

REFERENCE ARCHITECTURE

Dell EMC ScaleIO Ready Nodes for VDI

Integration of VMware Horizon with ScaleIO Ready Nodes.

Abstract

A Reference Architecture for integrating Dell EMC ScaleIO Ready Nodes with VMware vSphere and Horizon software to create virtual application and virtual desktop environments on 14th generation Dell EMC PowerEdge Servers.

Revisions

Date	Description
February 2018	Initial release

Acknowledgements

This paper was produced by the following members of the Dell EMC Ready Solutions team:

Author: Keith Keogh – Lead Architect

Peter Fine – Chief Architect

Support: Colin Byrne – Principal Systems Development Engineer

Rick Biedler – Engineering Director

David Hulama – Senior 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.

© 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.....	5
1 Introduction.....	6
1.1 Objective.....	6
1.2 What's new.....	6
2 Solution Architecture Overview.....	7
2.1 Introduction.....	7
2.2 Dell EMC SIO Ready Nodes.....	7
2.3 SIO Architecture.....	8
2.4 ScaleIO Ready Nodes.....	10
2.5 Dell EMC SIO Ready Nodes - VDI solution architecture.....	13
2.5.1 Networking.....	13
2.5.2 SIO Ready Nodes – Enterprise solution pods.....	17
3 Hardware components.....	19
3.1 Network.....	19
3.1.1 Dell Networking S3048 (1Gb ToR switch).....	19
3.1.2 Dell Networking S4048 (10Gb ToR switch).....	20
3.2 Dell EMC SIO Ready Nodes.....	21
3.2.1 Dell EMC SIO R640.....	23
3.2.2 Dell EMC SIO R740xd.....	24
3.3 NVIDIA Tesla GPUs.....	25
3.3.1 NVIDIA Tesla M10.....	25
3.3.2 NVIDIA Tesla M60.....	26
3.4 Dell Wyse Endpoints.....	27
3.4.1 Wyse 3040 Thin Client (ThinOS, ThinLinux).....	27
3.4.2 Wyse 5040 AIO Thin Client (ThinOS).....	27
3.4.3 Wyse 5060 Thin Client (ThinOS, ThinLinux, WES7P, WIE10).....	27
3.4.4 Wyse 7020 Thin Client (WES 7/7P/8, WIE10, ThinLinux).....	28
3.4.5 Wyse 7040 Thin Client (WES7P, WIE10).....	28
3.4.6 Latitude 3480 and 5280 Mobile Thin Clients (Win 10 IoT).....	28
4 Software components.....	29
4.1 VMware vSphere 6.5.....	29
4.2 VMware Horizon.....	30
4.2.1 VMware Horizon Key Features.....	31

4.2.2	VMware Horizon Apps.....	32
4.2.3	NUMA architecture considerations.....	33
4.3	NVIDIA GRID vGPU.....	35
4.3.1	vGPU profiles.....	35
5	Solution architecture for SIO with Horizon.....	41
5.1.1	RDSH VM Configuration.....	41
5.1.2	NVIDIA GRID License Server Requirements.....	42
5.1.3	SQL databases.....	42
5.1.4	DNS.....	42
5.2	Storage architecture overview.....	43
5.3	Virtual Networking.....	43
5.4	Scaling guidance.....	44
5.5	Solution high availability.....	46
5.6	Communication flow for Horizon.....	47
6	Solution performance and testing.....	48
6.1	Summary.....	48
6.2	Test and performance analysis methodology.....	49
6.2.1	Testing process.....	49
6.2.2	Resource monitoring.....	51
6.2.3	Resource utilization.....	51
6.3	Test configuration details.....	52
6.3.1	Compute VM configurations.....	53
6.3.2	Platform configurations.....	53
6.4	Test results and analysis.....	54
6.4.1	Task Worker, 230 users, ESXi 6.5, Horizon 7.x Instant Clones.....	55
6.4.2	Knowledge Worker, 140 users, ESXi 6.5, Horizon 7.x Instant Clones.....	60
6.4.3	Power Worker, 115 users, ESXi 6.5, Horizon 7.x Instant Clones.....	66
6.4.4	Power Worker (M60-1Q), 48 users, ESXi 6.5, Horizon 7.x Instant Clones.....	71
6.4.5	Power Worker (M10-1B), 64 users, ESXi 6.5, Horizon 7.x Instant Clones.....	78

Executive summary

This document provides the reference architecture for integrating Dell EMC ScaleIO (SIO) Ready Nodes and VMware Horizon software to create virtual desktop infrastructure environments.

The Dell EMC SIO Ready Nodes are a hyper-converged solution that combines storage, compute, networking, and virtualization using industry-proven Dell EMC PowerEdge™ server technology and Dell EMC ScaleIO™ software. By combining the hardware resources from each appliance into a shared-everything model for simplified operations, we deliver improved agility, greater flexibility and simple, cost-effective solutions for enterprise workloads.

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

1 Introduction

This document addresses the architecture design, configuration and implementation considerations for the key components required to deliver virtual desktops or shared sessions via VMware Horizon on the Dell EMC SIO Ready Nodes with vSphere™ 6.5.

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

- SIO Ready Nodes on 14th generation Dell EMC PowerEdge platforms

2 Solution Architecture Overview

2.1 Introduction

Dell EMC customers benefit in leveraging this integrated solution for their primary workload data protection needs. This integrated solution offers Virtual Machine (VM) deployment and lifecycle management for the combined solution offering as well as protection for newly deployed and existing VMs. Usage of policies and best practices and the consequent streamlining of the data protection workflow are the primary goals for this solution. This section will provide an overview of the products used to validate the solution.

2.2 Dell EMC SIO Ready Nodes

Dell EMC SIO ready nodes start with the proven 14th generation Dell EMC PowerEdge 14th generation server platform and incorporate many of the advanced software technologies that power leading web-scale and cloud infrastructures. Backed by Dell EMC global service and support, these 1U and 2U appliances are preconfigured for specific virtualized workloads, and are designed to maintain data availability in case of node and disk failure.

The SIO ready node infrastructure is a scalable cluster of high-performance servers, each running a standard hypervisor and containing processors, memory, and local storage consisting of solid state disk (SSD) drives for high performance. Each ready node runs virtual machines just like a standard hypervisor.

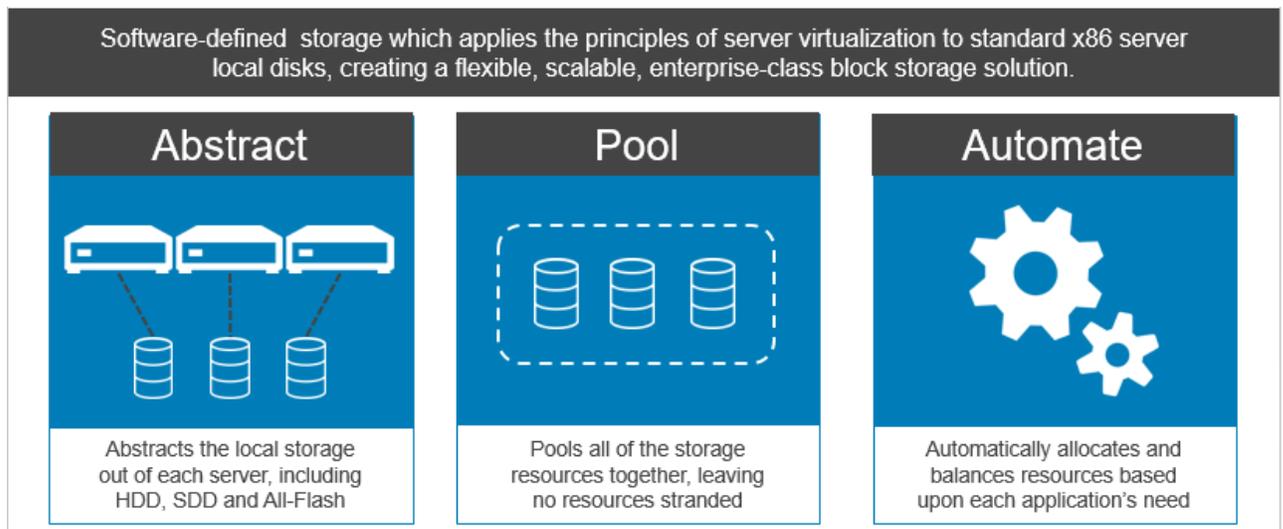


Figure 1 Dell EMC ScaleIO Ready Nodes overview

2.3 SIO Architecture

ScaleIO systems contain a number of elements including the SDC, SDS and MDM and are discussed in detail in the following sections.

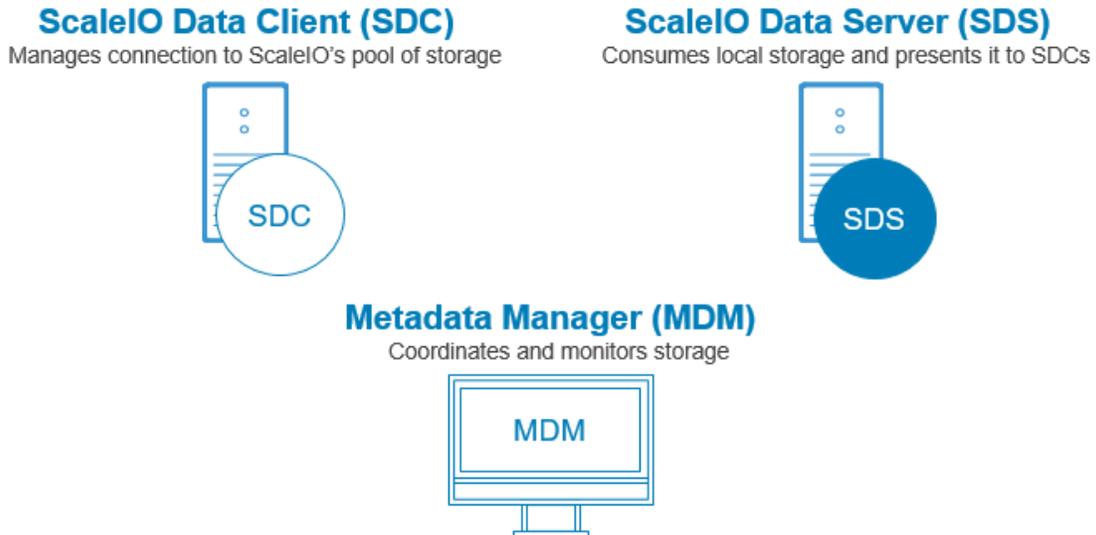


Figure 2 ScaleIO architecture overview

ScaleIO Data Client –SDC

The SDC is a lightweight block device driver that exposes ScaleIO shared block volumes to applications. The SDC runs on the Hypervisor and this enables the application to issue an IO request and the SDC fulfils it regardless of where the particular blocks physically reside. The SDC communicates with other nodes in the cluster over TCP/IP-based protocol, so it is fully routable.

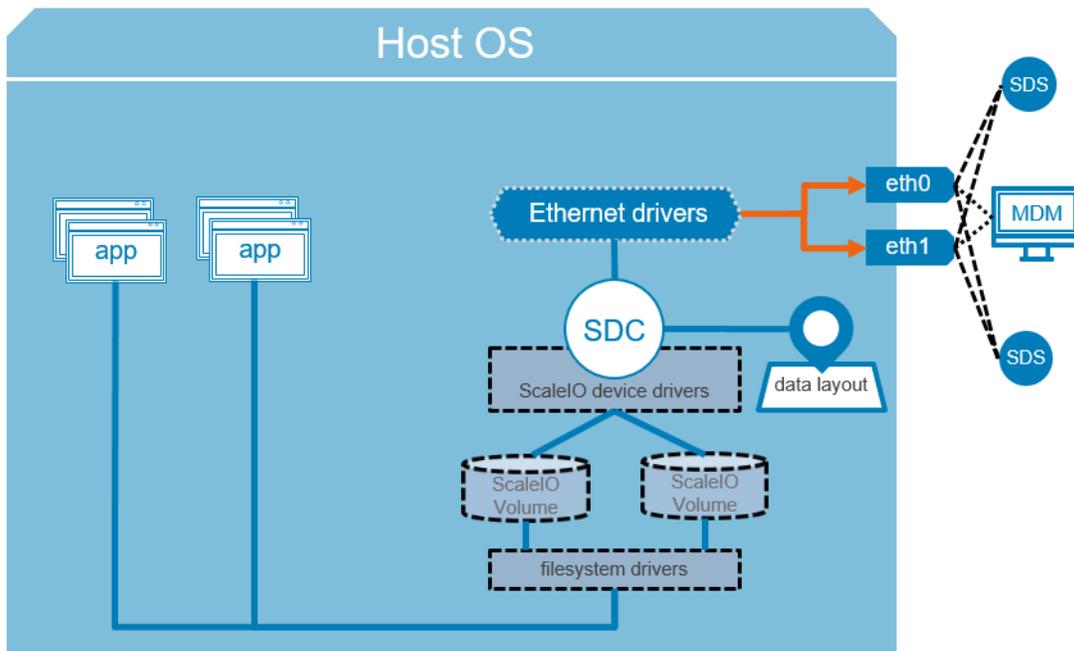


Figure 3 ScaleIO Data Client

ScaleIO Data Server –SDS

The SDS owns local storage that contributes to the ScaleIO Storage Pools. An instance of the SDS runs on every server that contributes some or all of its local storage space (HDDs, SSDs, PCIe, NVMe and flash cards) to the aggregated pool of storage within the ScaleIO datastore. The role of the SDS is to perform the Back-End IO operations as requested by an SDC. This is fully configurable in terms of how much memory is consumed, and which devices it's associated with.

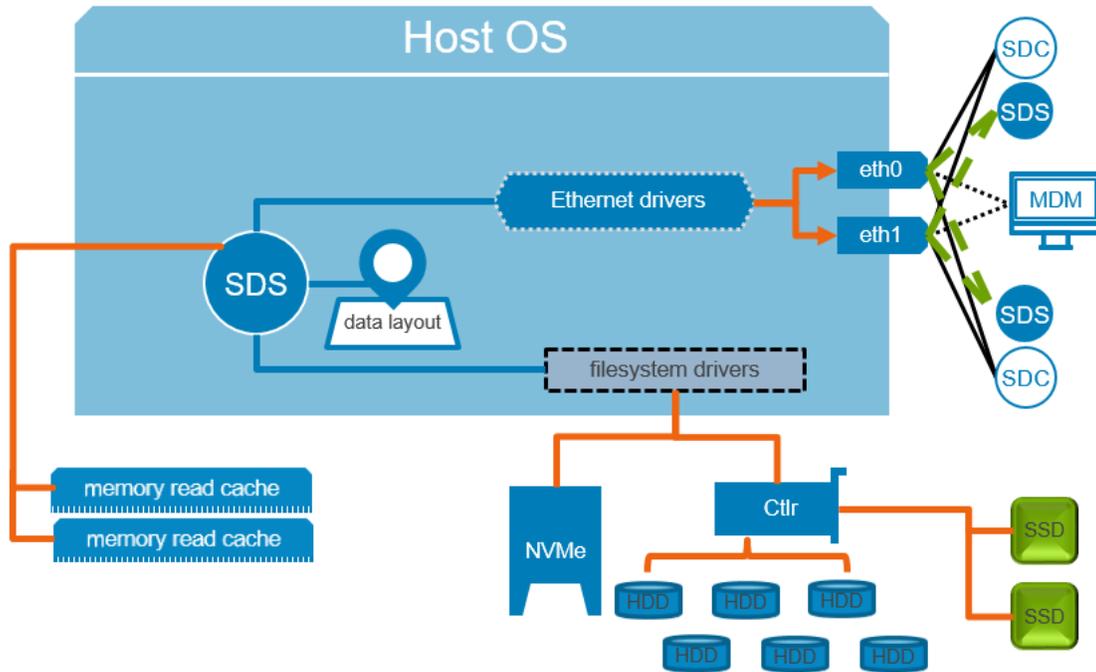


Figure 4 ScaleIO Data Server

Meta Data Manager –MDM

The Meta Data Manager (MDM) manages the ScaleIO system. The MDM contains all the metadata required for system operation, such as configuration changes. The MDM also provides monitoring capabilities to assist users with the most system management tasks.

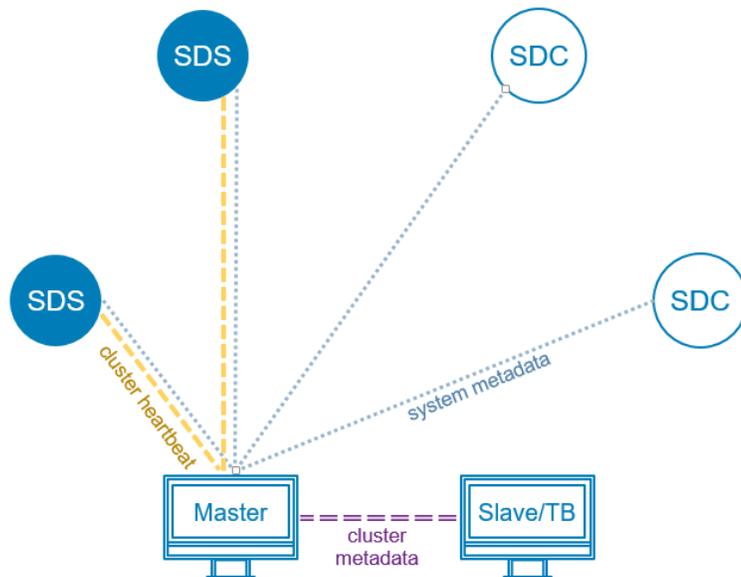


Figure 5 Meta Data Manager

The MDM manages the meta data, SDC, SDS, devices mapping, volumes, snapshots, system capacity including device allocations and/or release of capacity, RAID protection, errors and failures, and system rebuild tasks including rebalancing. In addition, all user interaction with the system is handled by the MDM. In a normal IO flow, the MDM is not part of the data path and user data does not pass through the MDM. Therefore, the MDM is never a performance bottleneck for IO operations. Currently, an MDM can manage up to 1024 servers. When several MDMs are present, an SDC may be managed by several MDMs, whereas, an SDS can only belong to one MDM. ScaleIO version 2.0 and later supports five MDMs (with a minimum of three) where we define a Master, Slave and Tie-breaker MDM.

The MDM is extremely lightweight and has an asynchronous interaction with the SDCs and SDSs. The MDM daemon produces a heartbeat where updates are performed every few seconds. If the MDM does not detect the heartbeat from an SDS it will initiate a rebuild. All ScaleIO commands are asynchronous with one exception. For consistency reasons, the unmap command is synchronous where the user must wait for the completion before continuing. Each SDC holds mapping information that is light-weight and efficient so it can be stored in real memory.

The SDS & MDM run on a standalone virtual machine called a ScaleIO Virtual Machine (SVM) in an ESXi environment with the SDC directly installed on the ESXi host.

2.4 ScaleIO Ready Nodes

The SIO hyper-converged infrastructure provides an ideal combination of both high-performance compute with localized storage to meet any demand. True to this capability, this reference architecture has been validated as optimized for the VDI use case.

A hyper-converged configuration involves the installation of both the SDC and the SDS on the same server. This is considered the best practice configuration as well. In this configuration, the applications and storage share the same compute resources and in many ways the storage is like another application running on the server.

The applications perform IO operations via the local SDC. All servers contribute some or all of their local storage to the ScaleIO system via their local SDS. Components communicate over the Local Area Network (LAN).

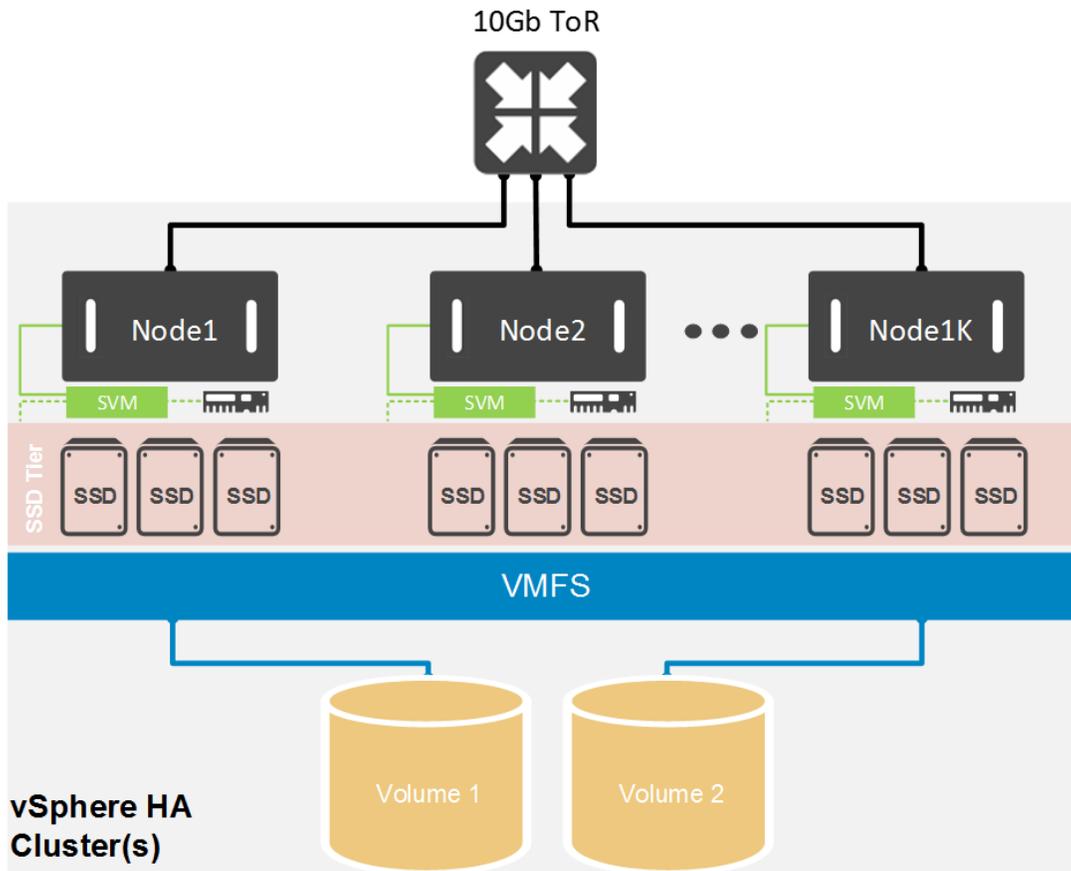


Figure 6 High-level example of the relationship between a SIO ready node, storage pool, container, pod and relative scale out.

This solution allows organizations to deliver virtualized or remote desktops and applications through a single platform and support end users with access to all of their desktops and applications in a single place.

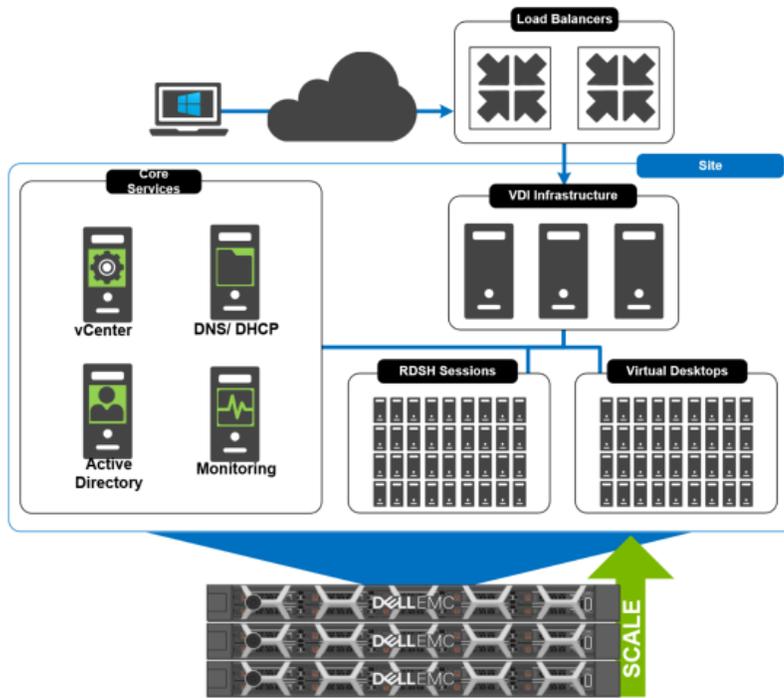


Figure 7 Delivery of remote desktops and applications from a single platform

2.5 Dell EMC SIO Ready Nodes - VDI solution architecture

2.5.1 Networking

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

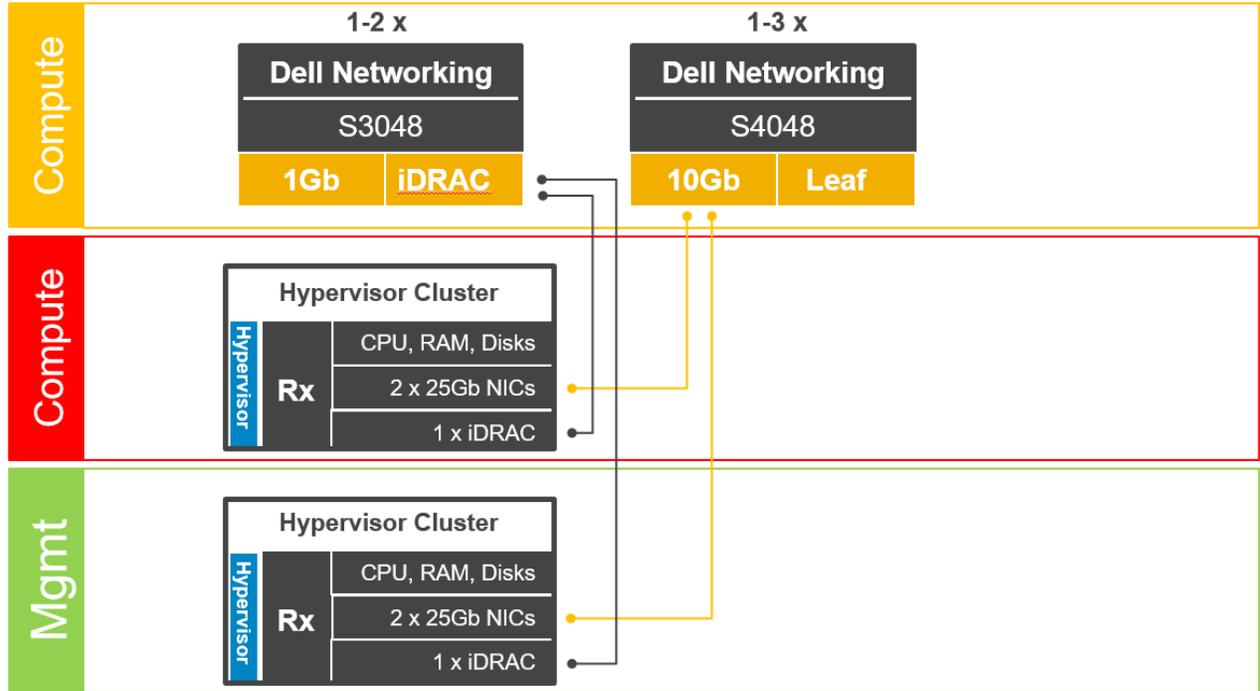


Figure 8 Networking architecture

Designed for true linear scaling, SIO ready nodes leverages a Leaf-Spine network architecture. A Leaf-Spine architecture consists of two network tiers: a 10Gb layer-2 (L2) Leaf segment and a layer-3 (L3) Spine segment based on 40GbE and non-blocking switches. This architecture maintains consistent performance without any throughput reduction due to a static maximum of three hops from any node in the network. The hosts are using 2 x 25Gb nics but auto-negotiated down to 10Gb.

