

VMware Horizon View 5.3, EMC ScaleIO and LSI Nytro WarpDrive™

Reference Architecture

Contents

Executive Summary	2
Overview	3
Tested Configuration	4
Server Configuration	5
Reference Configuration	5
VDI Tested Parameters.....	6
User Response Time	6
VDI Deployment	6
VDI Boot	6
Test Results	6
User Response Time	6
Timing Test – 1000 Linked Clone Desktop Provisioning	7
Timing Test – 1000 Linked Clone Desktop Boot Storm	8
Tested Configuration Summary.....	9
Conclusion	14
About EMC ScaleIO.....	14
About LSI Nytro	14
About Login VSI	14
About Mellanox.....	15

Executive Summary

When moving to virtual desktop infrastructure (VDI) or improving the performance of your existing VDI deployment, it is critical that you select the right combination of software, compute and storage resources to meet your current and future needs. It is extremely important to investigate and evaluate any solution’s initial cost, ease of deployment and future scalability. EMC ScaleIO software, combined with fast and reliable LSI Nytro WarpDrive PCIe Flash Accelerator Cards, maximize performance, reduce latency and scale compute resources predictably as you add more users.

This document discusses scale, density and performance testing of a VDI application using EMC ScaleIO supported by LSI Nytro WarpDrive Flash Accelerator Cards. With this solution, utilizing the Login VSI workload generator, 1,000 VMware Horizon View desktop sessions were supported utilizing just six Supermicro Ivy Bridge-based servers connected over Mellanox 40GE end-to-end solution. This generator creates the actual storage accesses that a typical office user would create.

EMC ScaleIO software provides an elastic storage infrastructure that scales both capacity and performance as business requirements change. EMC ScaleIO is a software solution that utilizes a server based PCIe NAND, local hard disk drives (HDDs) and LAN interconnects to implement a high speed virtual SAN. EMC ScaleIO also provides all the benefits of shared storage—but at a fraction of the cost and complexity. Because VDI can often be bottlenecked by HDD or SAN based storage solutions, the LSI Nytro WarpDrive is an integral part of this high performance VDI implementation. (Figure 1)



Figure 1: Login VSI results

Overview

The objective of this solution test was to verify the feasibility and performance of creating and deploying a 1,000 desktop environment. Typically administrators prefer to support around 150 desktops per node, therefore, six Supermicro-based Ivy Bridge based servers were deployed each containing two LSI Nytro WarpDrive 3.2TB PCIe accelerator cards and Mellanox ConnectX-3 40GE NICs, all connected over Mellanox SwitchX based 40GE Switch.

EMC ScaleIO creates a virtual pool of storage utilizing Nytro WarpDrives installed in each host. The storage volumes appear as a single block device available to each node within the cluster. Each local volume is broken into data chunks and distributed across the hosts in a fully balanced manner. This layout helps eliminate hot spots across the cluster and allows for the scaling of the overall I/O performance of the system through the addition of nodes or disks. The EMC ScaleIO architecture supports a distributed two-copy redundancy scheme. When a host or storage device fails, applications continue to access EMC ScaleIO volumes and data availability is maintained through the remaining mirrors.

PCIe Flash Reduces Storage Bottlenecks

VDI can be a demanding environment for storage. For example, when hundreds or thousands of users begin their day logging into their desktops, a “boot storm” occurs as the operating system and applications are simultaneously loaded stressing the virtual infrastructure. IT administrators are finding that HDD-based storage cannot support this demanding workload forcing users to wait far longer than their experience when booting directly from a laptop. LSI Nytro WarpDrive PCIe Flash is typically used to mitigate the excessive delays during a boot storm.

For this implementation, desktop linked clones and user data were pooled and protected by EMC ScaleIO utilizing LSI Nytro WarpDrive 3.2TB cards. These LSI Nytro WarpDrives use enterprise-grade eMLC NAND that offer more write cycles than competing PCIe based NAND solutions with average latencies well below 50 microseconds thus creating a solution far superior to SAN or HDD based solutions. Note: EMC ScaleIO is so flexible; user data is also available via HDDs for cost effectiveness.

