nexenta™

PERFORMANCE STUDY

NexentaConnect™ View Edition
Branch Office Solution

Nexenta Office of the CTO
Murat Karslioglu

Table of Contents

Desktop Virtualization for Small and Medium Sized Office

User experience continues to take on greater importance with today's knowledge workers having increasingly demanding expectations about the speed and reliability of their desktop sessions. Many knowledge workers have a choice and may elect to use desktop virtualization or optionally use a corporate laptop or a BYOD notebook option. For desktop virtualization to increase its share and acceptance among knowledge workers and power users, the end user experience must remain competitive which implies, for instance, better than 5 milliseconds (ms) IO latency.

Figure 1 illustrates widely accepted in the industry conventions as far as virtual desktop performance requirements for a variety of workloads, from light to heavy:

Figure 1: Types of desktop virtualization users and performance requirements

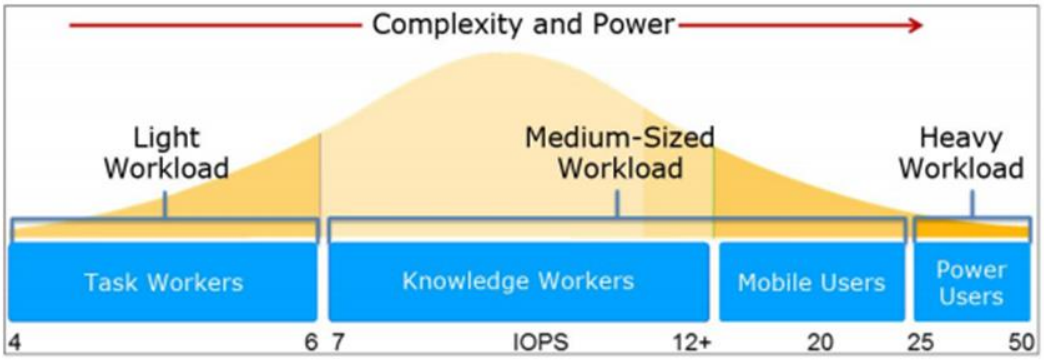

According to multiple third-party studies, expected number of IOPS (IO operations per second) for task workers (leftmost on the picture) is 4 to 6, for knowledge workers: 7 to 12, and for power users: 25 to 50. For mobile users the average IO performance is expected to be around 20 IOPS.

In addition to performance, resources and budget limitations make branch offices one of the most challenging environments for VDI. Some of the key problems that affect branch/small office desktop virtualization deployments are:

- Difficulty in measuring storage requirements prior to the deployment
- Long preparation and configuration times
- Long deployment times for full-clones
- Difficulty scaling up after POC
- Low storage performance
- High storage latency
- Monitoring performance after deployments
- Disaster Recovery strategy
- And finally, overall high VDI cost and TCO

NexentaConnect installed on Cisco UCS E server blades delivers a complete all-in-one desktop virtualization solution. That's because NexentaConnect offloads a significant portion of read and write IO from the HDD, removing the congestion for all peak (boot storm, login and write bursts) conditions while at the same time freeing more CPU cycles for utilization by running desktops.

Cisco UCS E Series and NexentaConnect

NexentaConnect combines a powerful GUI-based automation engine and enterprise grade Virtual Storage Appliance (VSA). Cisco UCS E-Series server blade is designed to host Cisco and third-party infrastructure services and mission-critical business applications.

NexentaConnect installed on Cisco UCS E server blades delivers a complete all-in-one desktop virtualization solution, including:

- Storage – Virtual Storage Appliance (VSA) creates storage pool and provides NFS share
- Networking – Dedicated local NFS network eliminates traffic flow over external network by keeping the storage and data traffic inside the vSphere vSwitch
- Compute – Desktop VMs and the business apps run on the same server blade

In this performance study we evaluate and measure branch office use-case and storage acceleration capabilities of the NexentaConnect product installed on the Cisco E series. Test results documented further in this paper show 47x (4700%) improvement with respect to IOPS, and 3x (300%) boost with respect to the desktop density. This in turn translates as significantly lower (per user) TCO and predictable measurable desktop virtualization SLAs.

All benchmark results in this white paper were independently reproduced by Cisco Data Center Solutions Engineering group, in the Cisco Interoperability Verification Testing (IVT) lab. The results can be separately obtained by Cisco Partners via Cisco CDN Partner Program.

Getting Started

Figure 2 depicts configuration used for all tests and benchmarks described in this paper.

One Cisco UCS E Series UCS-E160D Double Wide Server in Cisco UCS E Series 3945E Integrated Service Router chassis was installed in a VMware vSphere 5.1 cluster connected to the external network via the integrated router over 1GbE and Nexenta Virtual Storage Appliance (VSA) was running inside the same server blade using two local 10K RPM SAS drives (HDD#2 and HDD#3) to create the accelerated storage pool for the desktop virtualization workload.

