

# VMware View 5 and Tintri VMstore T540

Reference Architecture Testing Results  
(Version 1.1)

Testing performed at VMware labs in fall of 2012 by  
Tristan Todd, EUC Reference Architect, [ttodd@vmware.com](mailto:ttodd@vmware.com)  
Rex Walters, VP Technology, Strategic Alliances, [rex@tintri.com](mailto:rex@tintri.com)

### Introduction

Both VMware and Tintri are actively engaged at a number of large enterprise “virtual desktop infrastructure” (VDI) projects at a number of shared customer accounts. Both companies wanted to conduct formal performance tests, in a controlled lab environment, of VMware View with Tintri storage.

We wanted to determine the maximum number of desktops that can be simultaneously supported on a single T540 without running into performance issues. A secondary goal was to validate a minimal server hardware configuration (hosts and networking as well as the storage itself). Lastly, we wished to characterize the IO workloads presented by large numbers of desktop users, and validate the benefits Tintri technology provides to VMware View environments and VDI in general.

### VDI Fundamentals

Tintri believes quite strongly that a successful VDI project ultimately depends on three key characteristics:

- **Simplicity:** Administrators responsible for initially configuring and managing an environment are the most obvious beneficiaries of management simplicity.

At the risk of seeming hard-hearted, however, financially motivated executives care surprisingly little how hard someone has to work to perform their job. What they *do* care about, however, is the success or failure of new projects, how quickly they can be rolled out, and how costly they will be to manage.

Keeping things simple is the surest way to ensure quick and continued success.

- **Low cost:** While it’s generally a mistake to think that virtual desktops can ever be less costly than physical desktops, VDI projects are unlikely to reach even the pilot stage unless costs are kept within reason. Storage costs alone can often undermine a new VDI project (because of the surprising IO demands).

Just as an enterprise storage array is more costly than the collection of individual hard disk drives it contains, the enterprise-grade hardware infrastructure required to host virtual desktops will always be more costly than an equivalent number of physical desktops. The

former provides additional security, availability, and data management benefits that simply don’t come for free. Nonetheless, the financial justification for these benefits isn’t possible if the basic hardware infrastructure costs aren’t kept within reason.

- **Performance:** Put simply, users won’t be happy unless their virtual desktops provide at least the performance and overall user experience that they received from a physical desktop. The bar is continually being raised: these days users **expect** their desktops to perform like an SSD-enabled *ultrabook*, not the hard-disk based PCs of a few years ago.

Tintri believes that hard-disk-drive (HDD) based storage arrays are *incapable* of balancing these requirements for a VDI project—the random IO performance demands from virtual desktops completely overwhelm the cost and complexity of the design. With an HDD-based array it becomes: “Simplicity, low cost, or performance: pick any one.”

Used correctly, however, flash storage easily satisfies the performance demands of virtual desktops. The trick is to keep it simple and reasonably priced. Simply throwing a flash cache in front of an HDD-based array does not suffice. Invariably, a simple cache actually increases complexity and costs. Even sizing the cache correctly is a difficult and complex topic.

As the rest of this document will show, the Tintri VMstore T540 storage system provides an extremely simple and low cost storage system with sufficient performance for any VDI project with need of more than 200 simultaneously active desktops.

### Testing Methodology

Both VMware and Tintri agree that the only meaningful way to validate that a storage system is capable of meeting the demands of a VDI project is to actually simulate real user behavior.

A common but flawed approach to sizing is to create a brutally simple model of an “average desktop user”. A typical model would state that such a user generates, say, 20 IOPS on average during the course of their work. The belief then, is that any storage system that can support 20,000 IOPS would suffice to support 1,000 virtual desktop users.

Some vendors attempt to use (very) slightly more sophisti-

cated models with factors for read/write ratios, block sizes, and perhaps two or three different user types rather than a single “average desktop.” They fail to address the most important point, however, which is that *IO demands change over the course of a working day*.

We believe a more sophisticated approach than a simple model based on IOPS is to deploy some large number of virtual desktops, then simulate user activity within each one of the guest VMs. This ensures that the overall environment (hosts, network, as well as storage) provides adequate performance to meet the IO demands.

The basic methodology we used during each of the reference architecture tests was to provision a large pool of desktop VMs (typically 1,000 desktops which is the maximum number of VMs formally supported on a VMstore T540) then execute some number of iterations of the user activity simulator (View Planner) and evaluate the results

## View Planner

We chose to use VMware View Planner v2.1 to simulate user activity during our reference architecture testing. View Planner is a tool maintained by VMware expressly for VMware and VMware partners to test VDI performance.

VMware’s View Planner tool simulates application workloads for various user types (task workers, knowledge workers, and power users) by actually running applications typically used in a Windows desktop environment. During the execution of a workload, applications are randomly called to perform common desktop user operations, including open, save, close, minimize and maximize windows; view an HTML page, insert text, insert words and numbers, conduct a slideshow, view a video, send and receive email, and compress files.

View Planner then uses a patent pending watermark technique to quantify the user experience and measure application latency on a user client/remote machine.

