

# REFERENCE ARCHITECTURE

## Dell EMC Ready Bundle for VDI

Design, configuration and implementation of a Citrix XenDesktop environment with Dell EMC PowerEdge Servers and XtremIO Storage

### [Abstract](#)

A Reference Architecture for integrating Dell EMC PowerEdge servers and Citrix XenDesktop & XenApp brokering software on VMware ESXi and Microsoft Hyper-V 2016 hypervisors to create virtual application and virtual desktop environments on 14th generation Dell EMC PowerEdge Servers.

February 2018

## Revisions

Date	Description
April 2014	Initial release
March 2016	Document overhaul, new Broadwell CPUs, networking, servers
September 2016	Minor updates to Endpoints section
December 2016	Updated Citrix software section with latest XenDesktop 7.12 info and other minor updates
January 2017	Replaced Compellent and EqualLogic storage recommendations with Dell EMC storage
April 2017	Added test results to document
September 2017	Replaced server hardware with latest generation (14G) Dell EMC PowerEdge servers, Skylake CPUs and storage hardware with Dell EMC XtremIO X2
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.

© 2014 – 2018 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

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

# Table of contents

Revisions.....	2
Acknowledgements.....	2
Executive summary.....	6
<b>1 Introduction.....</b>	<b>7</b>
1.1 Objective.....	7
1.2 What's new .....	7
<b>2 Solution architecture overview .....</b>	<b>8</b>
2.1 Introduction .....	8
2.2 Physical architecture overview .....	8
2.3 Solution layers .....	9
2.3.1 Networking.....	10
2.3.2 Compute .....	10
2.3.3 Management.....	10
2.3.4 Storage .....	11
2.4 Local tier 1 .....	11
2.5 Shared tier 1 for rack servers .....	11
2.5.1 Shared tier 1 – network architecture.....	14
2.5.2 Shared tier 1 – rack cabling (HA) .....	15
2.5.3 Shared tier 1 – storage scaling guidance .....	15
2.6 Shared tier 1 for blade servers .....	16
2.6.1 Shared tier 1 – network architecture.....	17
2.6.2 Shared tier 1 – cabling.....	18
2.6.3 Shared tier 1 – scaling guidance .....	19
<b>3 Hardware components .....</b>	<b>20</b>
3.1 Network.....	20
3.1.1 Dell Networking S3048 (1Gb ToR switch).....	20
3.1.2 Dell Networking S4048 (10Gb ToR switch).....	21
3.1.3 Brocade 6510 (FC ToR switch) .....	21
3.1.4 Brocade M5424 (FC blade interconnect) .....	23
3.1.5 PowerEdge M I/O aggregator (10Gb blade interconnect).....	24
3.2 Servers .....	24
3.2.1 PowerEdge R640.....	25
3.2.2 PowerEdge R740.....	27
3.2.3 PowerEdge M630 .....	28
3.3 Storage .....	29

3.3.1 XtremIO X2 X-Brick – Combined Tier 1 and Tier 2 .....	29
3.4 GPUs .....	33
3.4.1 NVIDIA Tesla GPUs .....	33
3.5 Dell Wyse Endpoints .....	35
3.5.1 Wyse 3040 Thin Client (ThinOS, ThinLinux) .....	35
3.5.2 Wyse 5010 Thin Client (ThinOS).....	35
3.5.3 Wyse 5060 Thin Client (ThinOS, ThinLinux, WES7P, WIE10) .....	35
3.5.4 Wyse 7020 Thin Client (WES 7/7P, WIE10, ThinLinux).....	36
3.5.5 Wyse 7040 Thin Client (WES7P, WIE10) .....	36
3.5.6 Latitude 3480 and 5280 Mobile Thin Clients (Win 10 IoT) .....	36
3.5.7 Enhanced Security .....	37
<b>4 Software components.....</b>	<b>38</b>
4.1 Citrix.....	38
4.1.1 Citrix XenDesktop .....	38
4.1.2 Machine Creation Services (MCS) .....	40
4.1.3 Personal vDisk.....	41
4.1.4 AppDisks.....	41
4.1.5 HDX 3D Pro .....	41
4.1.6 Citrix Profile Manager .....	42
4.1.7 Citrix XenApp.....	43
4.1.8 Local Host Cache .....	46
4.1.9 Citrix NetScaler.....	48
4.2 Hypervisor platforms.....	50
4.2.1 VMware vSphere 6.x .....	50
4.2.2 Microsoft Windows Server Hyper-V.....	51
4.3 NVIDIA GRID vGPU .....	52
4.3.1 vGPU profiles.....	53
<b>5 Solution architecture for XenDesktop.....</b>	<b>60</b>
5.1 Management role configuration .....	60
5.1.1 vSphere .....	60
5.1.3 Hyper-V.....	61
5.1.4 XenApp VM configuration .....	61
5.1.5 SQL databases.....	62
5.1.6 DNS .....	62
5.2 Storage configuration overview .....	63
5.2.1 Local Tier 1 storage.....	63

5.2.2	Shared Tier 1 storage .....	63
5.2.3	Shared Tier 2 storage .....	63
5.2.4	Storage networking – XtremIO Fiber Channel (FC) .....	65
5.3	Virtual networking .....	66
5.3.1	Local Tier 1 .....	66
5.3.2	Shared Tier 1 – FC .....	67
5.4	Scaling guidance .....	69
5.5	Solution high availability .....	70
5.5.1	SQL Server high availability .....	72
5.6	Citrix XenDesktop communication flow .....	73
6	<b>Solution performance and testing</b> .....	<b>74</b>
6.1	Summary .....	74
6.2	Test and performance analysis methodology .....	74
6.2.1	Testing process .....	74
6.2.2	Resource monitoring .....	76
6.2.3	Resource utilization .....	77
6.3	Test configuration details .....	78
6.3.1	Compute VM Configurations .....	78
6.3.2	Platform Configurations .....	79
6.4	Test results and analysis .....	80
6.4.1	R740 Compute .....	81
7	<b>Related resources</b> .....	<b>89</b>

## Executive summary

This document provides the reference architecture for integrating Dell EMC PowerEdge servers, Dell EMC storage, and Citrix XenDesktop and XenApp software suite to create virtual application and virtual desktop environments. The available server choices include the PowerEdge R640, R740, and M630 servers. Shared storage for the solution is the Dell EMC XtremIO X2 X-Brick.

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

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

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

# 1 Introduction

This document addresses the architecture design, configuration and implementation considerations for the key components required to deliver virtual desktops or shared sessions via Citrix XenDesktop and XenApp on VMware vSphere 6.x or Microsoft Windows Server 2016 Hyper-V hypervisors. Proposed design choices include rack or blade servers, local disks or shared storage. Guidance contained within this document follows a building block methodology enabling the combination of several different components each with their own scaling capabilities.

## 1.1 Objective

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

- Define the detailed technical design for the solution.
- Define the hardware requirements to support the design.
- Define the constraints which are relevant to the design.
- Define relevant risks, issues, assumptions and concessions – referencing existing ones where possible.
- Provide a breakdown of the design into key elements such that the reader receives an incremental or modular explanation of the design.
- Provide solution scaling and component selection guidance.
- 

## 1.2 What's new

- Updated rack server hardware choices to the latest 14th generation Dell EMC PowerEdge servers with Skylake processors.
- Updated shared storage hardware to 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)
- Full Clone Virtual Desktops (Persistent)
- Shared XenApp Sessions (RDSH)

Additionally, our solution includes options for users who require:

- Graphics Acceleration
- Unified Communications

### 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.

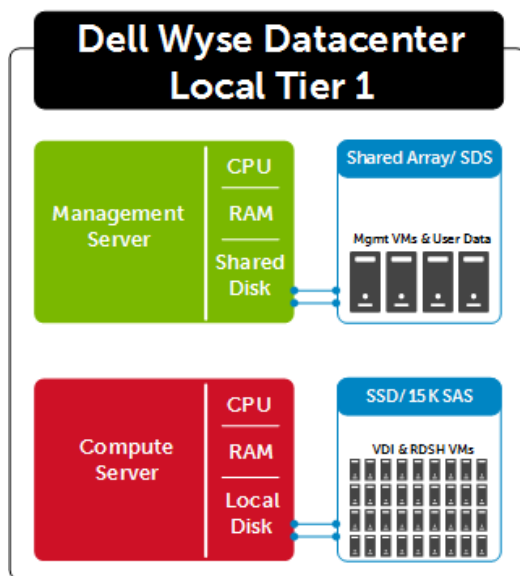


Figure 1 Local Tier 1

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.



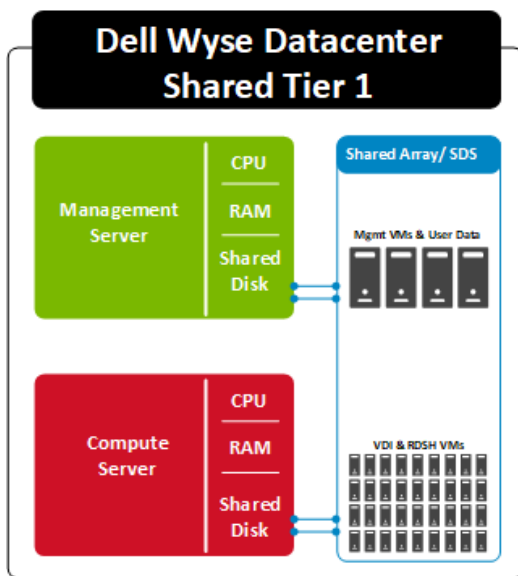


Figure 2 Shared Tier 1

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

**NOTE:** If using PVS, additional Management servers may be necessary.

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 (via file servers). Additional X-Bricks are added to the solution when scaling beyond 3,500 users.




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. SQL is required for System Center Virtual Machine Manager (SCVMM) when Hyper-V is the hypervisor.

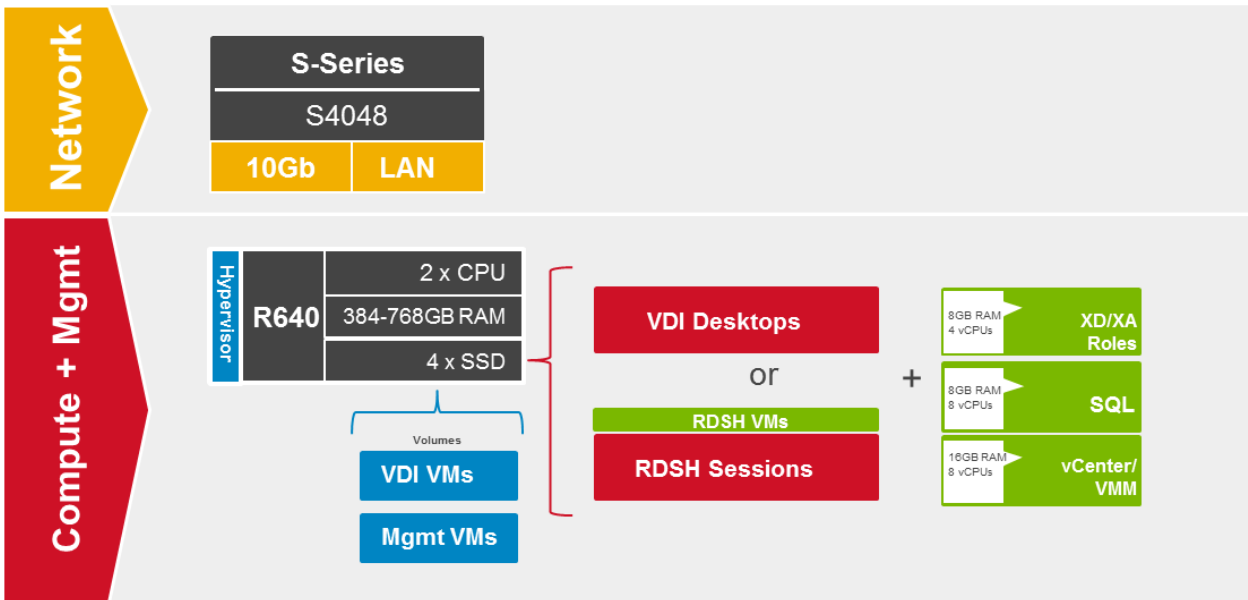


Figure 7 Local Tier 1 storage

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

2.5 Shared tier 1 for rack servers

This solution model provides a high-performance scalable rack-based configuration that incorporates shared T1 and T2 storage for execution of VDI sessions and management VMs. Since all VMs reside on the shared storage array, the servers are diskless and use a BOSS device (ESXi or Hyper-V) or SD cards (ESXi only) for the hypervisor/operating system. User data can either be stored on the same array (via a file server) as the

VMs or on another storage location. The figure below depicts the shared storage configuration with optional file server VM for user data.