The following figure shows a design of a scale-out Leaf-Spine network architecture that provides 20Gb active throughput from each node to its Leaf and scalable 80Gb active throughput from each Leaf to Spine switch providing scale from 3 SIO ready nodes to a thousand without any impact to available bandwidth:

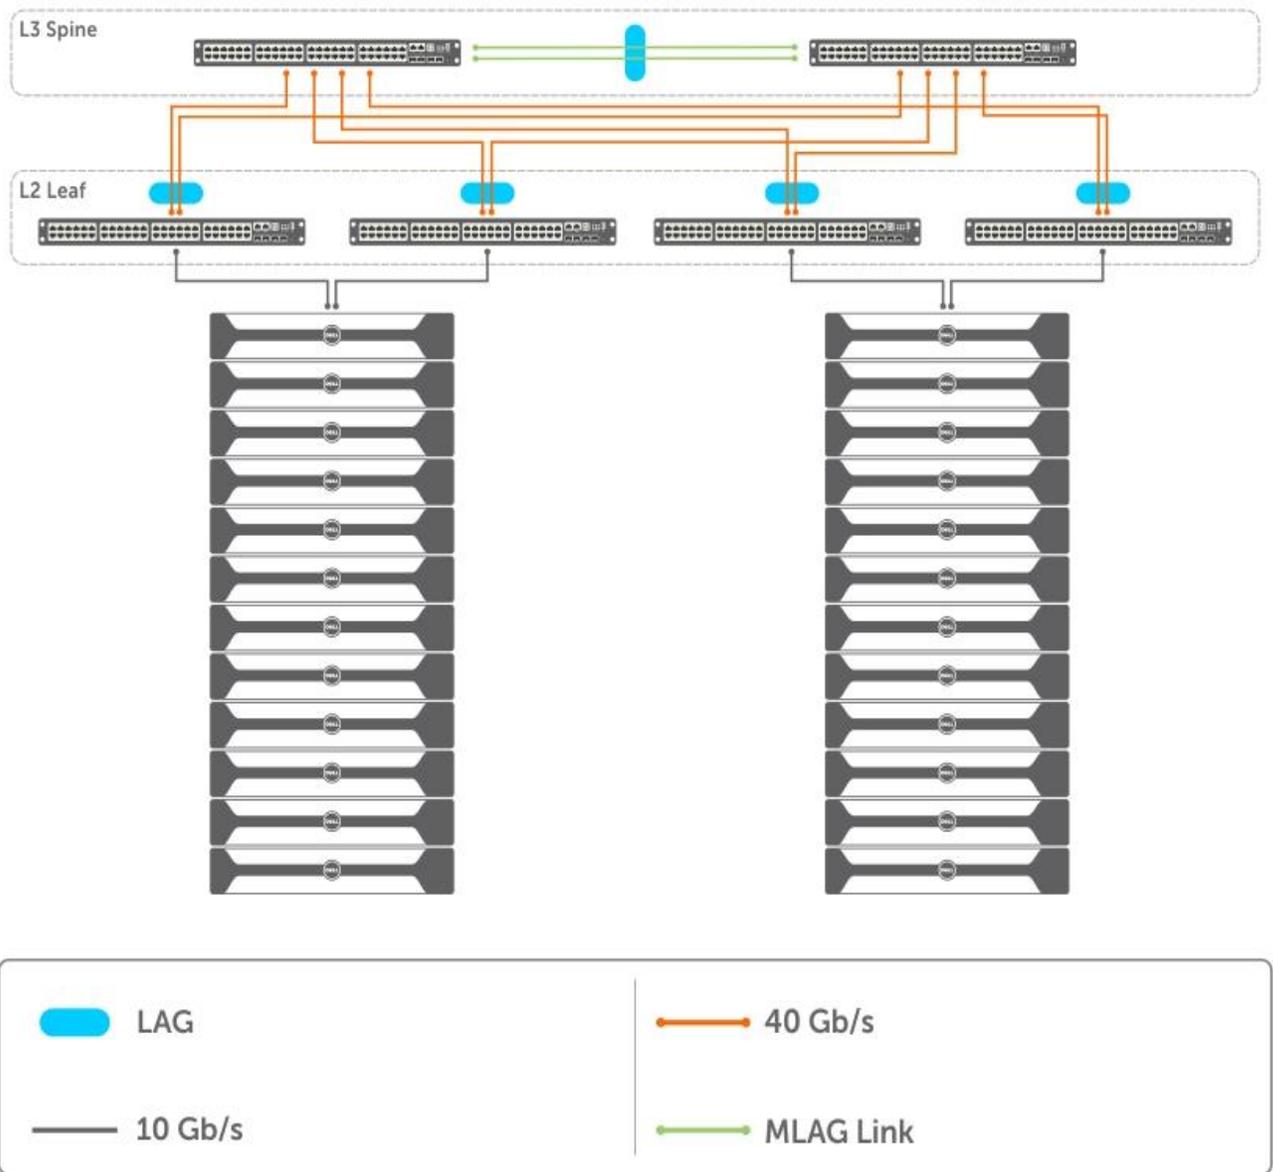


Figure 9 Scale out Leaf-Spine network architecture

2.5.1.1 SIO Traffic Types

The software components that make up SIO converse with each other in a predictable way. The image below shows SDC, SDS & MDM groups since a SIO Cluster will have multiples of each software component.

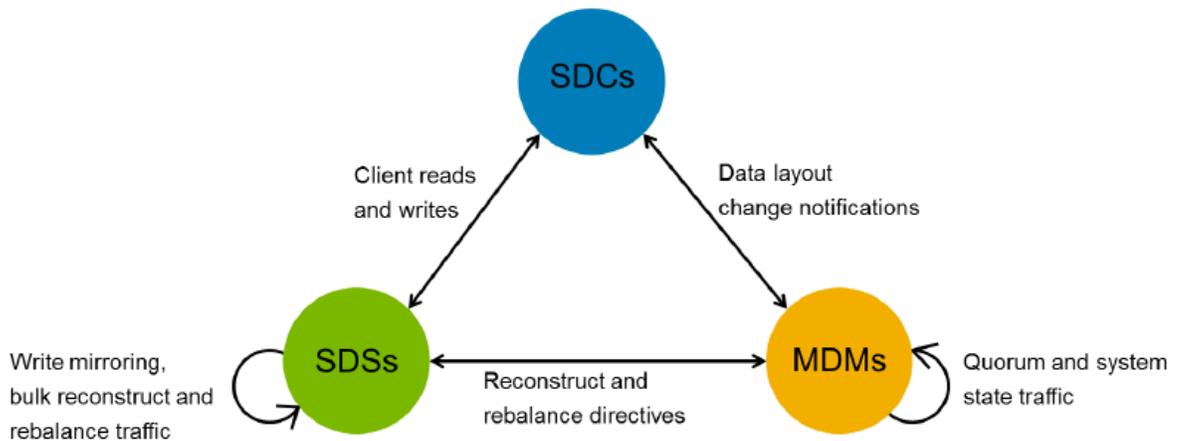


Figure 10 ScaleIO traffic types

ScaleIO Data Client (SDC) to ScaleIO Data Server (SDS)

Traffic between the SDCs and the SDSs forms the bulk of front end storage traffic. Front end storage traffic includes all read and write traffic arriving at or originating from a client. This network has a high throughput requirement. This network is used just for the ScaleIO read and writes operations. This network has a high throughput requirement. If there is a multitenancy requirement, ScaleIO SDC to SDS traffic can be isolated using VLANs and network firewalls.

ScaleIO Data Server (SDS) to ScaleIO Data Server (SDS)

Traffic between SDSs forms the bulk of back end storage traffic. Back end storage traffic includes writes that are mirrored between SDSs, rebalance traffic, and rebuild traffic. This network has a high throughput requirement. Although not required, there may be situations where isolating front-end and back-end traffic for the storage network may be ideal. This may be true in two-layer deployments where the storage and server teams act independently.

Meta Data Manager (MDM) to Meta Data Manager (MDM)

MDMs are used to coordinate operations inside the cluster. They issue directives to ScaleIO to rebalance, rebuild, and redirect traffic. MDMs are redundant, and must communicate with each other to maintain a shared understanding of data layout. MDMs also establish the notion of quorum in ScaleIO.

MDMs do not carry or directly interfere with I/O traffic. MDMs do not require the same level of throughput required for SDS or SDC traffic. MDM to MDM traffic requires a reliable, low latency network. MDM to MDM traffic is considered back end storage traffic.

Meta Data Manager (MDM) to ScaleIO Data Client (SDC)

The master MDM must communicate with SDCs if data layout changes. This can occur because the SDSs that host storage for the SDCs are added, removed, placed in maintenance mode, or go offline.

Communication between the Master MDM and the SDCs is asynchronous. MDM to SDC traffic requires a reliable, low latency network. MDM to SDC traffic is considered “front end” storage traffic.

Meta Data Manager (MDM) to ScaleIO Data Server (SDS)

The master MDM must communicate with SDSs to issue rebalance and rebuild directives. MDM to SDS traffic requires a reliable, low latency network. MDM to SDS traffic is considered “back end” storage traffic.

Management traffic

This network is optional but recommended. These IP addresses can be used to provide access to ScaleIO management applications like ScaleIO Gateway (REST Gateway, Installation Manager, and SNMP trap sender), traffic to and from the Light Installation Agent (LIA), and reporting or management traffic to the MDMs (such as syslog for reporting and LDAP for administrator authentication).

It should be noted that the management network does not come in the control plane data, which means that the ScaleIO system does not require nor use this network interface for any ScaleIO internal tasks.

For more information on networking best practices for SIO click [here](#)

2.5.2 SIO Ready Nodes – Enterprise solution pods

The compute, management and storage layers are converged into each SIO ready node in the cluster, hosting VMware vSphere. The recommended boundaries of an individual pod are based on number of nodes supported for a vSphere 6.5 cluster which is 64 nodes, although the SIO cluster can scale much larger, well beyond the boundaries of the hypervisor in use.

Dell EMC recommends that the VDI management infrastructure nodes be separated from the compute resources onto their own ready node cluster with a common VMFS namespace shared between them. One node for VDI management is required, minimally, and expanded based on size of the pod. The designations rdsh, compute, vgpu and mgmt as seen below are logical VMFS containers used to group VMs of a particular type.

Using distinct containers allows features and attributes, such as compression and deduplication, to be applied to groups of VMs that share similar characteristics. Compute hosts can be used interchangeably for Horizon Apps or Microsoft Remote Desktop Session Hosts (RDSH) as required. Distinct clusters should be built for management and compute hosts for HA, respectively, to plan predictable failover, scale and load across the pod. The VMFS namespace can be shared across multiple hypervisor clusters adding disk capacity and performance for each distinct cluster.

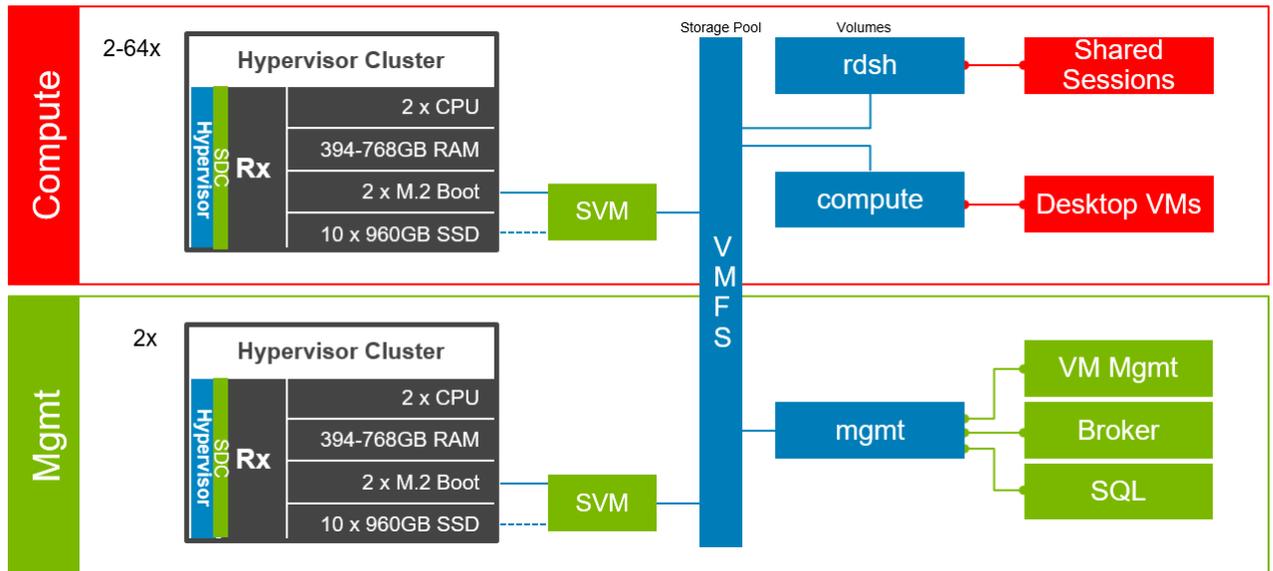


Figure 11 Shared VMFS namespace

High-performance graphics capabilities compliment the solution and can be added at any time to any new or existing SIO Ready Node vSphere-based deployment. Simply add the appropriate number of R740xd appliances to your SIO cluster and provide a superior user experience with vSphere 6 and NVIDIA GRID vGPU technology. The R740xd SIO Ready Node can be utilized for the non-graphics compute or management portions of this solution and vSphere will provide HA accordingly based on the type of VM.

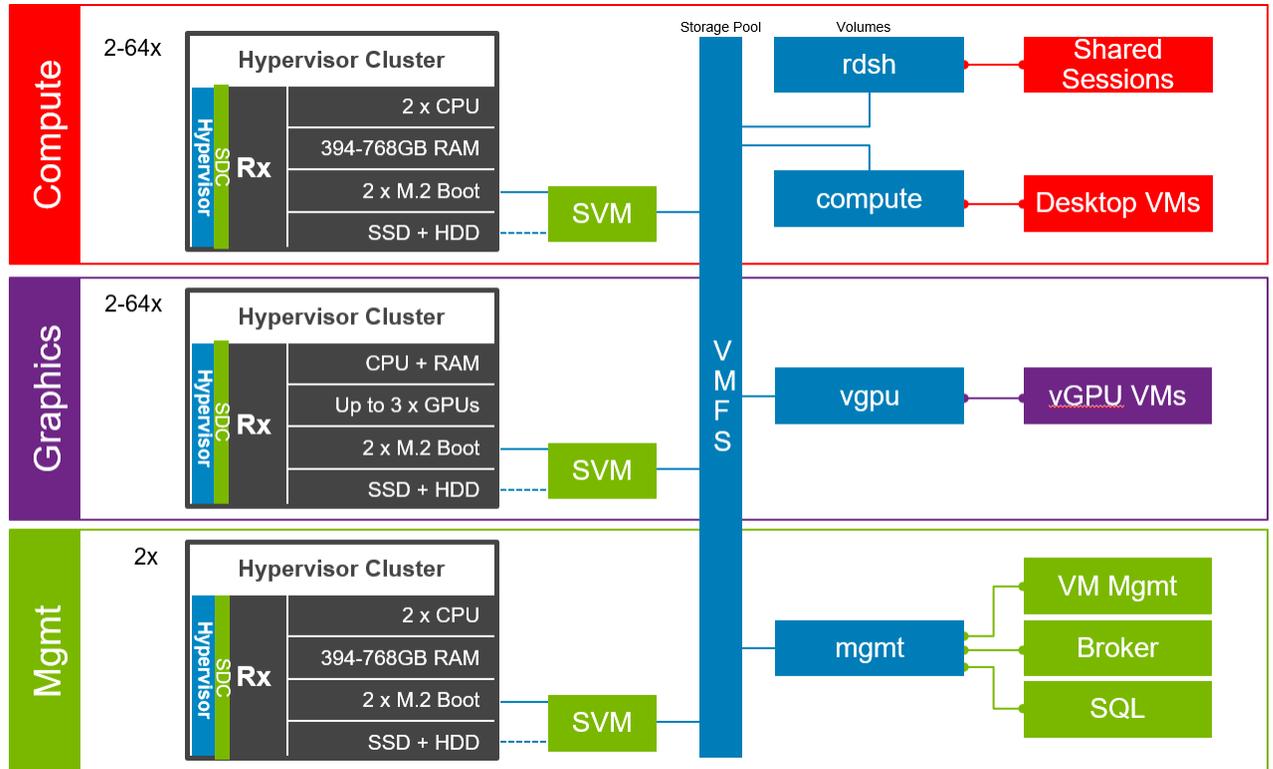


Figure 12 ScaleIO Ready Node deployment with vGPU integrated

3 Hardware components

3.1 Network

The following sections contain the core network components for the solution. General uplink cabling guidance to consider in all cases is that TwinAx or CAT6 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)

Accelerate applications in high-performance environments with 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 1 Dell Networking S3048 features

Model	Features	Options	Uses
<ul style="list-style-type: none">Dell Networking S3048-ON	<ul style="list-style-type: none">48 x 1000BaseT4 x 10Gb SFP+Non-blocking, line-rate performance260Gbps full-duplex bandwidth131 Mpps forwarding rate	<ul style="list-style-type: none">Redundant hot-swap PSUs & fansVRF-lite, Routed VLT, VLT Proxy GatewayUser port stacking (up to 6 switches)Open Networking Install Environment (ONIE)	<ul style="list-style-type: none">1Gb connectivity

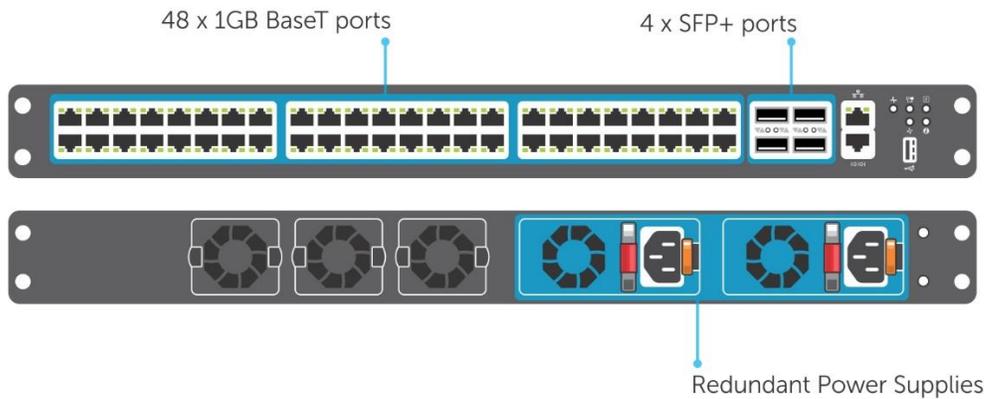


Figure 13 Dell Networking S3048 port diagram

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 10Gb SFP+ and 6 x 40Gb QSFP+ 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.

Table 2 Dell Networking S4048

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

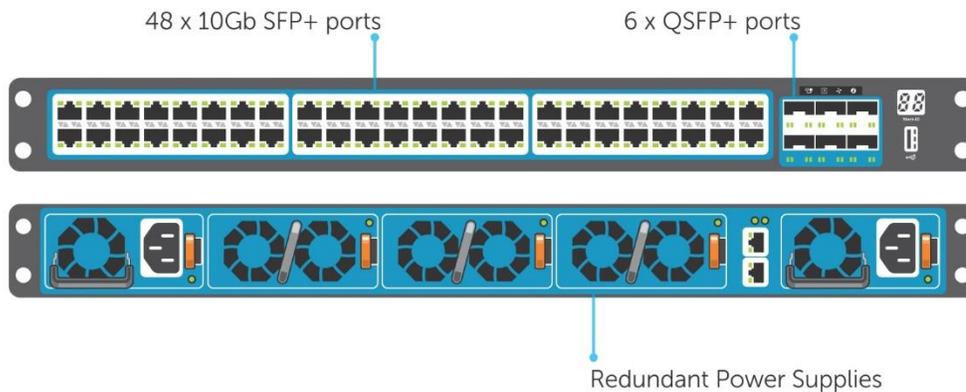


Figure 14 Dell Networking S4048 port diagram

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

3.2 Dell EMC SIO Ready Nodes

Dell EMC SIO Ready Nodes are based on the award-winning 14th generation of Dell EMC PowerEdge servers which offer many performance and feature enhancements. The table below outlines the hardware changes between generations.

Table 3 Hardware changes comparison

	R630	R640	R730xd	R740xd	730 -> 740 % Increase
CPU and chipset	Broadwell-EP	Skylake	Broadwell-EP	Skylake	
Front side bus	Intel QuickPath Interconnect @ 9.6 GT/s	Intel UltraPath Interconnect @ 11.2 GT/s	Intel QuickPath Interconnect @ 9.6 GT/s	Intel UltraPath Interconnect @ 10.4 GT/s	8%
Cores (max)	18 cores	28 cores	22 cores	28 cores	27%
TDP (max)	145 W	205 W	145 W	205 W	
Instruction set	AVX2	AVX2/ AVX-512	AVX2	AVX2/ AVX-512	
Max DP FLOPS / CLK	16 per core (w /AVX2)	32 per core (w / AVX-512)	16 per core (w /AVX2)	32 per core (w / AVX-512)	100%
Memory channels per socket	4 channels, DDR4	6 channels, DDR4	4 channels, DDR4	6 channels, DDR4	50%
Memory (max)	384 GB/ socket (768 GB total)	768 GB/ socket (1.5 TB total)	768 GB/ socket (1.5 TB total)	1.5 TB / socket (3 TB total)	100%
Memory speed (max)	2133 MT/s	2667 MT/s	2400 MT/s	2667 MT/s	11%
PCIe Lanes	40	48	40	48	20%

Consolidate compute and storage into a single chassis with SIO Ready Node, powered by Dell EMC SIO software. SIO Ready Nodes install quickly, integrate easily into any data center, and can be deployed for multiple virtualized workloads including desktop virtualization, test and development, and private cloud projects. For general purpose, virtual desktop and virtual application solutions, Dell EMC recommends the R640 and R740xd SIO Ready Nodes. For workloads requiring graphics the SIO R740xd with NVIDIA GRID vGPU can be integrated into any environment running any other SIO Ready Node Deployment.

The SIO ready node portfolio, optimized for VDI, has been designed and arranged in two top-level optimized configurations which apply to the available physical platforms showcased below.

- **B5** configuration is geared toward larger scale general purpose workloads, balancing performance and cost-effectiveness.
- **C7** is the premium configuration offering an abundance of high performance and tiered capacity where user density is maximized.

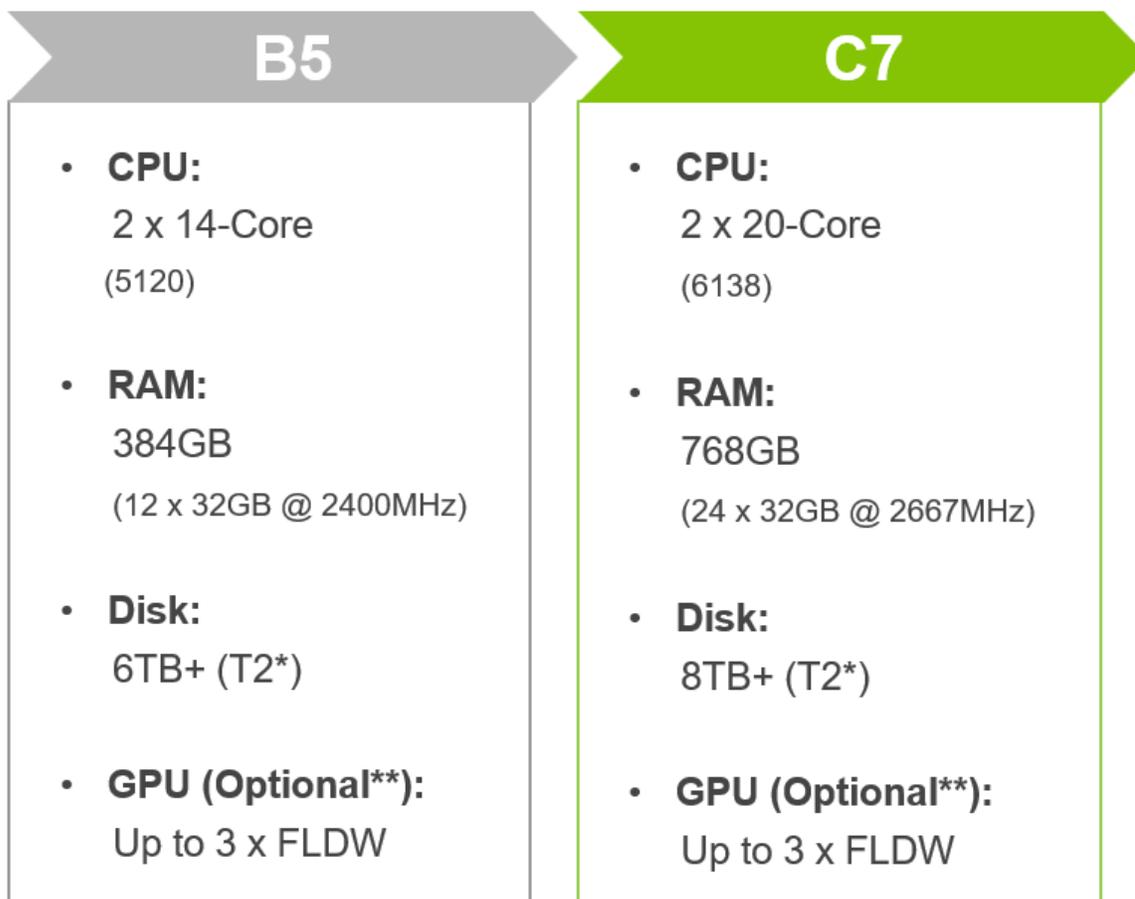


Figure 15 ScaleIO B5 and C7 configurations

3.2.1 Dell EMC SIO R640

The Dell EMC SIO R640 is a 10-disk 1U platform with a broad range of configuration options. Each appliance comes equipped with dual CPUs, 10 to 28 cores, and up to 1.5TB of high-performance RAM. A minimum of 5 x SSD are required in an All-Flash configuration. The M.2-based BOSS module boots the hypervisor and SIO Controller VM (SVM) while the PERC HBA330 connects the SVM to the SSDs. Each platform can be outfitted with SFP+ or BaseT NICs.