ScaleIO Software

EMC ScaleIO software provides an elastic storage infrastructure that scales both capacity and performance as business requirements change. EMC ScaleIO is a software solution that utilizes a server based PCIe Flash, local HDDs and LAN interconnects to implement a high speed virtual SAN. ScaleIO also provides all the benefits of shared storage—but at a fraction of the cost and complexity.

VMware Horizon View

VMware Horizon View provides unified access to virtual desktops and applications that run in a secure, centralized datacenter and are accessible from a wide variety of devices. Horizon View allows IT to manage desktops, applications, and data centrally while increasing flexibility and customization at the

endpoint for the user. It enables levels of availability and agility of desktop services unmatched by traditional PCs at about half the total cost of ownership (TCO) per desktop.

Unlike other desktop virtualization products, Horizon View is a tightly integrated, end-to-end solution built on the industry-leading virtualization platform, VMware vSphere®. Horizon View allows customers to extend business continuity and disaster recovery features to their desktops and to standardize on a common platform, from the desktop through the datacenter to the cloud.

Storage Components

The LSI PCIe Nytro flash accelerator cards offer high performance with exceptionally low latency, and a low CPU burden. Nytro WarpDrive cards with PCIe® technology are designed to be easy to install. The Nytro WarpDrive card uses VMware drivers native to vSphere ESXi. A single Nytro Flash Accelerator card can do the job of hundreds of spinning hard disk drives

Networking Components

Mellanox ConnectX-3 40GE Network Interface Cards (NIC) and SX6036 40GE Switch deliver high bandwidth and industry-leading connectivity for performance-driven server and storage applications in Enterprise Data Centers.

Running Horizon View over ESXi 5.5 with ConnectX-3 40GE inbox driver delivers significant throughput and latency improvements resulting in faster access, real-time response and increased number of virtual desktops per server.

Tested Configuration

Utilizing ESX version 5.5 that includes inbox Mellanox 40GE driver and Horizon View version 5.3, one thousand (1,000) desktops were supported using only six Supermicro servers (Table 1 & Figure 2) with two Nytro WarpDrive cards and ConnectX-3 40GE NIC in each connected by Mellanox SX6036 40GE Switch. Each desktop consisted of Win 7 (64bit), 2GB RAM and a 40GB HDD.

The ScaleIO Data Client (SDC) is a lightweight device driver that provides access to the ScaleIO virtual SAN block devices. The SDC exposes block devices representing the ScaleIO volumes that are currently mapped to that host.

The ScaleIO Data Server (SDS) is a lightweight software component that is situated in each host that contributes local storage to the central ScaleIO virtual SAN. In this tested configuration, both the SDC and SDS are co-located on each of the six servers.

Server Configuration

Item	Description	Quantity
SYS-2027R-72RFTP-OEM	SuperServer, 2U, WIO, Redundant 920W, 24 DIMM Slots, onboard 2 x 10GbE SFP+, 1GbE Management Port, onboard 2208 SAS2 controller, 1 x FHHL Gen 3 PCIe x16, 2 x FHHL Gen 3 PCIe x8, 2 x LPHL PCIe x 8	1
Dual Processor Xeon CPUs	Ivy Bridge 10C E5-2680V2 2.8G 25M 8G/s QPI	2
384GB Total Memory	MEM-DR316L-HL01-ER18, Hynix 16GB DDR3	24
3.2TB PCIe Flash Storage Card	Nytro WarpDrive BFH8-3200 (3.2TB)	2
SSD, 100GB, Enterprise e-MLC	Intel S3700 Series 100GB, SATA 6Gb/s, HET MLC 2.5" 7.0mm, 25nm (for boot support)	1
Add-on Card for 40GbE	Mellanox MCX314A-BCBT, ConnectX-3 EN NIC, 40/56GbE, dual-port QSFP, PCIe3.0x8, Tall Bracket	1