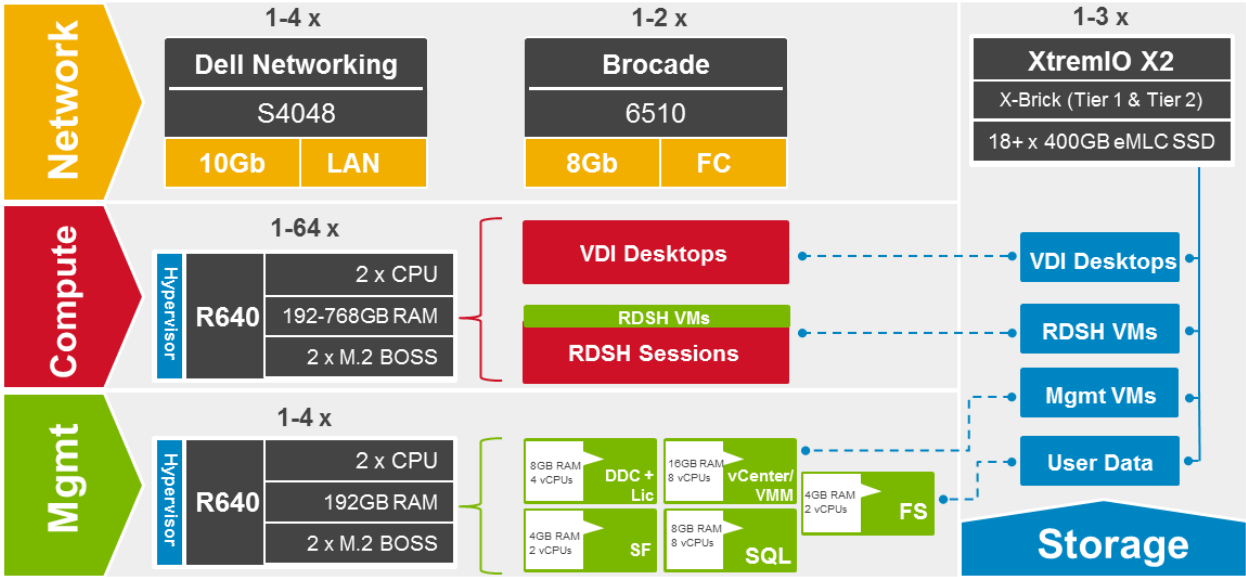


Figure 8 Shared Tier 1 storage - rack

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.

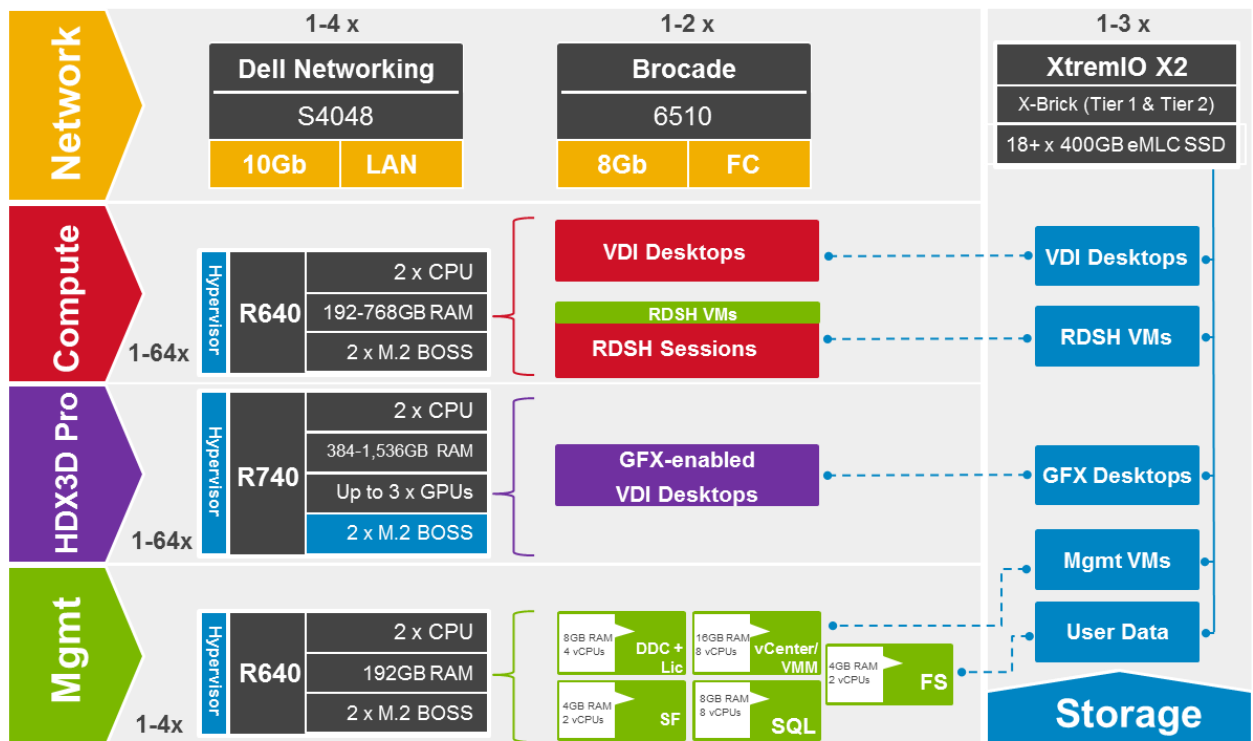


Figure 9 High performance graphics added

**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.

**NOTE:** If using PVS, additional Management servers may be necessary depending on requirements.

## 2.5.1 Shared tier 1 – network architecture

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




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 for ESXi

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

Table 2 Shared tier 1 storage scaling guidance for Hyper-V

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 – 18	Up to 720	Up to 14TB	1 x S4048	1 x 6510
3501 – 7000	2 x X-Brick	18 – 35	Up to 1,400	Up to 27TB	1-2 x S4048	1-2 x 6510
7001 – 10000	3 x X-Brick	35 – 50	Up to 2,000	Up to 38TB	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 drives (Hyper-V) or SD cards (ESXi only) for the operating system. User data can either be stored on the same array (via a file server) as the VMs or on another storage location. The figure below depicts the shared storage configuration with optional file server VM for user data.

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

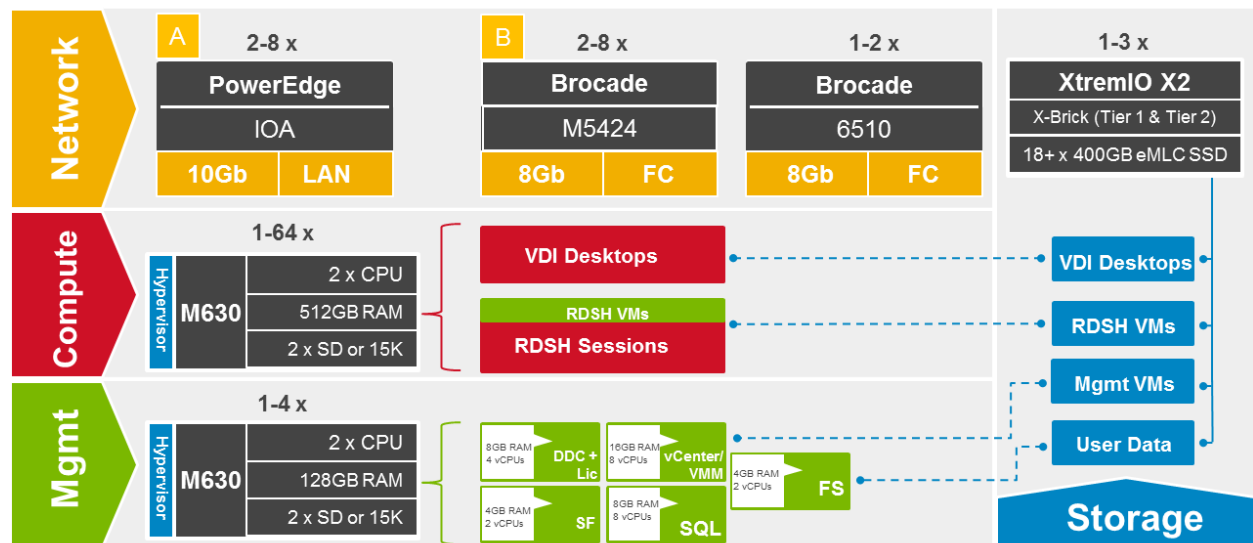


Figure 12 Shared Tier 1 – blade servers

**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.

**NOTE:** If using PVS, additional Management servers may be necessary depending on requirements.



## 2.6.1 Shared tier 1 – network architecture

In the Shared Tier 1 for blades 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.




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 3 Shared tier 1 scaling guidance ESXi

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

Table 4 Shared tier 1 scaling guidance ESXi

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 – 18	Up to 720	Up to 9TB	4 x IOA + 4 x M5424	1 x 6510
3501 – 7000	2 x X-Brick	18 – 35	Up to 1,400	Up to 18TB	6 x IOA + 6 x M5424	1 x 6510
7001 – 10000	3 x X-Brick	35 – 50	Up to 2,000	Up to 26TB	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 5 Dell Networking S3048 features

Model	Features	Options	Uses
Dell Networking S3048-ON	<ul style="list-style-type: none"><li>• 48 x 1000BaseT</li><li>• 4 x 10Gb SFP+</li><li>• Non-blocking, line-rate performance</li><li>• 260Gbps full-duplex bandwidth</li><li>• 131 Mbps forwarding rate</li></ul>	• Redundant hot-swap PSUs & fans	1Gb connectivity
		• VRF-lite, Routed VLT, VLT Proxy Gateway	
		• User port stacking (up to 6 switches)	
		• 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 6 Dell Networking S4048 features

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




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 7 Brocade 6510 features

Model	Features	Options	Uses
Brocade 6510	<ul style="list-style-type: none"><li>• 48 x 2/4/8/16Gb Fiber Channel</li><li>• Additional (optional) FlexIO module</li><li>• Up to 24 total ports (internal + external)</li></ul>	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 8 Brocade M5424 features

Model	Features	Options	Uses
Brocade M5424	<ul style="list-style-type: none"><li>• 16 x internal Fabric ports</li><li>• Up to 8 2/4/8Gb auto-sensing uplinks</li></ul>	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 9 PowerEdge M I/O aggregator features

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

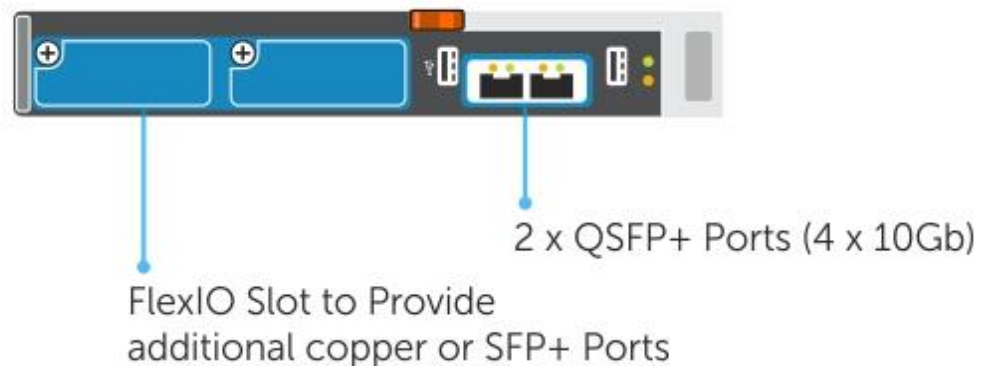


Figure 19 Table 7 PowerEdge M I/O aggregator

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

## 3.2 Servers

This reference architecture is built on the latest 14<sup>th</sup> generation Dell EMC PowerEdge R640 and R740 servers for rack installations and the M630 servers for blades. Optimized for VDI, the Dell EMC PowerEdge portfolio has been designed and arranged in three top-level overarching configurations which apply to the available physical platforms showcased below.

- **A3** configuration is perfect for small scale, POC or low density cost-conscious environments.
- **B5** configuration is geared toward larger scale general purpose workloads, balancing performance and cost-effectiveness.
- **C7** configuration is the premium configuration offering an abundance of high performance to maximize user density.