## View Planner Scoring

The standardized View Planner workload consists of nine applications performing a combined total of 44 user operations. These user operations are separated into three groups, shown in the table below. The operations in Group A are

Group A	Group B	Group C
AdobeReader-Browse	AdobeReader-Open	7zip-Compress
AdobeReader-Close	Excel_Sort-Open	PPTx-SaveAs
AdobeReader-Maximize	Excel_Sort-Save	Video-Play
AdobeReader-Minimize	Firefox-Open	
Excel_Sort-Close	IE_ApacheDoc-Open	
Excel_Sort-Compute	IE_WebAlbum-Open	
Excel_Sort-Entry	Outlook-Attachment-Save	
Excel_Sort-Maximize	Outlook-Open	
Excel_Sort-Minimize	PowerPoint-Open	
Firefox-Close	Video-Open	
IE_ApacheDoc-Browse	Word-Open	
IE_ApacheDoc-Close	Word-Save	
IE_WebAlbum-Browse		
IE_WebAlbum-Close		
Outlook-Close		
Outlook-Maximize		
Outlook-Read		
Outlook-Restore		
PPTx-AppendSlides		
PPTx-Close		
PPTx-Maximize		
PPTx-Minimize		
PPTx-ModifySlides		
PPTx-RunSlideShow		
Video-Close		
Word-Close		
Word-Maximize		
Word-Minimize		
Word-Modify		

Figure 1: View Planner Operations

used to determine the View Planner score, while the operations in Groups B and C are used to generate additional load.

The View Planner score is the 95th percentile value for application response time for Group A operations. This value represents a quality of service score based on operational latency experienced by users. A passing score is when the View Planner score is less than 1.5 seconds, which indicates an acceptable level of user experience at scale.

Tintri feels that while the 1.5 second threshold represents an adequate user experience, it is more indicative of the experience from a hard-disk based desktop. Ultrabook quality user experience would be indicated by View Planner scores well under the 1.5 second threshold.



Figure 2: Test hardware.

## Test Environment

The testing environment comprised a fairly modest amount of hardware: just four servers for desktop VMs, a small server for server VMs, a 10 gigabit network switch, a 1 gigabit network switch, and a Tintri VMstore T540. All hardware fit in approximately one half of a standard sized equipment rack.

Note that the VMstore contained all server VM images as well as the desktops themselves. We also performed our first tests (successfully) with just four hosts (without the dedicated host for the server VMs). View Planner run rules, however, require that View Planner and other server VMs run on separate server hardware.

All hosts were loaded with VMware vSphere ESXi 5.0 build 623860.

Single ESXi host for Management and Infrastructure VMs	
Dell PowerEdge R610	
CPU – 2 x Intel Xeon E5-2690 (2.9GHz) – 16 cores	
RAM – 128GB	
Storage – 500GB SAS Array	
Networking - 2 x 1GbE, 2 x 10GbE	
Four ESXi hosts for desktops	
Dell PowerEdge R720	
CPU - 2 x Intel Xeon E5-2690 (2.9GHz) - 16 cores	
RAM - 256GB	
Storage - 700GB SAS Array	
Networking - 2 x 1GbE, 2 x 10GbE	
Gigabit Switch	
Cisco Catalyst 4948 48-port Gigabit Switch	
10 Gigabit Switch	
Dell Force10 S2410 10 GbE Switch	
Shared Storage	
Tintri VMstore T540	

Table 1: Hardware configuration

## Network Configuration

Each host was configured with a traditional vSwitch that was used for the host management interface. This vSwitch was backed by dual 1GbE uplinks in a teamed configuration.

A dVSwitch was also deployed that was uplinked with redundant 10GbE adapters on each host. Trunk ports and VLAN tagging were utilized to properly segregate traffic (NFS, VM, and VMotion) traffic for security and performance reasons.

### Virtual Switch: vSwitch0

Console / Host Mgmt		Physical Adapters				
Port Group: Mgmt Interface		pNic	Type	Speed	Duplex	State
VLAN: 1		Vmnic0	Trad	1000	Full	Active
		Vmnic1	Trad	1000	Full	Active

### vSphere Distributed Switch

VMotion		Physical Adapters				
Port Group: dv_VMotion		pNic	Type	Speed	Duplex	State
VLAN: 977		Vmnic4	Trunk	10000	Full	Active
		Vmnic5	Trunk	10000	Full	Active
NFS Storage						
Port Group: dv_Storage						
VLAN: 3250						
VMs						
Port Group: dv_VirtualMachines						
VLAN: 3253						

Figure 3: Virtual network configuration.