Table 1: Solution Server Configuration

Reference Configuration




Figure 2. Benchmark Configuration of Workload Generators and Servers Supporting Desktops

1,000 VDI sessions were distributed across six servers each equipped with dual Xeon CPUs, an Ivy Bridge motherboard, 384GB of DDR3 memory, a Mellanox NIC and switch, and two 3.2TB Nytro WarpDrive flash cards. To support the mirroring of data across multiple servers, a 40GE network was installed. The workload generators communicated via a 10GE network. (Figure 2)

VDI Tested Parameters

There are a number of pain points that are important to be measured in a VDI environment.

User Response Time

This is the time required for users to wait until a file access is performed. To emulate a laptop experience for the user, response times should be in seconds.

VDI Deployment

This is the total time required to create the VDI sessions. This involves instantiating Windows 7 and an Office environment for each desktop. This can be a lengthy process when the number of desktops is in the hundreds.

VDI Boot

This is the total time to simultaneously boot the VDI sessions. This can be a lengthy process when desktop users arrive at work and sign on to their desktops at the same time of the day.

Test Results

User Response Time

The Login VSI benchmark creates the actual workload of a simulated user. A variety of typical desktop tasks are performed. Examples are loading a new document into Microsoft Word, compressing a document into a zip file (which briefly spikes CPU and disk I/O), opening a Web browser, editing a worksheet and loading a short video.

Maximizing the number of desktops sessions supported by a server (the VDI density) is a key metric in reducing CAPEX and OPEX. As more sessions are supported, user response times may lengthen as more I/O activity occurs. The benchmark will ramp the number of sessions while measuring the response times – with the goal of not exceeding a value known as the VSI Index Average.

