



Technical White Paper

AccelStor NeoSapphire VDI Reference Architecture for 500 users with VMware Horizon 7

Aug 2018

The information contained herein is subject to change without notice.
All trademarks and trade names are owned by their respective owners.

Table of Contents

Executive summary.....	5
Solution benefits and highlights.....	6
1. Introduction	7
2. AcceleStor: The best AFA solution for VDI	8
2.1. FlexiRemap Technology	8
2.2. Simplified Virtualization.....	8
2.3. True High Availability	8
3. Virtualization solution	9
3.1. VMware Horizon components	9
3.2. View Composer linked clone	9
4. Infrastructure and test configuration	10
4.1. Hardware Infrastructure	10
4.1.1. Storage architecture	10
4.1.2. Networking	12
4.1.3. Host configuration	12
4.2. Software components	13
4.3. Virtual desktop	14
5. Test methodology.....	15
5.1. Test objectives	15
5.2. Login VSI	15
5.2.1. Workloads.....	15
5.2.2. Workload resources usage.....	17
5.3. VDI working-day	17
Important notes:.....	17
5.4. Monitoring tools	18
5.4.1. VMware ESXTOP	18
5.4.2. vSphere client performance monitor	18
5.4.3. NeoSapphire performance monitor	18
6. Test validation and results	19
Considerations and notes for test validation	19
6.1. Provisioning 500 desktops	20
Storage array performance.....	20
Hypervisor performance.....	20
6.2. Booting 500 desktops	23
Storage array performance.....	23

Hypervisor performance.....	23
6.3. Login VSI test	26
6.3.1. Knowledge Worker results	26
6.3.2. Power Worker results	29
6.3.3. Login VSI extended steady-state	31
Knowledge Worker: 30 minutes steady-state	32
Power Worker: 30 minutes steady-state	33
6.3.4. Workloads comparison	35
6.4. Power-off and pool refresh	36
6.5. Storage capacity.....	36
7. Conclusion.....	37
8. AccelStor sizing criteria for VDI.....	38
Appendix.....	40
About Login VSI	40
References.....	40
Version history	40

List of Tables

Table 1: Common VDI operations of a Power Worker workload.	6
Table 2: VMware Horizon components	9
Table 3: Hardware components of server infrastructure	12
Table 4: Hardware components of desktops infrastructure	13
Table 5: Software components	13
Table 6: Virtual Machine (desktop) configuration	14
Table 7: Knowledge and Power Worker configuration.	16
Table 8: Login VSI 4.1 workloads resources usage.....	17
Table 9: vCenter Advanced Settings	18
Table 10: Knowledge Worker test overview	26
Table 11: Power Worker test overview	29
Table 12: 8 hours steady-state performance for a Power Worker.	31
Table 13: 30 minutes steady-state performance comparison between Knowledge and Power Worker.	35
Table 14: Desktop pool average values for power-off and pool refresh operations.....	36
Table 15: Booting and steady-state performance comparison.....	38

List of Figures

Figure 1: Overall test environment	10
Figure 2: Storage layout for desktops infrastructure	11
Figure 3: Storage layout for server infrastructure	11
Figure 4: Network switches	12
Figure 5: Array performance when provisioning 500 desktops	20
Figure 6: Array network traffic when provisioning 500 desktops	20
Figure 7: Hypervisor CPU utilization when provisioning 500 desktops	21
Figure 8: Read and write IOPS when provisioning 500 desktops	21
Figure 9: Throughput when provisioning 500 desktops	22
Figure 10: Average latency when provisioning 500 desktops	22
Figure 11: Array performance when booting 500 desktops	23
Figure 12: Array network traffic when booting 500 desktops	23
Figure 13: Hypervisor CPU utilization when booting 500 desktops	24
Figure 14: Read and write IOPS when booting 500 desktops	24
Figure 15: Throughput when booting 500 desktops	25
Figure 16: Average latency when booting 500 desktops	25
Figure 17: Knowledge Worker VSImax overview	27
Figure 18: Average CPU usage of ESXi cluster for a Knowledge Worker test	27
Figure 19: Read and write IOPS for a Knowledge Worker test	28
Figure 20: Average latency and total IOPS for a Knowledge Worker test	28
Figure 21: Power Worker VSImax overview	29
Figure 22: Average CPU usage of ESXi cluster for a Power Worker test	30
Figure 23: Read and write IOPS for a Power Worker test	30
Figure 24: Average latency and total IOPS for a Power Worker test	31
Figure 25: Average CPU usage of ESXi cluster for a steady-state Knowledge Worker	32
Figure 26: Read and write IOPS for a steady-state Knowledge Worker	32
Figure 27: Average latency and total IOPS for steady-state Knowledge Worker	33
Figure 28: Average CPU usage of ESXi cluster for a steady-state Power Worker	33
Figure 29: Read and write IOPS for a steady-state Power Worker	34
Figure 30: Average latency and total IOPS for steady-state Power Worker	34
Figure 31: VSImax 4.1 comparison between Knowledge and Power Workers	35
Figure 32: Capacity used after booting from NeoSapphire dashboard.	36

Executive summary

This reference architecture describes the solution infrastructure and benchmark validation of 500 linked-clones, the desktops are deployed with VMware vSphere 6.7 and Horizon View 7.3 on an AcceleStor NeoSapphire P710 All-Flash Array (AFA). The reference architecture also contains hardware and software configuration used through the validation with clear and comprehensive results of the storage performance tested. It also provides storage sizing guidance and best practices for efficiently designing, configuring and deploying Virtual Desktop Infrastructure (VDI) environments with AcceleStor NeoSapphire AFA and VMware.