A3 (low, entry level)	B5 (mid, balance cost/ perf)	C7 (high, max density, max perf)
<ul style="list-style-type: none"> <li><b>CPU:</b> 2 x 10-Core (4114)</li> <li><b>RAM:</b> 192GB (12 x 16GB @ 2133MHz)</li> <li><b>Disk (LT1 only):</b> 4 x 960GB SSD Or 10 x 600GB 15K SAS</li> </ul>	<ul style="list-style-type: none"> <li><b>CPU:</b> 2 x 14-Core (5120)</li> <li><b>RAM:</b> 384GB (12 x 32GB @ 2400MHz)</li> <li><b>Disk (LT1 only):</b> 4 x 960GB SSD Or 10 x 600GB 15K SAS</li> </ul>	<ul style="list-style-type: none"> <li><b>CPU:</b> 2 x 20-Core (6138)</li> <li><b>RAM:</b> 768GB (24 x 32GB @ 2667MHz)</li> <li><b>Disk (LT1 only):</b> 4 x 960GB SSD Or 10 x 600GB 15K SAS</li> </ul>

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

### 3.2.1 PowerEdge R640

The Dell EMC PowerEdge R640 is the ideal dual-socket, 1U platform for dense scale-out 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).



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

#### 3.2.1.1 Local Tier 1 Rack

For small deployments such as ROBO or POC/pilot setups, the Local Tier 1 model combines Compute and Management on the same server with VDI desktops (or XenApp sessions) and management role VMs executing from local storage. The OS can be installed to the BOSS device for vSphere or Hyper-V or to dual

SD cards (vSphere only). To provide sufficient capacity and IOPS, use at least 4 x SSDs for all-flash. Optionally, 10 x 15K SAS drives can be substituted for the SSDs.

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


### 3.2.1.2 Shared Tier 1 Rack (FC)

In the Shared Tier 1 model, VDI desktops or XenApp 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, while Hyper-V is installed to the BOSS device.


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

### 3.2.2 PowerEdge 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.



Up to 3 GPU Cards




R740 Shared T1	
	Compute
CPU	2 x Intel Gold 6138 (20C, 2.0GHz)
Memory	24 x 32GB 2667MT/s RDIMMs <b>Effective speed: 2667MT/s @ 768GB</b>
Storage Ctrlrs	PERC H330 – no RAID
Storage	2 x 120GB M.2 BOSS in RAID1 <b>(Hypervisor)</b>
GPUs	Up to 3 x FLDW GPU (NVIDIA)
Network	4 x 10Gb SFP+/BT 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 PowerEdge M630

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



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

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

The Shared Tier 1 blade server varies slightly from the rack server equivalent since the latest 14<sup>th</sup> generation blade servers are not yet available at the time of this writing. However, the processor cores are the same as the rack servers. Two network interconnect Fabrics are configured for the blades: the A-Fabric dedicated to 10Gb LAN traffic and the B-Fabric dedicated to 8Gb FC.

M630 Shared T1	Compute	Management	
	CPU	2 x E5-2698v4 (20C, 2.2GHz)	2 x E5-2660v4 (14C, 2.0GHz)
	Memory	16 x 32GB 2400MT/s RDIMMs 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.



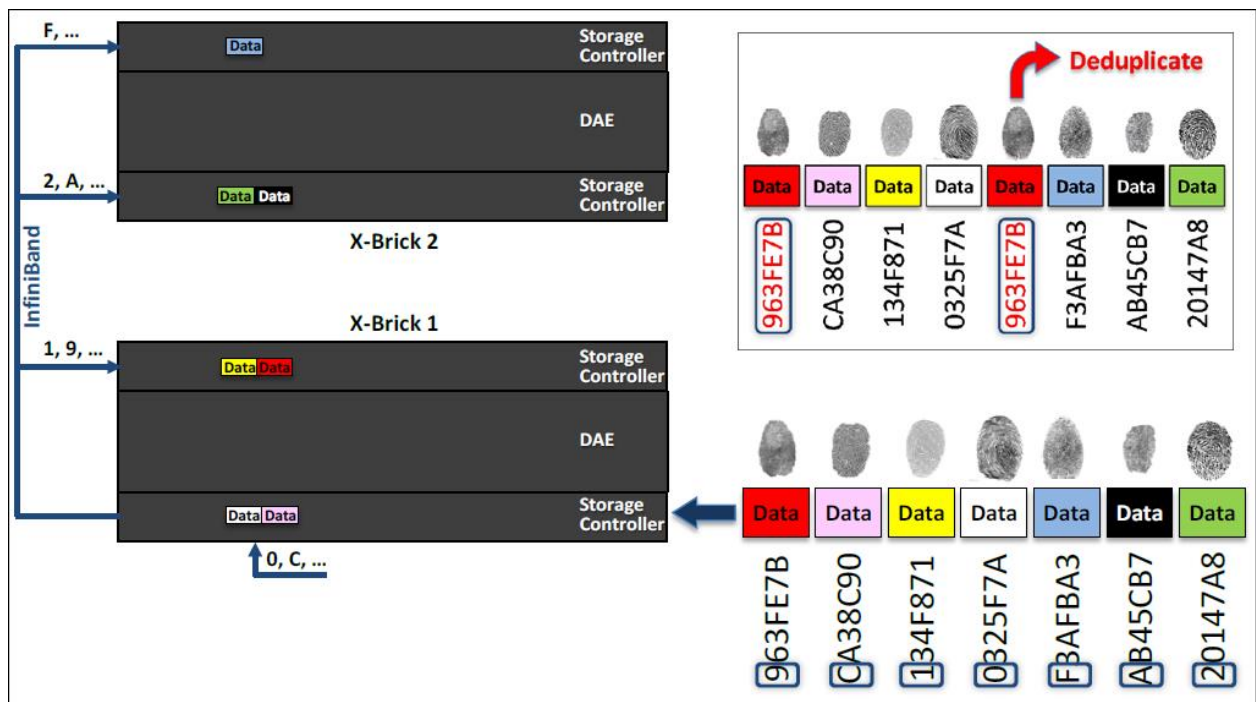
XtremIO leverages flash to deliver value across multiple dimensions:

- 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.



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



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


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

### 3.3.1.1 Architecture

An XtremIO X2 Storage System is comprised of a set of X-Bricks that form together a cluster. This is the basic building block of an XtremIO array. There are two types of X2 X-Bricks available: X2-S and X2-R. X2-S is for environments whose storage needs are more IO intensive than capacity intensive, as they use smaller SSDs and less RAM. An effective use of the X2-S is for environments that have high data reduction ratios (high compression ratio or a lot of duplicated data) which lower the capacity footprint of the data significantly. X2-R X-Bricks clusters are made for the capacity intensive environments, with bigger disks, more RAM and a bigger expansion potential in future releases. The two X-Brick types cannot be mixed together in a single system, so the decision as to which type is suitable for your environment must be made in advance. The X2-S is the recommended X-Brick for Dell EMC Ready Bundle for VDI solutions and is used to provide Shared Tier 1 and Tier 2 storage on the same array.

Each X-Brick is comprised of:

- Two 1U Storage Controllers (SCs) with:
  - Two dual socket Haswell CPUs
  - 346GB RAM (for X2-S) or 1TB RAM (for X2-R)
  - Two 1/10GbE iSCSI ports
  - Two user interface interchangeable ports (either 4/8/16Gb FC or 1/10GbE iSCSI)
  - Two 56Gb/s InfiniBand ports
  - One 100/1000/10000 Mb/s management port
  - One 1Gb/s IPMI port
  - Two redundant power supply units (PSUs)
- One 2U Disk Array Enclosure (DAE) containing:
  - Up to 72 SSDs of sizes 400GB (for X2-S) or 1.92TB (for X2-R)
  - Two redundant SAS interconnect modules
  - Two redundant power supply units (PSUs)





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

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



## 3.4 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. Plus, 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. GPUs can only be used with the Dell EMC XC730 platform.

#### 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.

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.



The NVIDIA® Tesla® M60 GPU accelerator works with NVIDIA GRID™ software to provide the industry's highest user performance for virtualized workstations, desktops, and applications. It allows enterprises to virtualize almost any application (including professional graphics applications) and deliver them to any device, anywhere.

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 Endpoints

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



### 3.5.1 Wyse 3040 Thin Client (ThinOS, ThinLinux)

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



### 3.5.2 Wyse 5010 Thin Client (ThinOS)



Designed for knowledge workers and power users, the Wyse 5010 is a high performance thin client based on Wyse ThinOS, the virus-resistant firmware base designed for optimal thin client security, performance, and ease-of-use. Highly secure, compact and powerful, it combines a dual-core AMD 1.4 GHz CPU with a revolutionary unified graphics engine for an outstanding user experience. It addresses the performance challenges of processing-intensive applications like computer-aided design, multimedia, HD video and 3D modelling. Scalable on premise or cloud-based management provides simple deployment, patching, and updates. Take a unit from box to productivity in minutes with auto configuration.

Delivering outstanding processing speed and power, security and display performance, the Wyse 5010 offers a unique combination of performance, efficiency, and affordability. For more information, please visit: [Link](#)

### 3.5.3 Wyse 5060 Thin Client (ThinOS, ThinLinux, WES7P, WIE10)

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



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

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



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

helps lower the total cost of ownership for large deployments and offers remote enterprise-wide management that scales from just a few to tens of thousands of cloud clients. For more information, please visit [Link](#)

### 3.5.5 Wyse 7040 Thin Client (WES7P, WIE10)

The Wyse 7040 is a high-powered, ultra-secure thin client running Windows Embedded Standard 7P (WES7P) or Windows 10 IoT Enterprise (WIE10) operating systems. Equipped with an Intel i5/i7 processors, it delivers extremely high graphical display performance (up to three displays via display-port daisy-chaining, with 4K resolution available on a single monitor) for seamless access to the most demanding applications. The Wyse 7040 is compatible with both data center hosted and client-side virtual desktop environments and is compliant with all relevant U.S. Federal security certifications including OPAL compliant hard-drive options, VPAT/Section 508, NIST BIOS, Energy-Star and EPEAT. Wyse enhanced WES7P OS provides additional security features such as BitLocker. The Wyse 7040 offers a high level of connectivity including dual NIC, 6 x USB3.0 ports and an optional second network port, with either copper or fiber SFP interface. Wyse 7040 devices are highly manageable through Intel vPRO, Wyse Device Manager (WDM), Microsoft System Center Configuration Manager (SCCM) and Dell Command Configure (DCC). For more information, please visit: [Link](#)



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

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

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

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



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

### 3.5.7 Enhanced Security

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

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

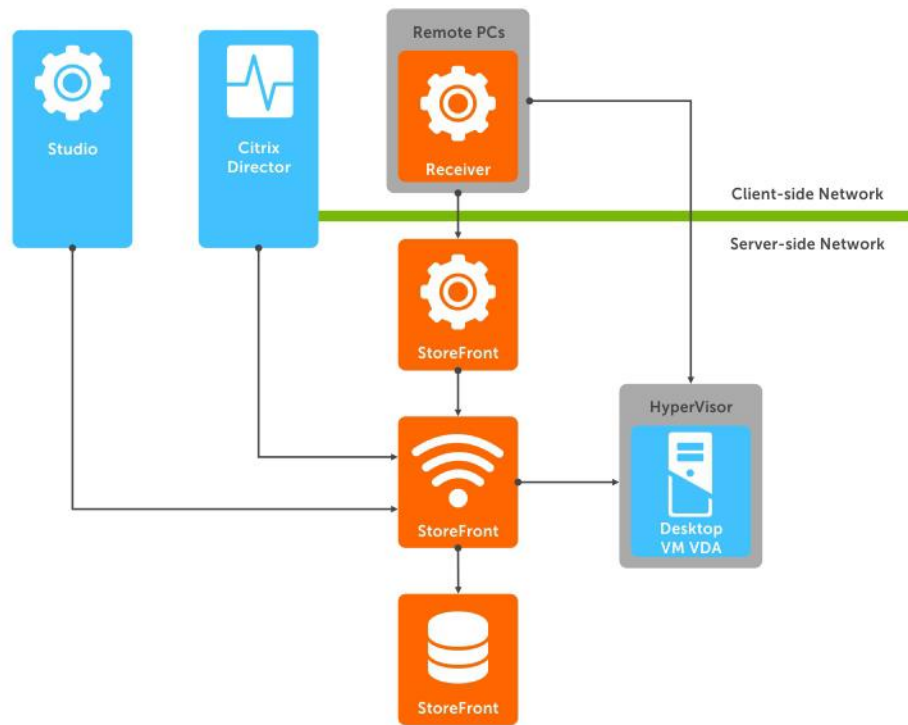
## 4 Software components

### 4.1 Citrix

#### 4.1.1 Citrix XenDesktop

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

Citrix XenDesktop provides a complete virtual desktop delivery system by integrating several distributed components with advanced configuration tools that simplify the creation and real-time management of the virtual desktop infrastructure.