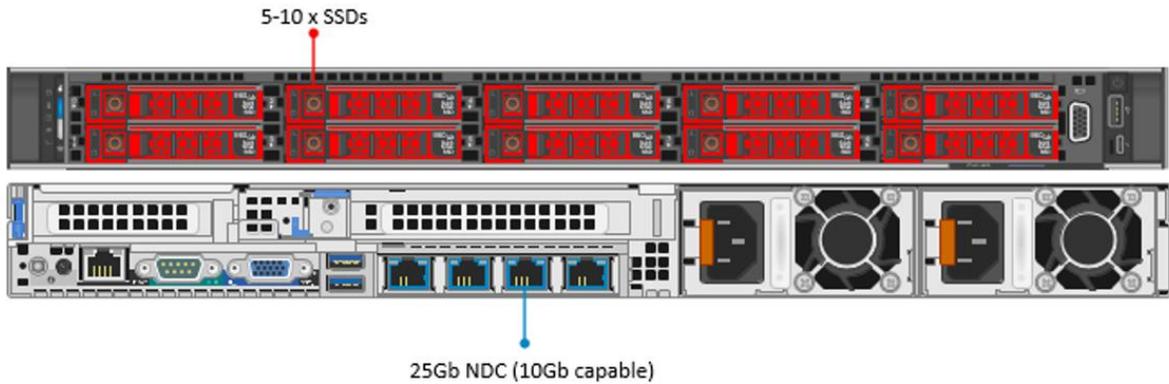


Figure 16 ScaleIO R640 port diagram

Table 4 R640 B5 and C7 configurations

R640 AF	B5	C7
CPU	2 x Intel Xeon Gold 5120 (14C, 2.2GHz)	2 x Intel Gold 6138 (20C, 2.0GHz)
Memory	12 x 32GB 2667MT/s RDIMMs Effective speed: 2400MT/s @ 384GB	24 x 32GB 2667MT/s RDIMMs Effective speed: 2667MT/s @ 768GB
Storage Ctrl	HBA330	
Storage	2 x 120GB M.2 BOSS in RAID1 8 x 960GB SSD 2.5" (T1)	2 x 120GB M.2 BOSS in RAID1 10 x 960GB SSD 2.5" (T1)
Network	4 x 10Gb/25Gb SFP+/ BT	
iDRAC	iDRAC9 Enterprise	
Power	2 x 1100W PSUs*	

3.2.2 Dell EMC SIO R740xd

The Dell EMC SIO 740xd is a 2U platform that can be configured with 24 x 2.5" disks to serve a broad range of capacity requirements. Each appliance comes equipped with dual CPUs, 10 to 28 cores, and up to 1.5TB of high-performance RAM. A minimum of 5 x SSD are required in an All-Flash configuration. The M.2-based BOSS module boots the hypervisor and SIO Controller VM (SVM) while the PERC HBA330 connects the SVM to the SSDs. Each platform can be outfitted with SFP+ or BaseT NICs. The R740xd can support up to 3 NVIDIA M60 or 2 x M10 GPU cards.

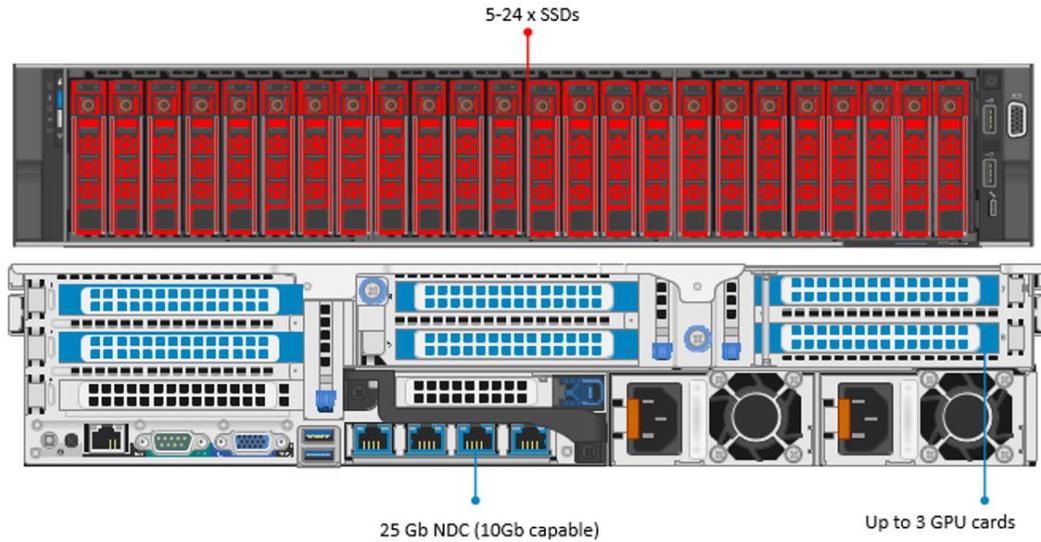


Figure 17 ScaleIO R740xd port diagram

Table 5 R740xd B5 & C7 configurations

R740xd AF	B5	C7
CPU	2 x Intel Xeon Gold 5120 (14C, 2.2GHz)	2 x Intel Gold 6138 (20C, 2.0GHz)
Memory	12 x 32GB 2667MT/s RDIMMs Effective speed: 2400MT/s @ 384GB	24 x 32GB 2667MT/s RDIMMs Effective speed: 2667MT/s @ 768GB
Storage Ctrl	HBA330	
Storage	2 x 120GB M.2 BOSS in RAID1 8 x 960GB SSD 2.5" (T1)	2 x 120GB M.2 BOSS in RAID1 10 x 960GB SSD 2.5" (T1)
GPU	2 x Tesla M10 or 3 x Tesla M60	
Network	4 x 10Gb/25Gb SFP+/ BT	
iDRAC	iDRAC9 Enterprise	
Power	2 x 1100W PSUs*	

*Higher wattage PSUs required when GPUs are in use

3.3 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 SIO R740XD platform.

3.3.1 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 using a 1GB framebuffer (up to 128 desktops per server) and gives businesses the power to deliver great graphics experiences to all their employees at an affordable cost.

Table 6 Tesla M10 specifications

Specs	Tesla M10
Number of GPUs/ card	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

3.3.2 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. M60 can support 3 cards in the R740xd providing 48 x Windows10 users assigned a 1GB framebuffer each.

Table 7 Tesla M60 specifications

Specs	Tesla M60
Number of GPUs/ card	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

3.4 Dell Wyse Endpoints

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

3.4.1 Wyse 3040 Thin Client (ThinOS, ThinLinux)

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



3.4.2 Wyse 5040 AIO Thin Client (ThinOS)



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

remote management from a simple file server, Wyse Device Manager (WDM), or Wyse Management Suite can help lower your total cost of ownership as you grow from just a few thin clients to tens of thousands. For more information, please visit: [Link](#)

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

The Wyse 5060 offers high performance and reliability, featuring all the security and management benefits of Dell thin clients. It comes with flexible OS options: ThinOS (with or without PCoIP), ThinLinux, Windows Embedded Standard 7P (WES7P) or Windows 10 IoT Enterprise (WIE10). Designed for knowledge workers demanding powerful virtual desktop performance, and support for unified communications solutions like Skype for Business, the Wyse 5060 thin client delivers the flexibility, efficiency and security organizations require for their cloud environments. It is powered by a quad-core AMD 2.4GHz processor, supports dual 4K (3840x2160) monitors and provides multiple connectivity options with six USB ports, two of which are USB 3.0 for high-speed peripherals, as well as two DisplayPort connectors, wired networking or wireless 802.11 a/b/g/n/ac. The Wyse 5060 can be monitored, maintained, and serviced remotely via Wyse Device Manager (WDM), cloud-based Wyse Management Suite or Microsoft SCCM (5060 with Windows versions). For more information, please visit: [Link](#)



3.4.4 Wyse 7020 Thin Client (WES 7/7P/8, 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/8 (WES), Windows 10 IoT Enterprise (WIE10), Wyse ThinLinux operating systems and it supports a broad range of fast, flexible connectivity options so that users can connect their favorite peripherals while working with processing-intensive, graphics-rich applications. This 64-bit thin client delivers a great user experience and support for local applications while ensuring security. Designed to provide a superior user experience, ThinLinux features broad broker support including Citrix Receiver, VMware Horizon and Amazon Workspace, and support for unified communication platforms including Skype for Business, Lync 2013 and Lync 2010. For additional security, ThinLinux also supports single sign-on and VPN. With a powerful quad core AMD G Series APU in a compact chassis with dual-HD monitor support, the Wyse 7020 thin client delivers stunning performance and display capabilities across 2D, 3D and HD video applications. Its silent diskless and fan less design helps reduce power usage to just a fraction (it only consumes about 15 watts) of that used in traditional desktops. Wyse Device Manager (WDM) helps lower the total cost of ownership for large deployments and offers remote enterprise-wide management that scales from just a few to tens of thousands of cloud clients. For more information, please visit [Link](#)

3.4.5 Wyse 7040 Thin Client (WES7P, WIE10)

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



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

Designed to securely deliver virtual desktops and applications to mobile users who want to connect a broad range of peripherals, the Latitude 3480 and 5280 mobile thin clients run **Windows 10 IoT Enterprise**. They support a wide variety of connection brokers including Citrix XenDesktop/XenApp, Microsoft RDS and VMware Horizon right out of the box, and are an ideal alternative to much less secure Chromebooks. The Latitude 3480 features an Intel dual core processor with integrated graphics for a rich multimedia experience, and delivers excellent 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](#) and [Latitude 5280](#).



4 Software components

4.1 VMware vSphere 6.5

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 Wyse Datacenter 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 with the latest 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 the Blog article located [here](#).

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

4.2 VMware Horizon

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

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

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

The core Horizon components include:

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

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

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

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

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

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

Horizon Transfer Server – Manages data transfers between the data center and the Horizon desktops that are checked out on the end users' desktops in offline mode. This Server is required to support desktops that run the Horizon client with Local Mode options. Replications and synchronizing are the functions it will perform with offline images.

4.2.1 VMware Horizon Key Features

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

4.2.1.1 Just in time delivery with Instant Clone Technology

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

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

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

4.2.1.2 Transformational user experience with Blast Extreme

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

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

4.2.1.3 Modernize application lifecycle management with App Volumes

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

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

4.2.1.4 Smart policies with streamlined access

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

Policy-Managed Client Features, which enables IT to use policy to define which specific security-impacting features, are accessible upon login. These include clipboard redirection, USB, printing, and client-drives. These can be enforced contextually, based on role, evaluated at logon/logoff, disconnect/reconnect and at pre-determined refresh intervals for consistent application of policy across the entirety of the user experience. For example, a user logging in from a network location consider unsecured, can be denied access to USB

and printing. Additionally, PCoIP bandwidth profile settings allow IT to customize the user experience based on user context and location.

True SSO streamlines secure access to a Horizon desktop when users authenticate via VMware Identity Manager. A short lived VMware Horizon virtual certificate is generated, enabling a password-free Windows login, bypassing the usual secondary login prompt users would encounter before getting to their desktop.

4.2.2 VMware Horizon Apps

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

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

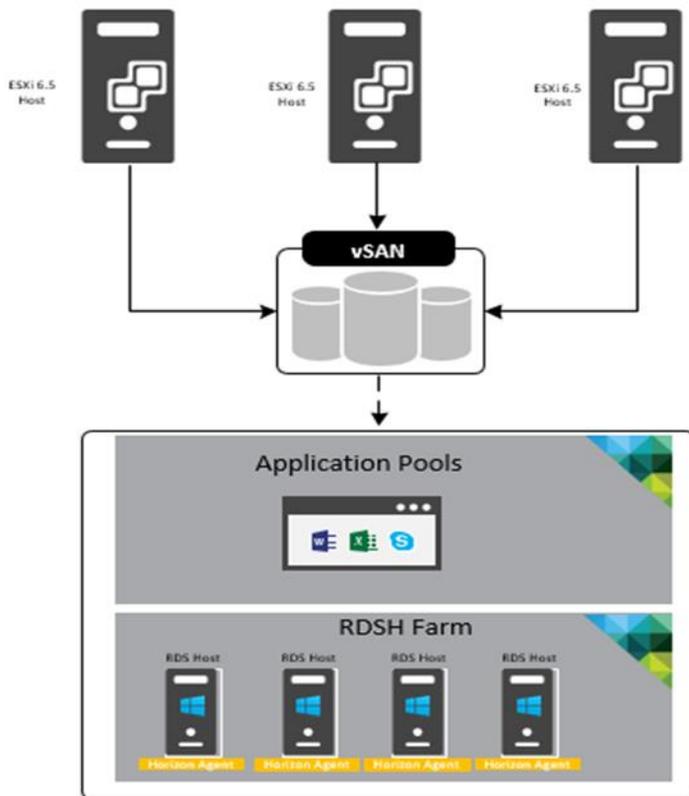


Figure 18 VMware Horizon Apps

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

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

Horizon JMP Apps are composed of:

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

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

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

4.2.3 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 UltraPath Interconnect (UPI). 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 ScaleIO Ready Nodes are listed in the below configurations. The SVM only consumes 2 cores but the image takes into account the possibility of one of the hosts on a SIO cluster also having the Gateway SVM which also consumes two cores.

4.2.3.1 SIO B5 NUMA alignment

14 physical cores per CPU in the SIO B5 configuration, 28 logical with Hyper-threading active, gives us a total of 56 consumable cores per appliance. The SVM only consumes 2 cores but the image take into account the possibility of one of the

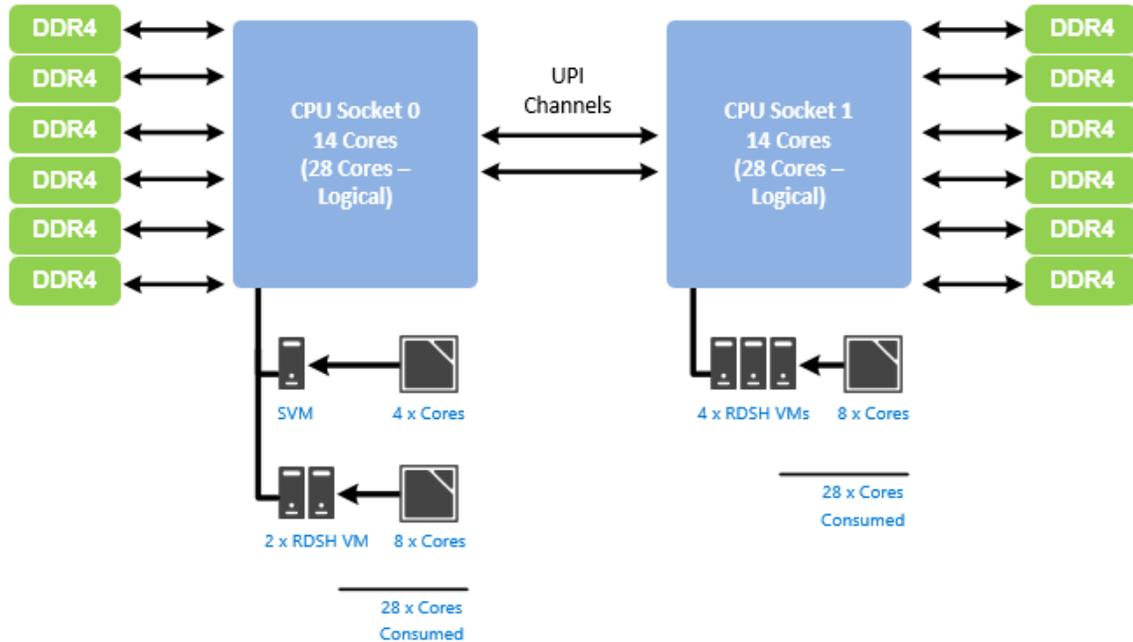


Figure 19 B5 NUMA alignment

4.2.3.2 SIO C7 NUMA alignment

20 physical cores per CPU in the SIO C7 configuration, 40 logical with Hyper-threading active, gives us a total of 80 consumable cores per appliance.

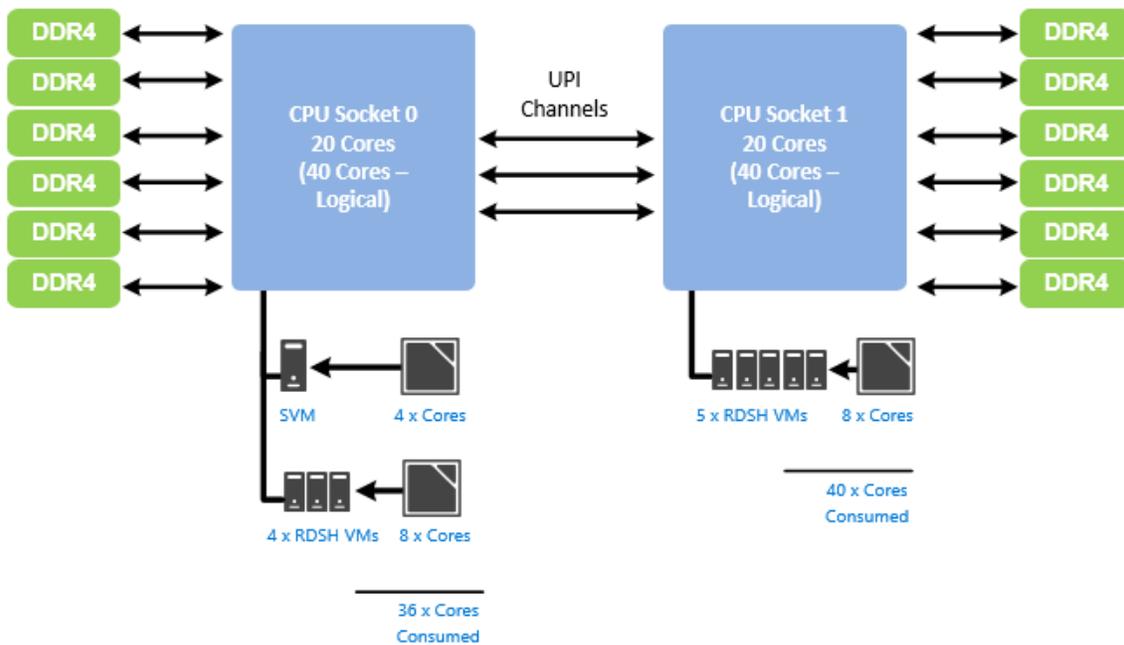


Figure 20 C7 NUMA alignment

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.

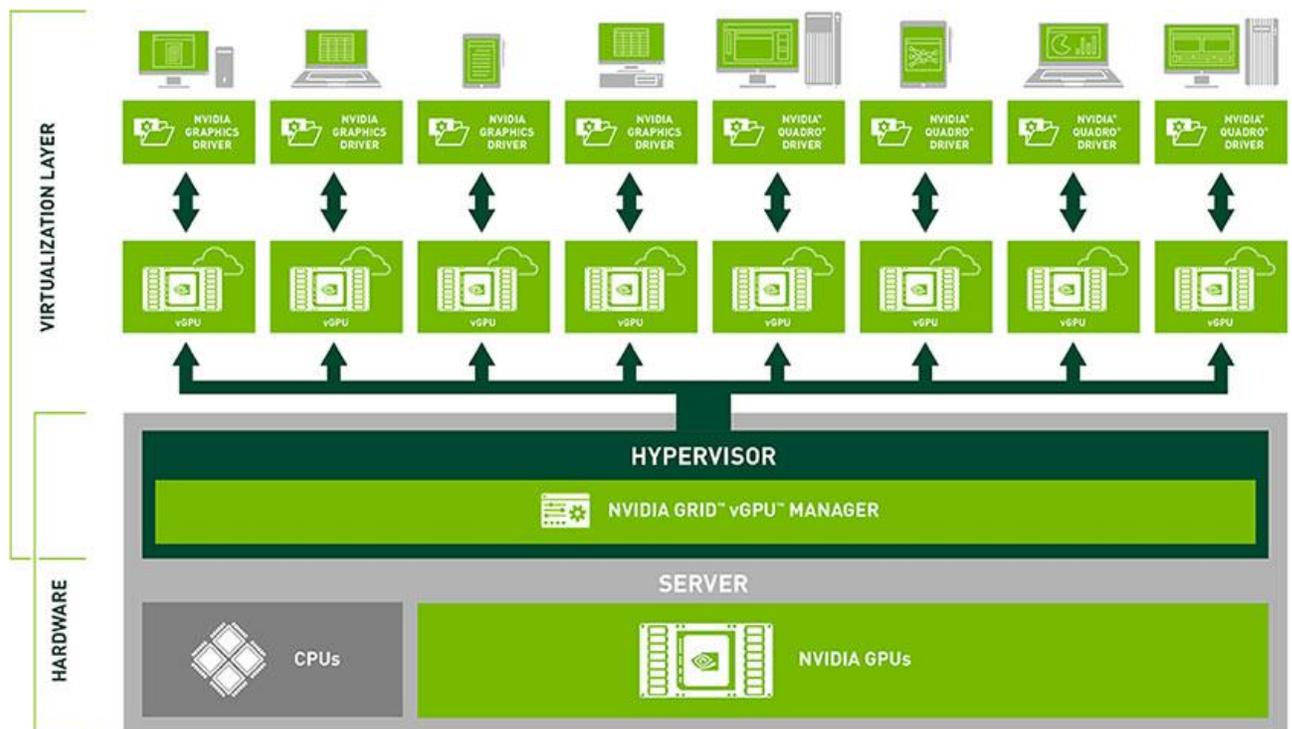


Figure 21 NVIDIA GRID vGPU

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 EMC servers, NVIDIA GRID vGPU technology and NVIDIA Tesla™ 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 2 x M10 or 3 x M60), vGPU profiles used for each GPU in a card, and GRID license type. The same profile must be used on a single GPU but profiles can differ across GPUs within a single card.

Table 8 NVIDIA® Tesla® M10 GRID vGPU Profiles:

Card	vGPU Profile	Graphics Memory (Frame Buffer)	Virtual Display Heads	Maximum Resolution	Maximum Graphics-Enabled VMs		
					Per GPU	Per Card	Per Server (2 cards)
Tesla M10	M10-8Q	8GB	4	4096x2160	1	4	8
	M10-4Q	4GB	4	4096x2160	2	8	16
	M10-2Q	2GB	4	4096x2160	4	16	32
	M10-1Q	1GB	2	4096x2160	8	32	64
	M10-0Q	512MB	2	2560x1600	16	64	128
	M10-1B	1GB	4	2560x1600	8	32	64
	M10-0B	512MB	2	2560x1600	16	64	128
	M10-8A	8GB	1	1280x1024	1	4	8
	M10-4A	4GB			2	8	16
	M10-2A	2GB			4	16	32
	M10-1A	1GB			8	32	64

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

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

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

Table 9 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
	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	24
	M60-1A	1GB			8	16	48

Card	vGPU Profile	Guest VM OS Supported*		License Required
		Win	64bit Linux	
Tesla M60	M60-8Q	●	●	NVIDIA Quadro Virtual Data Center 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.

Table 10 GRID vGPU licensing options



NVIDIA GRID Virtual Applications	NVIDIA GRID Virtual PC	NVIDIA Quadro Virtual Data Center Workstation
For organizations deploying 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.

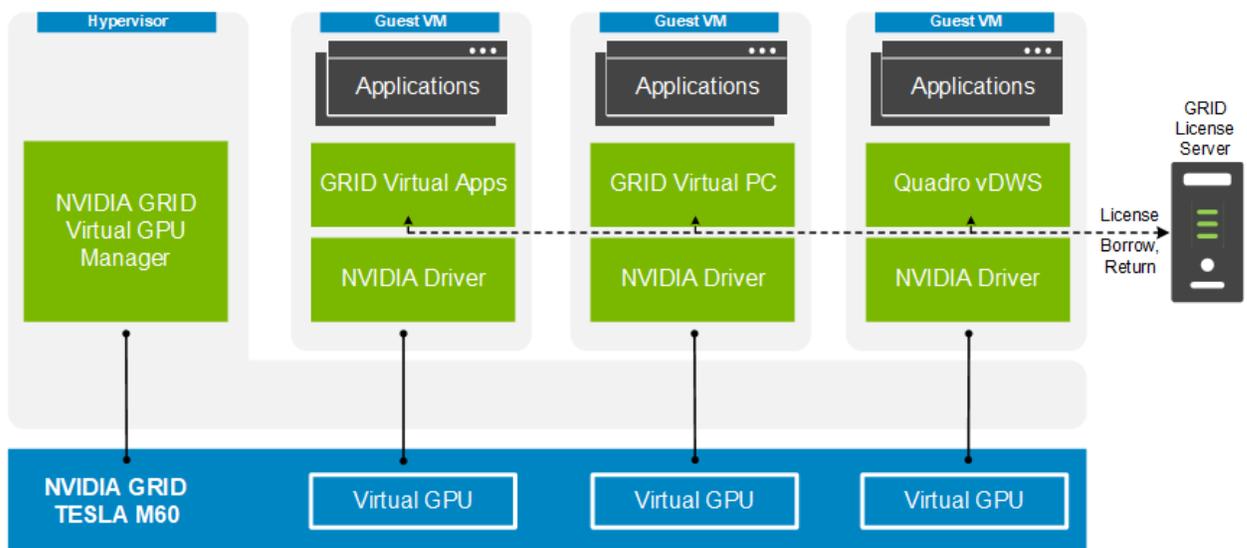


Figure 22 GRID vGPU Manager

5 Solution architecture for SIO with Horizon

There is the option to use an existing Virtual Center during the VxRail Appliance deployment but the sizing information below shows the details of the VC appliance and PSC that will be deployed during the factory install. This table represents all the management virtual machines in a three-node cluster so if the cluster was more than three nodes then the amount of SVM's will increase per node. There will be only one Gateway SVM per cluster.

