


Virtual Desktop Infrastructure on VMware Horizon 6* with View* using Maxta® Storage Platform

Maximize the Promise of Hyper-Convergence



Intel® Xeon® Processor E5-2600 v3
Product Family



Audience and Purpose

A growing number of businesses today are adopting virtualization technology in their data centers for consolidation, efficiency, ease of management, and greater mobility. Drawn by this success and the growing challenge of managing desktops, businesses are exploring Virtual Desktop Infrastructure (VDI) to extend the benefits of virtualization to their desktop environment.

The goal of this document is to provide a tested Hyper-Converged configuration of a desktop virtualization solution with Maxta® Storage Platform (MxSP™) on VMware Horizon* 6.0 with View*, running on Intel® servers, processors and PCIe SSDs. This solution will enable fast deployment with high storage efficiency and low cost. This MaxDeploy™ reference architecture is not intended to be a comprehensive guide to deployment and configuration for every aspect of this solution.




Table of Contents

Executive Summary 3

Introduction 3

The Solution: Tested and Validated Infrastructure for Desktop Virtualization 4

 Base Components 4

 Overview 4

 Key Benefits 5

 Architecture 6

 Maxta Metadev and SSD Caching 6

 Intel Components 7

 Intel® Server board S2600WT-based System 7

 Intel® Xeon® Processor E5-2600 v3 Product Family 7

 Intel® P3700 Solid-State Drive Data Center Family 7

 Intel developed NVMe driver for ESXi 5.5 7

 Mellanox* Ethernet Interconnects 8

 VMware vSphere 5.5* 8

 VMware Horizon 6.0 with View 8

 Login VSI 4.1* 9

 MxSP with VMware – Your best VDI solution 9

Architecture and Design of the Solution 10

Test Procedure 10

 Login VSI workload characteristics 10

 Login VSI measurements and graphs explanation 10

Test Results 11

Conclusion 11

Appendix A: Architecture and Design of the Solution 12

 Maxta Storage Platform 12

 Intel S2600WT Server Configuration 12

 VMware vSphere 5.5 Configuration 13

 VMware Horizon 6.0 with View Configuration 13

 Login VSI 4.1 Configuration 14

 Networking Configuration 15

Appendix B: Test Results 16

 Test 1: 100 Desktops – 2 Hyper-Converged Compute and Storage Nodes + 1 Compute-Only node 16

 Test 2: 300 Desktops – 3 Hyper-Converged Compute and Storage Nodes 17

 Test 3: 456 Desktops – 4 Hyper-Converged Compute and Storage Nodes 18

 Test 4: High Availability – Recover 100 desktops after single node failure 19

Executive Summary

Virtualized desktop environments have radically changed storage requirements for the IT environment. Setting up a virtual desktop infrastructure can be complicated, with the most vital and difficult part being the design of the supporting storage.

A MaxDeploy appliance with Maxta Storage Platform running on top of Intel servers helps to address these problems. The solution presented in this reference architecture reduces the cost of a VDI solution by aggregating storage resources across multiple servers and eliminating the need for storage arrays.

A well-designed and tested VDI configuration with MxSP can:





- Deliver a good end-user experience
- Scale the solution effortlessly as demand grows
- Simplify virtual desktop management
- Lower costs on VDI Deployments

Introduction

This reference architecture details a solution implementing Maxta storage software in a VMware Horizon 6.0 with View virtual desktop environment. The Intel® Server Board S2600WT family, using the Intel® Xeon® E5-2600 v3 processor family, with Intel® Solid-State Drive Data Center Family was used to form a cluster of hyper-converged nodes with capacity and flash optimized by Maxta Storage Platform. Mellanox* Ethernet adapters and switches were used as the cluster interconnect layer.

The configuration defined in this reference paper supports 456 virtual desktops, but Maxta Storage Platform (MxSP) provides the ability to scale to thousands of virtual desktops simply by adding MaxDeploy appliances to the environment. Metrics were collected for application response times as the system was being loaded with simulated users running active workloads. This reference architecture showcases how MxSP provides linear scalability without compromising performance.

Summary of results:

<p>VSImax v4.1 average: 995 ms 456 user sessions; VSIMax v4.1 not reached **Figure 10 in Appendix B</p>		<p>Faster application response times than industry average</p>
<p>8 minutes Recover 100 desktops after host failure **Test 4 in Appendix B</p>		<p>Support for HA minimizes system downtime</p>
<p>58.81% peak CPU usage 36 CPU x 2.294 GHz per node **Figure 11 in Appendix B</p>		<p>CPU is not limiting resource in scale up scenarios</p>
<p>33.97% peak Network usage 10 GbE private storage network **Figure 12 in Appendix B</p>		<p>Network bandwidth provides room to grow</p>

Maximize the Promise of Hyper-Convergence

The Solution: Tested and Validated Infrastructure for Desktop Virtualization

This reference architecture utilizes the latest data center technology, encompassing compute, storage, networking, and virtualization technologies.