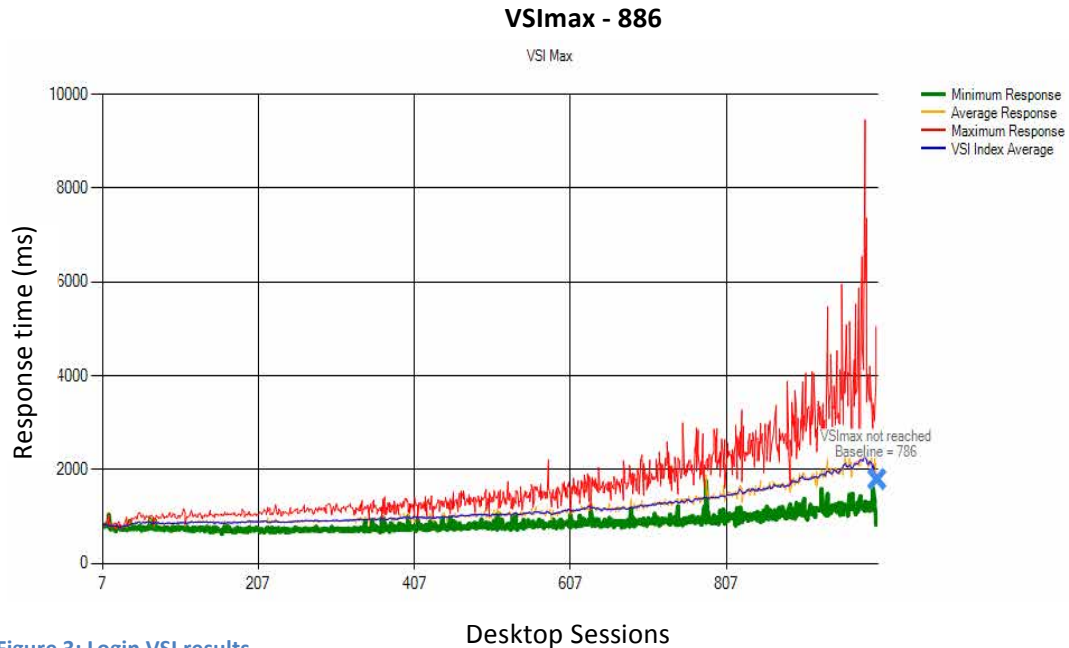


Figure 3: Login VSI results

The test results are used to calculate the maximum number of desktop sessions (VSI max) that may be supported while providing an acceptable response time. The baseline of 786 ms provides guidance on the best performance the system can give to an individual user. Note that the VSI max was not reached meaning that 1,000 sessions can be supported with an acceptable response time. (Figure 3)

Timing Test – 1000 Linked Clone Desktop Provisioning

In this tested configuration, a new desktop session was deployed in less than 0.09 seconds. This time is also representative of the time to redeploy sessions after a patch is implemented. Measured in the Horizon View Manager™ console, 1,000 desktops powered on in 86 minutes (Figures 4 and 5).