The core XenDesktop components include:

**Studio** - Studio is the management console that enables you to configure and manage your deployment, eliminating the need for separate management consoles for managing delivery of applications and desktops. Studio provides various wizards to guide you through the process of setting up your

environment, creating your workloads to host applications and desktops, and assigning applications and desktops to users.

**Delivery Controller (DC)** - Installed on servers in the data center, the controller authenticates users, manages the assembly of users' virtual desktop environments, and brokers connections between users and their virtual desktops. The Controller also manages the state of desktops, starting and stopping them based on demand and administrative configuration.

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

**Director** - Director is a web-based tool that enables IT support teams to monitor an environment, troubleshoot issues before they become system-critical, and perform support tasks for end users. You can also view and interact with a user's sessions using Microsoft Remote Assistance. Starting in version 7.12, Director now includes detailed descriptions for connection and machine failures, one month historical data (Enterprise edition), custom reporting, and notifications via SNMP traps.

**Receiver** - Installed on user devices, Citrix Receiver provides users with quick, secure, self-service access to documents, applications, and desktops from any of the user's devices including smartphones, tablets, and PCs. Receiver provides on-demand access to Windows, Web, and Software as a Service (SaaS) applications. For devices that cannot install the Receiver software, Citrix Receiver for HTML5 provides connectivity through a HTML5-compatible web browser.

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

**License Server** - The Citrix License Server is an essential component at any Citrix-based solution. Every Citrix product environment must have at least one shared or dedicated license server. License servers are computers that are either partly or completely dedicated to storing and managing licenses. Citrix products request licenses from a license server when users attempt to connect.



**Machine Creation Services (MCS)** - A collection of services that work together to create virtual servers and desktops from a master image on demand; optimizing storage utilization and providing a pristine virtual machine to users every time they log on. Machine Creation Services is fully integrated and administrated in Citrix Studio.

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

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

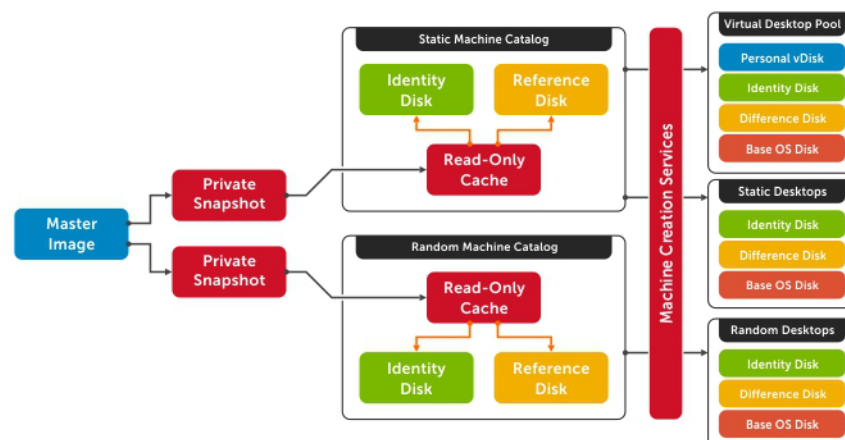
### 4.1.2 Machine Creation Services (MCS)

Citrix Machine Creation Services is the native provisioning mechanism within Citrix XenDesktop for virtual desktop image creation and management. Machine Creation Services uses the hypervisor APIs to create, start, stop, and delete virtual desktop images. Desktop images are organized in a Machine Catalog and within that catalog there are a number of options available to create and deploy virtual desktops:

**Random:** Virtual desktops are assigned randomly as users connect. When they logoff, the desktop is reset to its original state and made free for another user to login and use. Any changes made by the user are discarded at log off.

**Static:** Virtual desktops are assigned to the same user every time with three options for how to handle changes made to the desktop: Store on local vDisk, Personal vDisk, or discarded on user log off.

All the desktops in a random or static catalog are based off a master desktop template which is selected during the catalog creation process. MCS then takes snapshots of the master template and layers two additional virtual disks on top: an Identity vDisk and a Difference vDisk. The Identity vDisk includes all the specific desktop identity information such as host names and passwords. The Difference vDisk is where all the writes and changes to the desktop are stored. These Identity and Difference vDisks for each desktop are stored on the same data store as their related clone.



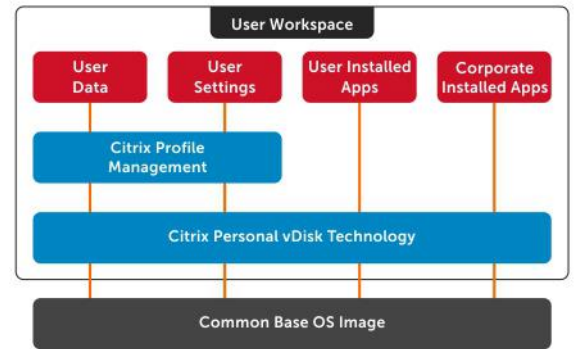
While traditionally used for small to medium sized XenDesktop deployments, MCS can bring along with it some substantial Tier 1 storage cost savings because of the snapshot/identity/difference disk methodology. The Tier 1 disk space requirements of the identity and difference disks when layered on top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



### 4.1.3 Personal vDisk

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

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

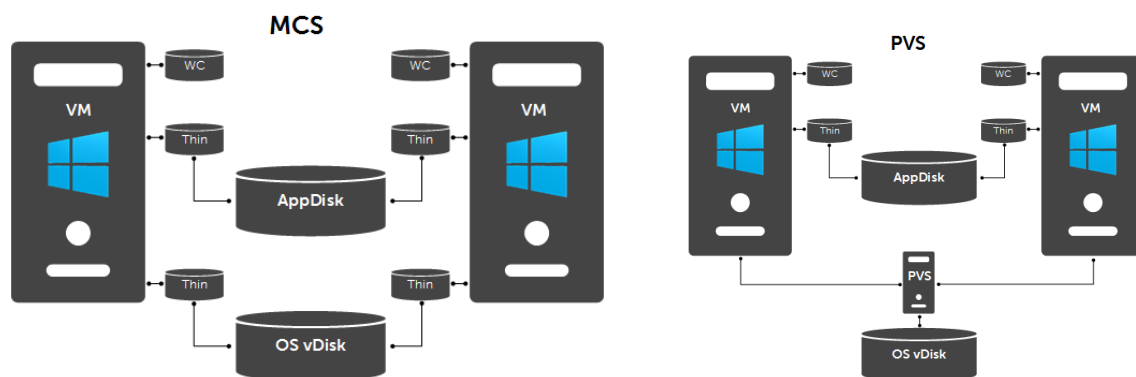


Personal vDisk provides the following benefits to XenDesktop:

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

### 4.1.4 AppDisks

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



### 4.1.5 HDX 3D Pro

XenDesktop with HDX 3D Pro is a desktop and app virtualization solution that supports high-end designers and engineers of 3D professional graphics applications as well as provides cost-effective support to viewers

and editors of 3D data. With XenDesktop, you can deliver a persistent user experience and leverage other virtualization benefits such as single-image management and improved data security.

Use HDX 3D Pro technologies with:

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

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

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

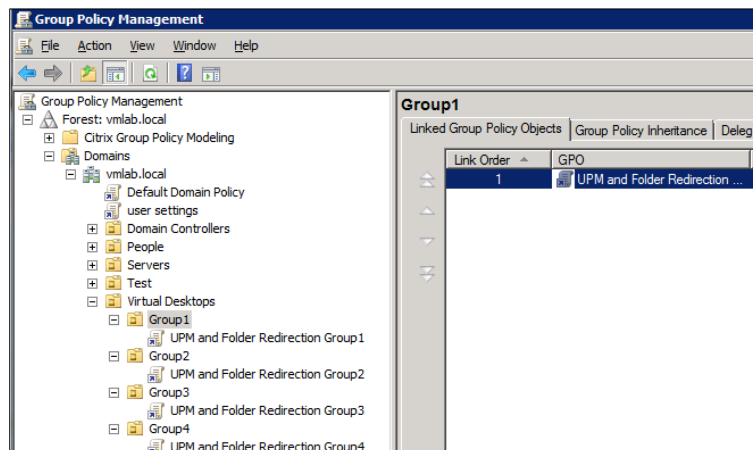
Framehawk is a display remoting technology implemented as an ICA virtual channel that optimizes delivery of virtual desktops and applications to users on broadband wireless connections where high packet loss or congestion occurs.

### 4.1.6 Citrix Profile Manager

Citrix Profile Management is a component of the XenDesktop suite which is used to manage user profiles and minimize many of the issues associated with traditional Windows roaming profiles in an environment where users may have their user profile open on multiple devices at the same time. The profile management toolset has two components: the profile management agent, installed on any device where the user profiles is managed, and a Group Policy Administrative Template, which is imported to a group policy.

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

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




- My Videos
- Desktop

### 4.1.7 Citrix XenApp

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

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



Benefits of hosted desktop sessions and applications:

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

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

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

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

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

#### 4.1.7.1 XenApp integration into Dell EMC Ready Bundle for VDI Architecture

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

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

#### 4.1.7.2 XenDesktop with XenApp and Personal vDisk integration

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

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

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

#### 4.1.7.3 NUMA architecture considerations

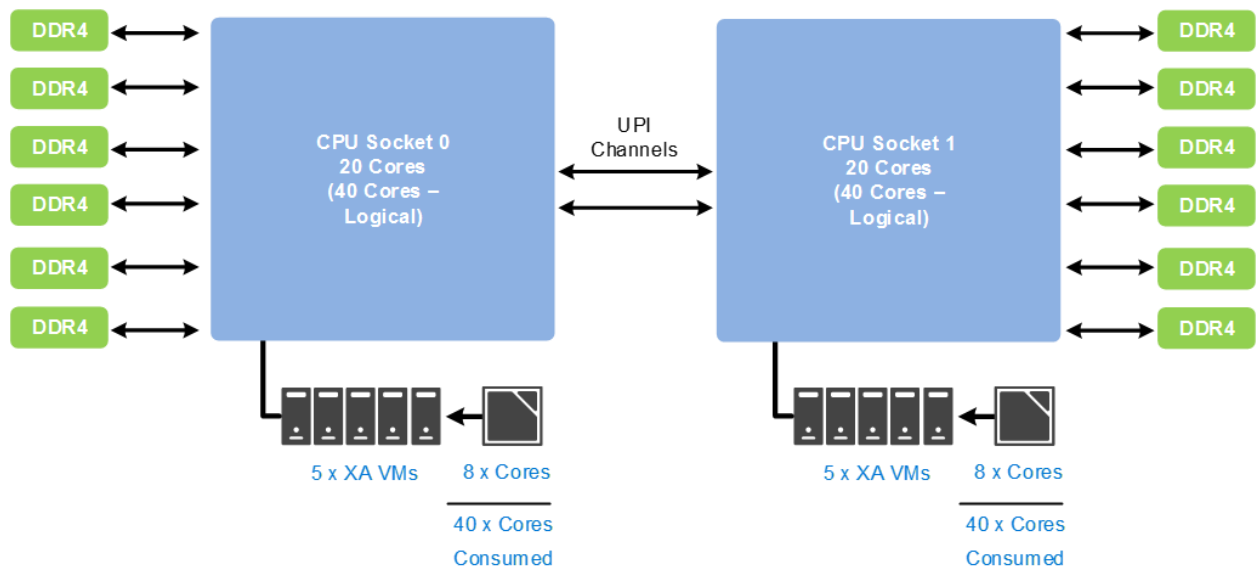
Best practices and testing has showed that aligning RDSH design to the physical Non-Uniform Memory Access (NUMA) architecture of the server CPUs results in increased and optimal performance. NUMA

alignment ensures that a CPU can access its own directly-connected RAM banks faster than those banks of the adjacent processor which are accessed via the Ultra Path Interconnect (UPI) which was formally the Quick Path Interconnect (QPI) with older processors. The same is true of VMs with large vCPU assignments, best performance will be achieved if your VMs receive their vCPU allotment from a single physical NUMA node. Ensuring that your virtual RDSH servers do not span physical NUMA nodes will ensure the greatest possible performance benefit.