Base Components

- Maxta Storage Platform version 2.4.1
- 4 Intel S2600WTT Servers
 - » 2 Intel Xeon E5-2699 v3 Processors per server
 - » 1 Intel SSD DC P3700 per server
- VMware vSphere* 5.5.2
- VMware Horizon 6.0 with View
- Mellanox SwitchX-2 SX1012 switch and ConnectX-3 Pro NIC
- Login VSI 4.1

Maxta Solutions

Overview

Maxta provides its solutions in two form factors: Maxta MaxDeploy™ Appliances and Maxta Storage Platform (MxSP™). MxSP is a highly resilient, scalable, distributed, software-defined VM storage platform that turns industry standard servers into a converged compute and storage solution, leveraging server-side flash and disk drives to

optimize performance and capacity. MxSP addresses the challenges related to storage for a virtual desktop infrastructure and simplifies it while delivering enterprise-class services, and CAPEX as well as OPEX savings. MxSP provides VM-level storage abstraction and full integration into the server virtualization management scheme and user interface. MxSP intelligently maps VMs to storage resources, optimizing data layout for virtual workloads and leverages SSDs for read/write caching.

The MxSP distributed architecture enables shared storage with enterprise-class data services such as snapshots, clones, thin provisioning, compression, de-duplication, replication and full scale-out without performance degradation. Additionally, Maxta supports advanced capabilities such as live migration of virtual machines, dynamic load balancing, high availability, data protection, and disaster recovery. This results in simplification of IT and significant cost savings by enabling the elimination of costly and complex storage arrays.

Maxta's MaxDeploy Appliances deliver an easy way of deploying hyper-converged solutions for the virtual data center by combining Maxta Storage Platform (MxSP) along with hardware and other infrastructure software platforms. This removes interoperability and performance guesswork and simplifies the ordering process. MaxDeploy appliances are available for all major server vendors and white box server platforms.

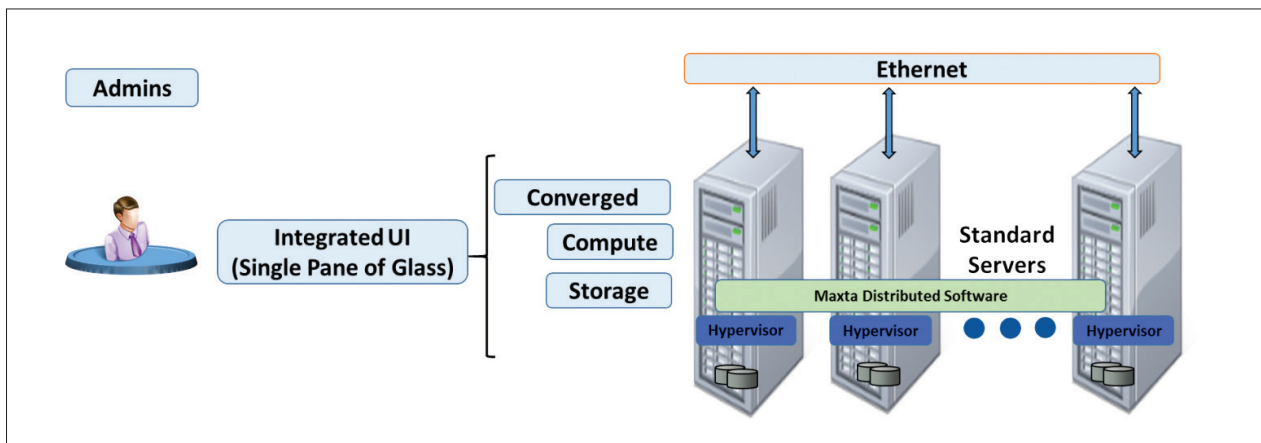


Figure 1: Maxta's simplified approach to storage in the virtual datacenter

The versatility of these solutions provides the ability to scale capacity and performance independently, on-demand, and without having to over-provision resources. Maxta solutions are hypervisor-agnostic and fully integrate with server virtualization at all levels from user interface to data management, while supporting all possible deployments of virtual data centers, including private, public, and hybrid clouds.

Key Benefits

Eliminate performance issues associated with traditional storage and deliver good end-user VDI experience

Delivering an end-user experience that is better than the traditional desktop deployment is critical for a successful VDI deployment. MxSP leverages SSDs as read/write cache to alleviate the performance issues commonly associated with VDI environments.

The boot image and frequently accessed information ("hot data") is cached on the SSDs and enables all of the desktops to boot from the SSDs without performance degradation. Additionally, updates are written and acknowledged from the SSDs to minimize write latency.

MxSP leverages HDDs to deliver capacity to the virtual desktops. This approach delivers the required performance and capacity at relatively low cost.

Simplify virtual desktop management

Administrators spend time and resources to manage, provision, patch, and update the thousands of desktop images. The ability to create thousands of snapshots or clones in minutes is paramount for VDI deployments. MxSP delivers an unlimited number of VM-level snapshots and zero-copy clones that can be used to provision the virtual desktops. These snapshots and clones are time, capacity, and performance efficient.

Upgrading or patching these virtual machines is as simple as incorporating the changes to the base or golden image and creating new snapshots or clones that can be used by the virtual desktops. MxSP provides flexibility and simplifies the ability to scale performance and capacity independently on-demand by adding industry standard servers.