Figure 2: Cisco UCS E Series and NexentaConnect test bed

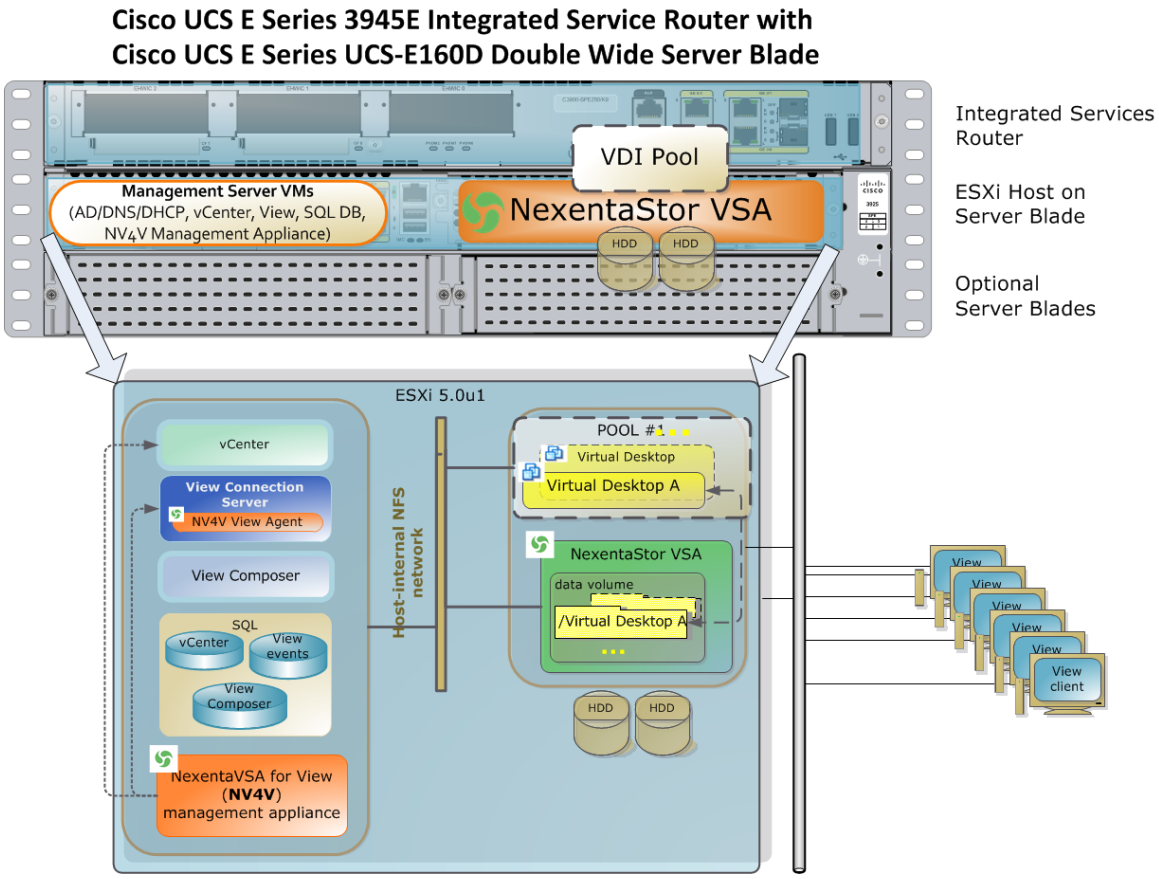

Notice (Figure 2) that the entire management infrastructure: the conventional DHCP/DNS/AD and SQL servers, VMware servers including vCenter and View, NexentaConnect management appliance, NAS server (the NexentaConnect VSA), and finally the desktops themselves – the entire suite of storage, networking and compute software to provide desktop virtualization service to end users – is hosted and runs on the same single server blade sharing the blade's memory, CPU and 3 local 10K RPM SAS HDDs.

Server Specs

- CPU – Intel Xeon processor E5-2428L (15-MB cache, 1.8 GHZ, and 6 cores)
- Memory – 48 GB
- Hard Disk – 3x 900 GB 10K RPM SAS HDD
- Firmware - Cisco IOS 15.2(4)M

Software

- NexentaConnect V 2.1.1
- VMware ESXi – v 5.0U1
- VMware Horizon View 5.2

Storage Specs

- Nexenta VSA (2 vCPU, 6GB Memory, local Virtual Storage Appliance) providing RAID0 volume over two 10K RPM SAS drives (900GB each / 1.8TB TB RAW space, no additional HDDs or SSDs used). The two SAS disks are used to store virtual desktop (linked clones and full clones) VMware ESXi – v 5.0U1
- 1 NexentaConnect volume exported as NFS to for the desktop virtualization deployment over a dedicated vSwitch local to the ESXi host (no external connection required)

LoginVSI and Iometer Testing by Nexenta and Cisco

We have used industry de-facto standard 3rd party tools to benchmark the performance and validate the improvements, namely:

- LoginVSI™ – Validates density and responsiveness of max number of desktops that can be run on the environment
- Iometer – Validates performance per desktop for the given desktop pool

VMware vSphere performance charts, LoginVSI graphics and NexentaConnect performance charts were used to capture the statistics.

Results of this study are independently published by Cisco (available for Cisco Partners) at:

- NexentaConnect on E-series: https://marketplace.cisco.com/catalog/products/3144
- Nexenta and Cisco Simplify IT Challenges White Papers: http://www.nexenta.com/corp/cisco
- NexentaConnect + UCS E-Series: https://marketplace.cisco.com/catalog/products/3509/files/55478

Certification resources:

- NexentaConnect + Cisco UCS – VMware Certification: https://marketplace.cisco.com/catalog/products/3509/files/55479
- VMware Horizon View Rapid Desktop Certified Appliances: http://www.vmware.com/resources/compatibility/vcl/poc.php

LoginVSI: with and without NexentaConnect

LoginVSI™ (www.loginvsi.com) is an industry de-facto standard desktop virtualization benchmarking tool validates the application response time on various predefined workload options.