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

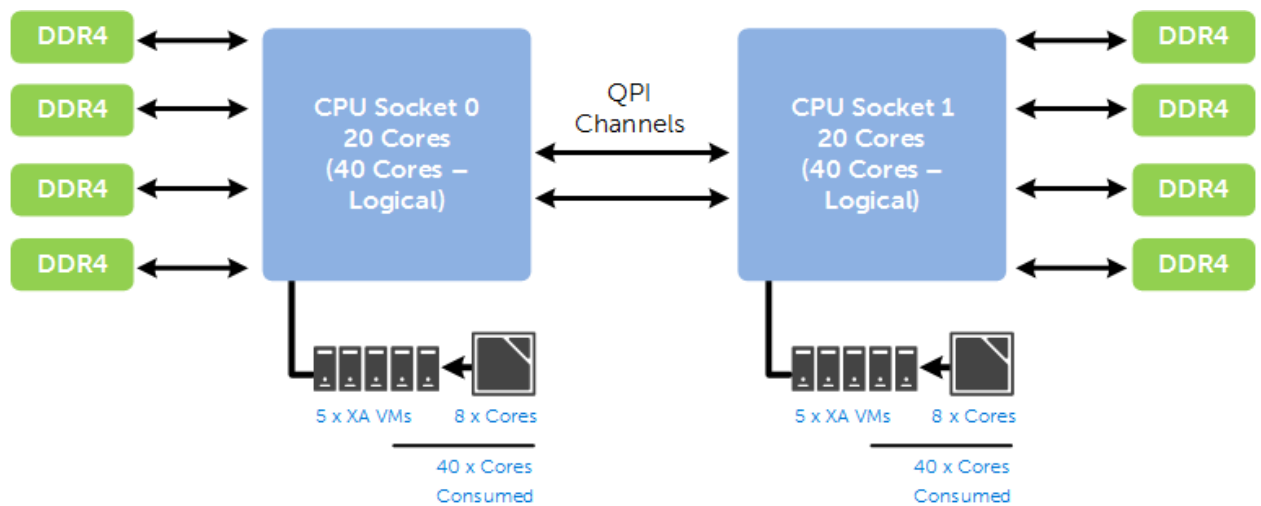
#### 4.1.7.4 NUMA alignment – Skylake processor (R640/R740)

20 physical cores per CPU given the Gold 6138 part, 40 logical cores with Hyper-Threading active, gives us a total of 80 consumable cores per compute node and falls in line with a 2x oversubscription rate. Configuring the XenApp VMs as shown below (as XA VMs) will ensure that no physical NUMA node spanning occurs which could lower performance for an effected VM.



#### 4.1.7.5 NUMA alignment – Broadwell processor (M630)

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



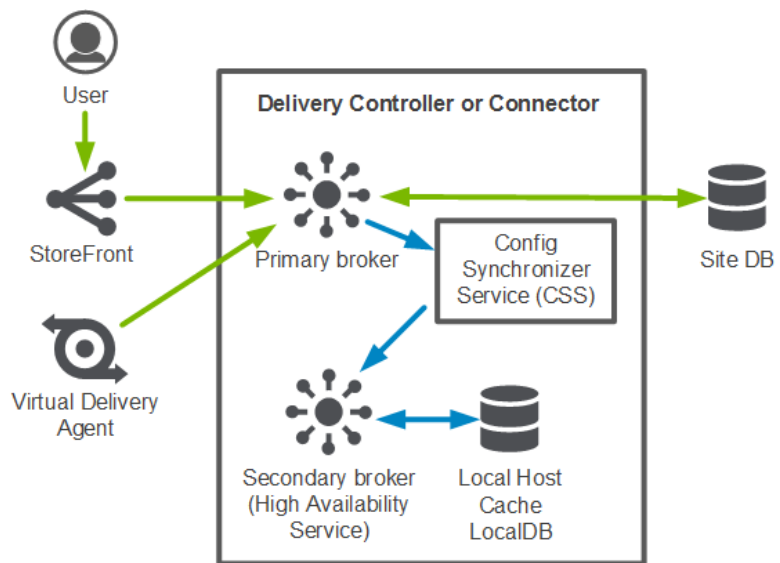
### 4.1.8 Local Host Cache

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

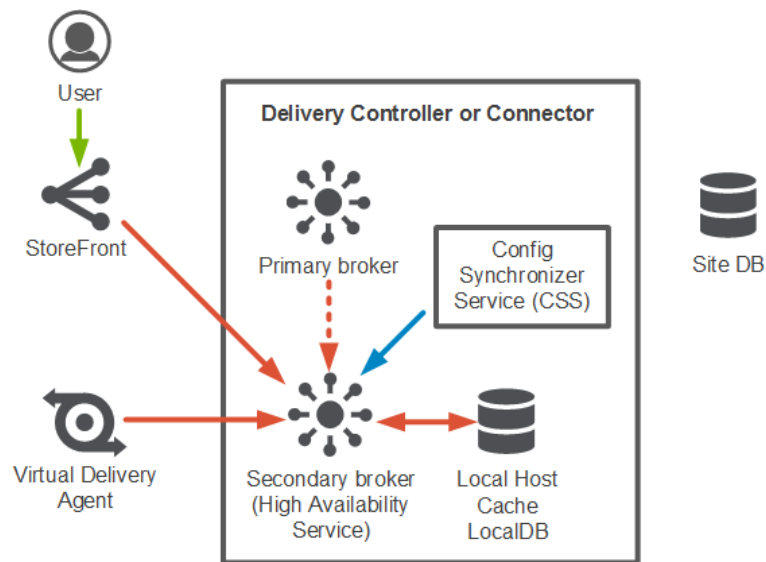
LHC replaces the connection leasing feature as the recommended XenApp and XenDesktop high availability solution. During an outage, LHC will support new users and existing users launching new resources, as well as users accessing pooled resources (shared desktops). Earlier versions of XenApp had a feature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

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

The following diagram shows the communication paths during normal operations.



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



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

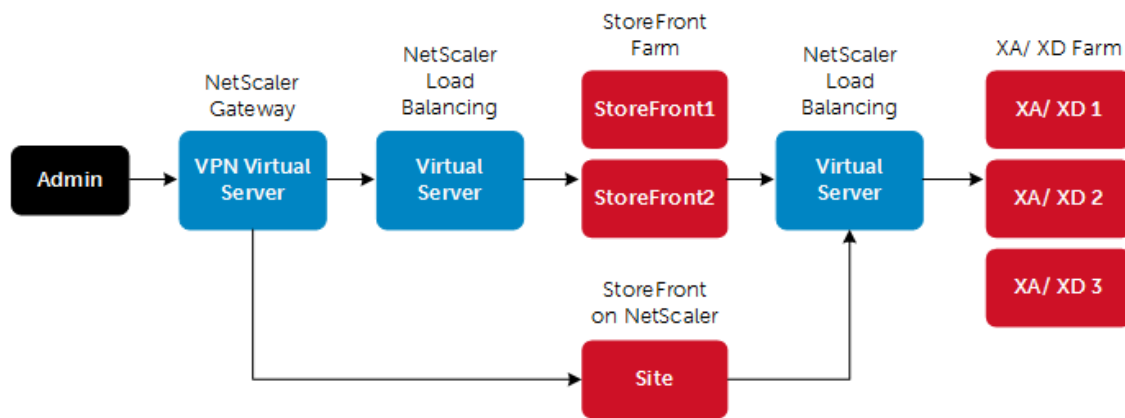
#### 4.1.9 Citrix NetScaler

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

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

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





## Global Server Load Balancing

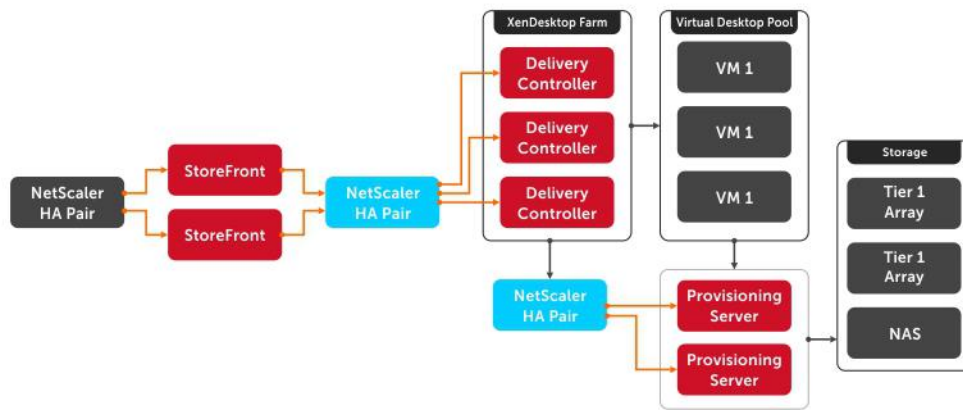
GSLB is an industry standard function. It is in widespread use to provide automatic distribution of user requests to an instance of an application hosted in the appropriate data center where multiple processing facilities exist. The intent is to seamlessly redistribute load on an as required basis, transparent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

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

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

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

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



## 4.2 Hypervisor platforms

### 4.2.1 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.

VMware vSphere® 6.x 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 HA Support for NVIDIA GRID vGPU, please visit this Blog article: [Link](#).

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

## **4.2.2 Microsoft Windows Server Hyper-V**

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

Hyper-V includes the following features:

### **Computing environment**

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

### **Disaster recovery and backup**

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

### **Optimization**

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

### **Portability**

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

### **Remote connectivity**

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

### **Security**

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

From a network management standpoint, virtual machines are much easier to manage than physical computers. To this end, Hyper-V includes many management features designed to make managing virtual machines simple and familiar, while enabling easy access to powerful VM-specific management functions. The primary management platform within a Hyper-V based XenDesktop virtualization environment is Microsoft Systems Center Virtual Machine Manager (SCVMM).

SCVMM provides centralized and powerful management, monitoring, and self-service provisioning for virtual machines. SCVMM host groups are a way to apply policies and to check for problems across several VMs at once. Groups are organized by owner, operating system, or by custom names such as “Development” or “Production”. The interface also incorporates Remote Desktop Protocol (RDP); double-click a VM to bring up the console for that VM—live and accessible from the management console.

## **4.3 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 exactly 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 outstanding shared virtualized graphics performance.

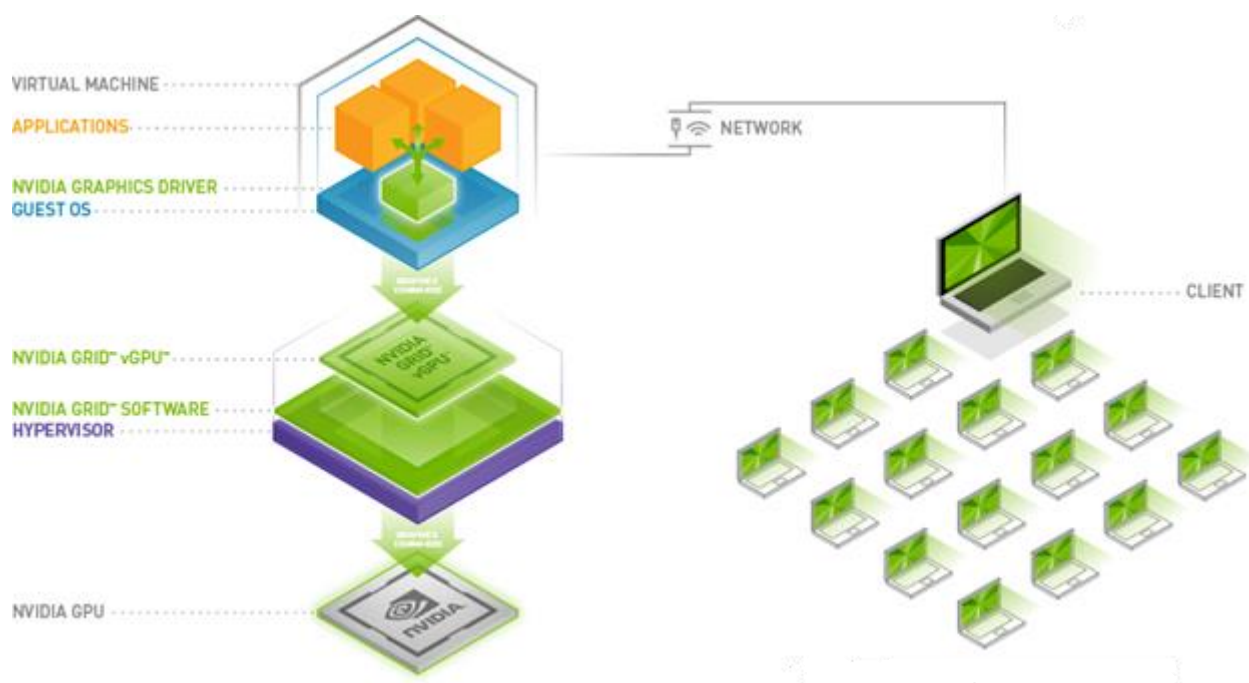


Image provided courtesy of NVIDIA Corporation, Copyright NVIDIA Corporation

### 4.3.1 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. Every virtual desktop has dedicated graphics memory so they always have the resources they need to launch and run their applications at full performance. 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 appliance is determined by the number of GPU cards in the system (max 3), vGPU profiles used for each GPU in a card (2-4 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.