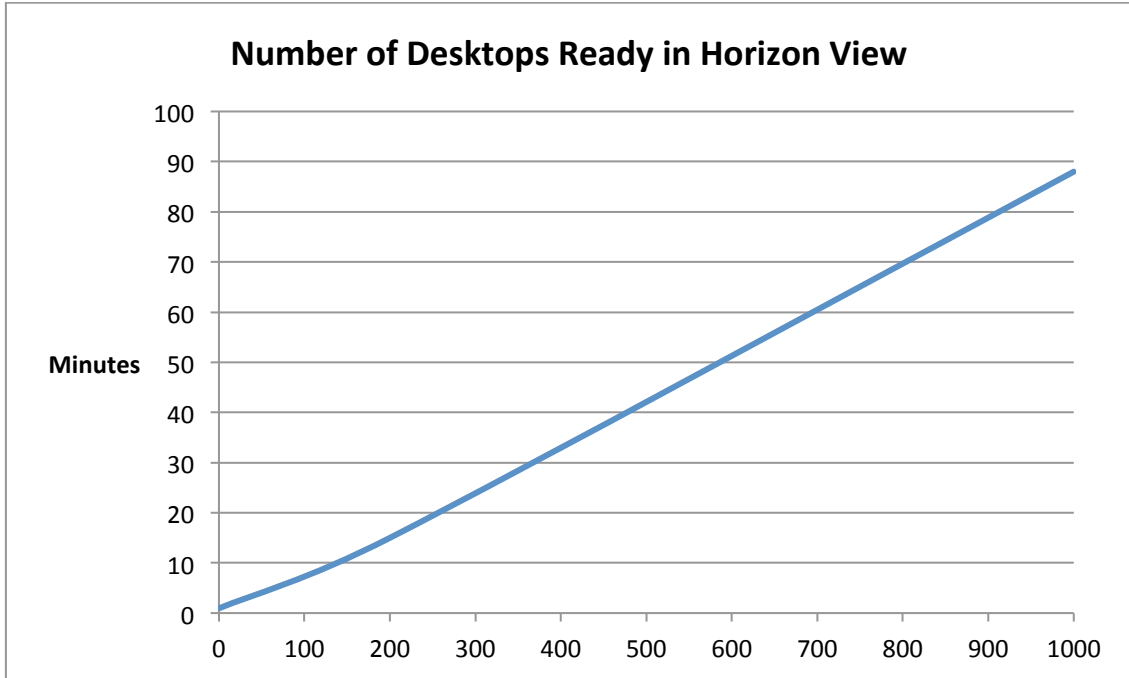


Figure 4: Desktop Provisioning

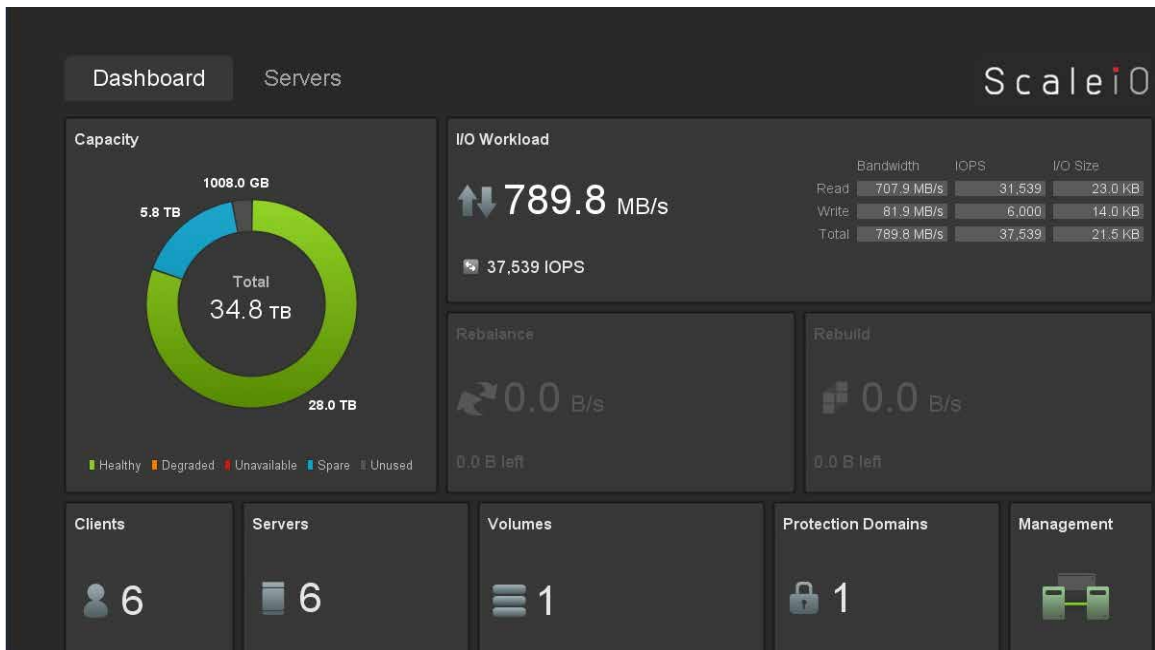


Figure 5: ScaleIO Performance during Desktop Provisioning

Timing Test – 1000 Linked Clone Desktop Boot Storm

When 1,000 desktops were booted simultaneously, 12 minutes were required for all to become usable. This is a worst case scenario as desktop users beginning their work day often may show up over a twenty minute period. In this more realistic situation, user’s desktop sessions will be up and running very rapidly (Figures 6 and 7).