MxSP further simplifies IT by aligning the virtual machine constructs with storage constructs and delivering all enterprise-class functionality at VM-level granularity, thus simplifying virtual desktop management. IT administrators

are able to manage the entire virtual desktop infrastructure through View and vCenter without the need for any storage management interface.

Maximize savings on VDI deployments

Storage has been the most expensive component in the data center, especially with networked storage solutions. MxSP delivers significant savings by leveraging industry standard components, while still delivering the required performance for VDI deployments. By using SSDs as read/write cache, MxSP delivers that performance without having to invest in an all-SSD appliance, significantly reducing the cost of VDI deployments.

Snapshots and clones do not consume any capacity when created and there is no upfront space reservation for configuring snapshots or clones. Additionally, creating snapshots or clones does not impact the performance of the primary virtual machine. Combining with capacity optimization features such as thin provisioning, compression, and de-duplication provides the ability to create thousands of virtual desktops with huge capacity savings.

In addition, MxSP provides significant OPEX savings in the area of virtual desktop management, eliminating the need for specialized skills for array-specific storage management and storage networking.

Maximize the Promise of Hyper-Convergence

Architecture

Below is a high level illustration of how the Maxta Storage Platform fits into virtualized environments. As shown in Figure 2, an instance of the Maxta Storage Platform software is installed on each of the servers that are part of a server virtualization cluster. MxSP creates a Maxta Global Namespace that all of the virtual machines in the cluster can access. Maxta Storage Platform supports two types of servers:

- Converged compute/storage servers have access to the Maxta Storage Platform storage pool and also contribute storage to it.
- Compute-only servers have access to Maxta Storage Platform storage pool but do not contribute storage to it.

This approach allows for simplistic scale of compute and memory resources when additional storage is not needed.

Maxta Storage Platform's management functionality is seamlessly integrated into the virtualization user interface, eliminating the need for a separate storage management interface. MxSP leverages any combination of SSDs and magnetic disk drives on standard servers to deliver flash performance and hard disk capacity, while meeting various price point targets.

MxSP intelligently optimizes the mapping of VMs to storage resources. The storage for a given virtual machine in most scenarios is local to the server where the VM is deployed. Under normal operation, both reads and writes are served

by the server hosting the VM. Intelligence is built into the software to minimize the probability that a VM and its associated storage are on separate servers. This eliminates the network round-trip latency associated with distributed storage systems that don't optimize the placement of VM data.

All data stored on Maxta storage is protected with strong checksums and is mirrored across at least two physical servers to provide high levels of data integrity and data availability. Additional levels of mirroring can be utilized within each physical server as well.

Maxta Metadev and SSD Caching

Maxta Metadev is a class of device in the Maxta file system. Metadev accelerates the read/write performance of metadata I/O's by storing only the file system metadata. Maxta deployments typically use small page sizes (4KB or 8KB) for space efficiency and performance. Smaller page sizes result in larger amounts of metadata and that creates random read/write I/O performance issues, since the metadata can't be cached entirely in memory.

Metadev is a way to separate metadata from data, allowing the metadata to be stored on high performance devices (like SSDs). This approach significantly accelerates metadata read and write operations. Metadev stores metadata such as file indirect blocks and file system's space allocation records, which provide key information to the actual data itself.

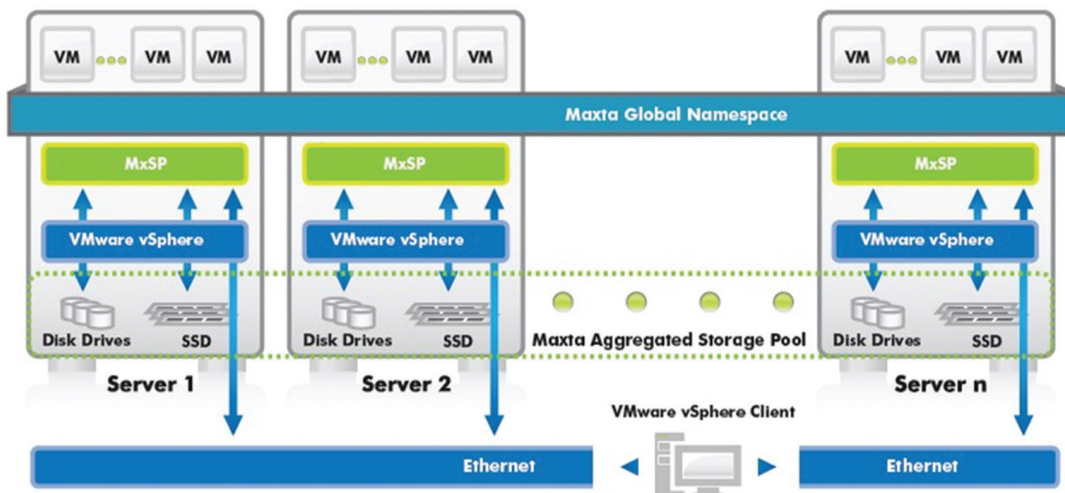


Figure 2: Maxta Storage Platform Architecture

Metadev places a requirement on the HDD to SSD ratio. The ratio depends specifically on the page size defined during the Maxta Storage Platform deployment. If the deployment is using a 4KB page size, the recommended Metadev capacity is at least 5% of the node's HDD capacity. If the deployment is using 8KB page size, then it's at least 3% of the node's capacity.