Table 1 below shows the CPU utilization for the medium LoginVSI workload on a Series E server blade with and without NexentaConnect. In the “NexentaConnect scenario, the VSA is residing on the same local datastore as the hypervisor and utilizes the remaining two local drives to provide NFS datastore for the virtual desktops (Figure 2 above).

As a best practice and a general guideline, in desktop virtualization deployments the ESXi host’s total CPU utilization must be kept below 90%. Table 1 demonstrates that up to 55 desktops consume less than 90% CPU while at the same time providing significant improvement in IO performance (demonstrated later in this Section, see Table 6).

Table 1: CPU Utilization for medium LoginVSI workload

# OF ACTIVE DESKTOPS	WITHOUT NexentaConnect	WITH NexentaConnect
10 Desktops	66%	45%
15 Desktops	81%	56%
18 Desktops	88%	64%
20 Desktops	90% (unresponsive)	68%
25 Desktops	100% (unresponsive)	75%
30 Desktops	N/A	80%
40 Desktops	N/A	81%
45 Desktops	N/A	82%
50 Desktops	N/A	84%
55 Desktops	N/A	88%

Using NexentaConnect, we achieved lowered IO latency and increased IO performance, reduced CPU utilization and increased desktop density – the latter from 18 (without NexentaConnect) to 55 desktops.

Pictures below include screenshots of the VMware vSphere Performance chart taken during the LoginVSI test. Figure 3 shows the CPU utilization during the medium LoginVSI workload ("without NexentaConnect" column of the Table 1). After the test reached 20 desktops density, LoginVSI RDP sessions became unresponsive. At 25 desktops RDP sessions started to drop, which can be seen on the Figure 3 below.

Figure 3: 55 linked-clone desktops starting medium workload on local disk

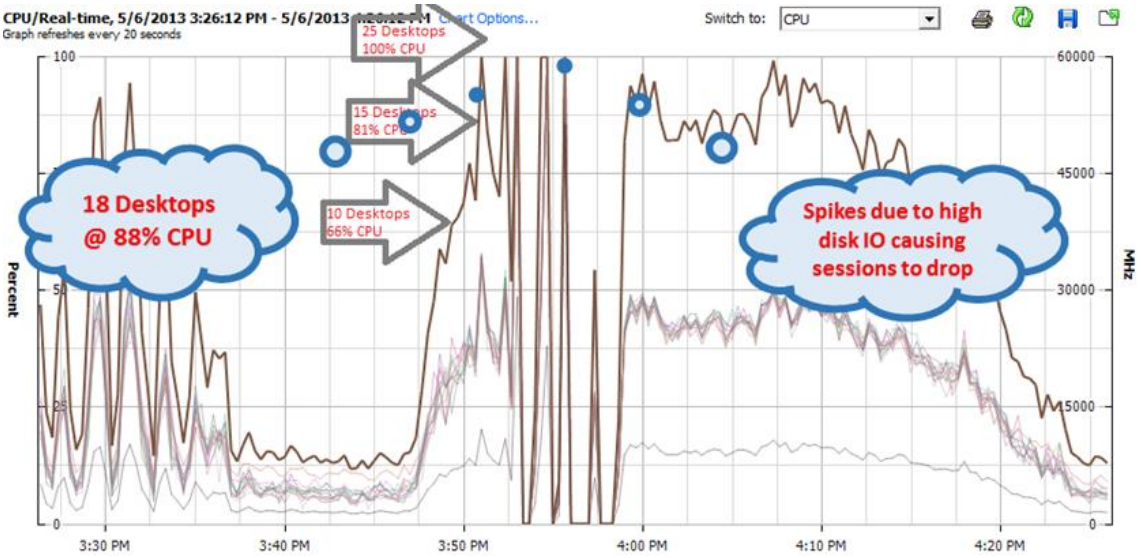

Figure 4 below shows the CPU utilization during the medium LoginVSI workload ("with NexentaConnect" column of the Table 1). At 25 active desktops CPU utilization was around 75%, when test reached to 55 active desktops CPU utilization reached to 88%, which was still below the acceptable 90% CPU utilization threshold.