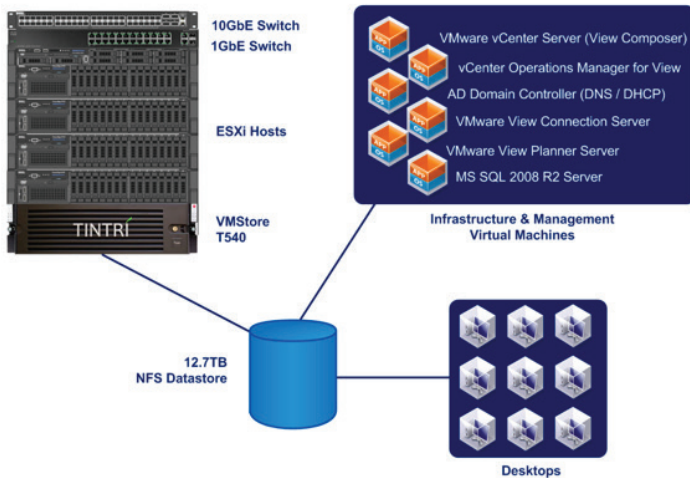


Figure 4: VMstore Contents

## VMstore Contents

The VMstore contained all VM images, both servers and desktops (refer to Figure 4).

Note that this document and the vSphere client displays the capacity of the T540 as "12 TB." This indicates  $12 \times 2^{10}$  bytes (or roughly 13.5 trillion bytes).

The virtual server VM configurations were as as indicated in table 2.

## VMware Software Configuration

We used a View Connection/Manager Server (v5.1, Build 704644).

A single vCenter server was deployed and the View Composer service was installed locally. The basic configuration of View Composer was as shown in table 3.

Description	View vCenter Server for Desktops
Connect using SSL	Yes
Port	443
Enable View Composer	Yes
Advanced Settings	40
•Maximum Concurrent vCenter Provisioning Operations	
•Maximum Concurrent Power Operations	40
•Maximum Concurrent View Composer Maintenance Operations	30
•Maximum Concurrent View Composer Provisioning Operations	30

Table 3: vCenter / View Composer configuration

Domain Controller VM	
OS – MS Windows Server 2008 64-bit	
vCPU – 2	
vRAM – 6GB	
Storage – 50GB via LSI SAS	
Networking – 1 x Vmxnet3	
vCenter Server VM (also running View Composer and MSSQL)	
OS – MS Windows Server 2008 64-bit	
vCPU – 4	
vRAM – 8GB	
Storage – 100GB via LSI SAS	
Networking – 1 x Vmxnet3	
View Connection Server VM	
OS – MS Windows Server 2008 64-bit	
vCPU – 4	
vRAM – 8GB	
Storage – 40GB via LSI SAS	
Networking – 1 x Vmxnet3	
View Planner Server VM	
OS – CentOS 6 32-bit	
vCPU – 1	
vRAM – 3GB	
Storage – 15GB via LSI SAS	
Networking – 1 x Vmxnet3	
vCenter Operations Manager – Analytics VM	
OS – SUSE Linux Ent 11 – 64-bit	
vCPU – 2	
vRAM – 9GB	
Storage – 200GB via LSI SAS	
Networking – 1 x Vmxnet3	
vCenter Operations Manager – UI VM	
OS – SUSE Linux Ent 11 – 64-bit	
vCPU – 2	
vRAM – 7GB	
Storage – 130GB via LSI SAS	
Networking – 1 x Vmxnet3	

Table 2: Server VM configuration.

Global Policies

Policy	Setting
USB access	Allow
Multimedia redirection (MMR)	Allow
Local Mode	Deny
PCoIP hardware acceleration	Allow – medium priority

**Note** The global policy can be overridden by a pool policy set at a pool level. Alternatively, it can be set at a user level.

Global Settings

Attribute	Specification
Session timeout	600 (10 hours)
View Administrator Session Timeout	120 mins
Automatic Status Messages	True
Display pre-join message	False
Display warning before logoff	False
Reauthenticate secure tunnel connections after network interruption	False
Use IPsec for Security Server Connections	True

Pool Configuration

Attribute	Configuration
Pool Type	Automated
User Assignment	Floating
State	Enabled
Connection Server Restrictions	None
Remote Power Policy	None
Automatically Logoff after Disconnect	No
Allow users to reset desktops	No
Display Protocol	PCoIP
Max number of monitors	2
Max Display Resolution	1920 x 1200
Adobe Flash Quality	n/a
Adobe Flash Throttling	n/a
Number of Desktops	1000
Number of spare, powered-on desktops	1000
Provision all desktops up-front	Yes
Replica Datastore	Tintri Data 1
Linked Clone Datastores	Tintri Data 1

Tables 4-6: View Manager configuration

## Desktop image configuration

The operating system image deployed for the desktop images uses Microsoft Windows 7 SP1 with all Important MS security updates installed. These images are considered to be the Gold images. Upon configuration and customization (according to View Planner specifications), these images were converted to virtual machine templates.

The virtual hardware configuration of the Master desktop virtual machine varies based on the computing resources required by the end user. The following tables define the default virtual hardware specifications. Exceptions to this default specification are handled on a pool-specific basis.

