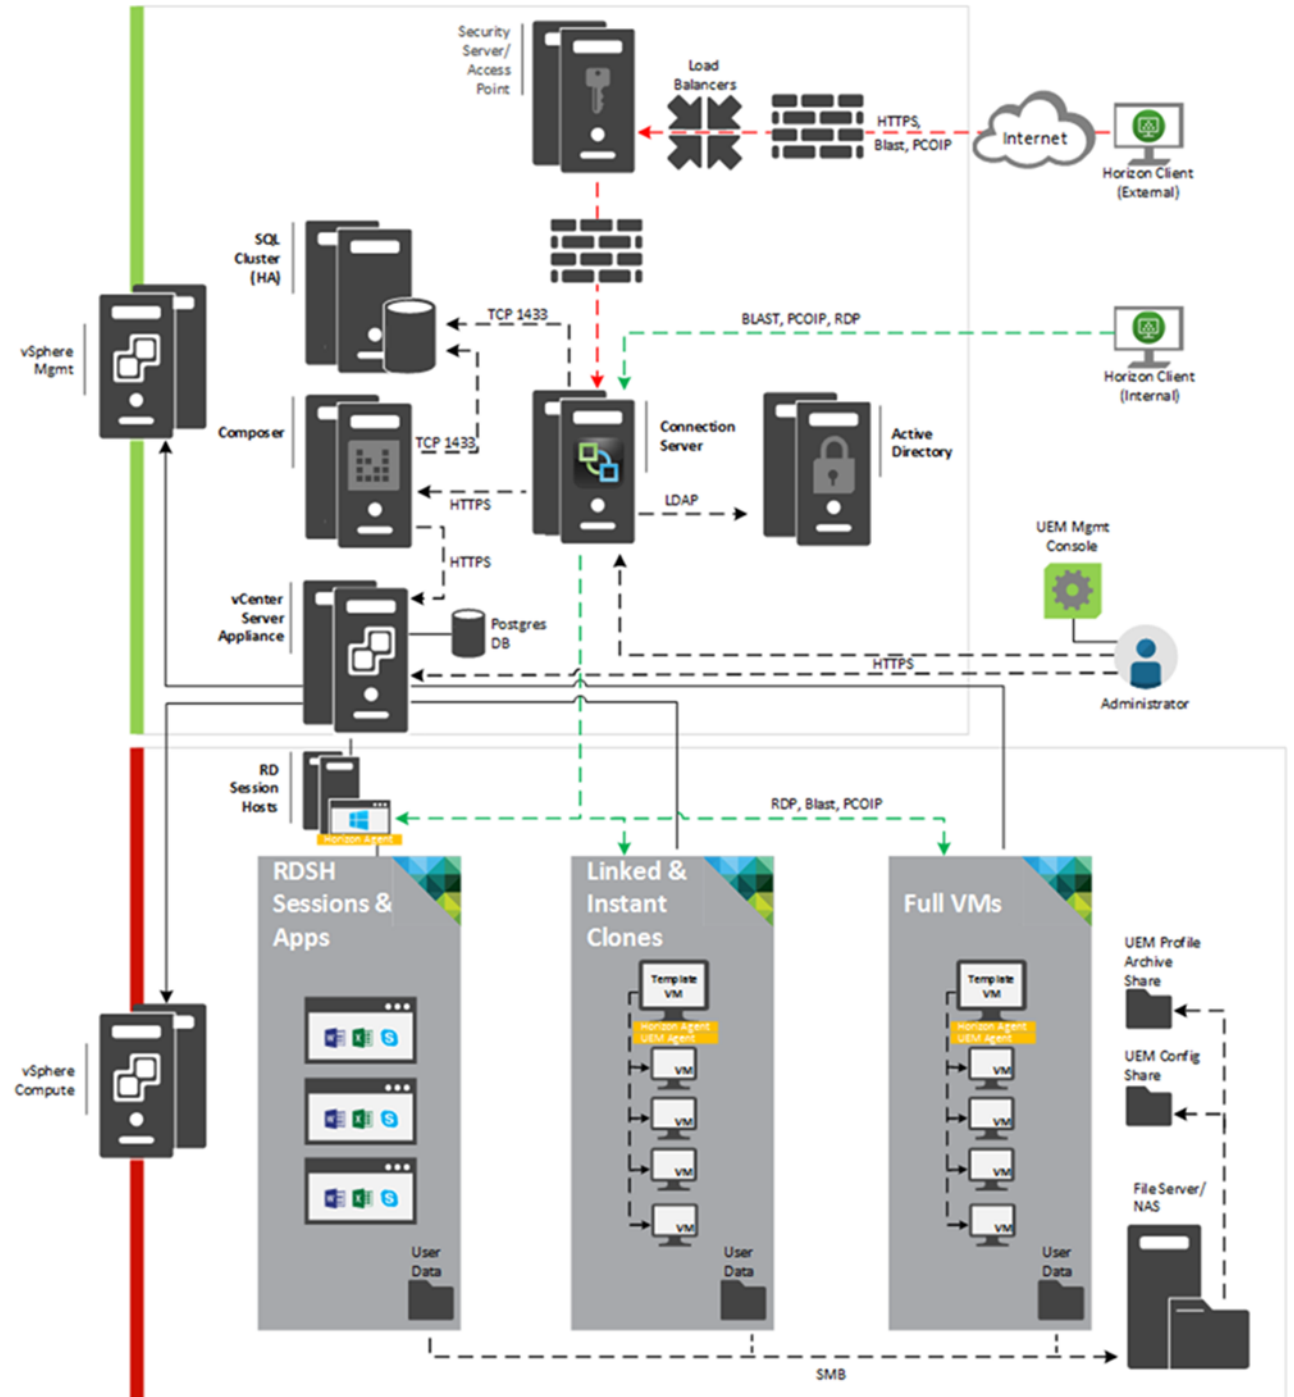
5.5.4 SQL Server High Availability

HA for SQL is provided via Always On 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 VMware Horizon communication flow



6 Solution Performance and Testing

At the time of publication, here are the available density recommendations per compute server. Please refer to [Section 3.2](#) for hardware specifications

Table 12 User density summary

Host Config	Hypervisor	Broker & Provisioning	Workload	Template	User Density
B5	ESXi 6.5	Horizon 7 & Instant Clone	Knowledge Worker	Windows 10 x64 & Office 2016	155

All tests above were performed with [LoginVSI](#) version 4.1, on VMware Horizon 7. For 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 four runs. First, a pilot run to validate that the infrastructure is functioning and valid data can be captured, and then, three 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 hour. 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. 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
Graphics LVSI Power + ProLibrary	Graphics - Login VSI Power worker with ProLibrary
Graphics LVSI Custom	Graphics – LVSI Custom

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.

Graphics - Login VSI Power Worker with ProLibrary workload

For lower performance graphics testing where lower amounts of graphics memory are allocated to each VM, the Power worker + Pro Library workload is used. The Login VSI Pro Library is an add-on for the Power worker workload which contains extra content and data files. The extra videos and web content of the Pro Library utilizes the GPU capabilities without overwhelming the lower frame buffer assigned to the desktops. This type of workload is typically used with high density vGPU and sVGA or other shared graphics configurations.