Figure 4: 55 linked-clone desktops starting medium workload on NexentaConnect

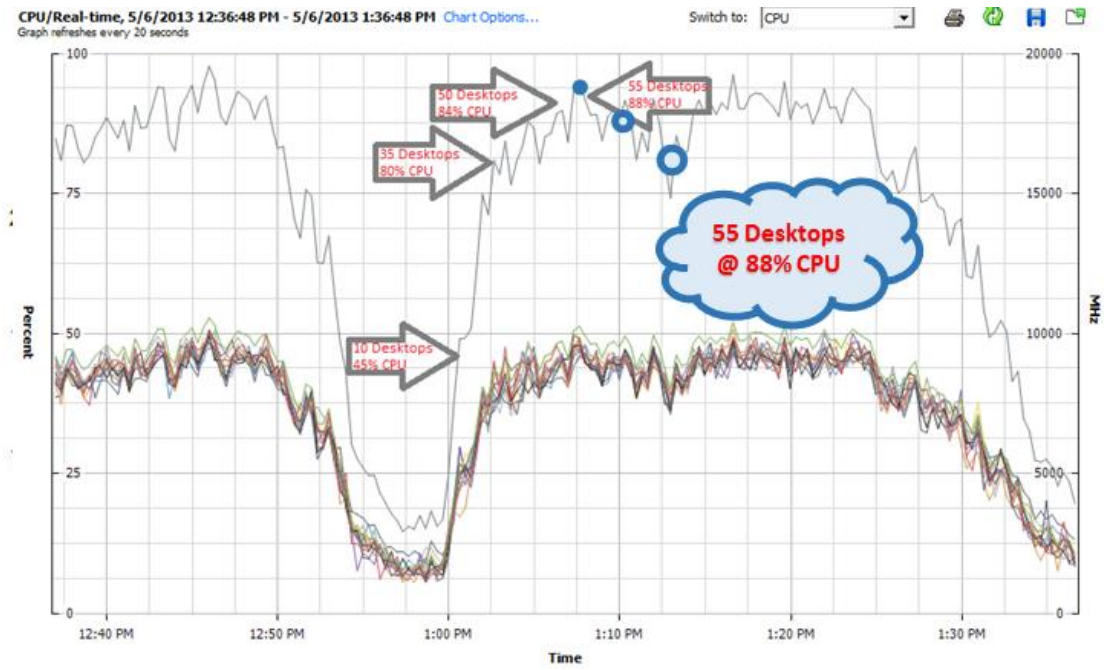

LoginVSI data (Table 2 below) documents desktop density with and without NexentaConnect. The first row in this table, baseline performance or VSIbase, is defined as performance of the system without desktop application running (on the virtual desktops). The number must be below 5000; the lower number indicates better response time (latency) of the desktop virtualization solution.

Table 2: Summary for the medium LoginVSI workload

PARAMETER	WITHOUT NEXENTACONNECT	WITH NEXENTACONNECT
LoginVSI Baseline	3419	2209
VSImax Score*	18 Desktops	55 Desktops, VSImax not reached*
Observations	After 25 desktops sessions dropped and desktops became unresponsive	Achieved 3x Higher Density *VSImax not reached. Test was able to reach up to 55 desktops due to 48GB maximum memory limitation on Cisco UCS E Series server platform. Used 20% more than physical memory via VMware memory overcommit.

* LoginVSI analyzer monitors the point in time when the target environment reaches its saturation point and the number of concurrent session and this point. The latter is called VSImax (www.loginvsi.com).

Cisco E Series UCS-E-160D server blades support up to 48GB physical RAM. In our tests we used VMware memory overcommit to provision up to 20% of more than physical available memory. Each virtual desktop had 1GB memory assigned to it, which means that 55GB was required to just run the 55 desktops. "VSImax not reached" on the Figure 6 shows that if we had more physically memory we could run more desktops without increasing the CPU and storage resources.

Desktop virtualization workloads generate bursts of small-block random reads and writes. Part of the NexentaConnect provided performance optimization can be attributed to:

Compression: NexentaConnect uses Lempel-Ziv family LZJB lossless data compression, which replaces repeated occurrences of data with references to a single copy and improve IO throughput to 1.9 times with linked clones (up to 4 times with full clones).

Read Caching: NexentaConnect VSA is configured to perform aggressive caching using the platform's Adaptive Replacement Cache (ARC) technology. At the same time, the virtual storage appliance strictly abides by the NFS semantics to store persistent data on stable storage without risking data integrity - often a trap for vendors trying to speed-up IO performance via data caching. For details on ARC see for instance Wikipedia at: en.wikipedia.org/wiki/Adaptive_replacement_cache

NFS: Note also that NFS RPC transactions are performed in the ESXi host memory, at the bandwidth that is not restricted by physical links.

Figure 5 shows aggregated response time (called "VSImax score") during the medium LoginVSI workload:

Figure 5: VSImax result for 55 linked-clone Desktops running medium workload on local disk

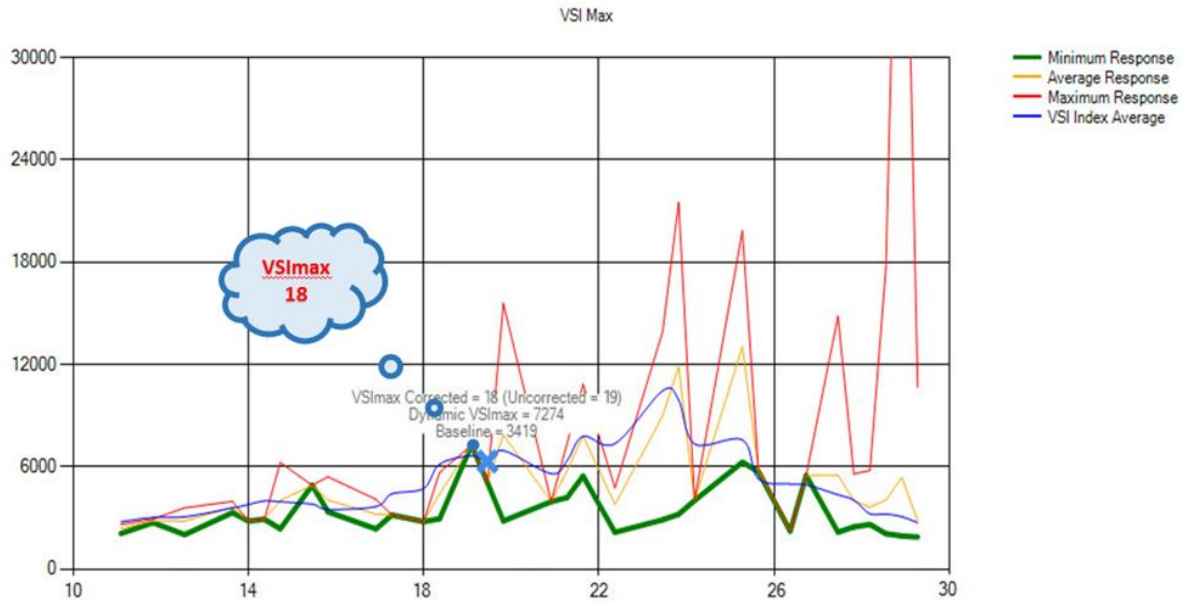

Figure 6 shows the VSImax score during the medium workload LoginVSI ("with NexentaConnect" column of the Table 2). Difference between Figure 5 and Figure 6 clearly shows the benefits of VSA. Without NexentaConnect V4V excessive read and write requests to the HDD result in increased latency and application responsive time. Application response time defines the user experience quality of the desktops.