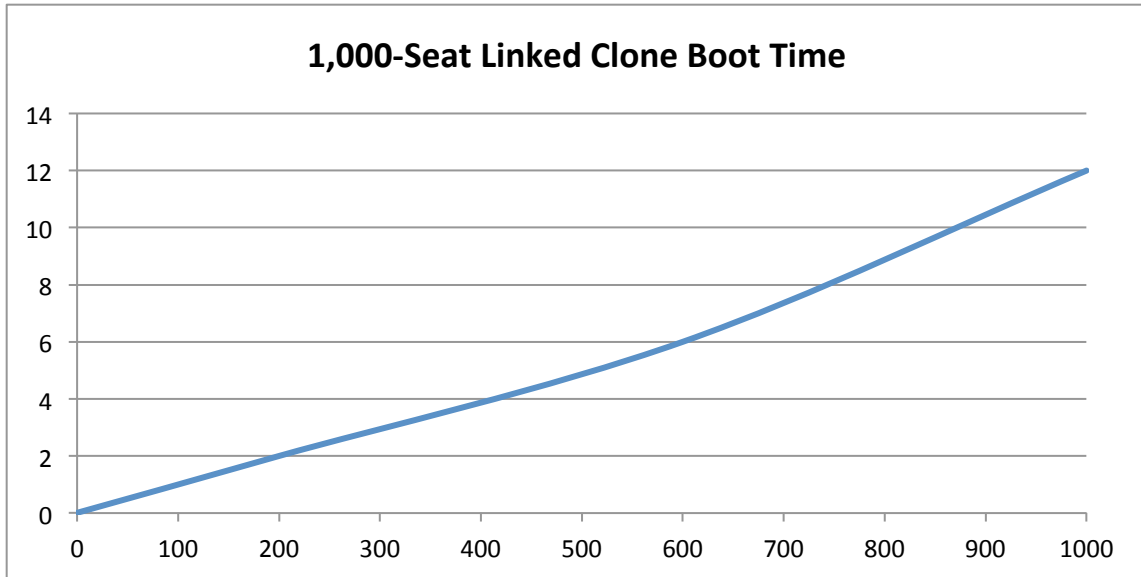


Figure 6: ScaleIO Linked Clone Boot Time

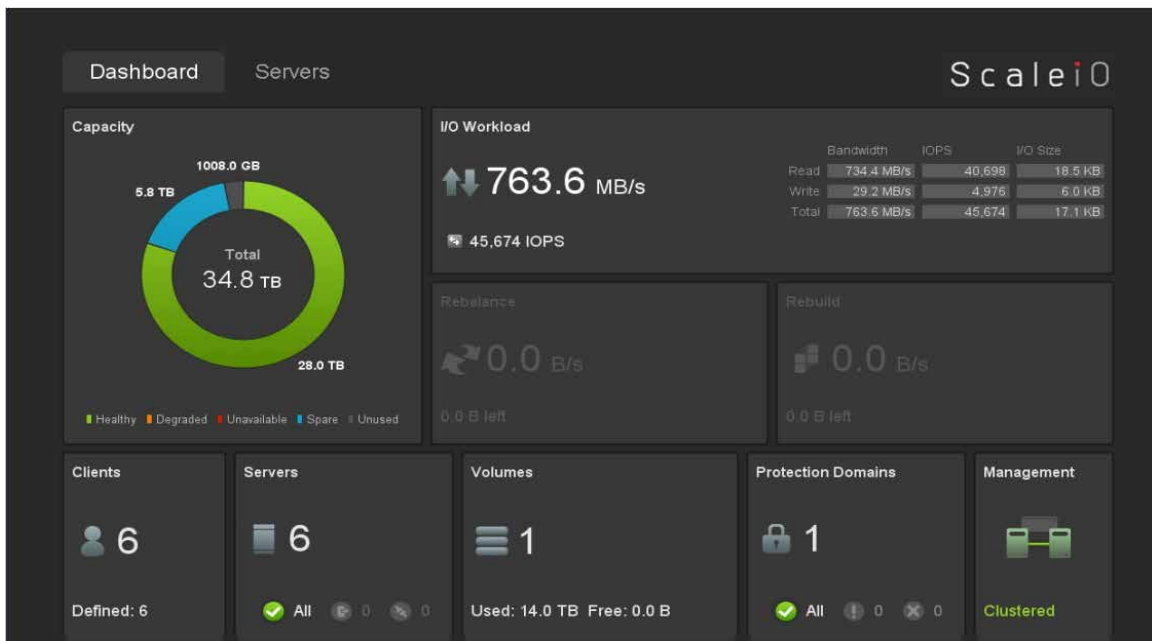


Figure 7: ScaleIO Performance during Desktop Boot

Note: As of Horizon View version 5.3, VMware supports managing up to 10,000 desktops with a single vCenter 5.5 server instance. This is a significant revision from prior versions.

All additional vCenter™ roles (inventory, SSO, vCenter) for the desktop vCenter were divested to separate servers to avoid any resource contention that might have resulted from combining roles on a busy vCenter Server.