There are two types of Metadev: Shared and Exclusive. Exclusive Metadev is allocated on SSDs that do not have read/write-back cache partitions. Shared Metadev can be created on SSDs that also host read/write-back cache partitions. Exclusive Metadev is available as an option to reduce the wear on the Metadev SSDs. eMCL/PCIe SSDs that support at least 10 drive writes per day for 5 years can be configured to use Shared Metadev.

Users cannot configure both Shared and Exclusive Metadev at the same time. During the Maxta Storage Platform installation, users can select either an Exclusive or a Shared Metadev configuration. If two SSDs are available and the user selects "Enable Metadev", an Exclusive Metadev will be configured with one SSD assigned to Exclusive Metadev and the other to read/write-back cache. If you cannot meet the minimum Metadev SSD ratio requirements, Maxta will still utilize the available SSDs automatically for read/write-back cache.

Intel Components

Intel® Server board S2600WT-based System

Intel® Server board S2600WT family of products are designed to support the performance and features of dual Intel® Xeon® processors E5-2600 v3, Intel® TXT Technology, Intel® Xeon Phi™ coprocessors and Intel® Solid State Drives. Systems include preconfigured 1 and 2U rack models, or disaggregate chassis and boards for greater flexibility. For more details, please refer to <http://ark.intel.com/products/82156/Intel-Server-BoardS2600WTT?q=2600WTT>.

Intel® Xeon® Processor E5-2600 v3 Product Family

The Intel® Xeon® processor E5-2600 v3 product family helps IT address the growing demands placed on infrastructure, from supporting business growth to enabling new services faster, delivering new applications in the enterprise, technical computing, communications, storage, and cloud segments. This new generation of processors enables powerful, agile data centers by supporting a software-defined infrastructure (SDI) to address the imminent need of greater flexibility with

higher levels of automation and orchestration. In addition, the Intel Xeon processor E5-2600 v3 product family delivers significant benefits in performance, power efficiency, virtualization, and security. Combining these benefits with a low total cost of ownership and Intel's acclaimed product quality, the Intel Xeon processor E5-2600 v3 product family enables compelling solutions for any organization. For more details, please refer to

<http://www.intel.com/content/www/us/en/processors/xeon/xeon-processor-e5-family.html?wapkw=Xeon%20E5&wapkgw=featured>

Intel® P3700 Solid-State Drive Data Center Family

Intel SSDs DC family for PCIe provides fast, low-latency, unwavering data streams directly to Intel® Xeon® processors, making server data transfers extremely efficient. Additionally, the new Non-Volatile Memory Express (NVMe) storage interface standard is engineered for current and future Non-Volatile Memory (NVM) technologies. NVMe overcomes SAS/SATA SSD performance limitations by optimizing hardware and software to take full advantage of the SSD's capabilities. Intel Xeon processors efficiently transfer data in fewer clock cycles with the NVMe optimized software stack compared to the legacy Advanced Host Controller Interface (AHCI) stack, reducing latency and overhead. Direct CPU connections also eliminate Host-Bus-Adapter (HBA) cards, further reducing system latency and CapEx. By combining SSD NAND management techniques and NAND silicon enhancements, High Endurance Technology (HET) enables the DC P3700 Series to achieve 10 total drive writes per day over a 5 year drive life. For more details, please refer to <http://www.intel.com/content/www/us/en/solid-state-drives/datacenter-family.html>.

Intel developed NVMe driver for VMware ESXi 5.5

The Intel developed NVMe driver compatible with ESXi 5.5 was used in this reference architecture. Intel is actively working to certify Intel PCIe SSDs and Intel NVMe drivers with ESXi 5.5 and future releases. Once certified, Intel SSD DC P3700 Series and others will be listed on the VMware Compatibility Guide and the NVMe driver will be publicly available.

Maximize the Promise of Hyper-Convergence

Mellanox* Ethernet Interconnects

The Mellanox SwitchX-2 SX1012 switch is a high performance top-of-rack networking solution in a half-width 1U form factor. The switch provides 12 QSFP interfaces that can operate at speeds of 1GbE, 10GbE, 40GbE, or 56GbE.

Additionally each QSFP interface can be used with a passive breakout cable to create four discrete SFP+ interfaces as 10GbE ports, creating a maximum of 48 10GbE ports.

The Mellanox switch provides a smaller footprint and more flexible configuration, which improves utilization of the hardware and lowers power consumption. This all comes together to provide a networking solution that significantly reduces capital and operational costs.

ConnectX-3 Pro adapter cards with 10/40/56 Gigabit Ethernet connectivity provide a high performing and flexible interconnect solution for PCI Express Gen3 servers used in public and private clouds, enterprise data centers, and high performance computing.