Figure 6: VSImax result for 55 linked-clone Desktops running medium workload on NexentaConnect

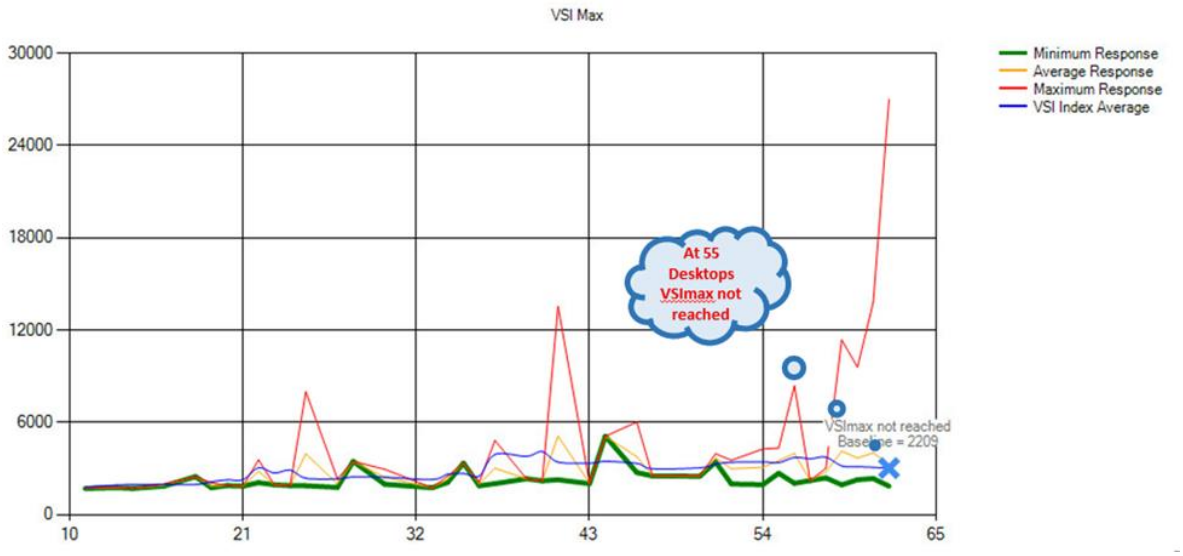

Table 3 shows the details of CPU utilization for the heavy LoginVSI workload on a Series E server blade, again with and without NexentaConnect.

Note that with NexentaConnect, up to 37 desktops utilize less than 90% of the CPU while at the same time providing significant improvement in the IO performance (this is further documented in the Table 6 below).

Table 3: CPU Utilization for heavy LoginVSI workload

# OF ACTIVE DESKTOPS	WITHOUT NEXENTACONNECT	WITH NEXENTACONNECT
10 Desktops	66%	47%
25 Desktops	93%	76%
30 Desktops	100%	79%
37 Desktops	N/A	85%
40 Desktops	N/A	90%
45 Desktops	N/A	92%
50 Desktops	N/A	94%

Similarly to the medium workload, NexentaConnect under heavy LoginVSI workload provides for lower IO latency and increased IO performance. The blade with NexentaConnect installed on it easily runs 37 desktops (Table 3) at average total CPU utilization below 90% threshold. Desktop density is thus getting increased more than 3 times (300+% improvement). The platform can be pushed to run more than 37 heavy-loaded desktops, as the data above shows.

Figure 7 shows CPU utilization during the heavy LoginVSI-generated workload represented in the "with NexentaConnect" column of the table 3 above.