All server resources were sized according to the current best practices from VMware. They are listed in the following table:

Infrastructure Server Role Configuration

SERVER ROLE	VCPU	RAM (GB)	STORAGE (GB)	OS
Domain Controller 1, DNS, DHCP	2	8	40	Windows Server 2008 64-bit R2
SQL Server 1	4	28	80, OS 350, Data root directory 50, Temp DB Directory 50, Backup Directory 50, User database log directory	Windows Server 2008 64-bit R2
SERVER ROLE	VCPU	RAM (GB)	STORAGE (GB)	OS
Horizon View Composer™ Server	2	18	40	Windows Server 2008 64-bit R2
vCenter Server Infrastructure and Desktop	4	28	50	Windows Server 2008 64-bit R2
Horizon View Connection Server 1	2	18	40	Windows Server 2008 64-bit R2

Table 2: Infrastructure Server Resource Sizing

Horizon View Configuration

The Horizon View installation included the following core systems:

- Four connection servers
- Dedicated vCenter with roles split to separate servers as follows: vCenter
 - vCenter single sign-on (SSO)
 - vCenter Inventory Service
- View Composer running on a separate server from vCenter

Note: Security servers were not used during this testing.

The basic Horizon View global settings are summarized below:

ATTRIBUTE	SPECIFICATION
GLOBAL POLICIES	
Multimedia redirection (MMR)	Allow
USB access	Allow
Remote mode	Allow
PCoIP hardware acceleration	Allow – medium priority

GLOBAL SETTINGS	
Session timeout	600 (10 hours)
SSO	Always Enabled
View Administrator session timeout	120 minutes
Automatic status updates	Enabled
Pre-login message	None
Warning before forced logoff	None
GLOBAL SECURITY SETTINGS	
Re-authenticate secure connections after network disrupt	Enabled
Message security mode	Enabled
Enable IPsec for Security Server pairing	Enabled

Table 3: Horizon View global settings

Desktop Pool Configurations

Linked-clone pool configurations conform to a typical knowledge worker profile with the exception of the storage overcommit policy.

ATTRIBUTE	SPECIFICATION
Pool type	Linked Clone
Persistence	Non Persistent
Pool ID	DesktopPool###
Display name	DesktopPool###
Folder ID	/
Separate datastores for replica and OS?	Not selected
State	Enabled
Connection Server restrictions	None
Remote desktop power policy	Take no power action
Auto logoff after disconnect	Never
User reset allowed	False
Multiple sessions per user allowed	False
Delete or refresh desktop on logoff	Never
Display protocol	PCoIP
Allow protocol override	False
Max number of monitors	2
Max resolution	1920 x 1200
HTML Access	Not selected
Flash quality level	Medium
Flash throttling level	Moderate
Enable provisioning	Yes
Stop provisioning on error	No
Provisioning timing	All up front
Use View Storage Accelerator	No
Reclaim virtual machine disk space	No
Storage overcommit policy	Aggressive*

Table 4: Desktop Pool Configurations

* Storage overcommit policy must be set to Aggressive with any storage platform that has an inline deduplication or data-reduction capability.

Cluster to Datastore Assignments

Each 1,000-seat pool was deployed across four datastores as summarized.

POOL NAME	DATASTORE NAME	DATASTORE SIZE
Cluster1 – Pool1		
1,000 desktops	Datastore 1	14TB

Table 5: Datastores

Test Image Configuration

We configured the virtual hardware of the master desktop virtual machine according to standard Login VSI specifications. It is important to note that in production deployments, virtual machine configurations vary based on individual use-case requirements.