VMware vSphere 5.5*

VMware vSphere is an industry-leading virtualization platform for building cloud infrastructures. It enables IT to meet SLAs for the most demanding business critical applications, at the lowest TCO.

vSphere can accelerate the shift to cloud computing for existing data centers and also underpins compatible public cloud offerings, forming the foundation for a hybrid cloud

model. With the support of more than 3,000 applications from more than 2,000 ISV partners, vSphere is a trusted platform for many applications.

VMware Horizon 6.0 with View

Horizon 6.0 allows IT to centrally manage virtual, physical, and BYO Windows images in order to streamline management, reduce costs, and maintain compliance. With Horizon 6.0, virtual or remote desktops and applications can be delivered through a single platform to end users. These desktop and application services— including RDS hosted apps, packaged apps with ThinApp, SaaS apps, and even virtualized apps from Citrix— can all be accessed from one unified workspace to provide end users with the resources they want, at the speed they expect, with the efficiency to meet business demands.

Feature Overview and What's New:

- Central Management of Virtual, Physical, and BYO Images
- Desktops and Applications Delivered Through a Single Platform
- Unified Workspace with User Experience
- Closed-Loop Management and Automation
- Optimized for the Software-Defined Data Center

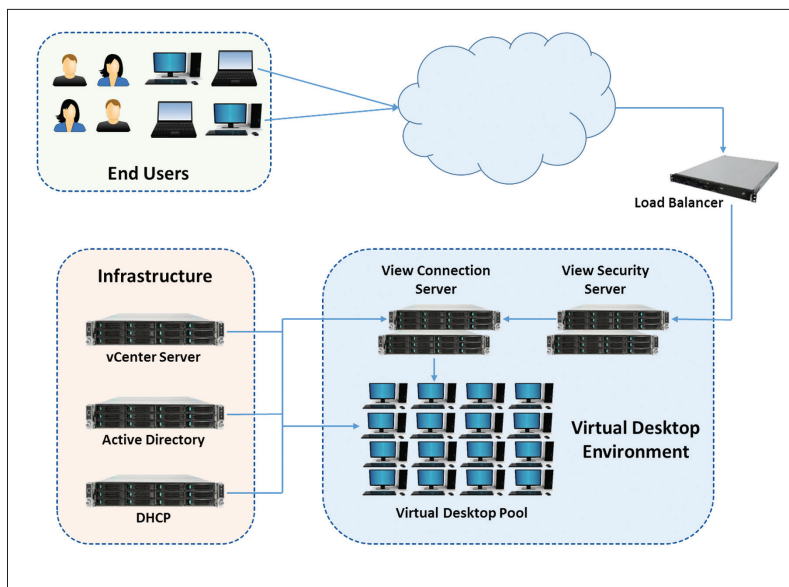


Figure 3: Virtual Desktop Infrastructure implementing VMware Horizon with View

Login VSI 4.1

Login VSI, Inc. (<http://www.loginvsi.com>) provides proactive performance management solutions for virtualized desktop and server environments. Enterprise IT departments use Login VSI products in all phases of their virtual desktop deployment— from planning to deployment to change management—for more predictable performance, higher availability and a more consistent end user experience. Virtualization vendors use the flagship product, Login VSI, to benchmark performance. With minimal configuration, Login VSI products works in VMware Horizon View, Citrix XenDesktop* and XenApp*, Microsoft* Remote Desktop Services (Terminal Services) and any other Windows-based virtual desktop solution.

For more information, download a trial at <http://www.loginvsi.com>.

MxSP with VMware

Maxta's VM-level clones make it fast and simple to provision virtual desktops. MxSP leverages SSDs as read/write cache to alleviate the performance bottlenecks common to VDI deployments. The innovative, peer-to-peer architecture aggregates storage resources from multiple industry standard servers, assimilating a global namespace and all of the storage functionality that a virtual desktop environment requires.

MxSP provides storage with VM-level granularity, which eliminates the need for storage management, allowing VDI administrators to focus entirely on vCenter and View management. This creates significant OPEX savings in VDI environments since there is no need for storage management and storage networking specialization.

The key benefits of MxSP for a VDI deployment are:

- **Eliminate** storage arrays and storage networking
- **Improve** performance compared to traditional storage
- **Deliver** good end-user experience
- **Simplify** virtual desktop management
- **Provide** VM-centric enterprise-class data services
- **Boost savings** on VDI deployments