Attribute	Specification
Desktop OS	Windows 7 Enterprise, SP1
Hardware	VMware Virtual Hardware version 8
vCPU	1
vMemory	2048MB
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 View Agent	VMware View Agent 5.1 build 704644
Applications Installed	MS Office 2010 (32-bit), Adobe Reader X, Firefox 6, Internet Explorer 9, 7Zip

Table 7: Windows 7 master image configuration

## Other VM Server Configuration

Microsoft SQL Server 2008 was installed on the same Windows 2008 virtual machine as vCenter and View Composer.

vCenter Operations Manager for View 1.1 was deployed in a standard configuration with appliances configured with sizing assumption of 100 hosts and 2000 virtual machines. VCOps was used for detail metrics gathering and to produce the performance graphs included in this document.

View Planner testing was conducted using “local” mode, where no remote simulated launchers were used to simulate PCoIP sessions. This is a standard configuration for conducting storage integration testing.

## Specific Best Practices

We carefully followed all relevant best practice published by VMware and Tintri.

All network traffic was carefully separated onto **separate vmkernel networks**:

- Management/service console (ESX only)
- VMkernel for IP storage
- VMkernel for vMotion
- VMkernel for FT
- Virtual machine network traffic

The Tintri VMstore was configured as a **single datastore** with a single logical data IP address on the 10 gigabit IP storage network, and was used by all esxi hosts. A separate IP address was used to configure and manage the T540.

The VMstore T540 was configured with **redundant physical 10 gigabit ports** to the data network to provide a failover path. Jumbo frames were **not** configured.

All network configuration was in accordance with the *Tintri NFS and vSphere Best Practices* guide. This includes **disabling Storage I/O Control (SIOC)** on every vSphere host.

All hosts used PCIe x8 or greater slots for their 10 gigabit network adapters.

**Note:** Though the published VMware tested maximum number of virtual machine objects per NFS volume in vSphere 5 is 250, it is possible to support 1000 virtual machines in a single VMstore T540.

We adjusted the concurrent power and provisioning settings in View Manager (max concurrent power ops, max concurrent view composer ops, and max concurrent view provisioning ops) for optimal performance during each test run.

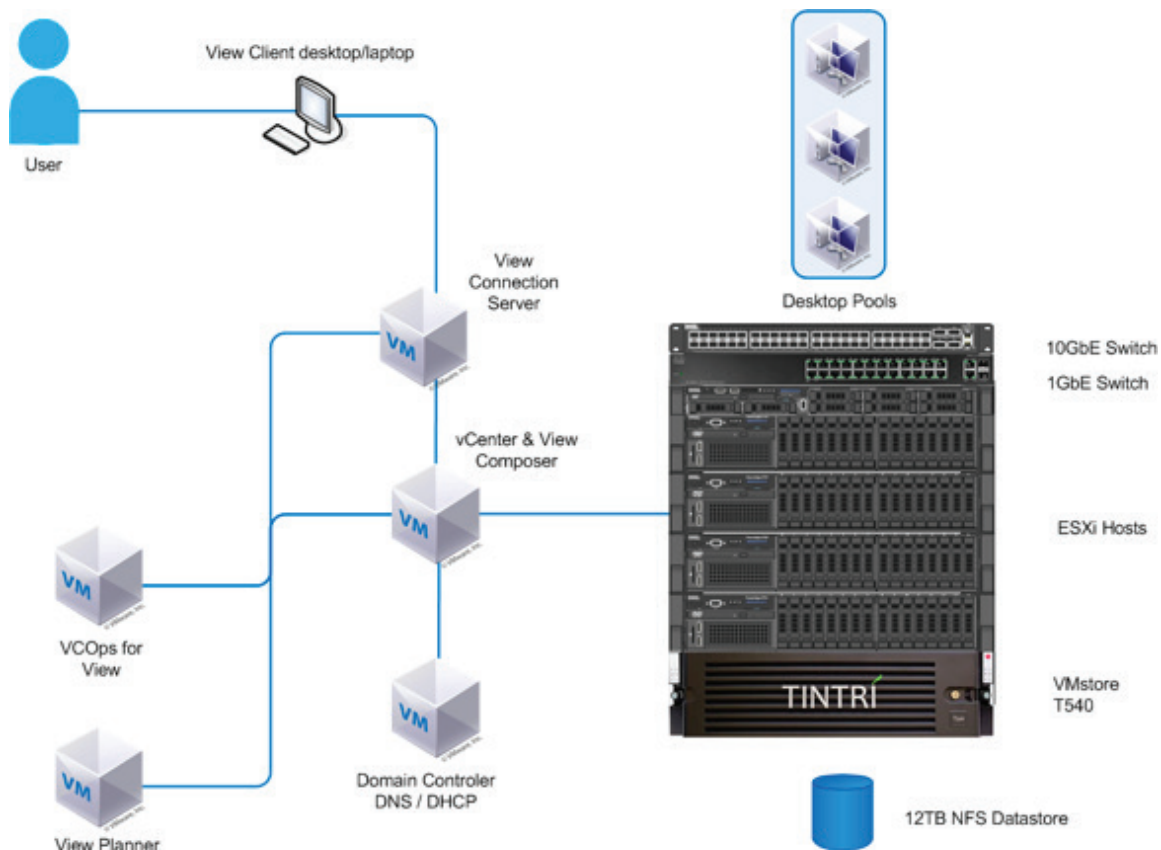


Figure 5: Logical overview.