ATTRIBUTE	SPECIFICATION
Desktop OS	Microsoft Windows 7, 64-bit
VMware virtual hardware	Version 9
VMware Tools version	1280544 (up to date)
Virtual CPU	1
Virtual memory	2048MB
OS pagefile	1.5GB starting and maximum
vNICs	1
Virtual network adapter 1	VMXNet3 Adapter
Virtual SCSI controller 0	LSI Logic SAS
Virtual Disk – VMDK	40GB
Virtual Floppy Drive 1	Removed
Virtual CD/DVD Drive 1	Removed
VMware Horizon View Agent	VMware Horizon View Agent 5.3 build 1427931
Installed applications (per Login VSI standards)	Microsoft Office 2010 Adobe Acrobat Reader 9.1 Adobe Flash Player 11

	Adobe Shockwave Player 10 Bullzip PDF Printer 9.1 FreeMind 0.8.1 Kid-Key-Lock
--	--

Table 5: Test Image Configuration

The master image used in this test environment underwent VDI optimization as defined in the VMware Horizon View Optimization Guide for Windows 7. VMware strongly recommends that the Windows image be optimized when master images are prepared for use with VMware Horizon View.

Conclusion

A VDI workload is often bottlenecked by HDD storage performance. In many cases, IT administrators use expensive all flash-based SANs and servers to solve this dilemma. EMC ScaleIO software eliminates the need for complicated and expensive infrastructures by using cost-effective server based storage instead. In this evaluation, an industry standard benchmark demonstrated that 1,000 VDI desktops are supported using only six servers connected over a Mellanox 40GE end-to-end solution. Creating a virtual pool of high performance storage with EMC ScaleIO and Nytro WarpDrive Flash Accelerator cards significantly accelerated VDI boot times, improved desktop performance and allowed for higher VDI density per node.

About EMC ScaleIO

EMC ScaleIO is the leader in elastic converged storage. The EMC ScaleIO solution enables data centers to cut storage costs by over 80% compared to similar configurations utilizing SAN storage. EMC ScaleIO now creates shared virtual storage pools using application servers and their locally attached storage over high speed IP data networks to easily create scale out high performance infrastructures.

About LSI Nytro

LSI PCIe Nytro Flash Accelerator cards accelerate server applications regardless of the mix of reads and writes. A single Nytro Flash Accelerator card can do the job of hundreds of spinning hard disk drives (HDDs) while significantly reducing power, cooling and physical space requirements for a lower TCO. Visit www.lsi.com for more information

About Login VSI

Login Virtual Session Indexer (Login VSI) is a vendor independent benchmarking tool to objectively test and measure the performance and scalability of Virtual Desktop Infrastructures and Server Based Computing environments by simulating unique user workloads.

Leading IT-analysts recognize and recommend Login VSI as the de-facto industry standard benchmarking tool for VDI and SBC. Login VSI can be used to test VMware Horizon View, Citrix XenDesktop and XenApp, Microsoft Remote Desktop Services or any other Windows based hosted desktop solution. Login VSI is the standard tool used in all tests that are executed in the internationally acclaimed Project Virtual Reality Check. Visit www.loginvsi.com for more information

About Mellanox

Mellanox Technologies is a leading supplier of end-to-end InfiniBand and Ethernet interconnect solutions and services for servers and storage. Mellanox interconnect solutions increase data center efficiency by providing the highest throughput and lowest latency, delivering data faster to applications and unlocking system performance capability. Mellanox offers a choice of fast interconnect products: adapters, switches, software, cables and silicon that accelerate application runtime and maximize business results for a wide range of markets including high performance computing, enterprise data centers, Web 2.0, cloud, storage and financial services. More information is available at www.mellanox.com.



EMC, the EMC logo, and ScaleIO are registered trademarks or trademarks of EMC Corporation in the United States and other countries. LSI Nytro is a registered trademark of LSI Corporation. Supermicro and SuperServer are registered trademarks of Supermicro Computer, Inc.

All other trademarks used herein are the property of their respective owners.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

Copyright © 2014 VMware, Inc. All rights reserved. This product is protected by U.S. and international copyright and intellectual property laws. VMware products are covered by one or more patents listed at

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies. VMware believes the information in this document is accurate as of its publication date. The information is subject to change without notice.