Figure 7: 50 linked-clone desktops under heavy workload on NexentaConnect

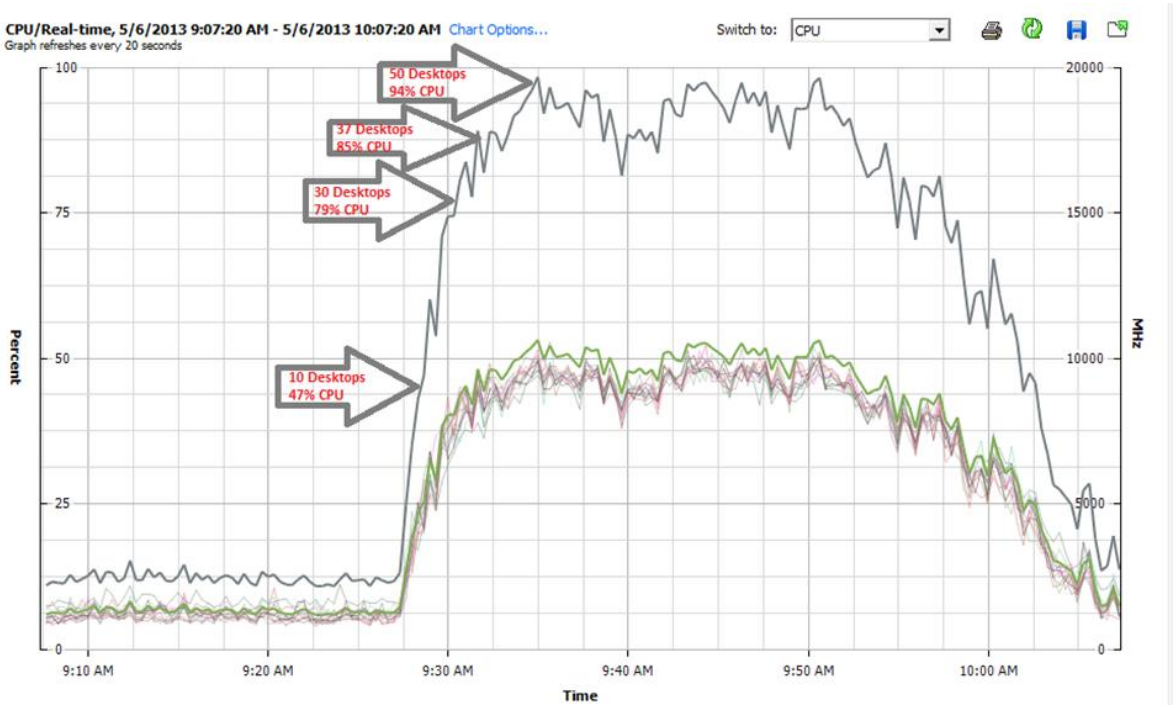

When using NexentaConnect, vSphere CPU utilization graph (Figure 7) for 50 linked-clones desktops under heavy desktop virtualization workload confirms again the Table 3 results: desktop density increased from 12 to 37 desktops - a 3x improvement.

Further, LoginVSI-generated VSImax score and the Baseline performance graph for the same benchmark on 50 heavily loaded desktops (Figure 8) shows and validates the maximum number of desktops the platform can handle (37, in this case).

LoginVSI Baseline improvement, from 5510 to 3561 reconfirms the NexentaConnect provided storage latency and compute response time improvements.

Figure 8: VSImax result for 50 linked-clone Desktops running heavy workload on Nexenta Connect

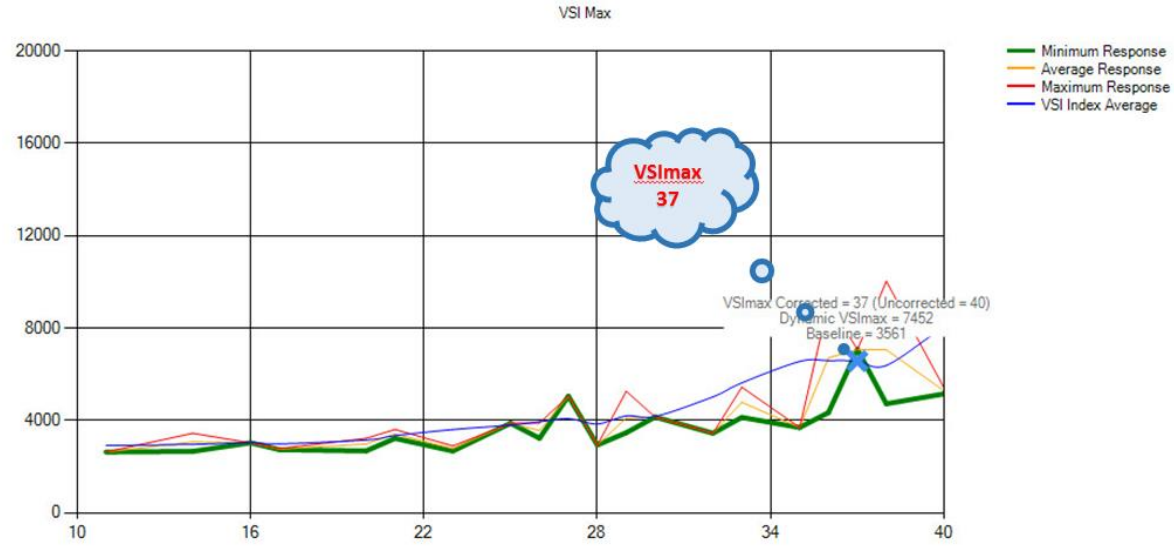

Compression Ratio

Compression ratio of the desktop virtualization persistent data varies based on the desktop pool type and installed desktop applications. With freshly deployed linked-clones pool of 55 desktops we've registered compression ratios 1.9x to 2.5x; compression ratio reached 4x for full-clone non-deduplicated desktops.

Figure 9: Data compression

```
root@NexentaStorVSA:/volumes# zfs get compressratio
NAME                              PROPERTY       VALUE  SOURCE
DESKTOP                           compressratio  1.91x  -
syspool                           compressratio  1.14x  -
syspool/dump                      compressratio  1.00x  -
syspool/rootfs-nmu-000            compressratio  1.38x  -
syspool/rootfs-nmu-000@initial    compressratio  1.37x  -
syspool/swap                      compressratio  1.00x  -
root@NexentaStorVSA:/volumes#
```

Iometer benchmarks – with and without NexentaConnect

NexentaConnect includes built-in Iometer benchmark, with the capabilities to easily schedule and run a user-selected workload (or a sequence of workloads with varying read/write ratios) simultaneously on each provisioned virtual desktop.

Iometer benchmark on 55 desktops (simultaneously), at 4K blocksize random, 25% read 75% write workload, measured a total of 52.3 IOPS – better than 50 IOPS required to support desktop virtualization power users (Figure 1).