#### NVIDIA® Tesla® M10 GRID vGPU Profiles:

Card					Maximum Graphics-Enabled VMs

	vGPU Profile	Graphics Memory (Frame Buffer)	Virtual Display	Maximum Resolution	Per GPU	Per Card	Per Server (3 cards)
<b>Tesla M10</b>	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	96
	M10-1A	1GB			8	32	192

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.

#### 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	48
	M60-1A	1GB			8	16	96



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.

#### 4.3.1.1 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.

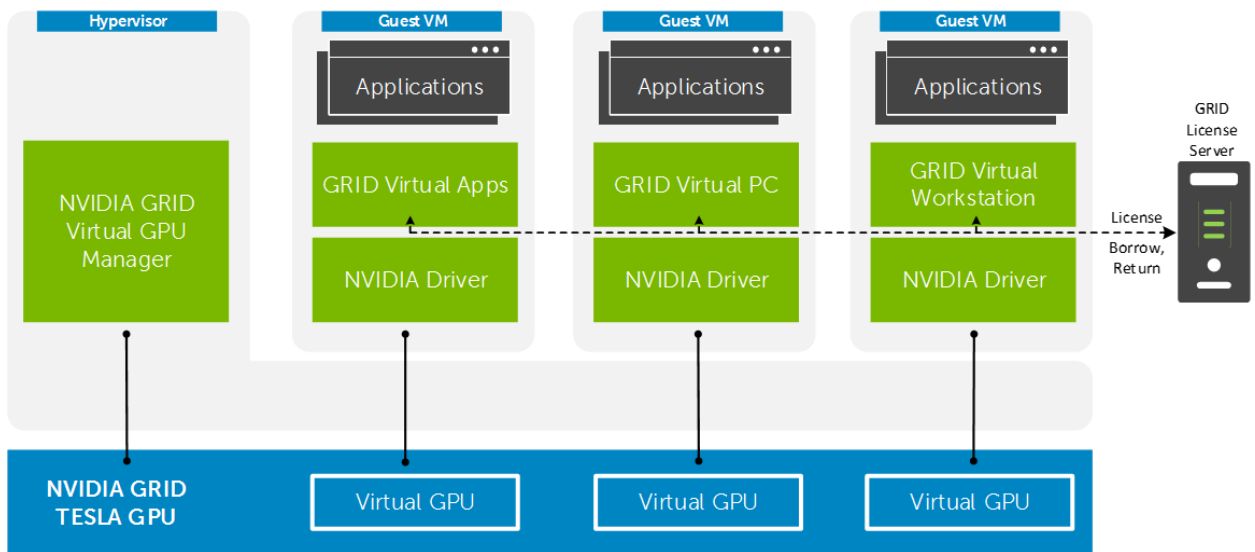


<b>NVIDIA GRID Virtual Applications</b>	<b>NVIDIA GRID Virtual PC</b>	<b>NVIDIA GRID Virtual Workstation</b>
For organizations deploying XenApp or other RDSH solutions. Designed to deliver Windows applications at full performance.	For users who want a virtual desktop, but also need a great user experience 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 XenDesktop

### 5.1 Management role configuration

The Management role recommendations for the base solution are summarized below. Use data disks for role-specific application files such as data, logs and IIS web files in the Management volume.

#### 5.1.1 vSphere

The Management role requirements for the base solution are summarized below and take into account the use of PVS. If MCS provisioning is desired remove this from your build plan. Depending on the scale of the deployment, the number of these services must be adjusted. Use data disks for role-specific application files such as data, logs and IIS web files in the Management volume.

Role	vCPU	RAM (GB)	NIC	OS + Data vDisk (GB)	Tier 2 Volume (GB)
DDC + Lic	4	8	1	60	-
Storefront	2	4	1	60	-
SQL Server	8	8	1	60	300 (VMDK)
vCenter	8	16	1	60	50 (VMDK)
PVS	4	24	1	60	100 (VMDK)
File Server	2	4	1	60	2048 (RDM)
<b>TOTALS</b>	<b>28</b>	<b>64</b>	<b>6</b>	<b>360</b>	<b>2,598</b>

#### 5.1.2 NVIDIA GRID License Server Requirements

When using NVIDIA Tesla M60 or M10 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	vCPU	RAM (GB)	NIC	OS + Data vDisk (GB)	Tier 2 Volume (GB)
NVIDIA GRID License Srv	2	4	1	40 + 5	-

At the time of this writing, 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 Hyper-V

Role	vCPU	Startup RAM (GB)	Dynamic Memory			NIC	OS + Data vDisk	
			Min Max	Buffer	Weight		Size (GB)	Tier2 Vol
DDC + Lic	4	8	4GB   10GB	20%	Med	1	60	-
Storefront	2	4	2GB   6GB	20%	Med	1	60	-
Primary SQL *	8	8	4GB   16GB	20%	Med	1	60	300 (VHDX)
SCVMM	8	8	4GB   16GB	20%	Med	1	60	50 (VHDX)
PVS	4	24	12GB   24GB	20%	Med	1	60	200 (VHDX)
File Server	2	4	2GB   6GB	20%	Med	1	60	2048 (PTM)
<b>Total</b>	<b>28</b>	<b>56GB</b>	<b>28GB   78GB</b>	<b>-</b>	<b>-</b>	<b>6</b>	<b>360GB</b>	<b>2598GB</b>

**\*NOTE:** Set memory to static (16GB or higher) if using SQL Standard version to avoid issue with memory not allocating properly.

### 5.1.4 XenApp VM configuration

The recommended number of XenApp VMs and their configurations on vSphere or Hyper-V are summarized below and take into account proper NUMA balancing assuming the CPU in use is the Intel Xeon Gold 6138 or E5-2698v4. For more information on NUMA please refer to the [NUMA Architecture Considerations](#) section.

XenApp VM configuration on Hyper-V

Role				Dynamic Memory	NIC	OS vDisk
------	--	--	--	----------------	-----	----------

	VMs per host	vCPUs per VM	Startup RAM (GB)	Min Max	Buffer	Weight			
								Size (GB)	Location
XenApp VM	10	8	16	8GB   48GB	20%	Med	1	80	Tier 1

XenApp VM configuration on vSphere

Role	VMs per host	vCPUs per VM	Memory	NIC	OS vDisk	
					Size (GB)	Location
XenApp VM	10	8	48GB	1	80	Tier 1

### 5.1.5 SQL databases

The Citrix, Microsoft and VMware databases are hosted by a single dedicated SQL 2012 R2 (or higher) Server VM in the Management layer. Use caution during database setup to ensure that SQL data, logs, and TempDB are properly separated onto their respective volumes. Create all Databases that are required for:

- Citrix XenDesktop
- vCenter or SCVMM

Initial placement of all databases into a single SQL instance is fine unless performance becomes an issue, in which case database need to be separated into separate named instances. Enable auto-growth for each DB.

Best practices defined by Citrix, Microsoft and VMware are to be adhered to, to ensure optimal database performance.

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

### 5.1.6 DNS




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

Pay consideration for eventual scaling, access to components that may live on one or more servers (SQL databases, Citrix services) during the initial deployment. Use CNAMEs and the round robin DNS mechanism to provide a front-end “mask” to the back-end server actually hosting the service or data source.

#### 5.1.6.1 DNS for SQL

To access the SQL data sources, either directly or via ODBC, a connection to the server name\ instance name must be used. To simplify this process, as well as protect for future scaling (HA), instead of connecting to server names directly, alias these connections in the form of DNS CNAMEs. So instead of connecting to SQLServer1\<instance name> for every device that needs access to SQL, the preferred approach is to connect to <CNAME>\<instance name>.

For example, the CNAME “VDISQL” is created to point to SQLServer1. If a failure scenario was to occur and SQLServer2 would need to start serving data, we would simply change the CNAME in DNS to point to SQLServer2. No infrastructure SQL client connections would need to be touched.

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

## 5.2 Storage configuration overview

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

### 5.2.1 Local Tier 1 storage

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

### 5.2.2 Shared Tier 1 storage

Selecting the Shared Tier 1 model means that the virtualization compute hosts are deployed without Tier 1 local storage and leverage shared storage hosted on a high performance array. In this model, shared storage is leveraged for Tier 1 and used for VDI execution and write cache. 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 cloning process (full clones or linked clones) as well as if thin or thick provisioning is used.

### 5.2.3 Shared Tier 2 storage

Tier 2 is shared file storage used to host the Management server VMs and user data. In this solution architecture, shared Tier 2 storage is using the same X-Brick array as the Tier 1 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. 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 or VHDX formats.

Volumes	Size (GB)	Storage Array	Purpose	File System
Management	410	Tier 2	vCenter/ SCVMM, XD roles, File & SQL	VMFS/ NTFS
PVS	200	Tier 2	PVS vDisks	VMFS/ NTFS
User Data	2048	Tier 2	File Server/ NAS	RDM/PTD
User Profiles	20	Tier 2	User profiles	VMFS/ NTFS
SQL DATA	100	Tier 2	SQL	VMFS/ NTFS
SQL LOGS	100	Tier 2	SQL	VMFS/ NTFS
TempDB Data	5	Tier 2	SQL	VMFS/ NTFS
TempDB Logs	5	Tier 2	SQL	VMFS/ NTFS
Templates/ISO	200	Tier 2	ISO storage (optional)	VMFS/ NTFS

---

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

---



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


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

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

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

The following QLogic HBA BIOS 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
- 



For Hyper-V deployments, configure Microsoft Multipath I/O (MPIO) with the least queue depth policy.

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

- 
- 
- 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. Round Robin or Fixed Paths are supported. Use a single-Target-per-single-Initiator (1:1) zoning scheme.

## 5.3 Virtual networking

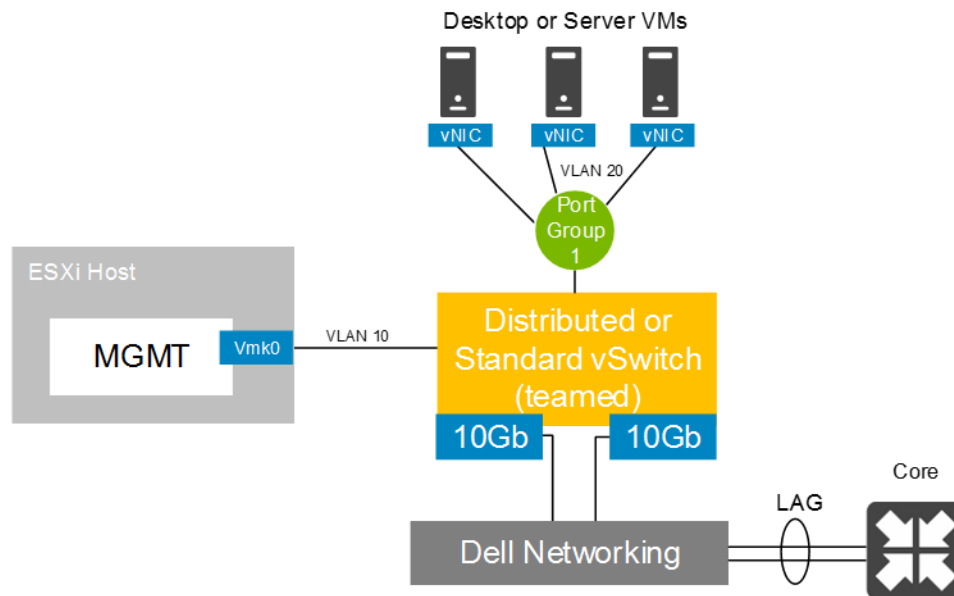
### 5.3.1 Local Tier 1

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

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

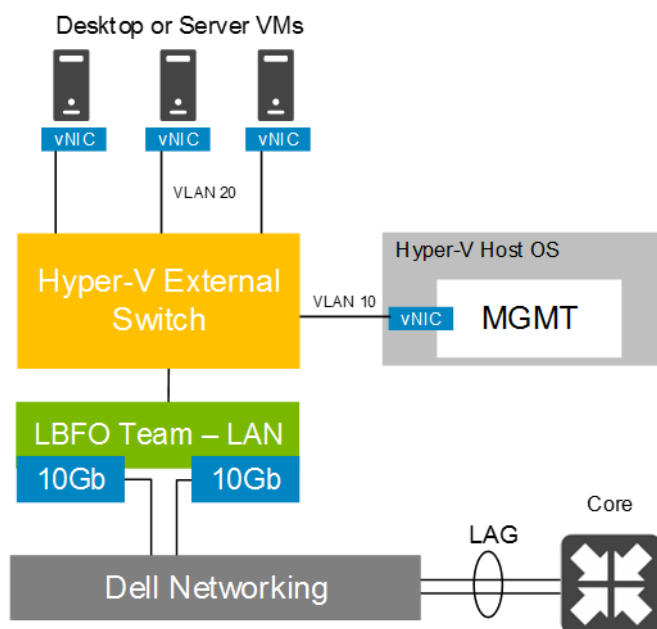
This traffic is combined within a single switch; however, VLANs are required for each traffic type to enable traffic separation. Configure the LAN traffic from the server to the ToR switch as a LAG.