Table 11 Management VMs in a three-node cluster

Role	vCPU	RAM (GB)	NIC	OS + Data vDisk (GB)	Tier 2 Volume (GB)
VMware vCenter Appliance	2	16	1	290	
Horizon Connection Server	2	8	1	40	-
SVM-GW	2	4	1	8	
SVM1	2	16	2	8	
SVM2	2	16	2	8	
SVM3	2	16	2	8	
SQL Server	4	8	1	40	210 (VMDK)
File Server	1	4	1	40	2048 (VMDK)
Total	17 vCPU	88GB	1 vNICs	442GB	2258GB

5.1.1 RDSH VM Configuration

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

Table 12 RDSH VM configuration on ESXi

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

5.1.2 NVIDIA GRID License Server Requirements

When using NVIDIA Tesla 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.

Table 13 NVIDIA GRID license server requirements

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

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

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

Additional license server requirements:

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

5.1.3 SQL databases

The VMware databases are hosted by a single dedicated SQL 2016 (or higher) Server VM in the Management layer. Use caution during database setup to ensure that SQL data, logs, and TempDB are properly separated onto their respective volumes. Create databases for:

- Horizon Connection Server

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

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

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

5.1.4 DNS

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

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

5.1.4.1 DNS for SQL

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

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

Figure 23 DNS CNAMEs

5.2 Storage architecture overview

SIO abstracts the local storage from each server, in this reference architecture we are focused on the All-Flash configuration so that would be SSDs, they are pooled together to provide a shared datastore across all Nodes in the SIO Cluster. The compute and management virtual machines are isolated into dedicated vmfs datastores.

5.3 Virtual Networking

The network configuration for the Dell EMC SIO Ready Nodes utilizes a 10Gb converged infrastructure model. All required VLANs will traverse 2 x 10Gb NICs configured in an active/active team. For larger scaling it is recommended to separate the infrastructure management VMs from the compute VMs to aid in predictable compute host scaling. The following outlines the suggested VLAN requirements for the Compute and Management hosts in this solution model:

- Compute hosts
 - Management VLAN: Configured for hypervisor infrastructure traffic – L3 routed via spine layer
 - Live Migration VLAN: Configured for Live Migration traffic – L2 switched via leaf layer
 - VDI VLAN: Configured for VDI session traffic – L3 routed via spine layer
 - Data VLAN: Configured for SIO data traffic- L2 switched via leaf layer
- Management hosts
 - Management VLAN: Configured for hypervisor Management traffic – L3 routed via spine layer
 - Live Migration VLAN: Configured for Live Migration traffic – L2 switched via leaf layer
 - VDI Management VLAN: Configured for VDI infrastructure traffic – L3 routed via spine layer
 - Data VLAN: Configured for SIO data traffic- L2 switched via leaf layer
- An iDRAC VLAN is configured for all hardware management traffic – L3 routed via spine layer

Both the compute and management host network configuration consists of a standard vSwitch teamed with 2 x 10Gb physical adapters assigned to VMNICs. The SVM connects to a private internal vSwitch to communicate directly with the hypervisor as well as the standard external vSwitch to communicate with other

SVMs in the cluster. All VDI infrastructure VMs connect through the primary port group on the external vSwitch.

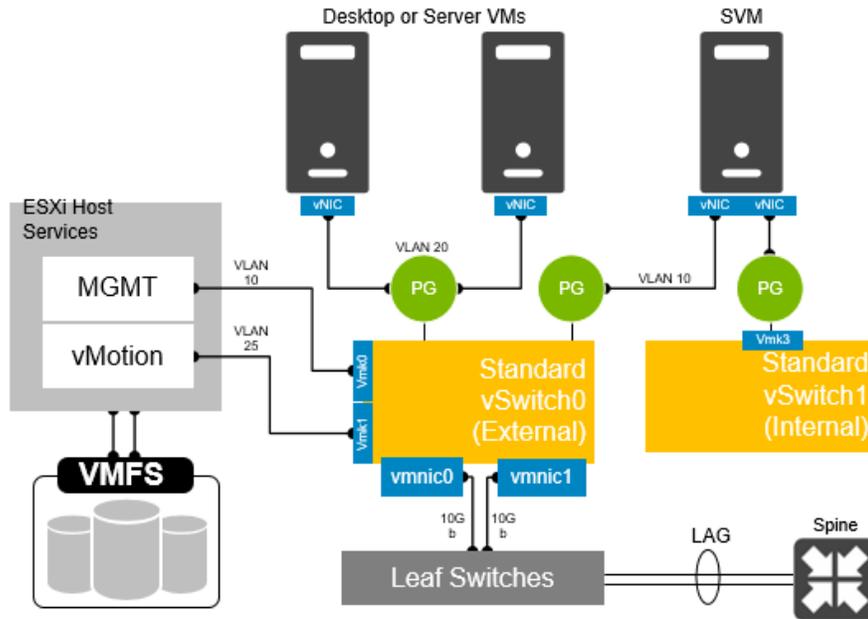


Figure 24 Virtual networking overview

5.4 Scaling guidance

Each component of the solution architecture scales independently according to the desired number of supported users. Additional appliance nodes can be added at any time to expand the SIO SDS pool in a modular fashion. While there is no scaling limit of the SIO architecture itself, practicality might suggest scaling pods based on the limits of hypervisor clusters which is 64 nodes for vSphere. Isolating management and compute to their own HA clusters provides more flexibility with regard to scaling and functional layer protection.

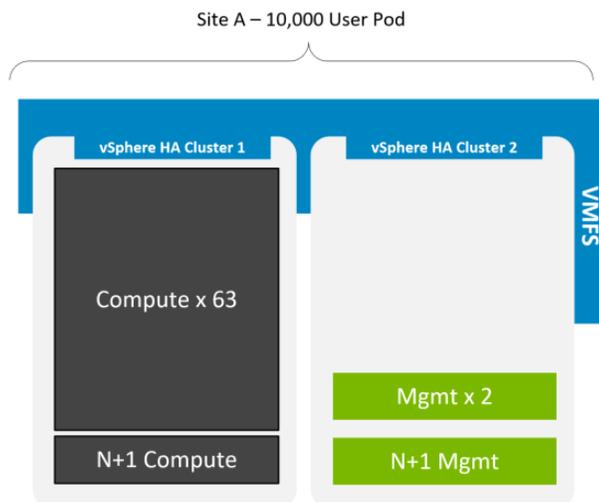


Figure 25 Scaling guidance – 10,000 user pod

Another option is to design a large single contiguous VMFS namespace with multiple hypervisor clusters within to provide single pane of glass management. For example, portrayed below is a large-scale user environment segmented by vSphere HA cluster and broker farm. Each farm compute instance is segmented into an HA cluster with a hot standby node providing N+1, served by a dedicated pair of management nodes per compute cluster in a separate HA cluster dedicated to mgmt. This provides multiple broker farms with separated HA protection while maintaining a single VMFS cluster across all nodes.

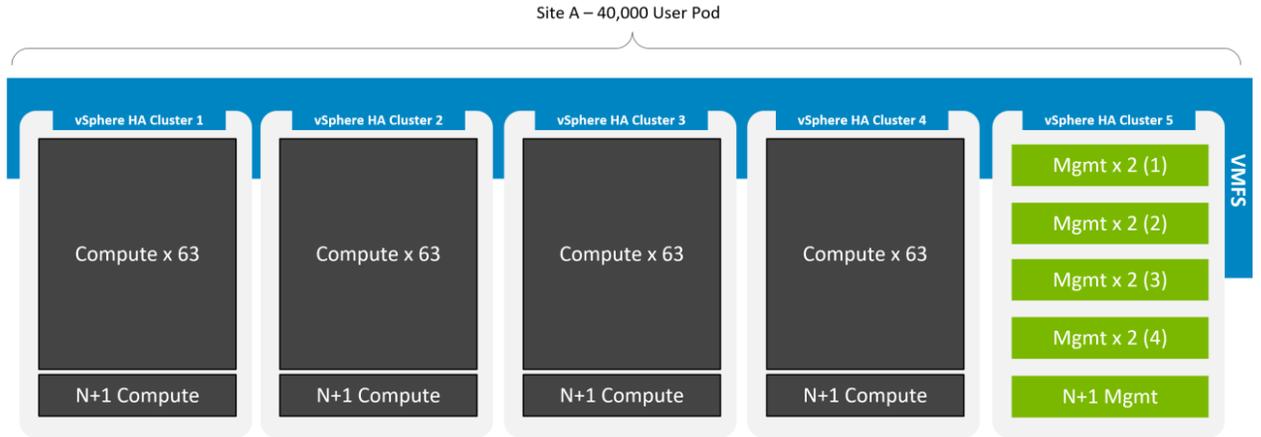


Figure 26 Scaling guidance – 40,000 user pod

- 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.

Table 14 Component scalability considerations

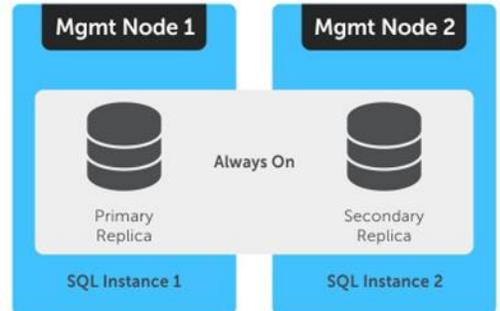
Component	Metric	Horizontal Scalability	Vertical Scalability
Virtual Desktop Host/Compute Servers	VMs per physical host	Additional hosts and clusters added as necessary	Additional RAM or CPU compute power
Composer	Desktops per instance	Additional physical servers added to the Management cluster to deal with additional management VMs.	Additional RAM or CPU compute power
Connection Servers	Desktops per instance	Additional physical servers added to the Management cluster to deal with additional management VMs.	Additional VCS Management VMs
VMware vCenter	VMs per physical host and/or ESX hosts per vCenter instance	Deploy additional servers and use linked mode to optimize management	Additional vCenter Management VMs
Database Services	Concurrent connections, responsiveness of reads/ writes	Migrate databases to a dedicated SQL server and increase the number of management nodes	Additional RAM and CPU for the management nodes
File Services	Concurrent connections, responsiveness of reads/ writes	Split user profiles and home directories between multiple file servers in the cluster. File services can also be	Additional RAM and CPU for the management nodes

		migrated to the optional NAS device to provide high availability.	
--	--	---	--

5.5 Solution 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: [SQL Server AlwaysOn Availability Groups](#) and [Windows Server Failover Clustering with SQL Server](#)



5.6 Communication flow for Horizon

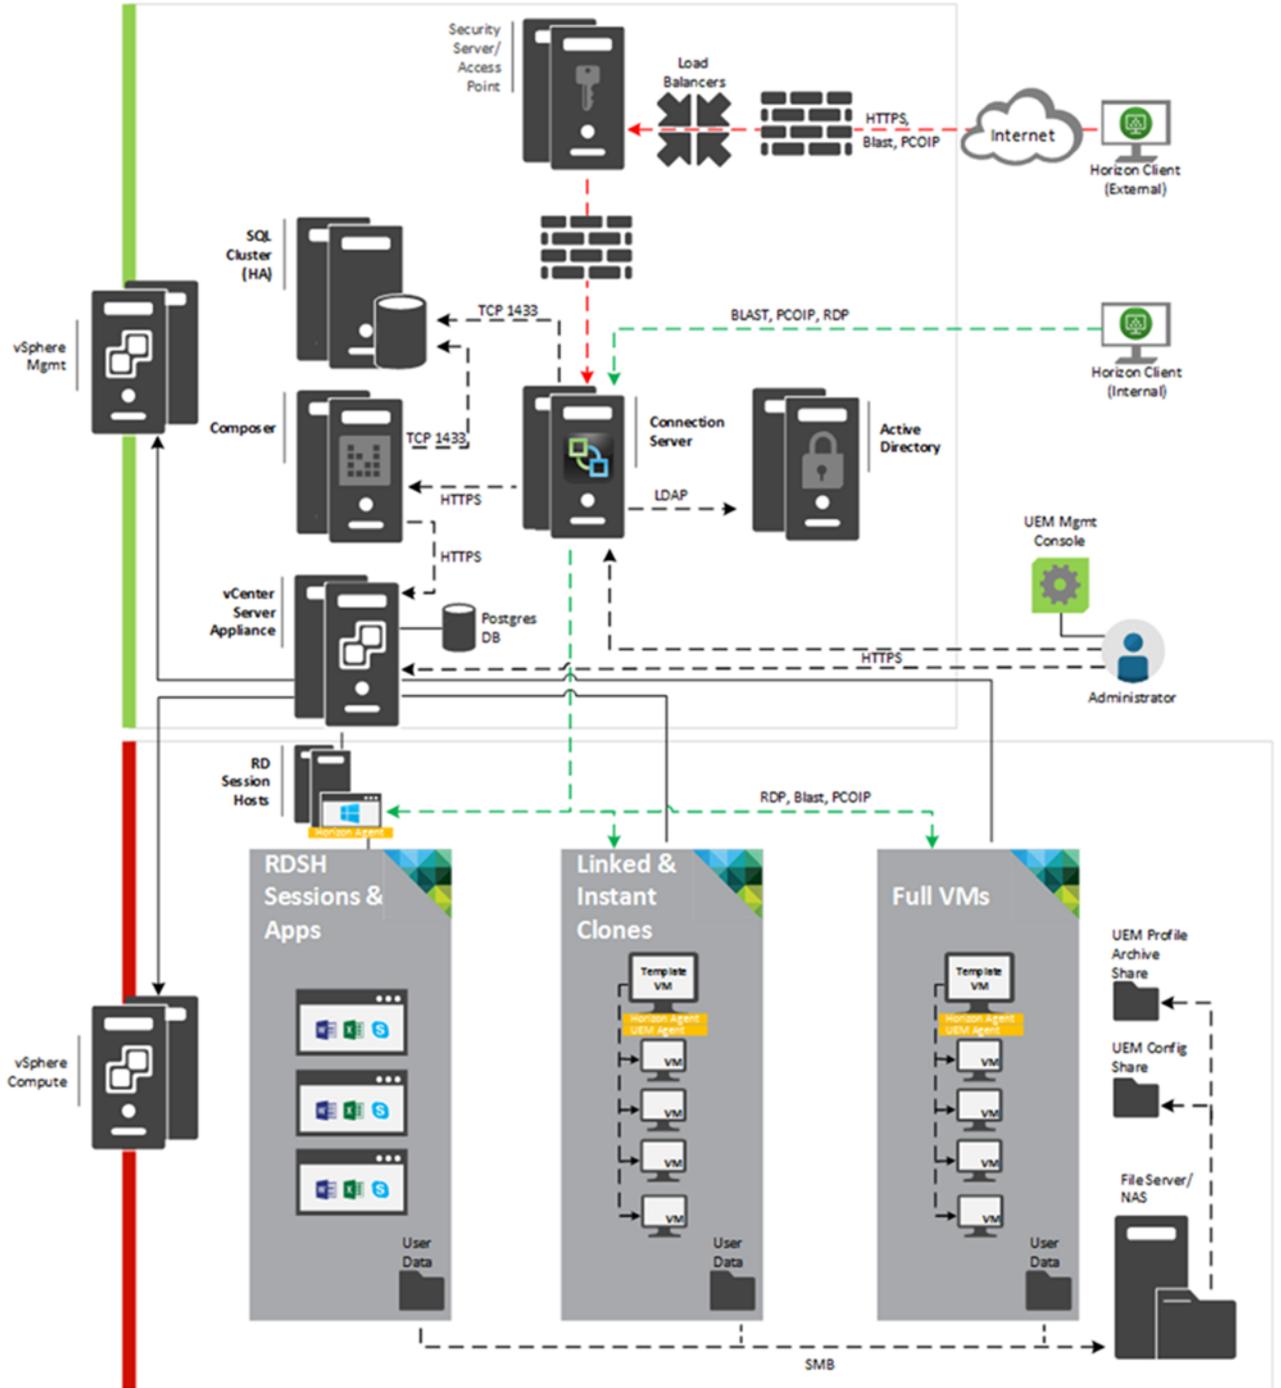


Figure 27 Communication flow for Horizon

6 Solution performance and testing

6.1 Summary

At the time of publication, these are the available density recommendations per ready node. Please refer to the [Platform Configurations](#) section for hardware specifications.

Table 15 User density summary

Host Config	Hypervisor	Broker & Provisioning	Workload	Template	User Density
SIORN R740XD-B5	ESXi 6.5	Horizon 7 & Instant Clone	Task Worker	Windows 10 & Office 2016	230
SIORN R740XD-B5	ESXi 6.5	Horizon 7 & Instant Clone	Knowledge Worker	Windows 10 & Office 2016	140
SIORN R740XD-B5	ESXi 6.5	Horizon 7 & Instant Clone	Power Worker	Windows 10 & Office 2016	115
SIORN R740XD-B5	ESXi 6.5	Horizon 7 & Instant Clone	M60-1Q	Windows 10 & Office 2016	48
SIORN R740XD-B5	ESXi 6.5	Horizon 7 & Instant Clone	M10-1Q	Windows 10 & Office 2016	64

Dell EMC is aware of the new side-channel analysis vulnerabilities, known as Meltdown and Spectre, affecting many modern microprocessors that were discovered and published by a team of security researchers on January 3, 2018. Further information is available at the following locations:

- https://emcservice--c.na55.visual.force.com/apex/KB_Security_KB?id=kA6f1000000FD0g
- <http://www.dell.com/support/article/SLN308588>
- <http://www.dell.com/support/article/SLN308587>

Horizon and vSphere management roles were also deployed on the cluster on a single host that also hosted desktops, this required a reduction in the amount of VMs that can be placed on the management host during the power user workload, 105 VMs were placed on the management host with 115 on the dedicated compute hosts. The user density numbers in the table are per host and the test information section 6.4 show the full cluster statistics.

6.2 Test and performance analysis methodology

6.2.1 Testing process

To ensure the optimal combination of end-user experience (EUE) and cost-per-user, performance analysis and characterization (PAAC) on Dell EMC 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 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 various 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 hour. Exception to this login timeframe occurs when testing low density solutions such as GPU/graphics based configurations. With those configurations, users are logged on every 10-15 seconds.
- All desktops are pre-booted in advance of logins commencing.
- All desktops run an industry-standard anti-virus solution. Windows Defender is used for Windows 10 due to issues implementing McAfee.

6.2.1.2 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 EMC, 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 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:

Table 16 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 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 multiples 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 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 RA.

Table 17 Resource utilization thresholds

Parameter	Pass/Fail Threshold
Physical Host CPU Utilization	85%
Physical Host Memory Utilization	85%
Network Throughput	85%
Storage IO Latency	20ms

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

6.3 Test configuration details

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

Table 18 SIO Series hardware and software test components

Component	Description/Version
Hardware platform(s)	Dell EMC SIORN 740xd B5
Hypervisor(s)	VMware vSphere ESXi 6.5
Broker technology	Horizon 7.3.2
Broker database	Microsoft SQL 2016
Management VM OS	Windows Server 2012 R2 (Connection Server & Database)
Virtual desktop OS	Microsoft Windows 10 64-bit
Office application suite	Microsoft Office 2016 Professional Plus
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.

Table 19 ESXi Desktop VM specifications

User Profile	vCPUs	ESXi Memory Configured	ESXi Memory Reservation	Screen Resolution	Operating System
Task Worker	2	2GB	1GB	1280 X 720	Windows 10 Enterprise 64-bit
Knowledge Worker	2	3GB	1.5GB	1920 X 1080	Windows 10 Enterprise 64-bit
Power Worker	2	4GB	2GB	1920 X 1080	Windows 10 Enterprise 64-bit

6.3.2 Platform configurations

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

Table 20 SIO R740XD B5 hardware configuration

Enterprise Platform	Platform Config	CPU	Memory	RAID Ctr	BOSS	HD Config	Network
R740XD	B5	Gold 5120 (14 Core, 2.2GHz)	384GB @2400 MT/s	HBA 330 Mini	2 x 120GB M.2	8x 960 GB SSD	Mellanox ConnectX-4 LX 25GbE SFP Rack NDC

Compute and Management resources were split out with the following configuration across a three node cluster and all test runs were completed with this configuration.

- Node 1 – PE-R740xd – Compute & ScaleIO Management. User & Management VMs.
- Node 2 – PE-R740xd – Dedicated Compute & ScaleIO Management. User VMs only.
- Node 3 – PE-R740xd – Dedicated Compute & ScaleIO Management. User VMs only.

10GB networking was used for all PAAC testing.

Instead of dedicated nodes, SIO SVMs and VDI management roles were deployed on the cluster with desktop VMs, which reduced maximum density on that node. Each compute node was loaded with desktops to its maximum density; no failover capacity was reserved.

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.

Table 21 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
SIO R740XD-B5	ESXi 6.5	Horizon 7 & Instant Clone	Task Worker	230	87%	345GB	168GB	7
SIO R740XD-B5	ESXi 6.5	Horizon 7 & Instant Clone	Knowledge Worker	140	84%	355GB	151GB	12
SIO R740XD-B5	ESXi 6.5	Horizon 7 & Instant Clone	Power Worker	115	86%	363GB	155GB	17
SIO R740XD-B5 (M60Q-1Q)	ESXi 6.5	Horizon 7 & Instant Clone	Power Worker	48	50%	222GB	209GB	15
SIO R740XD-B5 (M10-1B)	ESXi 6.5	Horizon 7 & Instant Clone	Power Worker	64	66%	288GB	273GB	16

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

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

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

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

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

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

6.4.1 Task Worker, 230 users, ESXi 6.5, Horizon 7.x Instant Clones

The below graph shows the performance data for 230 user sessions per host or 690 sessions per cluster. The CPU reaches a steady state average of 87 % across the three compute hosts during the test cycle when 230 users were logged on to each compute host. The CPU usage was approximately 11% on the hosts before the start of test.

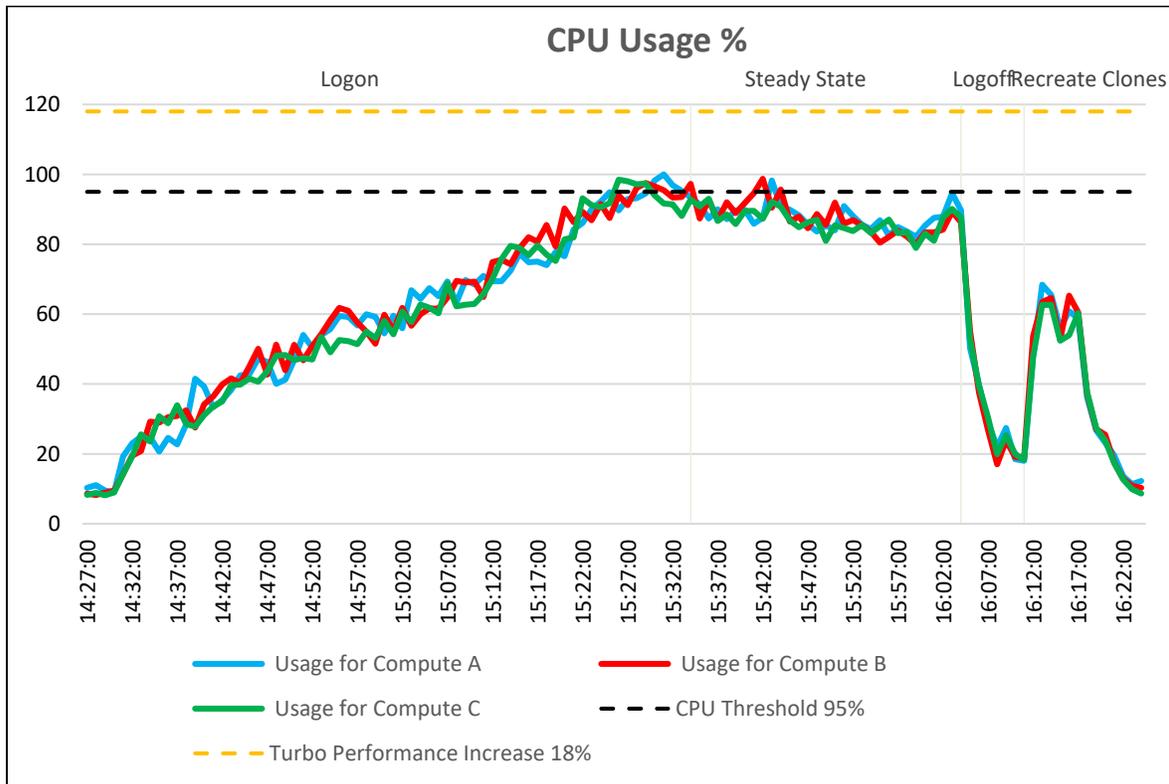


Figure 28 CPU usage

In regard to memory consumption for the cluster, out of a total of 384 GB available memory per node, memory usage was not pushed close to its maximum usage. The compute hosts reached a maximum memory consumption of 317 GB with active memory usage reaching a max of 260 GB during the recreation of the instant clones. A small amount of ballooning ~ 2 GB on ESXi A host and ballooning of ~ 1 GB host took place on all hosts for approximately 2 minutes prior to steady state. Some memory swapping occurred on ESX A where the ScaleIO Gateway SVM resides.