Iometer Workload:

- Number of desktops: 55
- Duration: 600 seconds
- Workload: 100% random
- Block size: 4K
- Workers = 1 (equals to desktop vCPU)

Table 6 below shows Iometer read and write IOPS on a Series E server blade, at a variety of read/write ratio from 100% write (leftmost) to 100% read (rightmost column).

Table 4: Each of the 55 full-clones desktops running Iometer: with and without NexentaConnect

	R:0% W: 100% w/o NC	R:0% W:100% w/ NC	R:25% W:75% w/o NC	R:25% W:75% w/ NC	R:50% W:50% w/o NC	R:50% W:50% w/ NC	R:75% W:25% w/o NC	R:75% W:25% w/ NC	R:100% W:0% w/o NC	R:100% W:0% w/ NC
Write Ops IOPS	6.25	74.50	3.5	39.2	3.2	61.1	6	45.6	0	0
Read Ops IOPS	0	0	4	12.9	3.4	62.0	3.2	137.3	6	292.1
Avg. Write KB/s	30.25	297.98	20.46	157	15.7	244.7	15.4	182.4	0	0
Avg. Read KB/s	0	0	10.2	51.7	18.4	248.0	18.3	549.5	28	1.1
Write Latency ms	12	3.6	7	0.3	5.5	2.1	3.4	0.9	0	0
Read Latency ms	0.0	0.0	3.5	1.3	6.2	1.5	8.9	2.6	12.1	3.6

NexentaConnect achieves lower latency and increased performance, reduced CPU utilization and increased desktop density (from 18 to 55 desktops).

Figure 10 shows disk level IO statistics collected from VSA. “r/s” and “w/r” columns contain read and write operations per second at the NexentaConnect VSA disk level. Second and third lines (Device: C1t1d0 and C2t1d0) represent the two 900GB 10K SAS HDDs used by the VSA to store desktop pool.

Figure 10: VSA disk-level IO

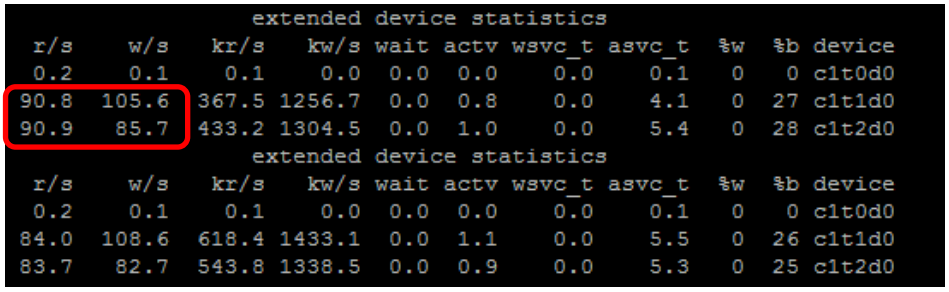

Performance improvement can be further explained by comparing the disk level IO from the Figure 10 versus Iometer results from Table 6.

Iometer measured 137.5x55 = 7562.5 read IOPS vs. 90.9+90.8 = 181.7 read IOPS from HDDs

Iometer measured 182.4x55 = 10032 write IOPS vs. 105.6+85.7 = 191.3 write IOPS to HDDs

Total average IOPS measured during this Iometer test: 17594.5 vs. 373 IOPS without NexentaConnect. Note that the latter number (373 IOPS) is within the 10K RPM SAS drive specifications.

NexentaConnect significantly accelerates storage performance, a boost from 373 IOPS to 17594 IOPS, or a 47x (4700%) improvement.

Figure 11: CPU utilization during 1 hour Iometer benchmark

Performance Chart Legend

Key	Object	Measurement	Rollup	Units	Latest	Maximum	Minimum	Average
	0	Usage in MHz	Average	MHz	876	1547	132	1053.139
	3	Usage in MHz	Average	MHz	863	1506	95	1033.922
	1	Usage in MHz	Average	MHz	852	1646	99	1074.867
	NV4V_VSA_3142BA42_7022	Usage	Average	Percent	46.15	84.21	6.14	56.724
	NV4V_VSA_3142BA42_7022	Usage in MHz	Average	MHz	3322	6063	442	4083.672
	2	Usage in MHz	Average	MHz	725	1426	97	917.039

Figure 12 demonstrates an overall read cache performance and cache efficiency – notice that almost all reads (99% in this case, to be exact) are satisfied from the cache.

63% of the prefetched data is read – a very good result for a random IO workload.

The underlying NexentaConnect technology for read caching is called Adaptive Replacement Cache, or ARC (see, for instance, http://en.wikipedia.org/wiki/Adaptive_replacement_cache). Figure 12 (and all the rest benchmarks performed on the E series server blade with NexentaConnect) proves in part that NexentaConnect ARC is extremely efficient in the desktop virtualization environment.