Graphics – LVSI Custom workload

This is a custom Login VSI workload specifically for higher performance, intensive graphics testing. For this workload, SPECwpc benchmark application is installed to the client VMs. During testing, a script is started that launches SPECwpc, which executes the Maya and sw-03 modules for high performance tests and module sw-03 only for high density tests. The usual activities such as Office application execution are not performed with this workload. This type of workload is typically used for lower density/high performance pass-through, vGPU, and other dedicated, multi-user GPU configurations.

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 GPU resources

ESXi hosts

For gathering of GPU related resource usage, a script is executed on the ESXi host before starting the test run and stopped when the test is completed. The script contains NVIDIA System Management Interface commands to query each GPU and log GPU utilization and GPU memory utilization into a .csv file.

ESXi 6.5 and above includes the collection of this data in the vSphere Client/Monitor section. GPU processor utilization, GPU temperature, and GPU memory utilization can be collected the same was as host CPU, host memory, host Network, etc.

6.1.2.2 VMware vCenter

VMware vCenter is used for VMware vSphere-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 (AHV & ESXi hypervisors)*	100%
Physical Host CPU Utilization (Hyper-V)	85%
Physical Host Memory Utilization	85%
Network Throughput	85%
Storage IO Latency	20ms

*Turbo mode is enabled; therefore, the CPU threshold is increased, as it will be reported as over 100% utilization when running with turbo.