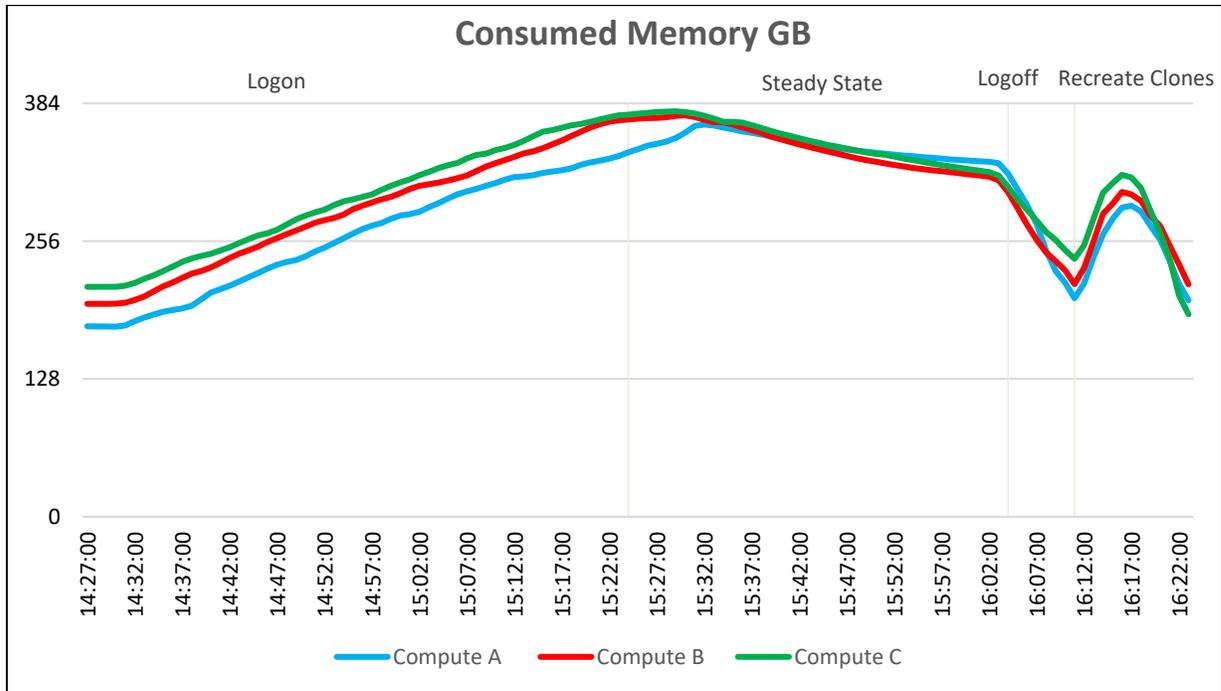


Figure 29 Consumed memory

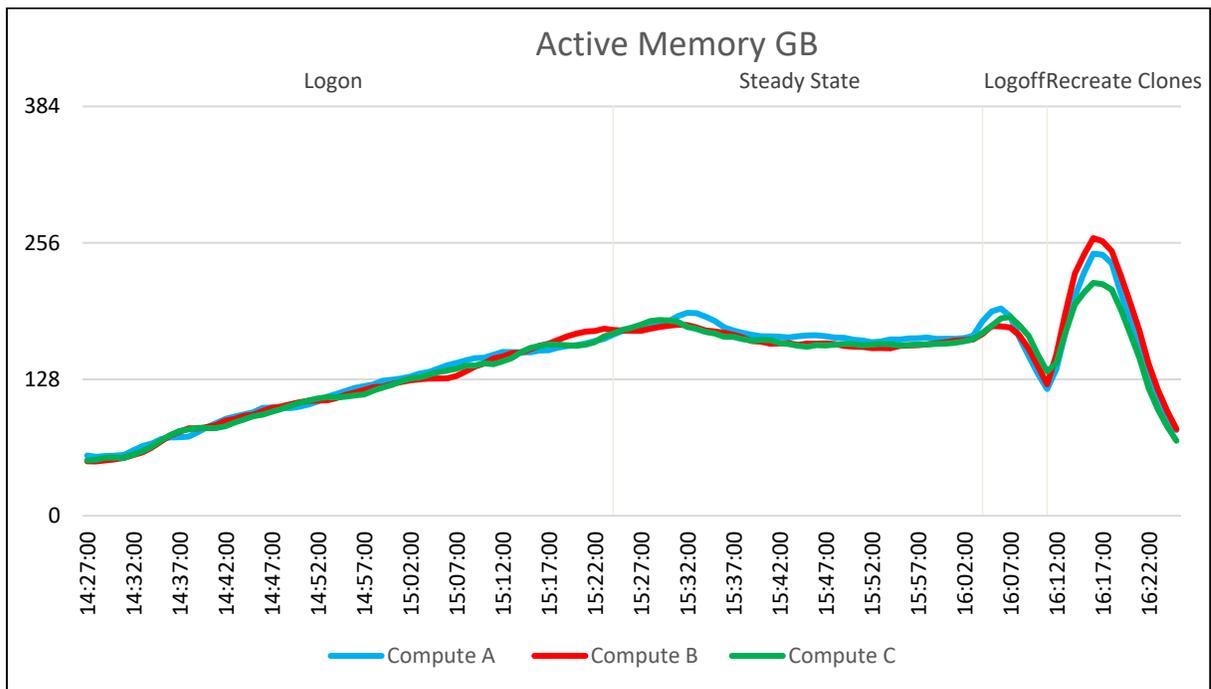


Figure 30 Active memory

Network bandwidth is not an issue on this test run with a steady state peak of approximately 1,561 Mbps. The busiest period for network traffic was during the recreation of the instant clones after testing had completed. One of the hosts reached a peak of 3,352 Mbps during the deletion and recreation of the instant clones.

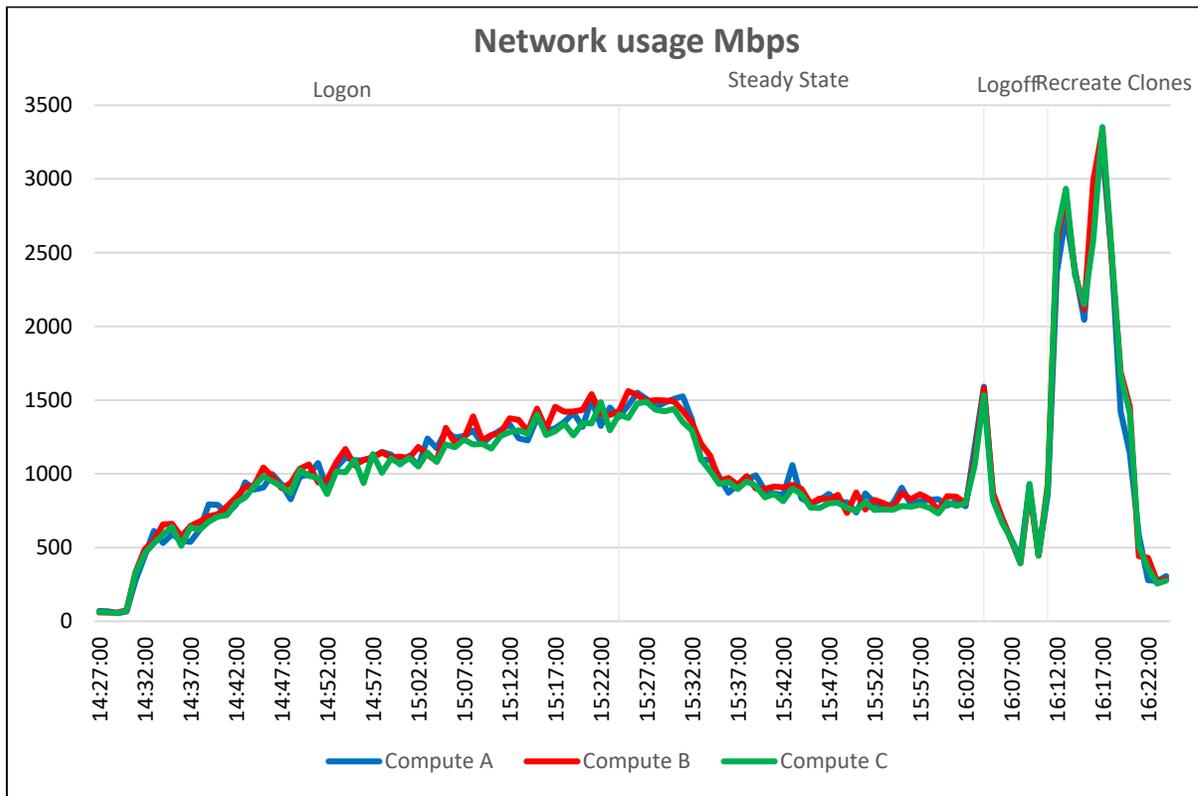


Figure 31 Network usage

The IOPS graphs and IOPS numbers are taken from the vCenter web console and the graphs clearly display the initial logon of the desktops, the steady state and logoff phases and finally the recreation of the desktops after testing was complete. The graph displays the Disk IOPS figure for the ScaleIO cluster.

The cluster reached a maximum of 26,900 Disk IOPS during the Recreate Clones period after testing and 10,724 IOPS at the end of steady state.

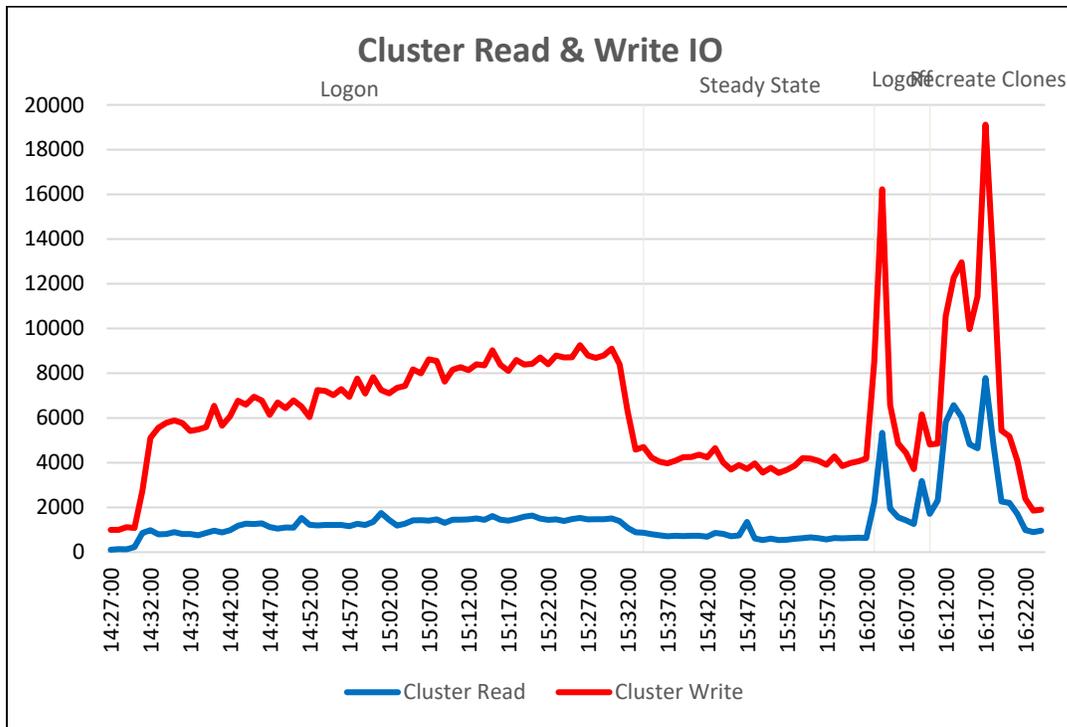


Figure 32 Cluster read & write IO

Prolonged Disk I/O Latency was not an issue during the Login VSI testing period of this test run. The maximum latency reached was approximately 29 ms spike during the Recreate Clones state. Steady state average for Read IO was 0.1 ms, Write IO was 4 ms and combined Read & Write IO average was 4.1 ms, median 3 ms. this was well below the 20 ms threshold that is regarded as becoming potentially troublesome.

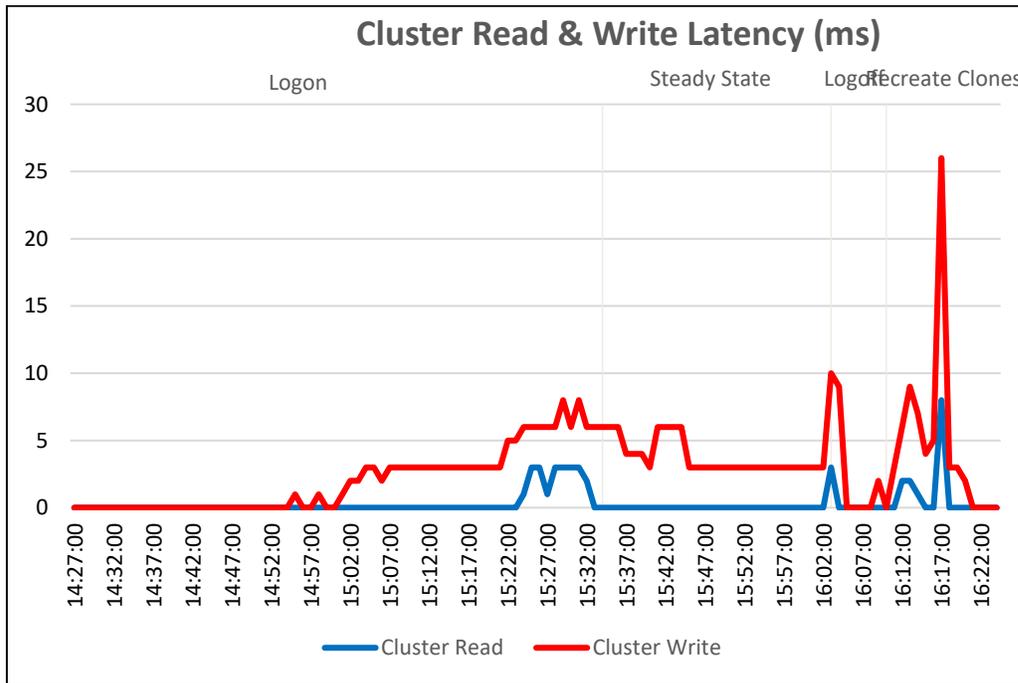


Figure 33 Cluster read & write latency

The Login VSI Max user experience score shown below for this test was not reached indicating that the number of users tested is appropriate for this configuration. Manually interacting with the desktop sessions showed mouse response and switching between windows to be quick and smooth and video playback was of good quality.

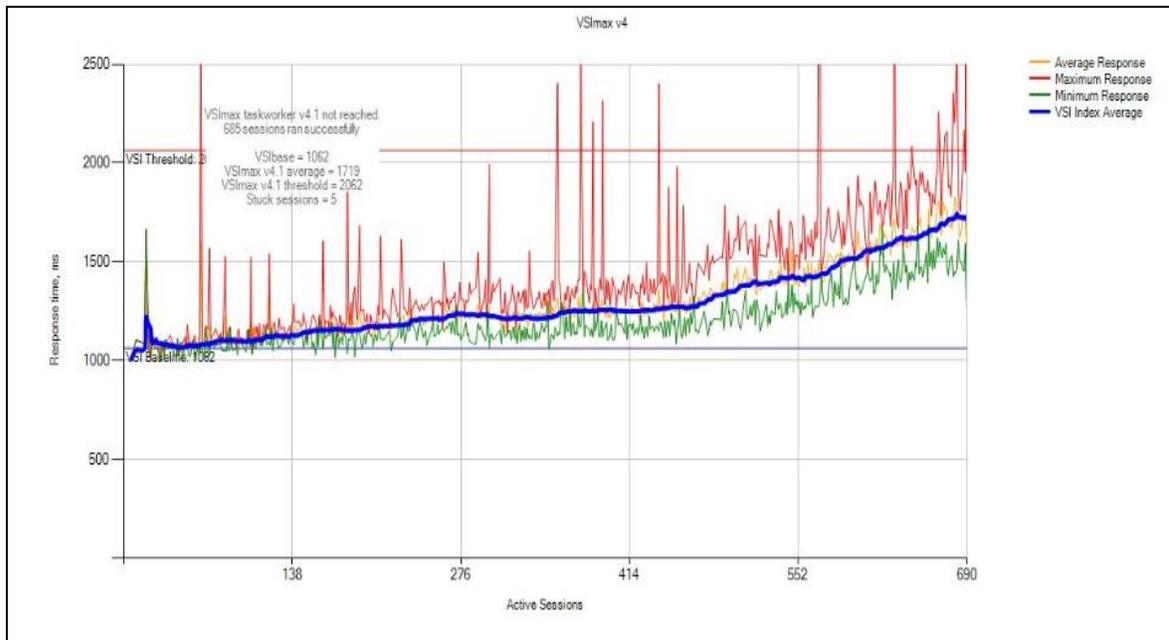


Figure 34 VSI max

6.4.2 Knowledge Worker, 140 users, ESXi 6.5, Horizon 7.x Instant Clones

The below graph shows the performance data for 140 user sessions per host or 420 sessions per cluster. The CPU reaches a steady state average of 84 % across the three compute hosts during the test cycle when 140 users were logged on to each compute host. The CPU usage was approximately 5% on the dedicated compute hosts before starting test.

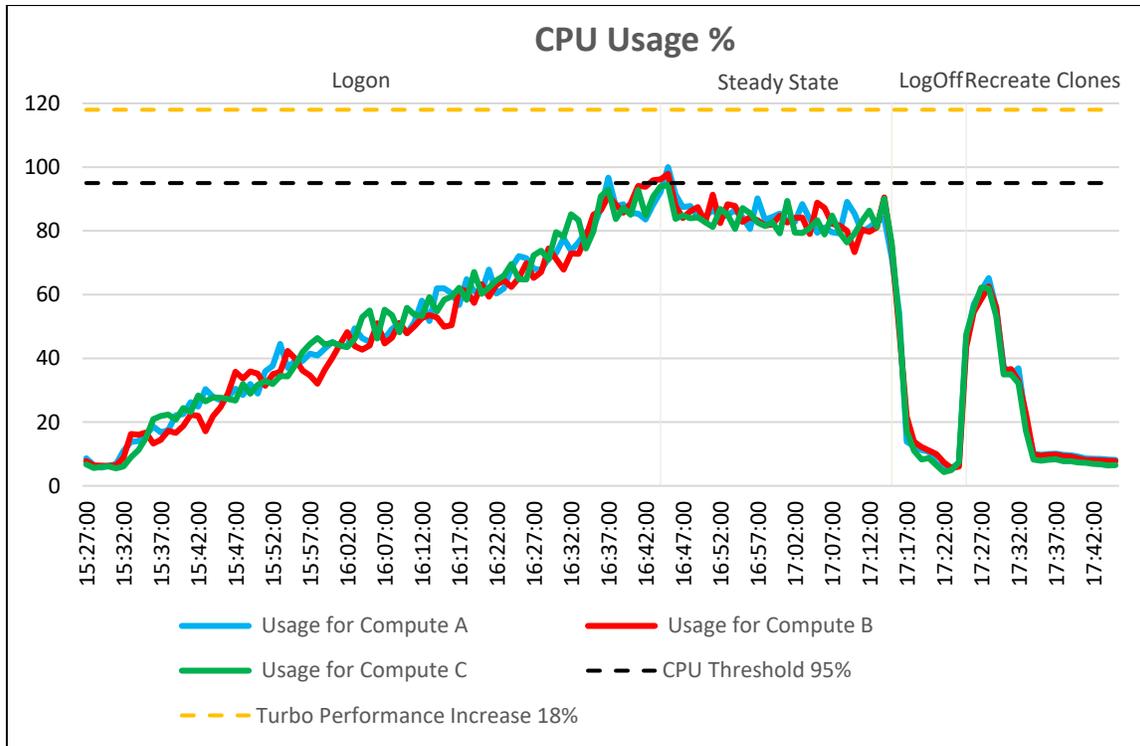


Figure 35 CPU usage

In regard to memory consumption for the cluster, out of a total of 384 GB available memory per node, memory usage was not pushed close to its maximum usage. The compute hosts reached a maximum memory consumption of 192 GB with active memory usage reaching a max of 399 GB during the recreation of the instant clones. A small amount of ballooning ~ 10GB per host took place on all hosts for approximately 15 minutes prior to steady state. Some memory swapping occurred on ESX1 where the ScaleIO Gateway SVM resides.

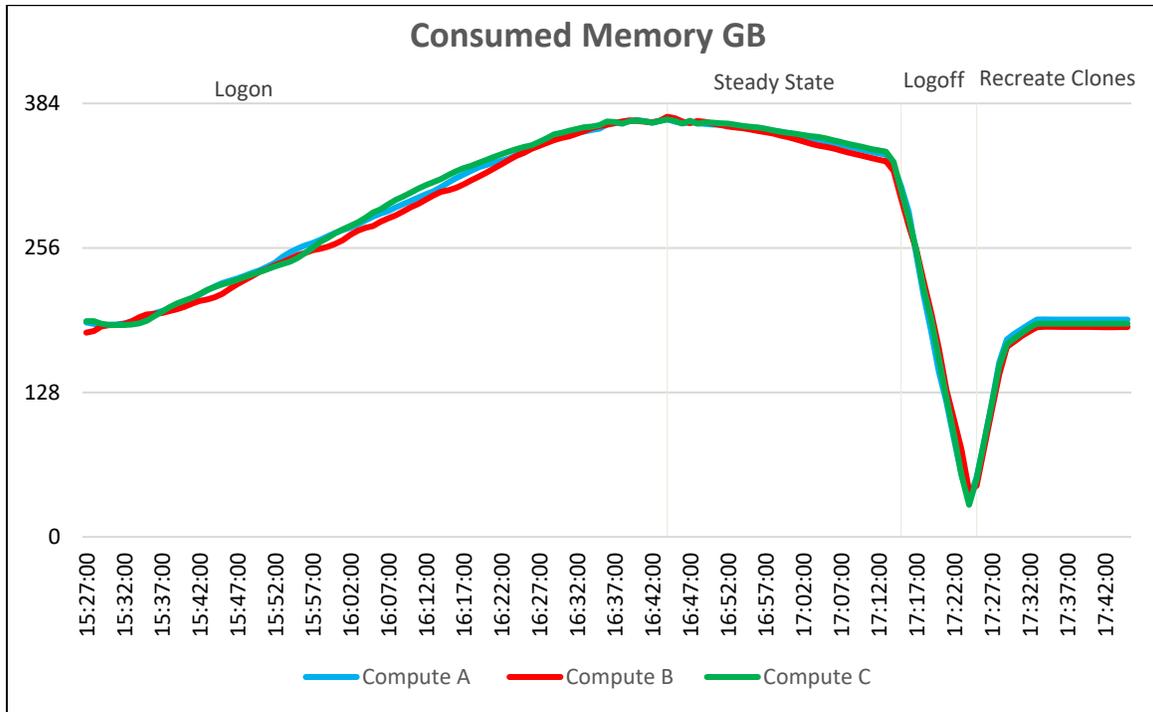


Figure 36 Consumed memory

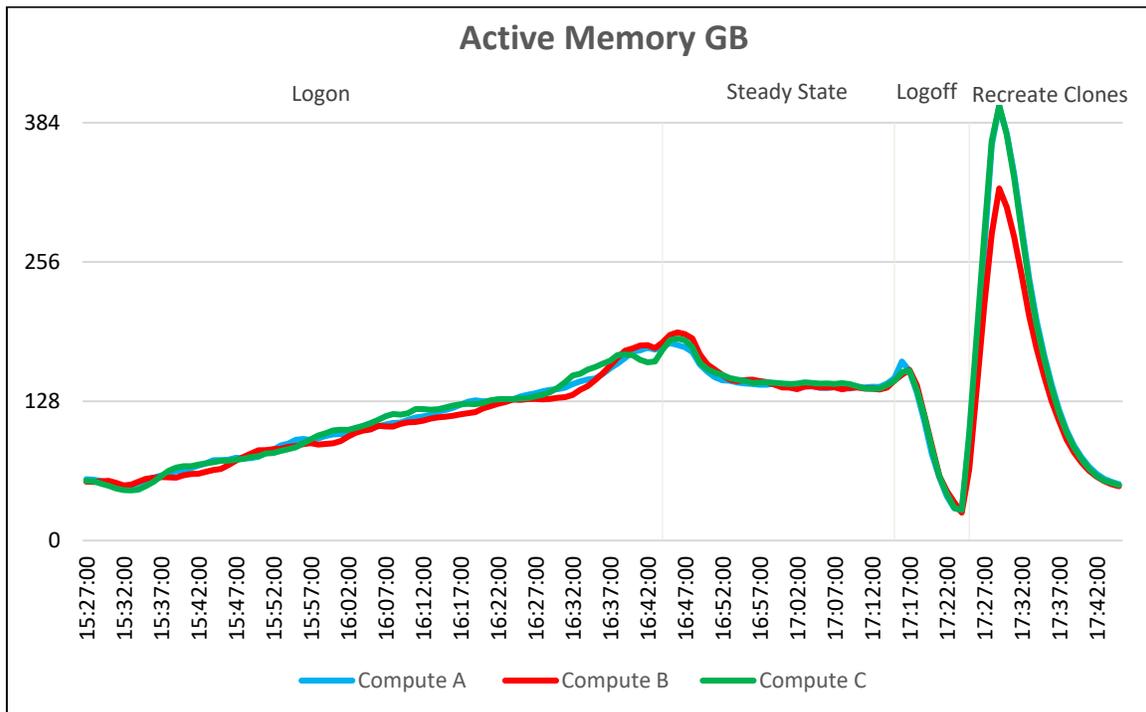


Figure 37 Active memory

Network bandwidth is not an issue on this test run with a steady state peak of approximately 1,728 Mbps. The busiest period for network traffic was during the recreation of the instant clones after testing had completed. One of the hosts reached a peak of 3,234 Mbps during the deletion and recreation of the instant clones.