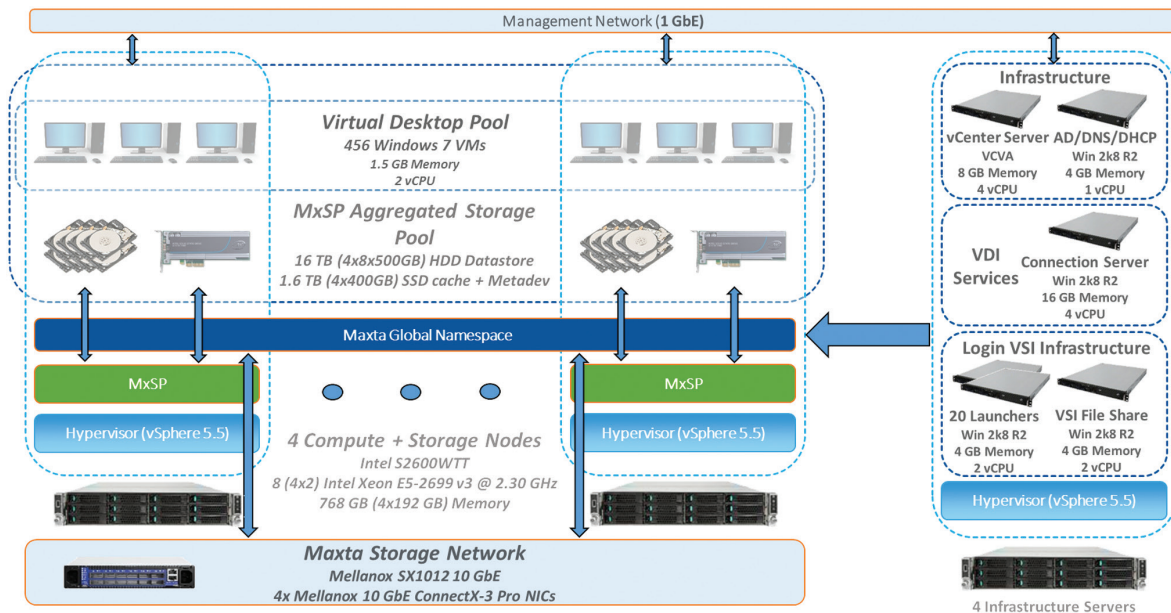


Figure 4: Solution Topology

Architecture and Design of the Solution

Maxta labs hosted the VDI environment under an isolated 1GbE management network. Four hyper-converged Intel S2600WTT servers were configured for this targeted 456 desktop test. Each server was equipped with two Intel Xeon E5-2699 v3 processors and 192GB memory. Four additional servers were utilized to support the infrastructure.

The MxSP aggregated storage pool consisted of one Intel SSD DC P3700 400GB SSD and eight 500GB HDDs per server, providing a total capacity of 17.6TBs. Shared Metadev was enabled while all other MxSP features were left as default, including allocation of 4 vCPU and 8GB memory for the MxSP controller VM. The storage pool hosted all 456 Windows 7 VMs, each with 1.5GB and 2 vCPU to support the Login VSI Knowledge Worker.

The Maxta Storage Network was configured on a Mellanox SwitchX-2 SX1012 switch with each server having a Mellanox 40 GbE ConnectX-3 Pro NIC. For purposes of this test, the private storage network was configured to have a maximum bandwidth of 10GbE.

The infrastructure nodes hosted one Login VSI File Share and 20 Launcher VMs in addition to one vCenter Server Appliance, one View Connection Server, and one virtual machine running Active Directory, DNS, and DHCP services. Infrastructure nodes were connected to the hyper-converged nodes over the isolated management network.

In a production environment, additional switches, View Connection servers, and other resources should be utilized to support all high availability and failover scenarios. For purposes of this document, HA functionality has been limited to that which is provided by default via MxSP replicas and vSphere clustering.

Full details of the MaxDeploy testbed configuration can be found in Appendix A.

Test Procedure

Login VSI workload characteristics

The tests covered in this document were executed using a Login VSI knowledge worker workload. This workload is designed to get accurate, consistent measurements for stress and performance testing in virtual desktop environments.

All Login VSI workloads are divided into four 12-minute segments, with each segment simulating normal user activities such as multimedia, productivity, and information consumption tasks. The end of every segment consists

of a maximum two minute idle buffer, which generally decreases as the workload slows down due to performance degradation. Workloads are compatible with all operating systems later than Windows XP and all Microsoft Office versions more recent than 2003. Documents and applications are closed and cleaned up to ensure that workloads can run continuously.

The knowledge worker is derived from the medium workload used in previous versions of Login VSI. This new workload is more stable and smooth while still generating statistics that are comparable to those of previous Login VSI versions. The knowledge worker runs 8.5 IOPS with 5-9 applications open, reaching a CPU usage of 100%. This workload requires a minimum of 1.5GB memory and 2 vCPU.

Login VSI measurements and graphs explanation

Login VSI performs benchmarks by loading the system with simulated user workloads, gradually increasing the amount of simulated users until the system is saturated. At this saturation point, there will be a significant increase in response time of the applications. The latency experienced as we get closer to overloading the system makes it possible to find out what the true maximum user capacity is.

The following statistics are collected during every knowledge worker workload:

- **VSI_{max} v4.1:** Amount of active sessions reached before the system is saturated.
- **VSI_{base}:** Best performance of the system during a test.
- **VSI_{max} v4.1 threshold:** Point at which the environment's saturation point is reached.
- **Minimum Response:** Minimum response time for all measurements taken when the number of sessions indicated on the X axis were active.
- **Average Response:** Average response time for all measurements taken when the number of sessions indicated on the X axis were active.
- **Maximum Response:** Maximum response time for all measurements taken when the number of sessions indicated on the X axis were active.
- **VSI Index Average:** Average response value calculated by VSI. Applies statistical rules to avoid spikes from influencing the average too much.