Figure 12: VSA ARC Cache Hits

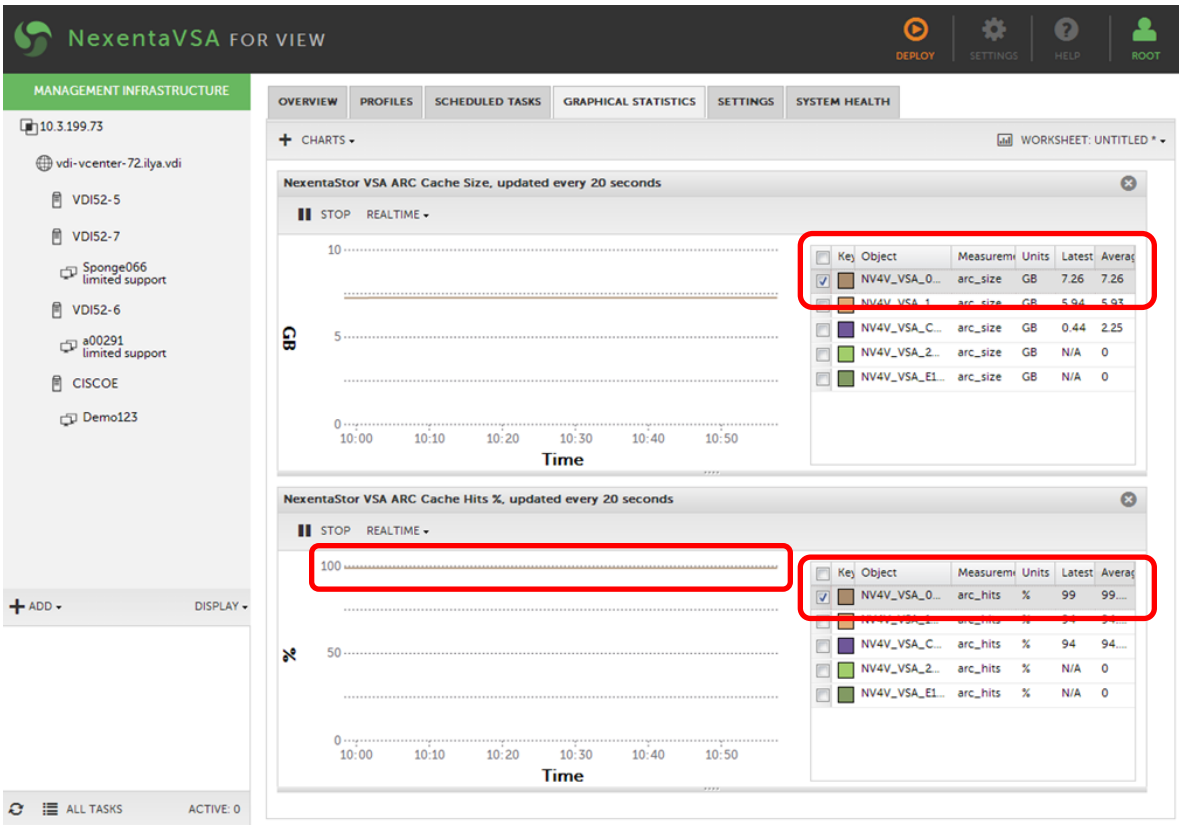

Extra details on ARC utilization can be obtained from the VSA's console.

```
Data Demand   Efficiency:    99%
Data Prefetch Efficiency:    63%
```

Full-Clone Desktops: Optimizing Times and Used Space

NexentaConnect includes NAS VAAI license feature that implements all the 4 NAS VAAI primitives in compliance with the VMware vSphere 5 API:

- Full File Clone
- Native Snapshot
- Extended Statistics
- Reserve Space

A combination of the NAS VAAI (NexentaConnect) capability and inline deduplication of the NexentaConnect VSA provides two immediate benefits: it optimizes (speeds-up) full-clone deployment and effectively reduces physical space dedicated to the full clone desktop pool.

In the following test, a pool of 24 full-clone desktops was deployed with and without NexentaConnect NAS VAAI license feature.

Table 7 shows that without NAS VAAI deployment of 24 full-clone desktops on Cisco UCS E series blade server took 7 hours 28mins including customization. When NAS VAAI is used the average clone time drops from 13mins to 2mins; total deployment time including desktops customization drops to 2hours and 38mins – an overall speed-up of 5.4 times (540%).

Table 5: Full clone deployment: with and without NAS VAAI + deduplication

	WITHOUT NAS VAAI	WITH NAS VAAI
"Pure" cloning time	13 min 36 sec	2 min 36 sec
Cloning + Customization	18 min 10 sec	4 min 38 sec
Total Deployment Time	7 hours 28 min	2 hours 38 min

Table 8 documents 8.5x storage space saving for the virtual full-clone desktops:

Table 6: Full clone capacity

	WITHOUT NAS VAAI	WITH NAS VAAI
Used Space	408GB	48GB

Results and Conclusions

This paper describes the following benchmarks:

- 3x Increased Density (LoginVSI)
- 47x Increased Performance (Iometer)
- 5.4x Faster Deployment Time (including customization)
- 8.5x Capacity saving for the full-clone desktops (via NAS VAAI + dedup)

Final conclusion and observations:

- For LoginVSI benchmarks we have achieved 3x higher density (from 18 to 55 desktops for medium workload, and from 12 to 37 desktops under heavy workload)
- Inline compression reduced write (KB/s) bandwidth to the disks 1.9 times to 4 times
- Iometer measured read performance improved from 370 IOPS to 17700 total IOPS at 75% write 25% read random 4K blocksize workloads
- A combination of reduced I/O to the disk and increased storage performance has in turn contributed to reduced CPU utilization by the storage subsystem – and therefore more CPU being available for the business logic (that is, the apps of the virtual desktops)
- For the full-clone desktops, we have achieved both significant reduction in the deployment time and significant savings of the storage capacity

All results of this study are independently verified by Cisco (the following links are available for Cisco partners):

- NexentaConnect on E-series: https://marketplace.cisco.com/catalog/products/3144
- NexentaConnect View + UCS E-Series: https://marketplace.cisco.com/catalog/products/3509/files/55478
- NexentaConnect + Cisco UCS – VMware Certification: https://marketplace.cisco.com/catalog/products/3509/files/55479