## Testing Result Summary

During the course of the testing we validated the following:

### General

- 1000 desktops (full clone or linked clone) easily fit and perform well on a single Tintri VMstore T540.
- The 5 host, one T540 architecture is a scalable unit for any number of virtual desktops (a half rack of equipment for each 1,000 desktops).
- All management, monitoring, infrastructure servers can be housed on a single Tintri VMstore T540
- Achieved View Planner benchmark certification for delivering low application latency (and thus good user experience).

### Simplicity

- Simple architecture, excellent documentation.
- 8 minutes until first VM deployed.

- No tuning or special optimization required.
- Innovative per-VM management and monitoring interface.

### Performance

- Fast desktop pool deployments (1000 desktops in a little over 2 hours)
- Excellent View Planner benchmark results
- Predictable performance and low storage latency even during workload bursts (including deploy or re-compose operations).

### Cost-Effective

- No complicated sizing exercises required.
- Simple, intuitive administration (no need for training or specialized skills).
- Modest acquisition costs.

### Detailed Results

This section documents the details of each test performed.

#### Overall Simplicity

The Tintri VMstore is, as both authors agreed, one of the easiest shared storage solutions to deploy and integrate into a vSphere environment. Because the Tintri VMstore is presented as a single NFS mount of 13.5 TB, the administrator only has to manage a single large datastore. The tested architecture relied on 10 GbE networking which lends itself well to simple architecture. Cables were minimized and bandwidth was efficiently allocated using distributed virtual switches and VLAN tagging.

No complicated workload or sizing estimation was required prior to deploying the Tintri array. The VMstore is designed to handle data tiering and SSD placement automatically on the back end. Therefore, little effort is required to tune or optimize the storage environment through storage placement or tiering.

Both authors believe that Tintri deployment documentation is better than that provided by most storage vendors. The VMware lab teams were able to easily rack and configure the VMstore equipment; using excellent deployment prerequisite lists and deployment checklists provided by Tintri.

Bringing the Tintri storage online was fast and painless. With the Tintri array powered on, from logging into vCenter until the vCenter VM was migrated to the NFS store took approximately 8 minutes. Minimal vSphere setting changes were required to adhere to standard VMware NFS and Tintri best practices.

#### Pool Deployment Operations

During testing, desktop pool deployments were fast and consistent. The deployment of linked clones and full clones was tested at scale to 1000 seats with very few stability issues. Desktop pool deployments were fast, in fact some of the fastest seen in the technical marketing labs. Some key performance milestones were:

- 500 linked clones deployed per hour
- 1,000 linked clones deployed in 2 hours and 30 minutes.
- 1,000 full clones deployed in 2 hours and 22 minutes

(with the Tintri VAAI plugin installed on the esxi hosts).

- 1,000 desktops booted (cold start) in 16 minutes.
- Individual desktop recompose in 1 minute, 15 seconds.
- 1,000 linked clone pool recompose in 2 hours and 45 minutes.

### View Planner Results

All View Planner testing was carried out in strict accordance with the View Planner benchmark rules as defined in the View Planner Workload Usage Rules (revised 21 October 2011). All View Planner test results were reviewed and approved by the VMware Performance Benchmark Team in early October 2012.

View Planner Test (1000 seat)	View Planner Score ( <b>&lt; 1.5s to pass</b> )
Linked Clone 5 Iteration Test	0.97s
Linked Clone 15 Iteration Test	0.78s
Full Clone 5 Iteration Test	0.57s
Full Clone 15 Iteration Test	0.52s

Table 8: View Planner results

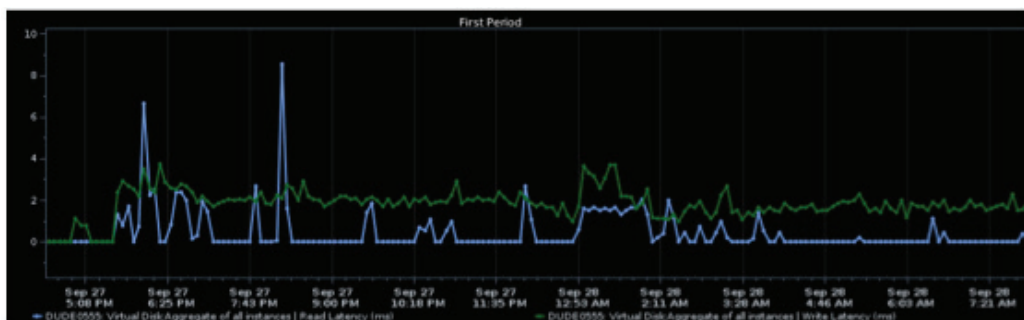


Figure 6: Desktop vDisk latency using linked clones.