Test Results

A summary of the Login VSI measurements, CPU usage, and other key metrics is provided in this section. Detailed results and graphs can be found in Appendix B.

Three different scenarios were tested to represent varying levels of scale and resource consumption. The minimal configuration included 100 users total across two servers running a knowledge worker workload. The medium scale configuration consisted of 300 total users running a knowledge worker workload across three nodes. The large scale configuration had 456 total users across four nodes, also running a knowledge worker workload.

As the table above shows, Maxta Storage Platform maintains linear scalability as nodes and desktops are added to the VDI environment. Response times remain nearly constant, showing MxSP’s ability to support thousands of desktops simply by adding more nodes. Intel S2600WT servers also maintain low CPU usage percentages even as scale increases.

Conclusion

Maxta solutions provide significant cost, management, and performance improvements in the VMware View VDI environment. Rapid creation of Maxta clones into a virtual desktop pool makes the deployment process quick and

straightforward. The 4-node MaxDeploy configuration with Intel S2600WT servers supports 456 virtual desktops while maintaining high levels of performance according to industry standard metrics.

Maxta solutions provide organizations the choice to deploy hyper-convergence on any x86 server, and use any hypervisor and any combination of devices, allowing those organizations to exploit the potential of VDI deployments. The simplicity of Maxta’s VM-centric solution reduces IT management to further maximize cost savings. Hyper-scale, enterprise-level data services and capacity optimization empower organizations to hyper-converge, eliminating the need for SAN or NAS devices. Maxta solutions can be purchased as a MaxDeploy fully configured hyper-converged appliance or as MxSP software for flexible configurations. Maxta MxSP and MaxDeploy appliances deliver a compelling end-user experience, simplify virtual desktop management, and significantly reduce capital and operational costs.

For more information on each product,

Maxta: <http://www.maxta.com/>

Intel: <http://www.intel.com/storage>

Maxta Storage Platform (MxSP)			
Number of Desktops	100	300	456
Number of Nodes (ESXi hosts)	2+1*	3	4
Number of stuck or unresponsive desktops (Sessions)	0	0	0
Controller VM configuration (CVM)	4vCPU/8GB Memory		
Performance Metrics			
Minimum response time (VSIbase)	731 ms	797 ms	843 ms
Average response time (VSImax Average)	823 ms	929 ms	995 ms
Maximum response time (VSI Max Threshold)	1731 ms	1798 ms	1844 ms
CPU Usage			
Peak	27.01 %	28.54%	58.81%
Average	18.42%	19.37%	40.27%

Figure 5: Performance Metrics * 2 Hyper-Converged Compute and Storage Nodes + 1 Compute-Only Node

Appendix A: Architecture and Design of the Solution

Maxta Storage Platform

Maxta Storage Platform	
MxSP version	2.4.1
# of nodes	4
Maxta Controller VM	
CPU	4
Memory	8GB
MxSP settings	
Page size	4K
Metadev	Enabled Shared
Local copies	Disabled
Rack awareness	Disabled
Disk usage	Hybrid
Total capacity	17.6TB

Intel S2600WT Server Configuration

Intel servers	
Number of servers	4
Server type	Intel S2600WTT
CPU	2 Intel Xeon E5-2699 v3 @ 2.30 GHz
Memory	192GB
SSD	1 NVMe 400GB
HDD	8 SATA 500GB
ESXi version	5.5.2
Public management network	Intel X540-AT2 1GbE
Maxta private storage network	Mellanox ConnectX-3 Pro 10GbE

VMware vSphere 5.5 Configuration

VMware vSphere	
ESXi version	5.5.2 build 2068190
vCenter Server version	5.5.0 build 1476327
Cluster Settings	
HA	Enabled
DRS	Enabled
EVC	Disabled
Swapfile location	Store the swapfile in the same directory as the virtual machine (recommended)
vSphere HA	
Host monitoring	Enabled
Admission control	Disabled
Virtual machine options -> VM restart priority	Medium
Virtual machine options -> Host isolation response	Leave powered on
VM monitoring status	VM monitoring only
Datastore heartbeating	Select any of the cluster datastores taking into account my preferences
vSphere DRS	
Automation level	Partially automated
Power management	Off

VMware Horizon 6.0 with View Configuration

VMware Horizon 6.0 with View	
Connection server	1 VM 4 vCPU 16GB memory
Windows 7 Desktops	456 target VMs 2 vCPU 1.5GB memory
Desktop Pool Settings	
Type	Manual
User assignment	Floating
ID	LoginVSI
Display name	LoginVSI
Machine source	vCenter
Access group	/
State	Enabled
View storage accelerator	Disabled

Maximize the Promise of Hyper-Convergence

Desktop Pool Settings continued	
Connection server restrictions	None
Remote machine power policy	Take no power action
Automatic logoff after disconnect	Never
Allow users to reset their machine	No
Allow multiple sessions per user	No
Default display protocol	PCoIP
Allow users to choose protocol	Yes
Max number of monitors	2
Max resolution of any one monitor	1920x1200
HTML access	Disabled
3D renderer	Disabled
Adobe flash quality	Do not control
Adobe flash throttling	Disabled
Override global Mirage settings	No