#### 5.3.1.1 vSphere



#### 5.3.1.2 Hyper-V

The Hyper-V configuration, while identical in core requirements and hardware, is executed differently due to how Hyper-V and Windows Server implement networking and virtual switches. As shown in the diagram below, native Windows Server NIC Teaming is utilized to load balance and provide resiliency for network connections. A single LBFO NIC team is configured to connect to a Hyper-V switch for external traffic. All vNICs associated with the Management OS connect directly to the external Hyper-V switch.



### 5.3.2 Shared Tier 1 – FC

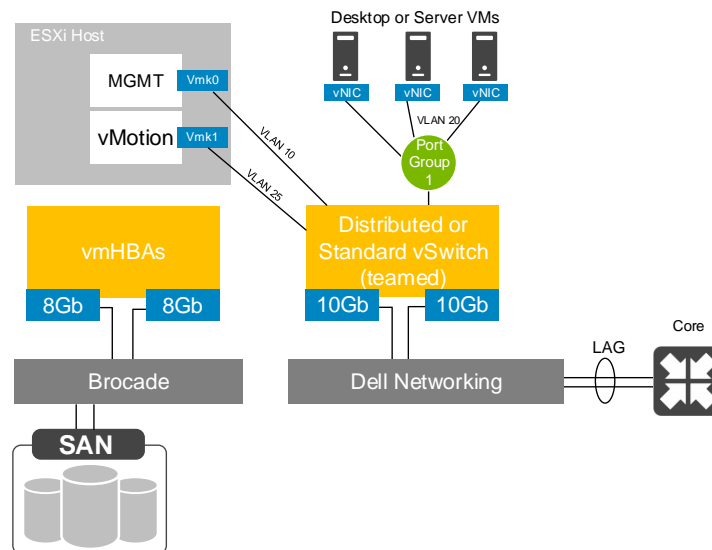
Using Fiber Channel based storage requires additional storage fabrics to be built out in the network stack. The network configuration in this model is identical between the Compute and Management hosts. The benefits of shared storage are available to all hosts such as Live Migration and HA. The following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

- Compute hosts (Shared Tier 1)
  - Management VLAN: Configured for hypervisor Management traffic – L3 routed via core switch
  - Live Migration VLAN: Configured for Live Migration traffic – L2 switched only, trunked from Core
  - Failover Cluster VLAN: Configured for Cluster and Cluster Shared Volume traffic – L2 switched only, trunked from core (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.

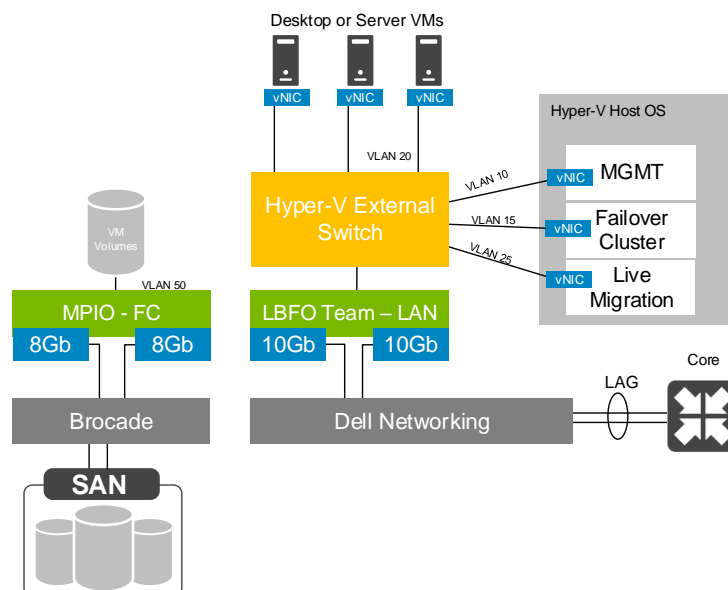
### 5.3.2.1 vSphere

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.



### 5.3.2.2 Hyper-V

The Hyper-V configuration, while identical in core requirements and hardware, is executed differently due to how Hyper-V and Windows Server implement networking and virtual switches. As shown in the diagram below, native Windows Server NIC Teaming is utilized to load balance and provide resiliency for network connections. For the compute or management hosts in this scenario, a single LBFO NIC team is configured to connect to a Hyper-V switch for external traffic. All vNICs associated with the Management OS connect directly to the external Hyper-V switch with MPIO used to connect to shared storage.



## 5.4 Scaling guidance

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

Component	Metric	Horizontal scalability	Vertical scalability
Compute Servers	Desktop VMs per physical host based on available CPU	Additional hosts and clusters added as necessary	Additional RAM or CPU compute power
Mgmt. Servers	Number of server VMs per host	Add additional hosts	Add RAM or network adapters
Provisioning Servers	Desktops per instance	Additional servers added to the Provisioning Server farm	Additional network and I/O capacity added to the servers
Desktop Delivery Servers	Desktops per instance (dependent on SQL performance as well)	Additional servers added to the XenDesktop Site	Additional virtual machine resources (RAM and CPU)
XenApp Servers	Desktops per instance	Additional virtual servers added to the XenDesktop Site	Additional physical servers to host virtual XenApp servers.
Storefront Servers	Logons/ minute	Additional servers added to the Storefront environment	Additional virtual machine resources (RAM and CPU)
Database Services	Concurrent connections, responsiveness of reads/ writes	Migrate databases to a dedicated SQL server and increase the number of management nodes	Additional RAM and CPU for the management nodes
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

### XtremIO X2 Scaling

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

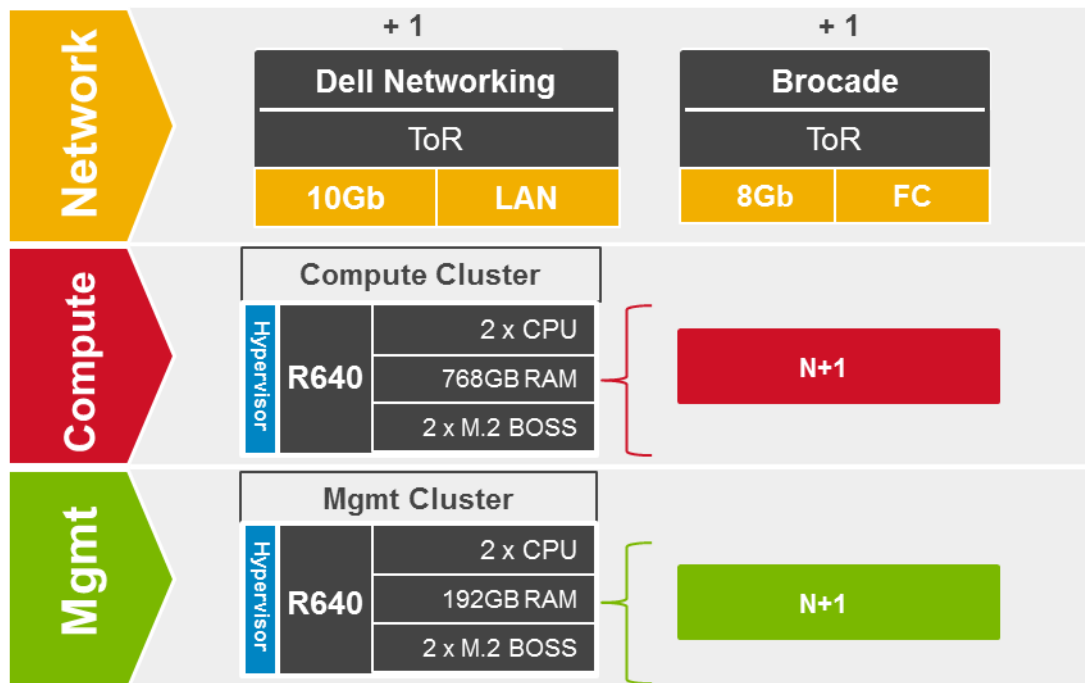
Scale-up of an XtremIO cluster is implemented by adding SSDs to existing DAEs in the cluster. Adding SSDs to existing DAEs to scale-up an XtremIO cluster is intended for environments that currently grow in capacity needs without the need for extra performance. This occurs, for example, when the same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

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

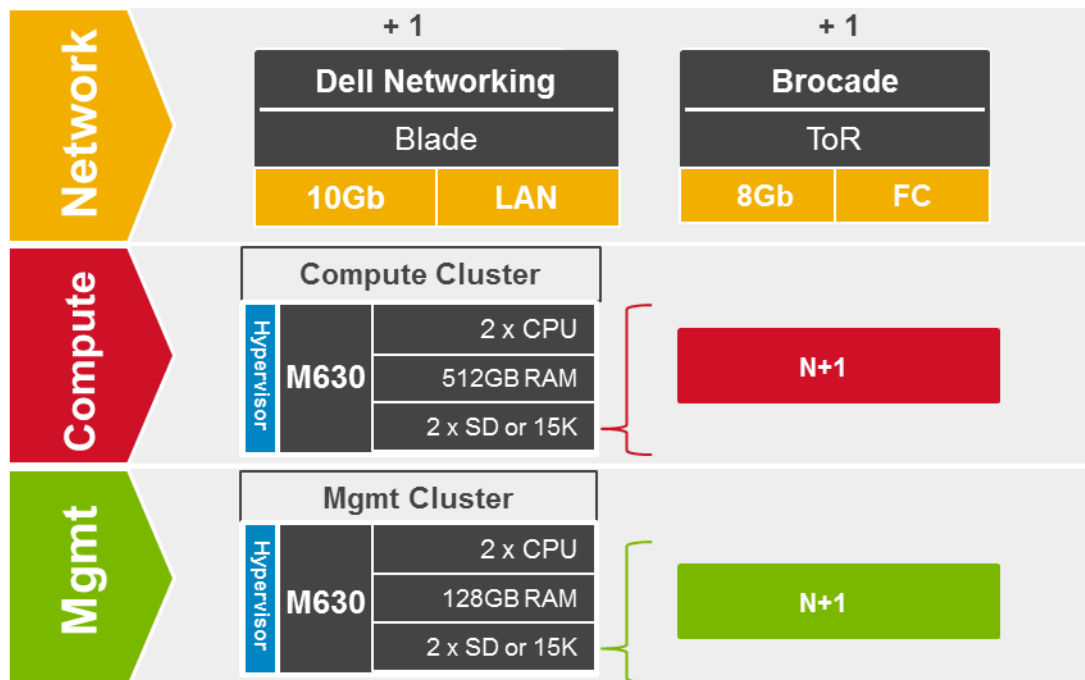
## 5.5 Solution high availability

High availability (HA) is offered to protect each architecture solution layer, individually if desired. Following the N+1 model, additional ToR switches are added to the Network layer and stacked to provide redundancy as required, additional compute and management hosts are added to their respective layers, vSphere or Hyper-V clustering is introduced in both the management and compute layers, SQL is configured with AlwaysOn and NetScaler is leveraged for load balancing. Storage protocol switch stacks and NAS selection will vary based on chosen solution architecture.