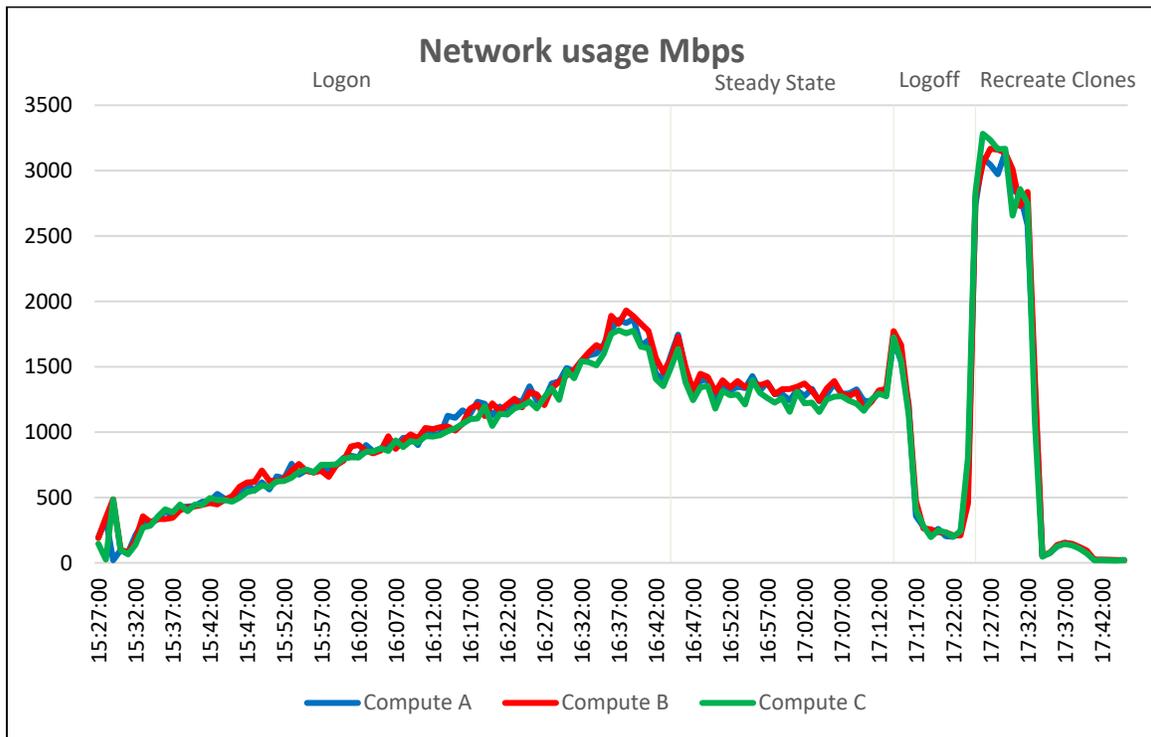


Figure 38 Network usage

The IOPS graphs and IOPS numbers are taken from the vCenter web console and the graphs clearly display the initial logon of the desktops, the steady state and logoff phases and finally the recreation of the desktops after testing was complete. The graph displays the Disk IOPS figure for the ScaleIO cluster.

The cluster reached a maximum of 24,254 Disk IOPS during the Logoff period after testing and 8,562 IOPS at the start of steady state.

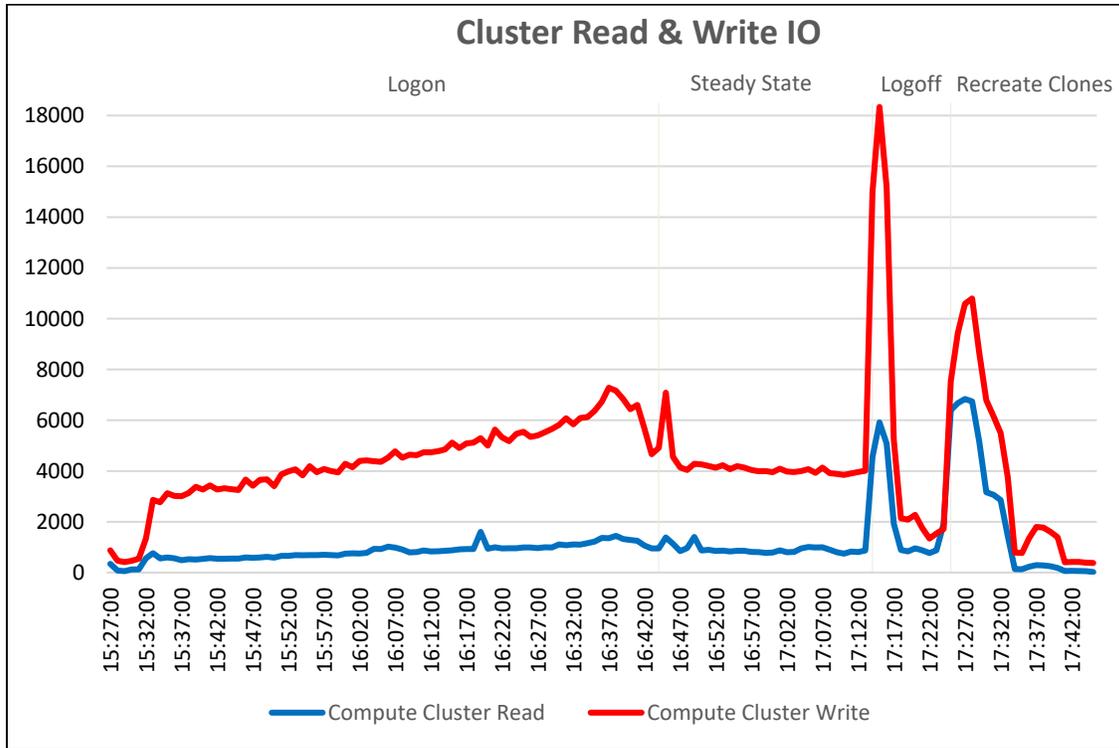


Figure 39 Cluster read & write IO

Prolonged Disk I/O Latency was not an issue during the Login VSI testing period of this test run. The maximum latency reached was approximately 118 ms spike at the beginning of steady state. Steady state average for Read IO was 1.2 ms, Write IO was 9 ms and combined Read & Write IO was 10 ms, median 3 ms. this was well below the 20 ms threshold that is regarded as becoming potentially troublesome.

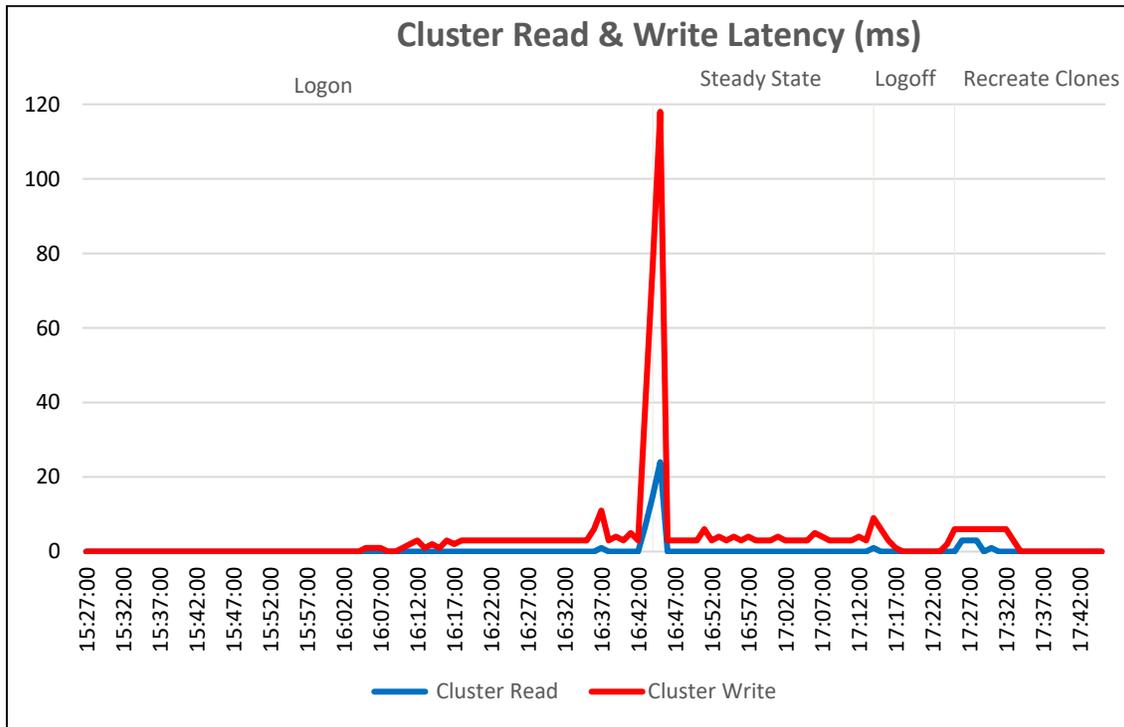


Figure 40 Cluster read & write latency

The Login VSI Max user experience score shown below for this test was not reached indicating that the number of users tested is appropriate for this configuration. Manually interacting with the desktop sessions showed mouse response and switching between windows to be quick and smooth and video playback was of good quality.

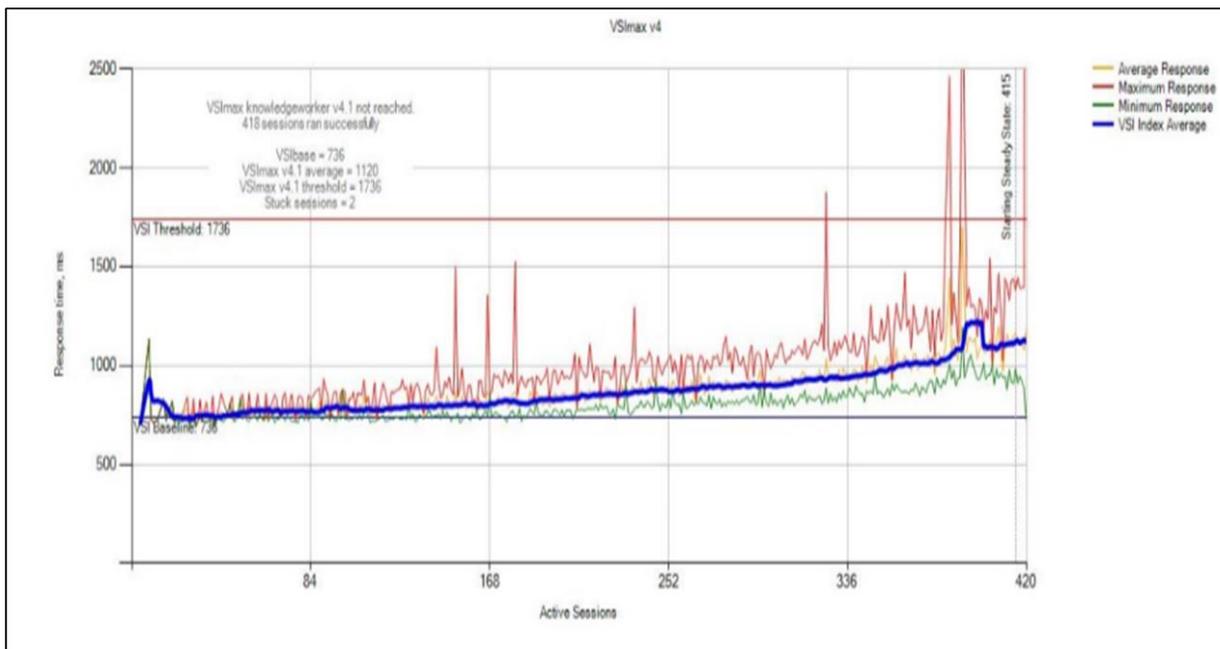


Figure 41 VSImax

6.4.3 Power Worker, 115 users, ESXi 6.5, Horizon 7.x Instant Clones

The below graph shows the performance data for the 105 management/115 compute user sessions per host or 335 sessions per cluster. The CPU reaches a steady state average of 86 % across the three compute hosts during the test cycle when 105/115 users were logged on to each compute host. The CPU usage was approximately 11% on the hosts before the start of test.

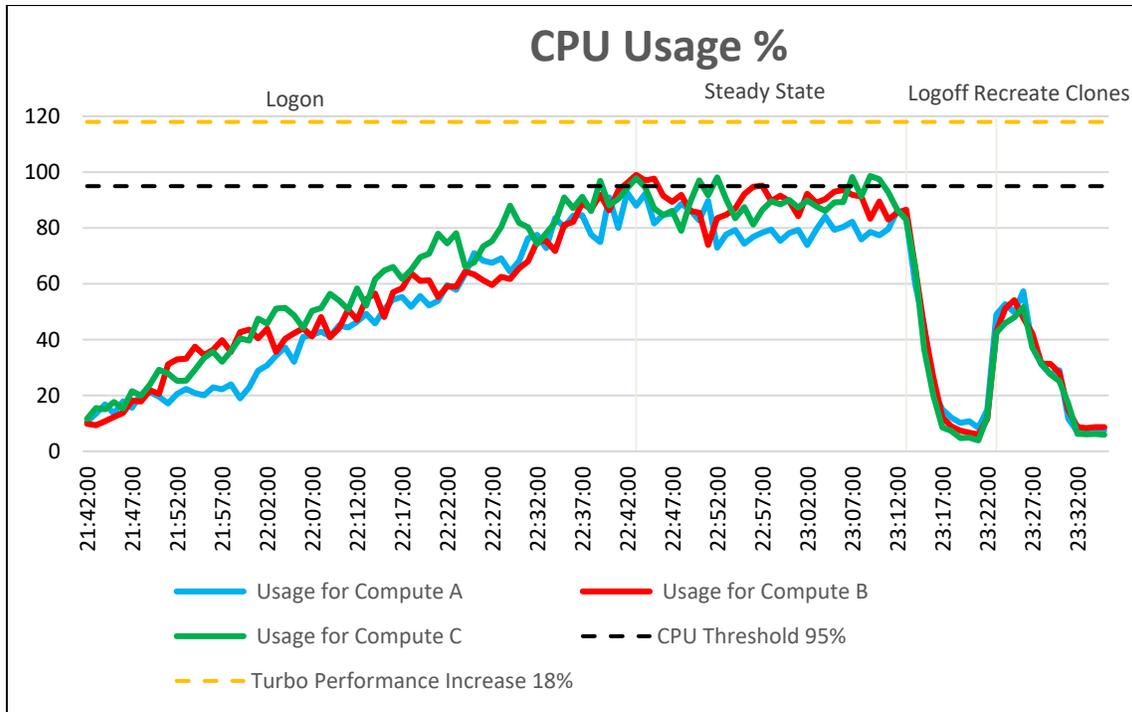


Figure 42 CPU usage

In regard to memory consumption for the cluster, out of a total of 384 GB available memory per node, memory usage was not pushed close to its maximum usage. The compute hosts reached a maximum memory consumption of 180 GB with active memory usage reaching a max of 414 GB during the recreation of the instant clones. A small amount of ballooning ~ 15 GB per host took place on all hosts for approximately 15 minutes prior to steady state. Some memory swapping occurred on ESX A where the ScaleIO Gateway SVM resides.

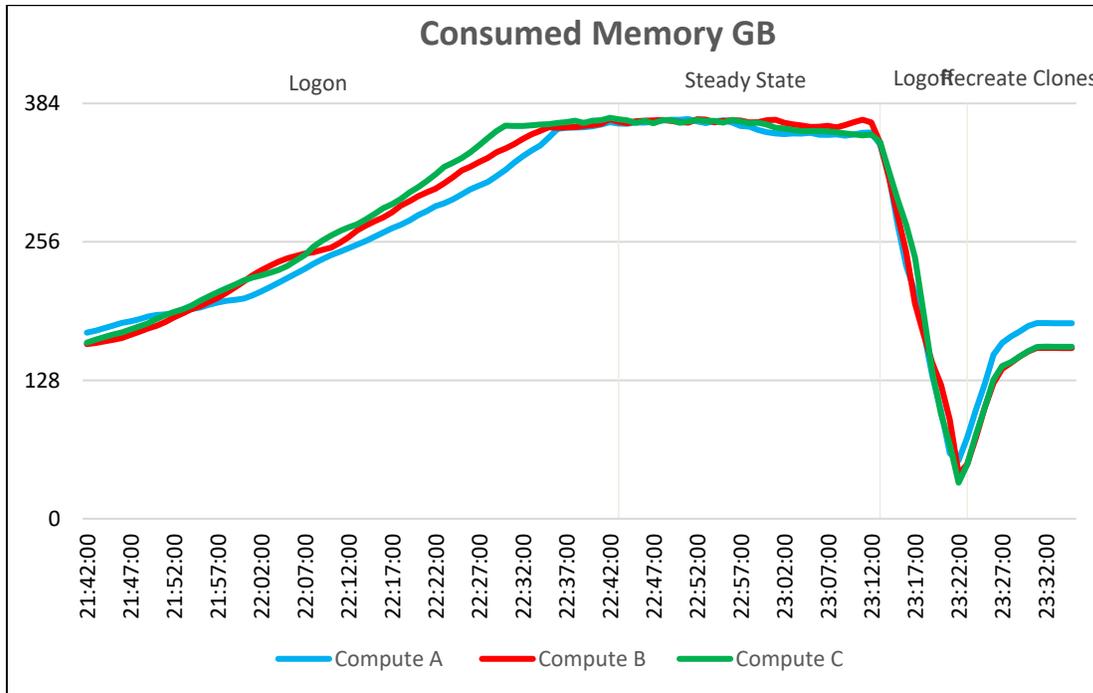


Figure 43 Consumed memory

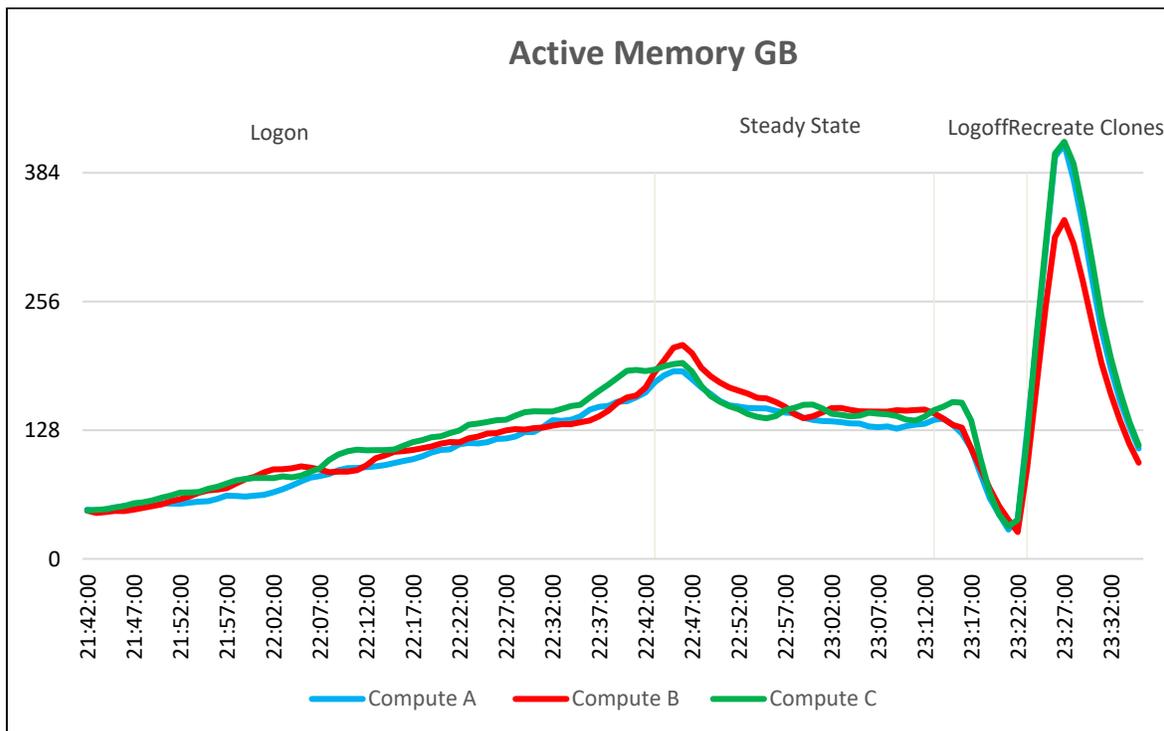


Figure 44 Active memory

Network bandwidth is not an issue on this test run with a steady state peak of approximately 1,736 Mbps. The busiest period for network traffic was during the recreation of the instant clones after testing had completed. One of the hosts reached a peak of 3,443 Mbps during the deletion and recreation of the instant clones.

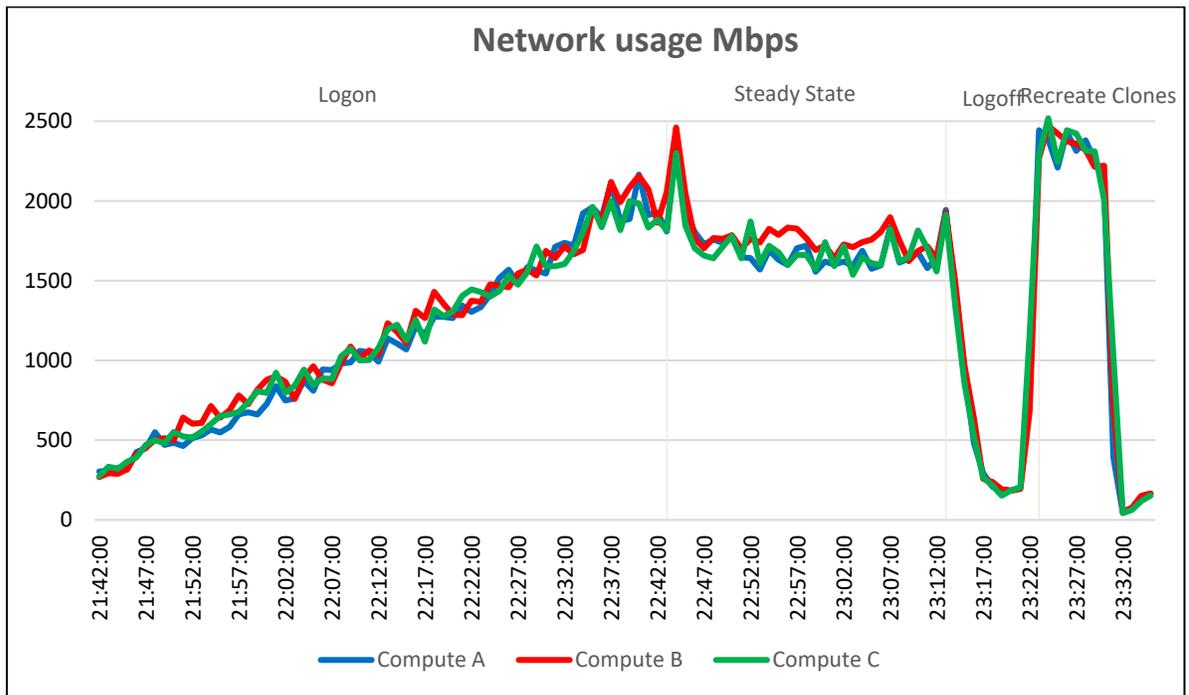


Figure 45 Network usage

The IOPS graphs and IOPS numbers are taken from the vCenter web console and the graphs clearly display the initial logon of the desktops, the steady state and logoff phases and finally the recreation of the desktops after testing was complete. The graph displays the Disk IOPS figure for the ScaleIO cluster.

The cluster reached a maximum of 22,458 Disk IOPS during the Logoff period after testing and 8,582 IOPS at the start of steady state.

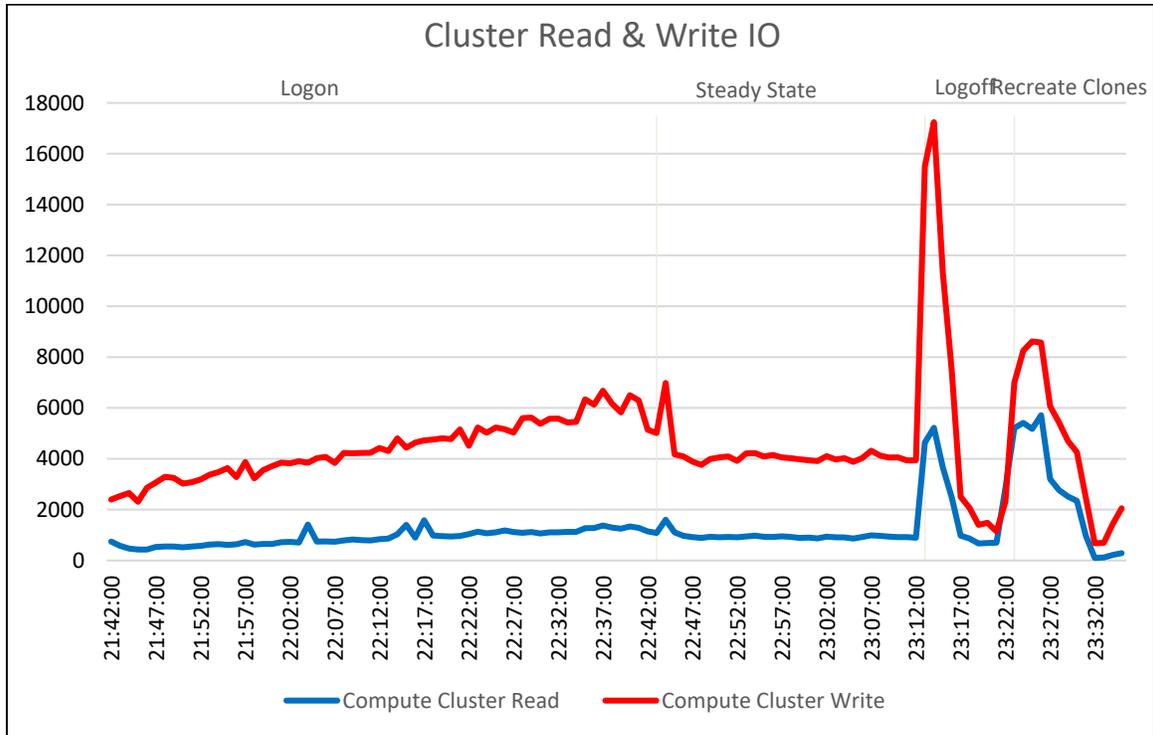


Figure 46 Cluster read & write IO

Prolonged Disk I/O Latency was not an issue during the Login VSI testing period of this test run. The maximum latency reached was approximately 64 ms spike at the beginning of steady state. Steady state average for Read IO was .9 ms, Write IO was 8 ms and combined Read & Write IO average was 8.6 ms, median 5.5 ms. this was well below the 20 ms threshold that is regarded as becoming potentially troublesome.