Login VSI 4.1 Configuration

Login VSI	
File share	Windows Server 2008 R2 64-bit 2 vCPU 4GB memory
Launcher	20 Launcher VMs Windows Server 2008 R2 64-bit 2 vCPU 4GB memory
Targets	456 Windows 7 32-bit
Users	456 Active Directory domain users
Workload	Knowledge worker
Virtual Desktop Settings	
CPU	2
Memory	1.5GB
NICs	1
Hard Disk 1	32GB (thin provision)
Applications	Adobe Flash Player 11 ActiveX Adobe Reader XI Doro 1.82 Microsoft Office Professional Plus 2010 Internet Explorer 8
VM hardware version	8
VMware Tools version	9.4.10, build-2068191
VMware View Agent version	6.0.0, build 1884979

Networking Configuration

Maxta Storage Network	
Switch	Mellanox SwitchX-2 SX1012
Adapter card	Mellanox ConnectX-3 Pro
Ethernet speed	10GbE
Management Network	
Switch	HP 1810G-24
Adapter card	Intel X540-AT2
Ethernet speed	1GbE

Appendix B: Test Results

Test 1: 100 Desktops – 2 Hyper-Converged Compute and Storage Nodes + 1 Compute-Only node

The knowledge worker workload test ran 100 active sessions successfully and did not reach VSImax. Response time at the saturation point (threshold) was 1731ms, best performance of the system (baseline) was 731ms, and the VSI index average was 823ms.




Figure 6: VSImax v4.1 results graph for 100 Desktops

(X axis = Number of Active Sessions; Y axis = Response time in milliseconds)

The CPU usage peaked at 27.01% of 82.584 GHz (36 CPU x 2.294 GHz) across the two servers, with an average usage of 18.42%.




Figure 7: CPU usage for 100 Desktops

Test 2: 300 Desktops – 3 Hyper-Converged Compute and Storage Nodes

The knowledge worker workload test ran 300 active sessions successfully and did not reach VSImax. Response time at the saturation point (threshold) was 1798ms, best performance of the system (baseline) was 797ms, and the VSI index average was 929ms.




Figure 8: VSIMax v4.1 results graph for 300 Desktops

(X axis = Number of Active Sessions; Y axis = Response time in milliseconds)

The CPU usage peaked at 28.54% of 82.584 GHz (36 CPU x 2.294 GHz) across the three servers, with an average usage of 19.37%.




Figure 9: CPU usage for 300 Desktops

Maximize the Promise of Hyper-Convergence

Test 3: 456 Desktops – 4 Hyper-Converged Compute and Storage Nodes

The knowledge worker workload test ran 456 active sessions successfully and did not reach VSImax. Response time at the saturation point (threshold) was 1844ms, best performance of the system (baseline) was 843ms, and the VSI index average was 995ms.




Figure 10: VSImax v4.1 results graph for 456 Desktops

(X axis = Number of Active Sessions; Y axis = Response time in milliseconds)

The CPU usage peaked at 58.81% of 82.584 GHz (36 CPU x 2.294 GHz) across the four servers, with an average usage of 40.27%.




Figure 11: CPU usage for 456 Desktops

Network usage statistics across all four servers were collected from the vSphere client performance chart for this test. Total bandwidth consumed across all four hosts peaked at 424.68 MBps, which accounted for a peak network usage of 33.97% on the 10 GBps storage network infrastructure.

VDI Network Usage					
Peak network usage per host					
ESX-1	ESX-2	ESX-3	ESX-4	Total network bandwidth	Total peak network usage
122807 KBps	105544 KBps	99264 KBps	97064 KBps	1250 MBps	424.68 MBps (33.97%)

Figure 12: Network usage for 456 Desktops

Test 4: High Availability – Recover 100 desktops after single node failure

One server in the four node cluster was brought down to simulate a disaster recovery scenario. The server was hosting 100 virtual desktops running the Login VSI simulated workload, all of which were recovered to other nodes. It took eight minutes total from the time when the node was brought down to the time that all 100 desktops became available on the other three servers in the cluster.

Disclaimers

© 2015 Maxta. All rights reserved. Maxta is a trademarks of Maxta Inc. All other trademark are the property of their respective owners. MDS-031301

Testing data is provided by the software provider and is not verified by Intel, nor does Intel guarantee that the information is error-free. All information is provided for informational purposes only. For more info, see www.maxta.com and/or contact the software provider for details.

Δ Intel processor numbers are not a measure of performance. Processor numbers differentiate features within each processor family, not across different processor families. See www.intel.com/products/processor_number for details.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

*Other names and brands may be claimed as the property of others.

By using this document, in addition to any agreements you have with Intel, you accept the terms set forth below.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to: <http://www.intel.com/design/literature.htm>.

Copyright © 2015 Intel Corporation. All rights reserved. Intel, the Intel logo, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