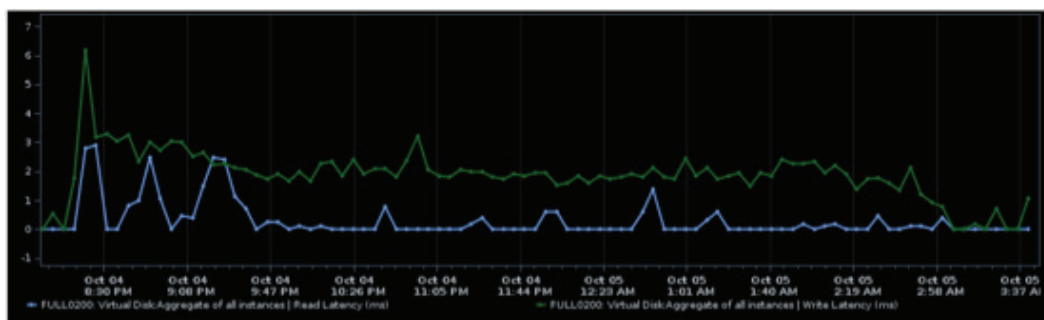


Figure 7: Desktop vDisk Latency using full clones (with VAAI).

### Desktop In-Guest VM Latency

Desktop Virtual Disk latency (vDisk Latency) is considered to be one of the key metrics associated with application response time and thus VDI user experience. This metric directly shows user perceived, end-to-end latency inclusive of host, network, and storage delays. During benchmark testing, this is a key metric to watch as density or workload intensity are increased.

vDisk performance during all Tintri testing was quite good and exhibited latency characteristics similar to those seen on high performance ultrabooks. During testing, both read and write latency levels were easily better than physical PCs and laptops with 7200-rpm or 8400-rpm hard drives. Read and write latency levels of 5ms or lower are considered quite good in the world of hosted desktops.

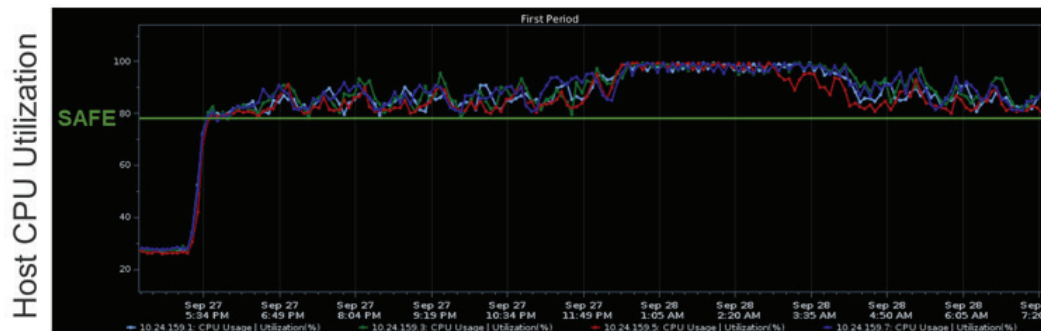


Figure 8: Host CPU utilization with 4 esxi hosts.

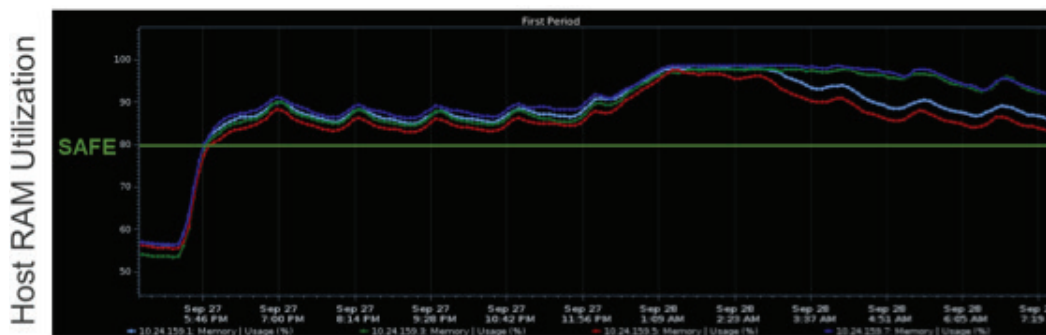


Figure 9: Host RAM utilization with 4 esxi hosts.

### vSphere Host Utilization

With 250 desktops on each host, CPU and RAM resources were pushed well beyond the limits considered by many to be best-practice maximums.

It is strongly advised not to run at this level of utilization in production. It is always a good idea to operate at no more than 80% host CPU and RAM utilization.

The tested architecture, while it can run 1000 task-worker desktops in a stable manner, is ideally suited for no more than 80% concurrency, where only 800 of the desktops are active with a logged-on user at one time. To scale beyond 800 active sessions would require additional host and network resources, beyond what was deployed for the initial testing.