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 R730
Hypervisor(s)	ESXi 6.5

Broker technology	Horizon 7
Broker database	Microsoft SQL 2016
Management VM OS	Windows Server 2012 R2 (Connection Server & Database)
Virtual desktop OS	Windows 10 Enterprise
Office application suite	Office Professional 2016
Login VSI test suite	Version 4.1

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	ESXi Memory Configured	ESXi Memory Reservation	Screen Resolution	Operating System
Task Worker	1	2GB	1GB	1280 X 720	Windows 10 Enterprise 64-bit
Knowledge Worker	2	3GB	1.5GB	1920 X 1080	Windows 10 Enterprise 64-bit
Power Worker	2	4GB	2GB	1920 X 1080	Windows 10 Enterprise 64-bit
Graphics LVSI Power + ProLibrary	2	4 GB	4GB	1920 X 1080	Windows 10 Enterprise 64-bit
Graphics LVSI Custom – Density	2	4 GB	4GB	1920 X 1080	Windows 10 Enterprise 64-bit
Graphics LVSI Custom - Performance	4	8GB	8GB	1920 X 1080	Windows 10 Enterprise 64-bit

6.2.2 Platform Configuration

Please refer to [Section 3.2](#) for hardware specifications

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.

Platform Config	Hypervisor	Broker & Provisioning	Login VSI Workload	Density Per Host	Avg CPU	Avg Mem Consumed	Avg Mem Active	Avg IOPS / User
B5	ESXi 6.5	Horizon 7 & Instant Clones	Knowledge Worker	155	95%	330 GB	170 GB	7.4

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. On the latest Intel series processors, the ESXi host CPU metrics will exceed the rated 100% for the host if Turbo Boost is enabled (by default). An additional 35% of CPU is available from the Turbo Boost feature but this additional CPU headroom is not reflected in the VMware vSphere metrics where the performance data is gathered. Therefore, CPU usage for ESXi hosts is adjusted and a line indicating the potential performance headroom provided by Turbo boost is included in each CPU graph.

Avg Consumed Memory: Consumed memory is the amount of host physical memory consumed by a virtual machine, host, or cluster. For clusters, this is the average consumed memory across all compute hosts over the steady state period.

Avg Mem Active: For ESXi hosts, active memory is the amount of memory that is actively used, as estimated by VMkernel based on recently touched memory pages. For clusters, this is the average amount of guest “physical” memory actively used 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.

6.3.1 B5 Compute


Please refer to [Section 3.2](#) for hardware specifications

6.3.1.1 ESXi 6.5, Horizon 7, 155 per Host, Knowledge Worker Workload

CPU Usage


The single compute host was populated with 155 virtual machines. With all user virtual machines powered on during test, the CPU usage was approximately 7%.

The below graph shows the performance data for 155 user sessions per host. The CPU reaches 100% usage during the logon phase but settles to about 95% usage during the steady state of the test when 155 users are logged on to the host.




Memory Utilization

Memory usage is monitored on ESXi host, memory usage monitored are consumed, active, balloon and swap used. Sustained levels of swap and ballooning usage would indicate host memory reached saturation point and the VM performance may start to deteriorate. All tests were carried out on host with 352GB physical memory installed. No ballooning or swapping was measured during the test even when the Instant clones were being Consumed memory in each test reached the maximum of 342GB but active memory only reached 183GB during the Login VSi test just as the steady state phase starts. Active memory briefly reached 443GB after logoff during the recreation of the Instant clones above 300GB during the task worker workload.




Network bandwidth is not an issue on this test run with a steady state peak of approximately 50 MBps. Network usage is affected by the Login VSI content and profile and home folder redirection Network traffic on the Fibre Channel network connecting the XtremIO storage is not measured from VSphere. The datastore bandwidth utilization is shown in the XtremIO graphs below.




The IOPS graphs and IOPS numbers are taken from the XtremIO Administration console and the graphs clearly display the initial logon of the desktops, the steady state and logoff phases, and finally the

recreation of the desktops after testing is complete. The graph shows the Disk IOPS, bandwidth and latency for the XtremIO cluster.

The cluster reached a maximum of 6,000 Disk IOPS during the instant clone recreation period after testing and 5,700 IOPS during the steady state. Latency remained at zero (0) milliseconds throughout the test, when measured from the XIO and also from vCenter. The XtremIO storage was not expected to be stressed in any way by the test.



The overall performance graph from XtremIO is shown below



The Login VSI VSIMAX score was not reached during the test with a good baseline of 925. The maximum VSI response during the test was about 1450.