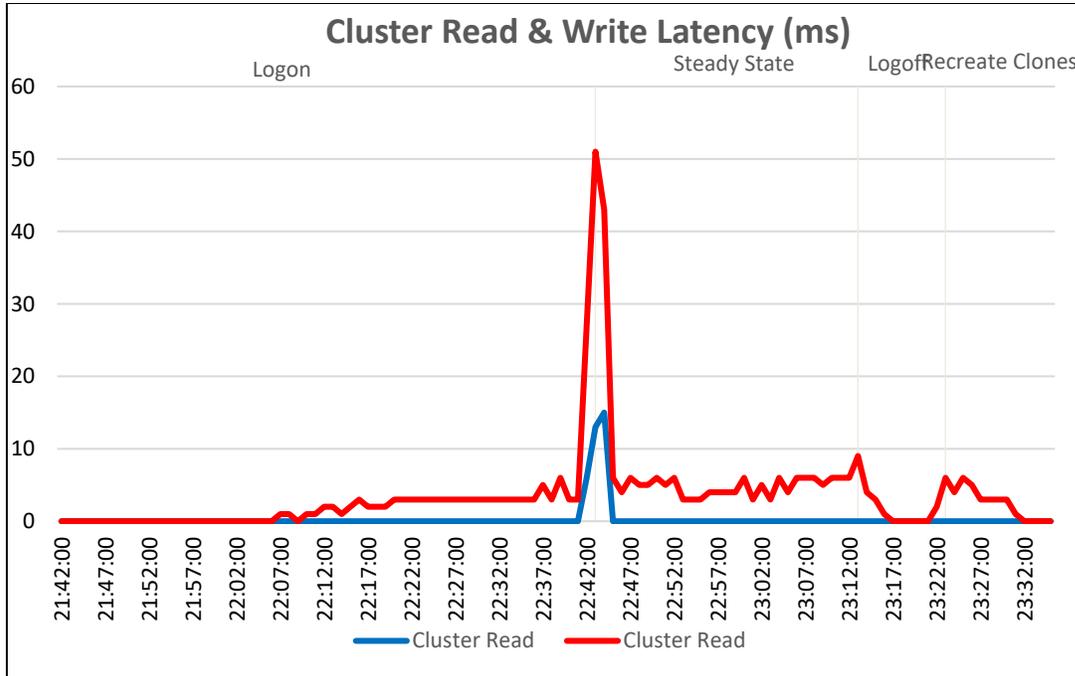


Figure 47 Cluster read & write latency

The Login VSI Max user experience score shown below for this test was not reached indicating that the number of users tested is appropriate for this configuration. Manually interacting with the desktop sessions showed mouse response and switching between windows to be quick and smooth and video playback was of good quality.

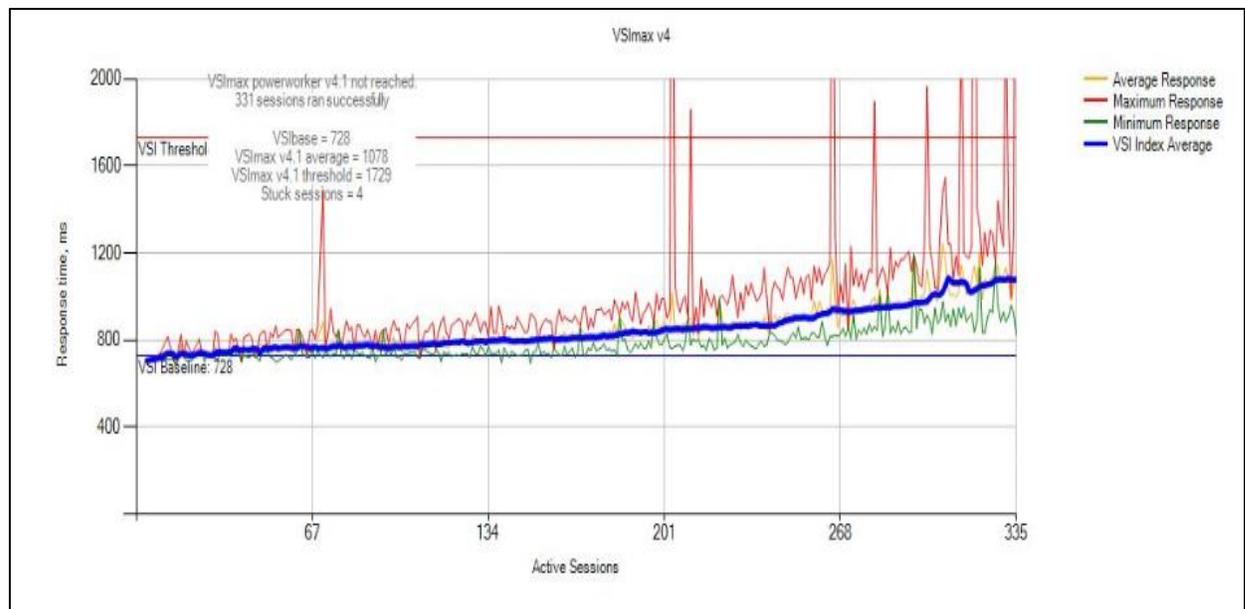


Figure 48 VSImax

6.4.4 Power Worker (M60-1Q), 48 users, ESXi 6.5, Horizon 7.x Instant Clones

With all the virtual machines restarted before starting test, the CPU usage reached a max of 62% on Compute Host C during the Boot Storm stage.

The below graph shows the performance data for 16 user sessions per M60 GPU Card or 48 sessions on Compute Host C. The CPU reaches a steady state average of 50 % on Compute Host C during the test cycle when 48 Graphic users were logged on.

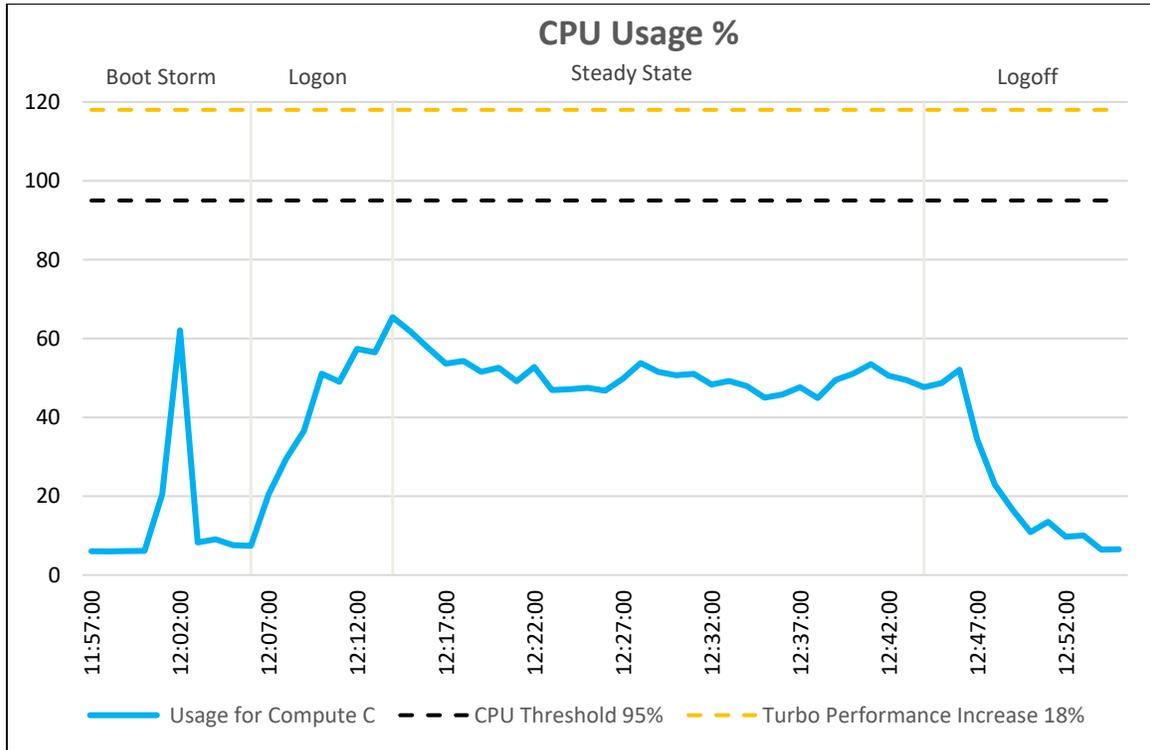


Figure 49 CPU usage

The below graph shows the GPU performance data for 3 X M60 Nvidia Telsa GPU Cards. Each GPU Card contains two GPU Cores. GPU core steady state average of 34% and Max of 35%, during the test cycle when 48 Graphic users were logged on.

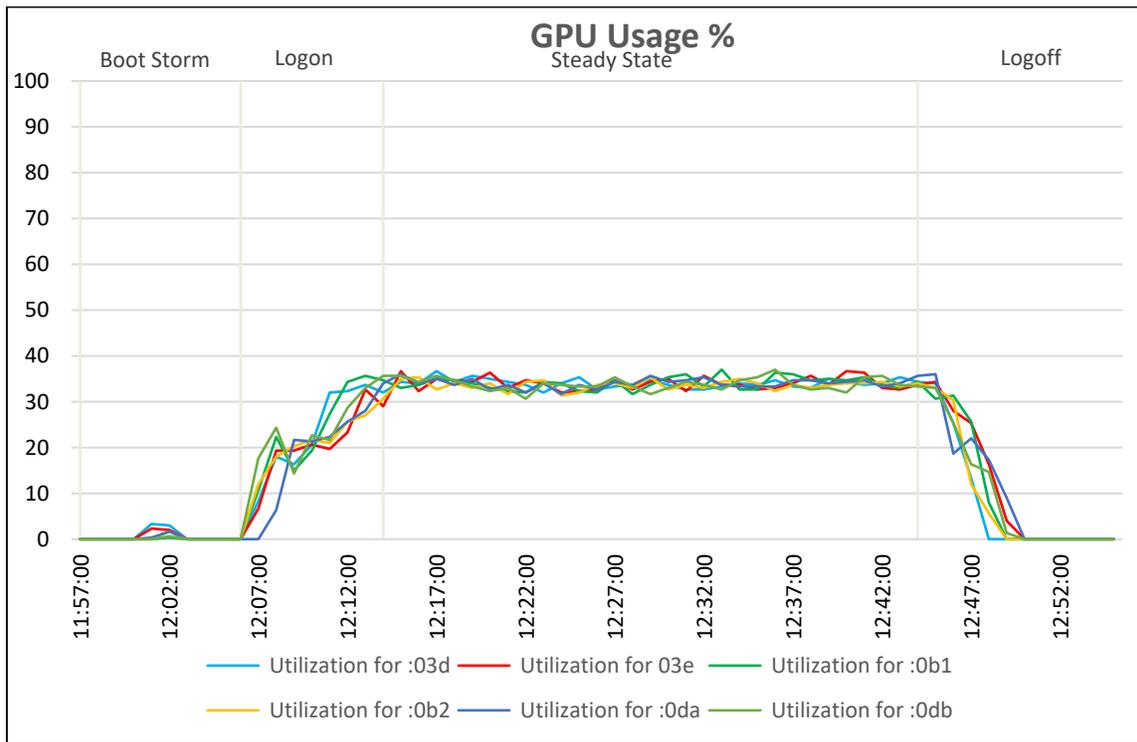


Figure 50 GPU usage

In regard to memory consumption for the cluster, out of a total of 384 GB available memory per node. Compute host C memory consumption remained steady at 222 GB with active memory remaining steady at 208 GB during the test period.

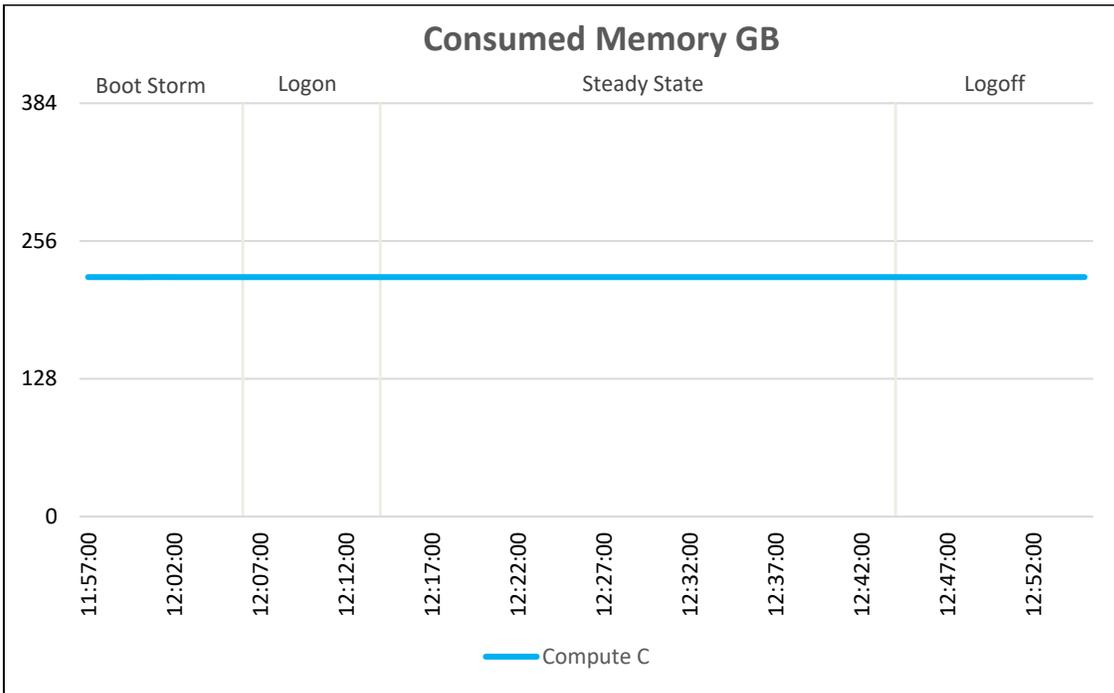


Figure 51 Consumed memory

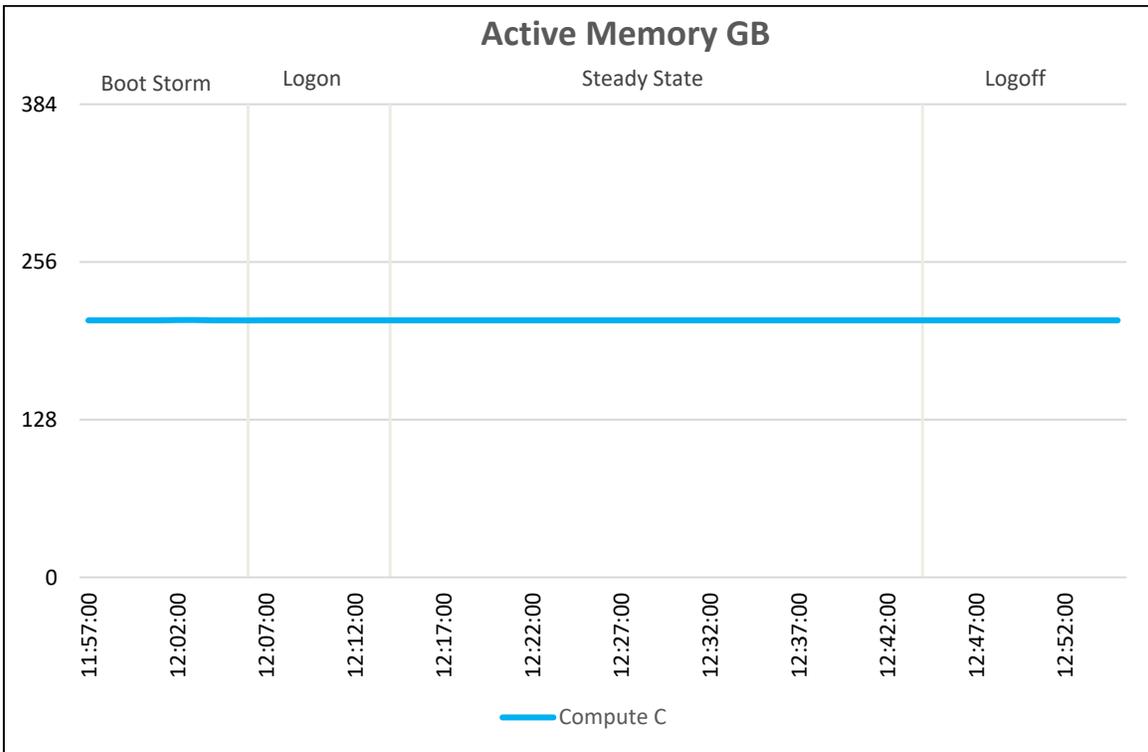


Figure 52 Active memory

Network bandwidth is not an issue on this test run with a Logon peak of approximately 864 Mbps. The busiest period for network traffic was during the Boot storm and Logon stages, averaging at 336 Mbps.

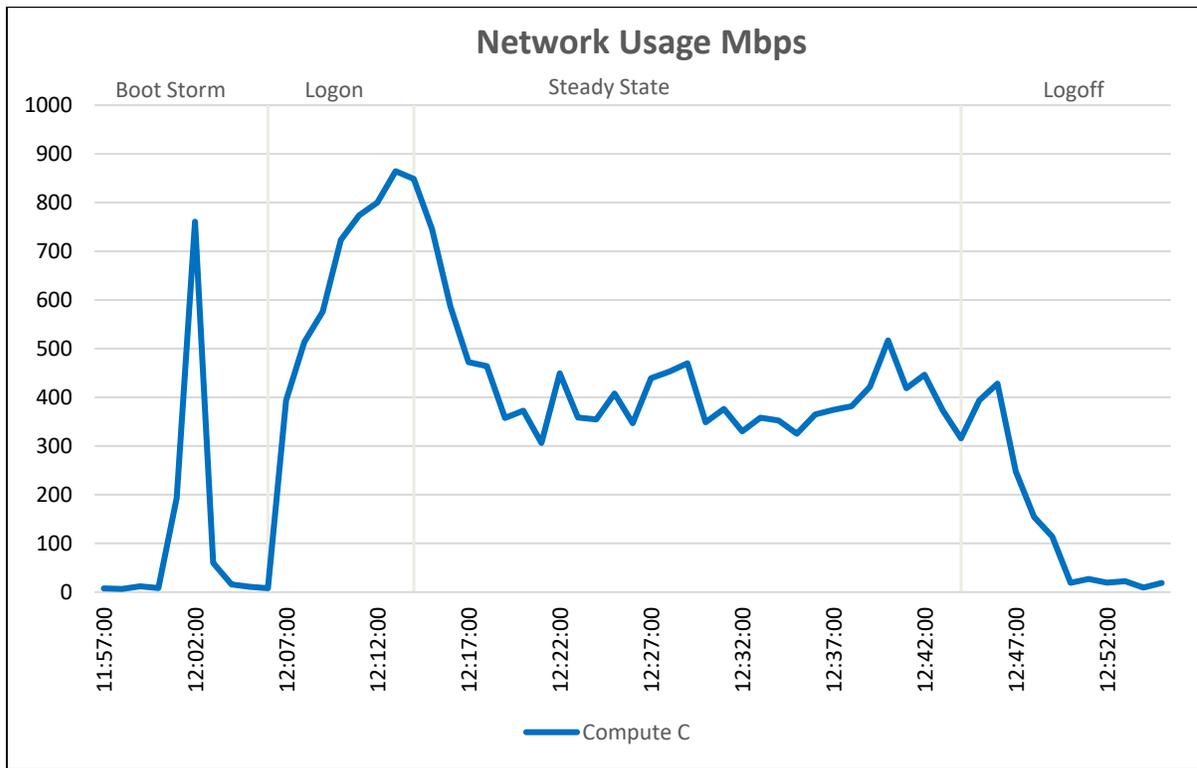


Figure 53 Network usage

The IOPS graphs and IOPS numbers are taken from the vCenter web console and the graphs clearly display the initial logon if the desktops, the steady state and logoff phases and finally the recreation of the desktops after testing was complete. The graph displays the Disk IOPS figure for the ScaleIO cluster.

The cluster reached a maximum of 6,133 Disk IOPS during the Boot Storm period before testing and 2151 IOPS during Logoff state.

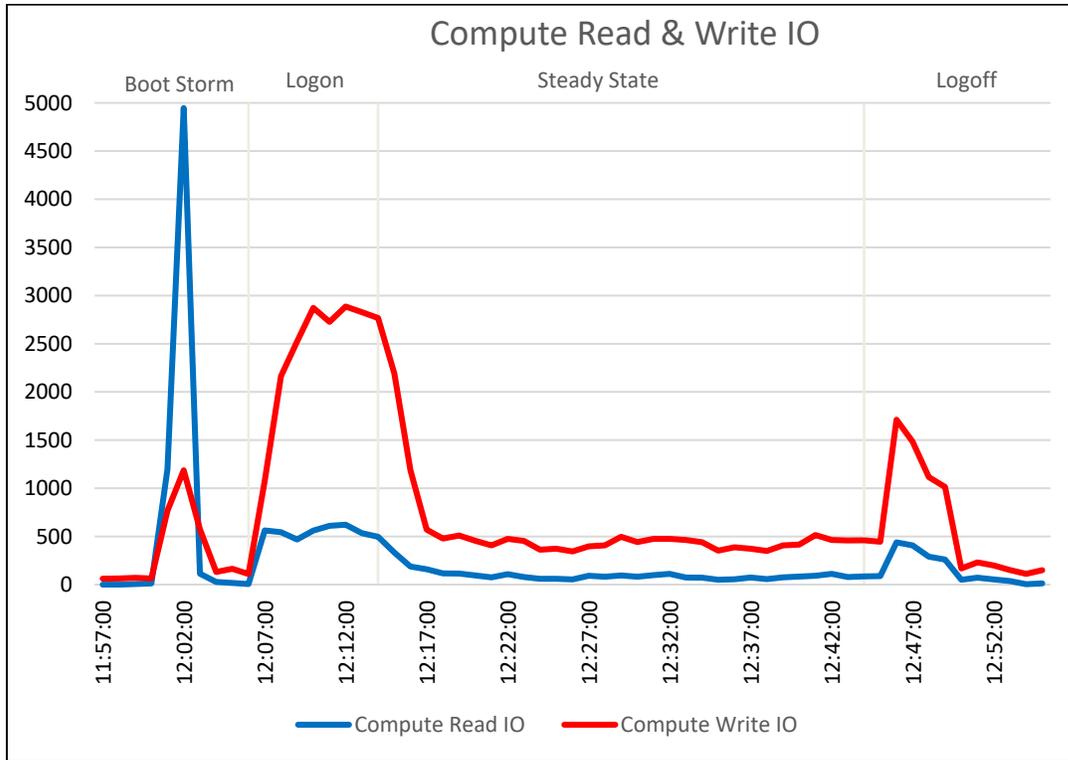


Figure 54 Compute read & write IO

Prolonged Disk I/O Latency was not an issue during the Login VSI testing period of this test run. The maximum latency reached was approximately 5 ms spike during the Boot Storm state. Steady state average for Read IO was 0.4 ms, Write IO was 0.4 ms and combined Read & Write IO average was 0.8 ms. This was well below the 20 ms threshold that is regarded as becoming potentially troublesome.

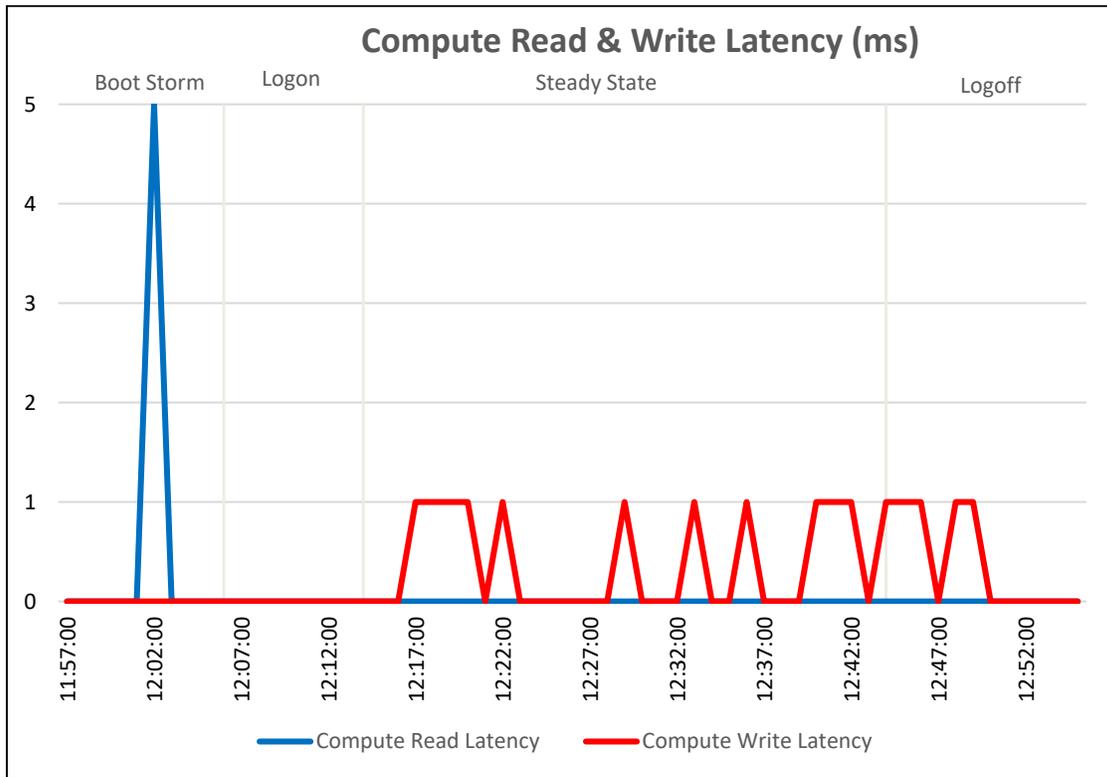


Figure 55 Compute read & write latency

The Login VSI Max user experience score shown below for this test was not reached indicating that the number of users tested is appropriate for this configuration. Manually interacting with the desktop sessions showed mouse response and switching between windows to be quick and smooth and video playback was of good quality.