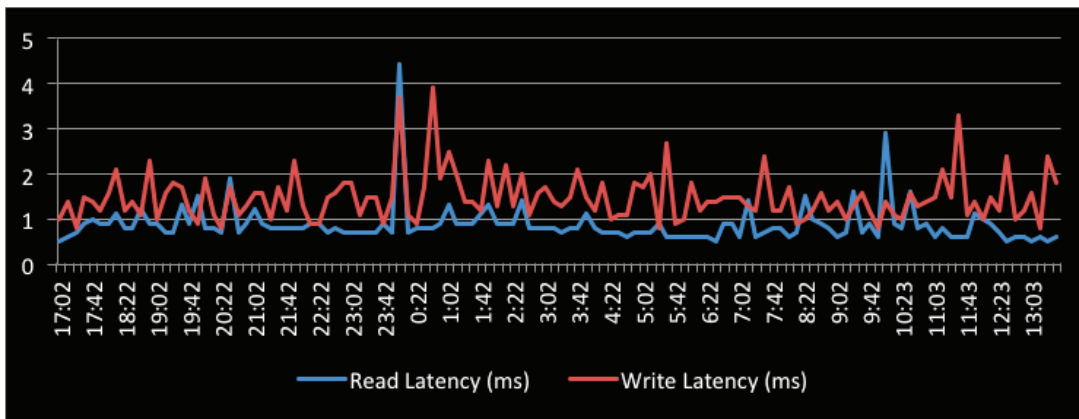


Figure 10: Datastore latency as measured by VCOPS

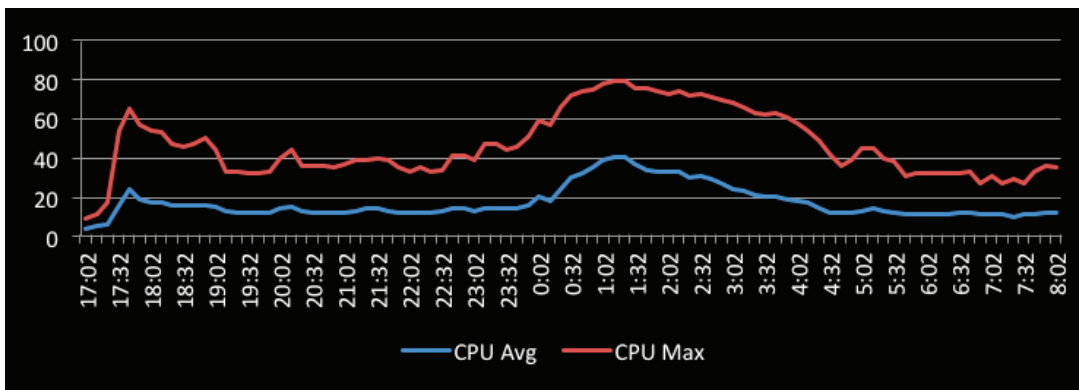


Figure 11: Datastore CPU as measured by VCOPS

### Storage Performance (VCOPS View)

Storage performance measured at cluster level (front end) in vCenter Operations Manager was quite good. Overall the front-end latency was low and the VMstore was able to maintain low latency even during periodic IO bursts.

- Average read latency of 1.7 ms
- Peak read latency of 7.2 ms
- Average write latency of 6.29 ms
- Peak write latency of 30.5 ms (host bottleneck during data collection)
- Peak throughput (during deployment) of 400 MB/s
- Peak IOPS (during deployment) of 26K IOPS

### Pool Operations

In typical VDI environments it is ill-advised to perform storage intensive operations such as pool deployments when shared infrastructure (host and especially storage) is in use with production workload. During testing, a 500-seat linked clone pool was deployed on the VMstore. During a View Planner run, another 500-seat pool was composed and deployed in the background on the same storage array (and same hosts). As follows, you can see IO-burst activity associated with the pool deployment (at approximately 0030 on 10/5). The second graph shows that datastore latency levels remained relatively consistent despite the background operations.