## Rack Servers



## Blade Servers




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

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

### 5.5.1 SQL Server high availability

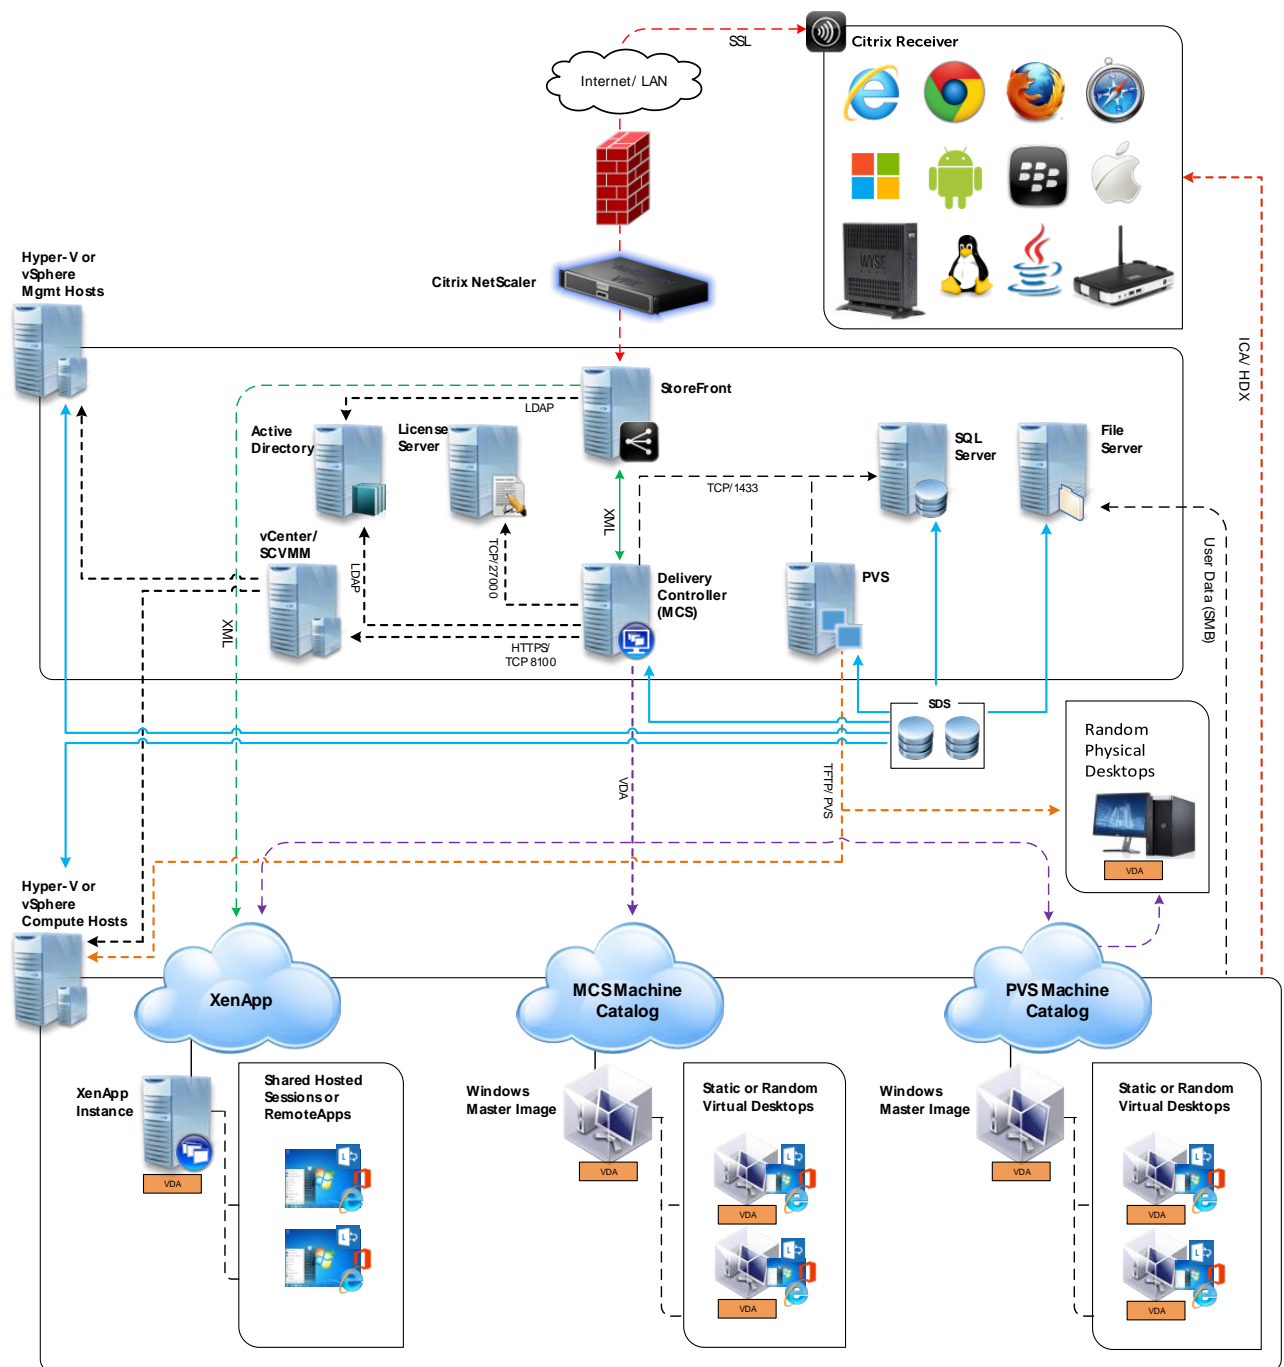
HA for SQL is provided via AlwaysOn using either Failover Cluster Instances or Availability Groups. This configuration protects all critical data stored within the database from physical server as well as virtual server problems. DNS is used to control access to the primary SQL instance. Place the principal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



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



## 5.6 Citrix XenDesktop communication flow



## 6 Solution performance and testing

### 6.1 Summary

At the time of publication, these are the available density recommendations per appliance/node. Please refer to the [Servers](#) and [Platform Configurations](#) sections for hardware specifications.

User density summary

Host Config	Hypervisor	Broker & Provisioning	Workload	Template	User Density
B5	ESXi 6.5	XenDesktop 7.15 MCS	Knowledge Worker	Windows 10 x64 & Office 2016	145
B5	Hyper-V 2016	XenDesktop 7.15 MCS	Knowledge Worker	Windows 10 x64 & Office 2016	155

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

### 6.2 Test and performance analysis methodology

#### 6.2.1 Testing process

In order to ensure the optimal combination of end-user experience (EUE) and cost-per-user, performance analysis and characterization (PAAC) on Dell EMC Ready Bundle for VDI solutions is carried out using a carefully designed, holistic methodology that monitors both hardware resource utilization parameters and EUE during load-testing.

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

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

##### 6.2.1.1 Load generation

Login VSI by Login Consultants is the de-facto industry standard tool for testing VDI environments and server-based computing (RDSH environments). It installs a standard collection of desktop application software (e.g. Microsoft Office, Adobe Acrobat Reader) on each VDI desktop; it then uses launcher systems to connect a specified number of users to available desktops within the environment. Once the user is connected, the workload is started via a logon script which starts the test script once the user environment is configured by the login script. Each launcher system can launch connections to a number of ‘target’ machines (i.e. VDI desktops). The launchers and Login VSI environment are configured and managed by a centralized management console.

Additionally, the following login and boot paradigm is used:

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

### 6.2.1.2 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.

## **6.2.2 Resource monitoring**

The following sections explain respective component monitoring used across all Dell EMC Ready Bundle for VDI solutions where applicable.

### 6.2.2.1 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.

### 6.2.2.2 Microsoft Performance Monitor

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

### 6.2.2.3 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.

## 6.2.3 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 (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, for vSphere specifically.

## 6.3 Test configuration details

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

Hardware and software test components

Component	Description/Version
Hardware platform(s)	PowerEdge R740
Hypervisor(s)	VMware 6.5 and Hyper-V 2016
Broker technology	XenDesktop 7.15
Broker database	Microsoft SQL 2016
Management VM OS	Windows Server 2016
Virtual desktop OS	Windows 10 x64 Enterprise
Office application suite	Office Professional 2016
Login VSI test suite	4.1.25

### 6.3.1 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	Hyper-V Start-up Memory	Hyper-V Min   Max Dynamic	Operating System
Task Worker	1	2GB	1GB	1GB	1GB   2GB	Windows 10 Enterprise 64-bit
Knowledge Worker	2	3GB	1.5GB	1.5GB	1GB   3GB	Windows 10 Enterprise 64-bit
Power Worker	2	4GB	2GB	2GB	1GB   4GB	Windows 10 Enterprise 64-bit
Graphics LVSI Power + ProLibrary	2	4 GB	4GB			Windows 10 Enterprise 64-bit
Graphics LVSI Custom – Density	2	4 GB	4GB			Windows 10 Enterprise 64-bit

Graphics LVSI Custom - Performance	4	8GB	8GB			Windows 10 Enterprise 64-bit
---------------------------------------	---	-----	-----	--	--	---------------------------------

#### Screen resolutions

User Profile	Screen Resolution
Task Worker	1280 X 720
Knowledge Worker	1920 X 1080
Power Worker	1920 X 1080
Graphics LVSI Power + ProLibrary	1920 X 1080
Graphics LVSI Custom – Density	1920 X 1080
Graphics LVSI Custom - Performance	1920 X 1080

#### XenApp VM specifications

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

## 6.3.2 Platform Configurations

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

#### Compute hardware configuration

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

#### Shared storage hardware configuration

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

## 6.4 Test results and analysis

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

Test result summary

Platform Config	Hypervisor	Broker & Provisioning	Login VSI Workload	Density per Host	Avg CPU	Avg Mem Consumed	Avg Mem Active	Avg IOPS / User	Avg Net Mbps / User
B5	ESXi 6.5	XenDesktop 7.15 MCS	Knowledge	145	95%	344GB		11	2
B5	Server 2016 Hyper-V	XenDesktop 7.15 MCS	Knowledge	155	80%	344GB		11.5	1.3

**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 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:** ESXi consumed memory is the amount of host physical memory consumed by a host. 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.




## 6.4.1 R740 Compute

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


### 6.4.1.1 Knowledge Worker, 145 Users, ESXi 6.5, XenDesktop 7.15 MCS

The single compute host was populated with 145 virtual machines. With all user virtual machines powered on and before starting test, the CPU usage was approximately 16%.


The below graph shows the performance data for 145 user sessions. The CPU reaches 100% usage during the logon phase but settles to about 95% usage during the steady state of the test when all users are logged on to the host.



Memory usage shown below is for consumed and active. Some ballooning and swapping occurred during the test from late in the logon phase until the completion of the test but was within acceptable limits. Consumed memory reached the maximum of 350GB but active memory reached 224GB during the Login VSI test just as the steady state phase starts. Active memory briefly reached 284GB after logoff.




Network bandwidth is not an issue on this test run with a steady state peak of approximately 330 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 graphs below are taken from the XtremIO Administration console and show the Disk IOPS, bandwidth and latency for the XtremIO cluster.


The cluster reached a maximum of 25,000 Disk IOPS during the boot process and 5,000 IOPS during the steady state and again at the logoff. Latency remained at one millisecond or less 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 The VSI baseline was good at 960. The maximum VSI response during the test was about 1500.




### 6.4.1.2 Knowledge Worker, 155 Users, Hyper-V 2016, XenDesktop 7.15 MCS


The below graph shows the CPU performance data. The CPU reaches a steady state average of 80% during the test cycle when 155 users are logged on to the compute host.



Memory consumption reached just over 90% for a system with 384GB. Nearly all the memory was consumed by the VDI VMs. It can be seen that although CPU and IO latency are low the system is limited by the memory available.





Network bandwidth is not an issue with a steady state peak of approximately 200 Mbps. Network usage is affected by the Login VSI content and profile and home folder redirection. The output queues on the network adapters were monitored during the testing and remained at zero (0).




Citrix temporary memory caching for XenDesktop MCS desktops was not used because of the limited memory available to the compute host. All temporary data was stored on the shared storage.

The host reached a maximum of 17000 IOPs during the boot phase, 5000 during the login phase and dropped to 1800 during the steady state resulting in about 11 IOPS per user in the later state of the test. In all cases the latency remained below 1 ms during the test as measured from the Hyper-V host. The X-Brick Performance Chart is shown following the Hyper-V chart. This shows that the latency was observed as zero (0) during the test



The Login VSI Max user experience score shown below for this test was not reached. However login times started to increase towards the end of the logon period. The VSI index is increasing and it can be seen that it would probably soon reach the threshold for VSIMAX.

As the CPU and disk latency remain well within thresholds it is assumed that the low available memory in the system is causing the increased response times.





## Related resources

See the following referenced or recommended resources:

- The [Dell EMC Cloud-Client Computing Solutions for Citrix Tech Center page](#) which includes this RA and other Citrix XenDesktop based RAs.
- [EMC XtremIO Storage Array Host Configuration Guide](#)
- [Microsoft Software Assurance for Thin Clients with Windows 10 IoT Enterprise](#)