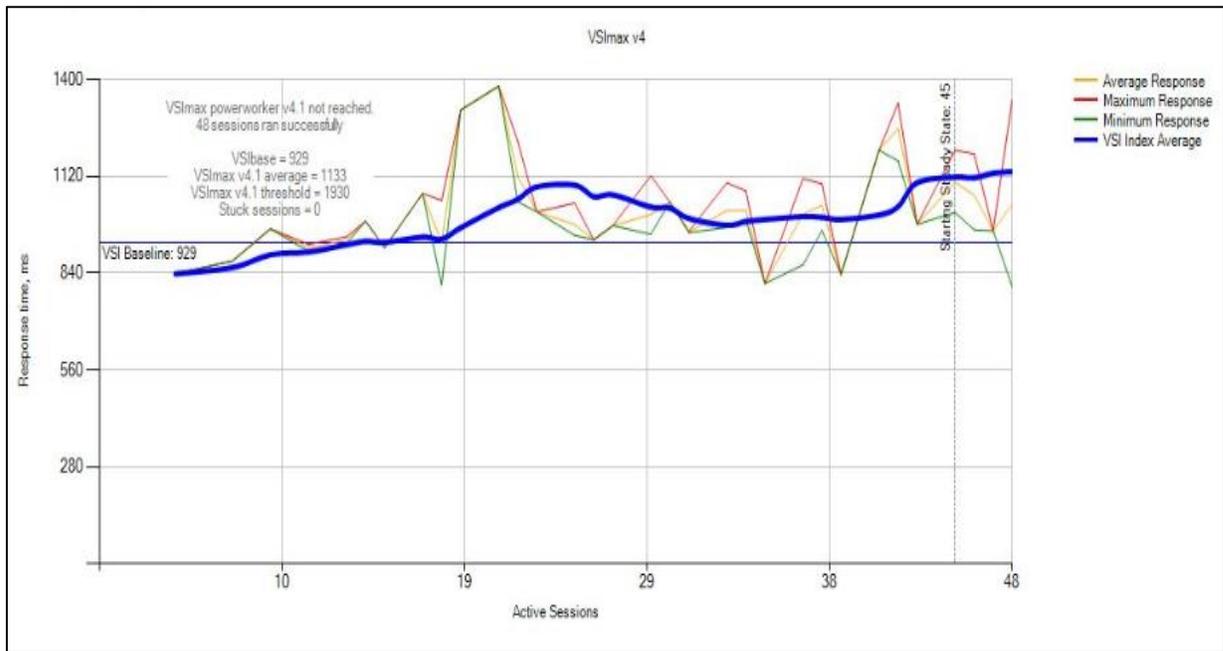


Figure 56 VSImax

6.4.5 Power Worker (M10-1B), 64 users, ESXi 6.5, Horizon 7.x Instant Clones

The below graph shows the performance data for 32 user sessions per M60 GPU Card or 64 sessions on Compute Host C. The CPU reaches a steady state average of 66 % on Compute Host C during the test cycle when 64 Graphic users were logged on.

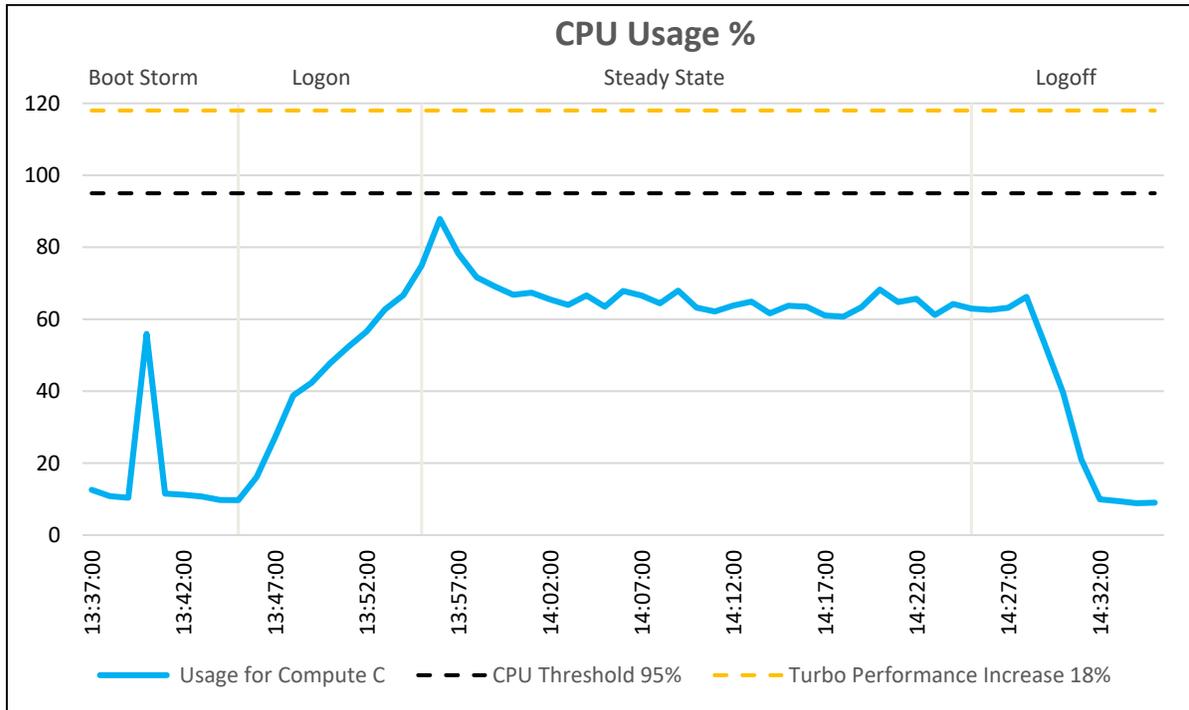


Figure 57 CPU usage

The below graph shows the GPU performance data for 2 X M10 Nvidia Telsa GPU Cards. Each GPU Card contains four GPU Cores. GPU core steady state average of 41% and Max of 42%, during the test cycle when 64 Graphic users were logged on.

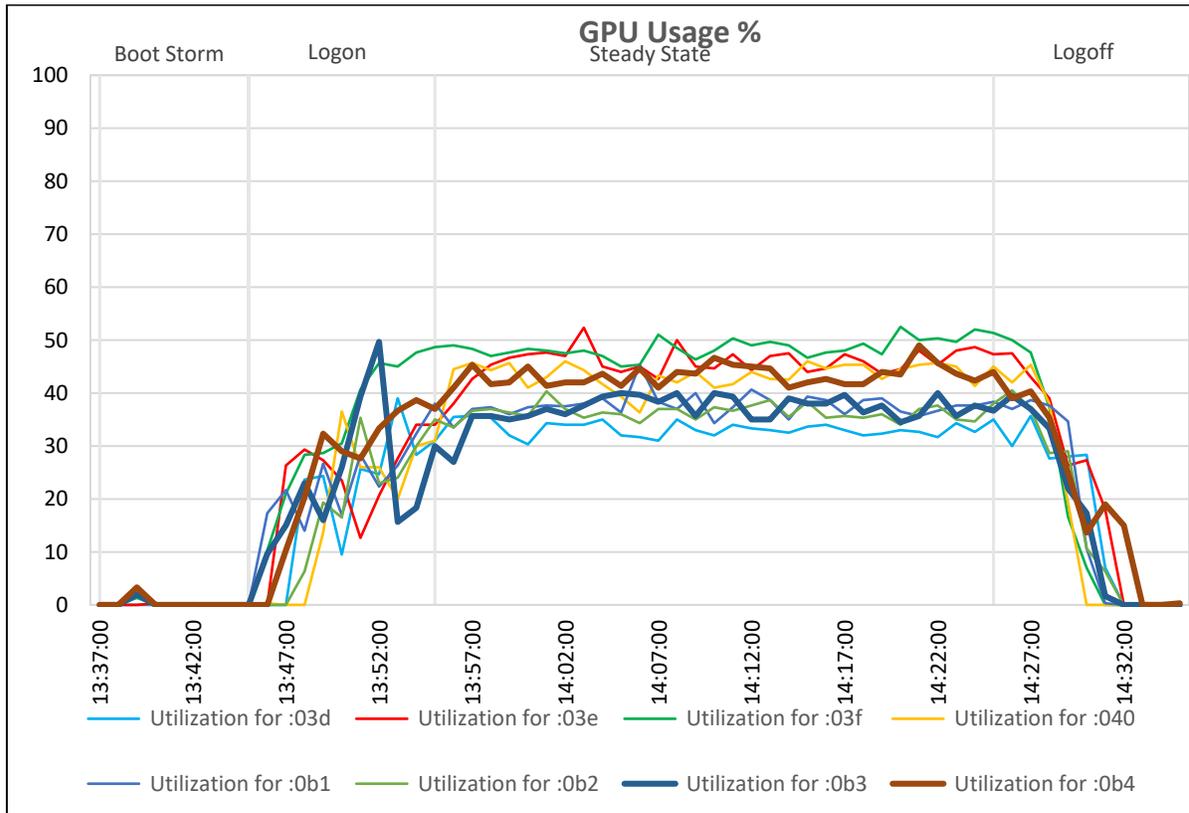


Figure 58 GPU usage

In regard to memory consumption for the cluster, out of a total of 384 GB available memory per node. Compute host C memory consumption remained steady at 288 GB with active memory remaining steady at 273 GB during the test period. Zero memory ballooning occurred on Compute Host C during testing.

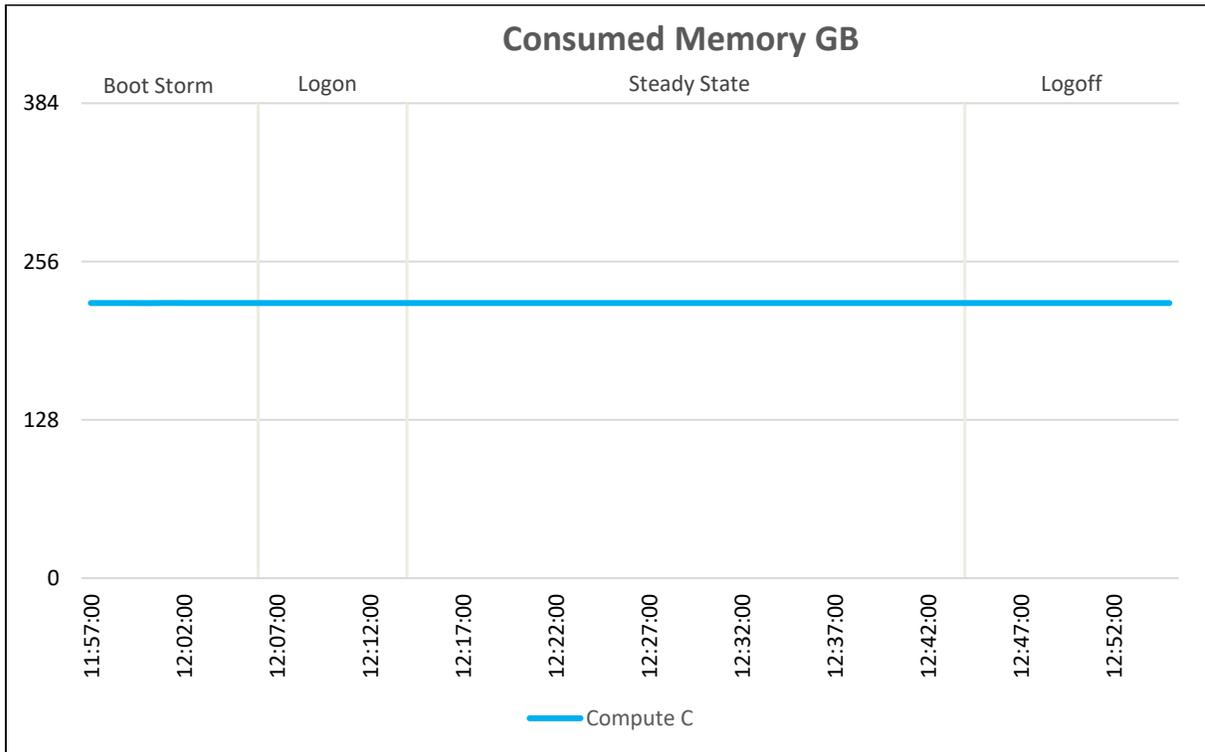


Figure 59 Consumed memory

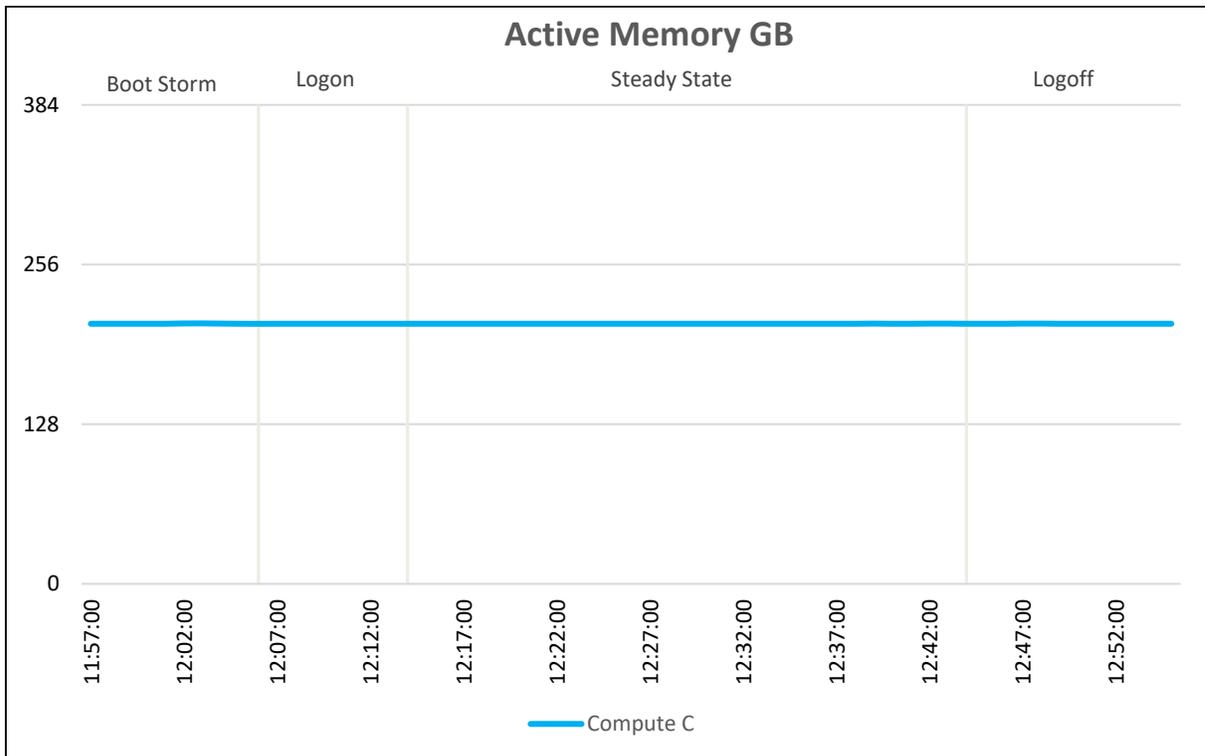


Figure 60 Active memory

Network bandwidth is not an issue on this test run with a Logon peak of approximately 968 Mbps. The busiest period for network traffic was during the Logon and Steady State, averaging at 562 Mbps.

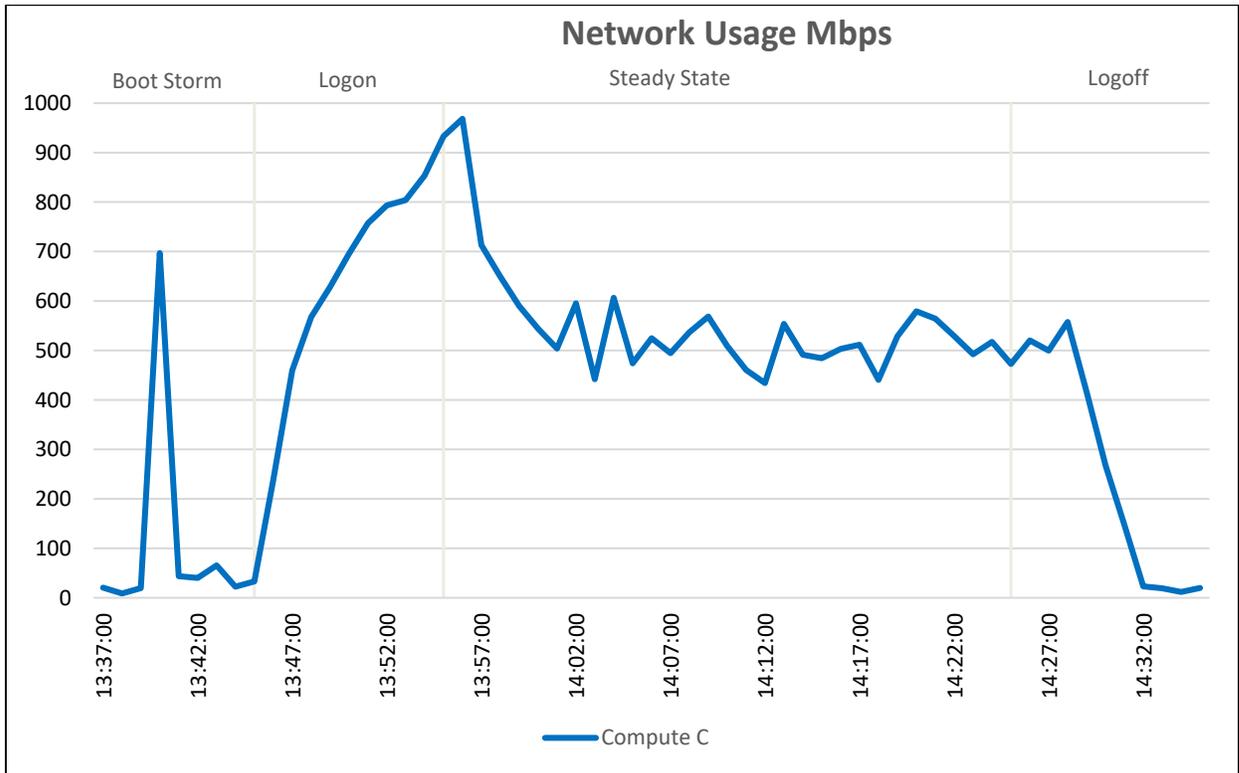


Figure 61 Network usage

The IOPS graphs and IOPS numbers are taken from the vCenter web console and the graphs clearly display the initial logon if the desktops, the steady state and logoff phases and finally the recreation of the desktops after testing was complete. The graph displays the Disk IOPS figure for the ScaleIO cluster.

The cluster reached a maximum of 5,245 Disk IOPS during the Boot Storm period before testing and 2,758 IOPS during Logoff state.

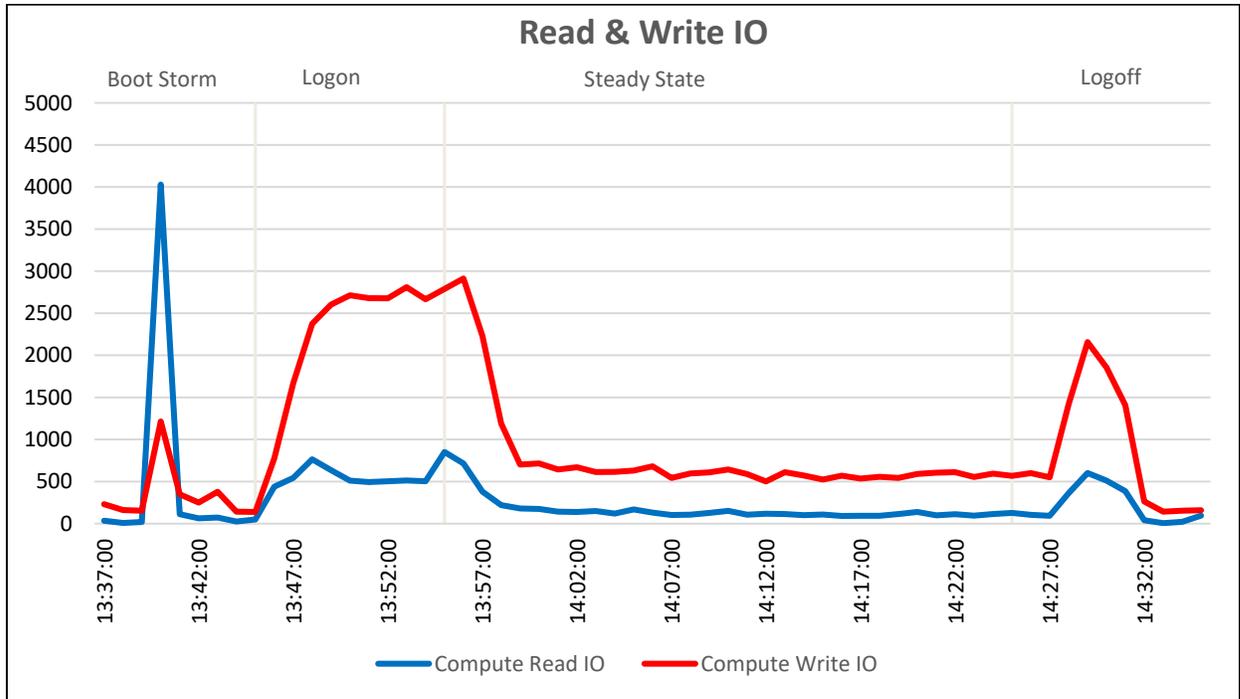


Figure 62 Read & write IO

Prolonged Disk I/O Latency was not an issue during the Login VSI testing period of this test run. The maximum latency reached was approximately 1 ms spike during the Boot Storm state. Steady state average for Read IO was 0.0 ms, Write IO was 0.8 ms and combined Read & Write IO average was 0.8 ms. This was well below the 20 ms threshold that is regarded as becoming potentially troublesome.

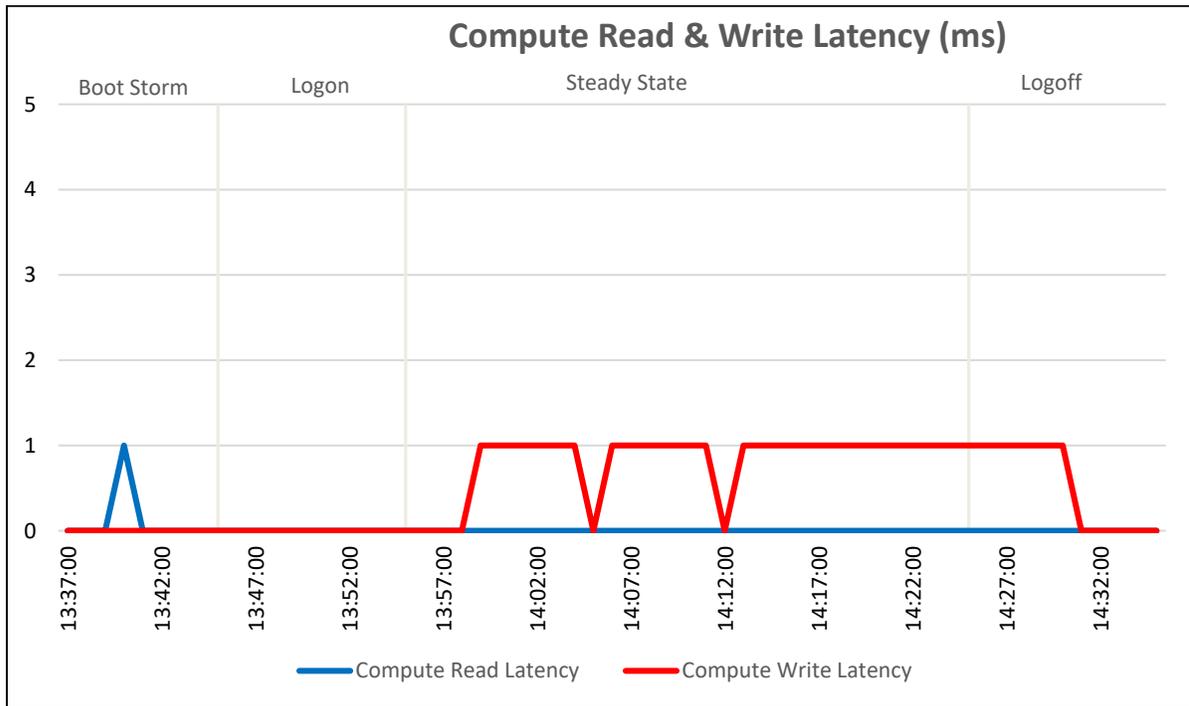


Figure 63 Compute read & write latency

The Login VSI Max user experience score shown below for this test was not reached indicating that the number of users tested is appropriate for this configuration. Manually interacting with the desktop sessions showed mouse response and switching between windows to be quick and smooth and video playback was of good quality.

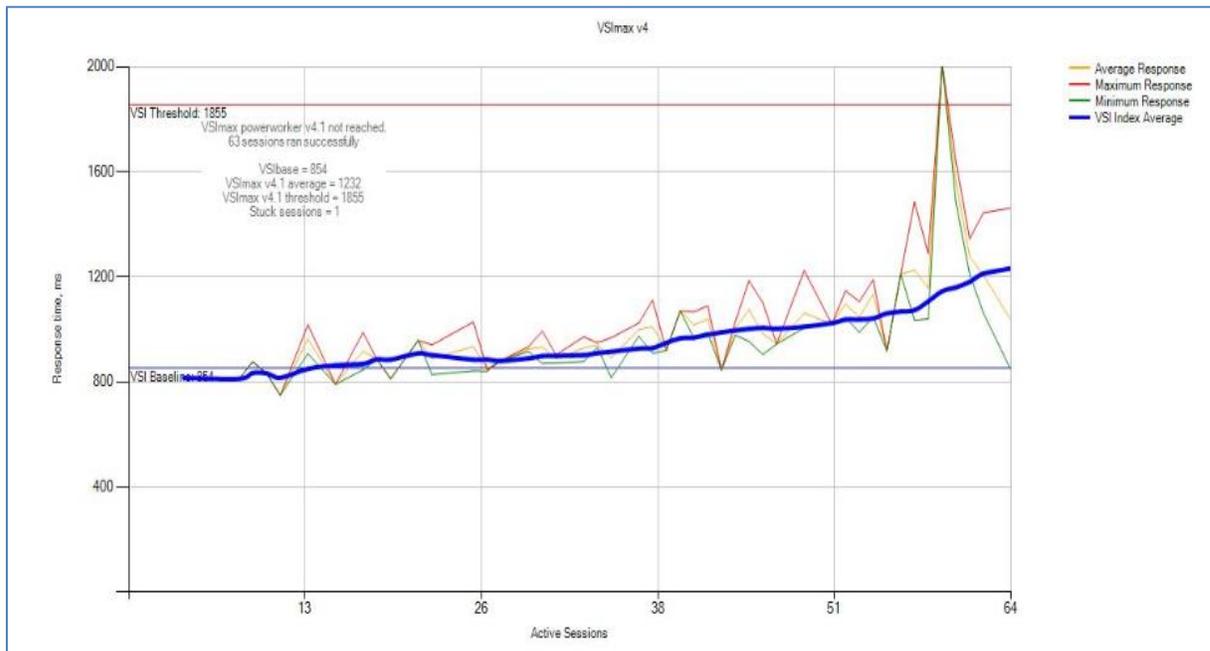


Figure 64 VSImax