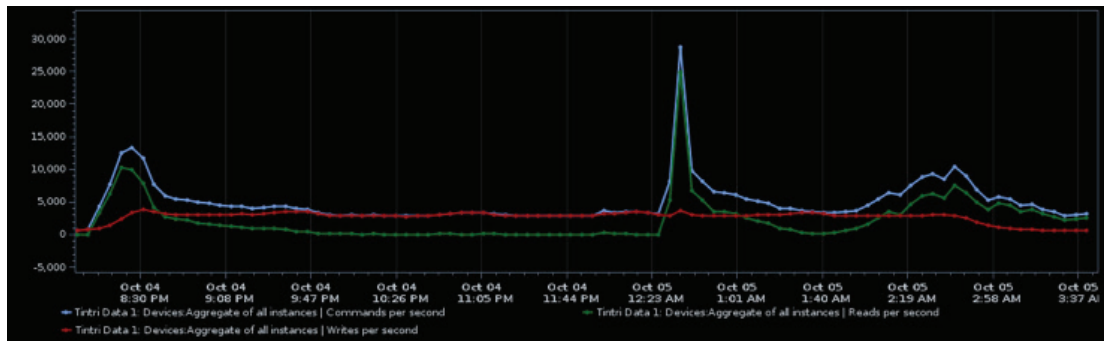


Figure 12: IOPS with deploy operation during View Planner run

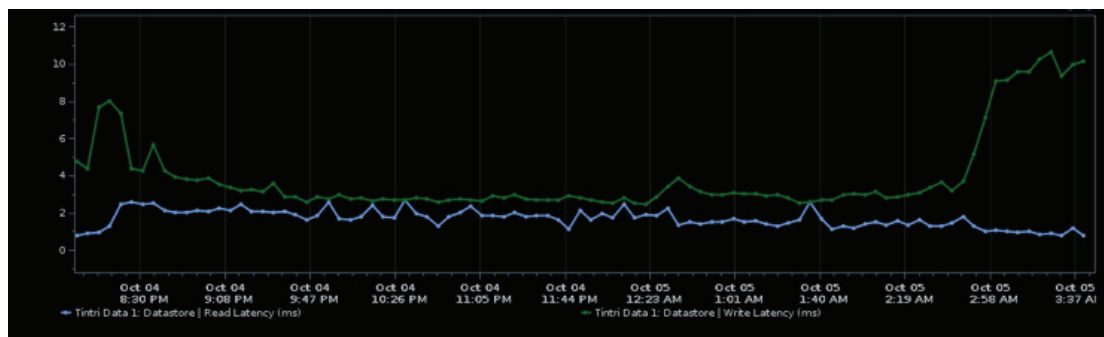


Figure 13: Latency with deploy operation during View Planner run.

NOTE: Write latency levels increases slightly at the tail-end of the test period. This increase is associated with View Planner log files being written out to the View Planner test harness.

### Boot and Login Storm

During boot storm the Tintri VMstore demonstrated the ability to serve up over 30K IOPs during a startup of 1000 desktops. Total duration to bring all desktops online was approximately 16 minutes (measured at View Manager).

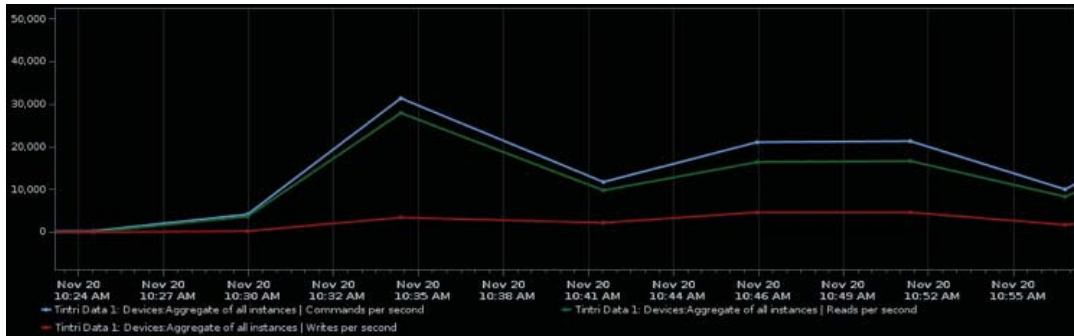


Figure 14: 1,000 user boot storm.

During a login storm the Tintri VMstore demonstrated the ability to burst to nearly 30K IOPs during a 2-hour, 1000 user login window

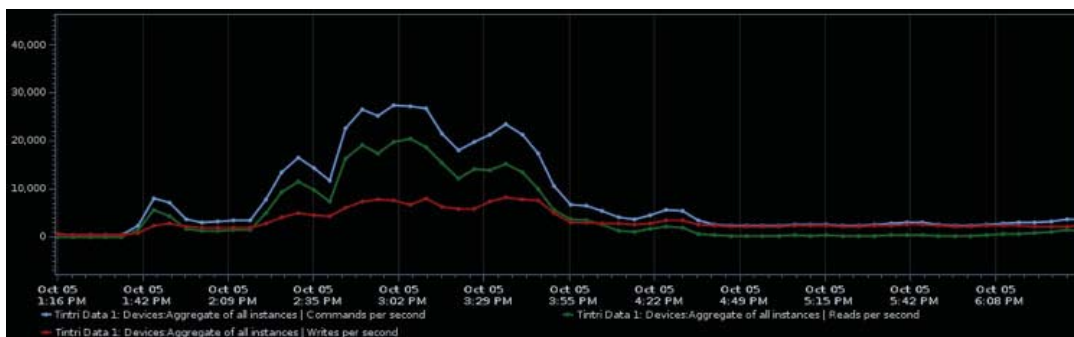


Figure 15: Monday morning login storm (1,000 users in 2 hours)

### In Summary

The results of this testing are significant for a number of reasons. Together, VMware and Tintri were able to define a compact and easy-to-deploy VDI infrastructure pod that goes beyond the conventional thinking for throughput, client support, performance and cost per desktop. Beyond that, through Tintri's unique VM-aware architecture and per-VM QoS, workloads that would otherwise be incompatible with one another, were proven to coexist successfully in the same datastore thereby further driving down the cost of the infrastructure. And finally, because of Tintri's ease of setup and administration, new benchmarks for initial deployment and desktop provisioning were established that were previously unseen in the VMware labs.