The solution demonstrates the storage performance capabilities of VDI's heaviest and most resource-intensive workloads, this is proved by using only "knowledge" and "power workers" desktop configurations to stress the whole system infrastructure. Later, the results are compared to validate that no matter the workload, the user experience, or desktop application response-time will remain unaffected.

The overall solution cost is also considered in this report, that is a key criteria for the success of a VDI project. VDI helps the lower total long-term cost of ownership, but for some solutions, the upfront solution cost could be expensive. AcceleStor's FlexiRemap technology combined with other storage techniques such as thin-provisioning and deduplication, cut-down the storage cost as low as US\$30 per power-desktop for all NeoSapphire Professional models when full desktop density is used in the storage arrays. In addition to the inexpensive cost per desktop, all data service fees and helpdesk support are inclusive in the AcceleStor solution.

Solution benefits and highlights

The highlights and benefits found when using AcceleStor NeoSapphire for VDI solutions are summarized below:

- **FlexiRemap:** With the Award-Winning Technology, FlexiRemap, storage I/O performance is no longer an issue for VDI environments. Power Worker workload test shows average steady-state of 3,900 total IOPS while keeping latency under 0.6ms.
- **Simple configuration:** A pre-configured and simple to manage system. Forget about other complex solutions such as traditional RAID, SSD cache, automated storage tiering and write-back cache. With AcceleStor NeoSapphire, easily create volumes, present them to the hosts and the storage configuration is done.
- **Lower overall cost:** With license-free data services, AcceleStor is leading in Price-Performance on SCP-1, helping to reduce TCA and TCO dramatically. At full desktop density, for all AcceleStor NeoSapphire models, the VDI storage cost average as low as US\$30 per power-desktop.
- **Shared nothing architecture:** Full redundancy hardware design with 99.9999% reliability and dedicated storage pool per node for a real no single point of failure VDI storage solution.
- **Data reduction:** With the fastest data reduction that does not compromise VDI performance, using both inline and post-process deduplication increases desktop density from 5x up to 15x if full clones are used. With the use of linked clones, capacity efficiency save is proven to be over 98%.

The key performance results found for common VDI operations of a Power Worker workload, are listed in table 1.

Table 1: Common VDI operations of a Power Worker workload.

Pool of 500 Linked Clones	Time to complete	Avg. write IOPS	Avg. read IOPS	Avg. Latency
Provisioning	~32 min	4,800	3,800	0.59ms
Bootimg	~2 min	9,400	11,700	0.68ms
Login VSI full test (Login)	50 min	2,200	289	0.46ms
30 minute steady-state	30 min	3,200	377	0.56ms
Power off	~7 min	1,500	652	0.45ms
Pool Refresh	~14 min	492	860	0.27ms

1. Introduction

In a VDI environment, user desktops are hosted as virtual machines (VMs) in a centralized infrastructure and delivered over the network to an end-user client device. Organizations have much to gain from this end-user computing model, which has proven to be an effective method for reducing costs and complexities of client networks while increasing IT management efficiency. With the boom of BYOD (Bring your own device), the modern worker now can expand their environment including a variety of end-point devices such as zero-clients, notebooks, tablets, and mobiles, while better utilizing resources through the VDI practice.

When opting to deploy VDI, many factors should be considered. Desktop pool deployments require high capacity and performance demands from the storage array. When large amounts of desktops are deployed into the centralized infrastructure, an enormous I/O stresses is placed on the storage network. But performance demands are not determined only by the number of I/O operations per second (IOPS), that could be served easily by any modern storage array; latency is one of the key metrics to analyze during the desktop operations of the whole VDI-day. Operations such as system boot, login and steady-state, must be sustained with enough IOPS capacity while keeping ultra-low latencies.

Traditional storage arrays have failed to fulfil the demands for VDI performance. Complicated storage layouts that require near-obsolete RAID configurations, poor storage I/O and latency performance and costly scale-out options, translate into common storage bottlenecks for VDI projects. Often these legacy systems do not even passes the POC stage due to the poor end-user experience. Issues with poor performance forces IT departments to abandon VDI and move back towards traditional desktops.

When storage performance really matters, an all-flash array built on a flash-optimized architecture is the best possible choice. AcceleStor's award-winning FlexiRemap software technology is the most innovative solution to empower the all-flash array system and keep VDI workloads and applications running smoothly.

2. AccelStor: The best AFA solution for VDI

NeoSapphire Professional AFAs deliver stunning performance, ranging from 360K IOPS to 700K IOPS at 4Kb random write, in a 2U platform with 24 SSDs. These AFAs support mainstream connectivity protocols, like 10GbE, 16G Fibre Channel for enterprises, and even 56G FDR Infiniband for lab or research environments. Plus, with the FlexiRemap software technology, NeoSapphire Professional models maintain ultra-high IOPS and low latency even under intensive workloads, making them optimal solutions for handling huge amounts of high-value data for enterprise-level businesses.

2.1. FlexiRemap Technology

AccelStor's FlexiRemap technology is software that is based on the fundamentals of flash memory. It has been awarded Best-of-Show Technology Innovation at Flash Memory Summit 2016. This proprietary algorithm enables AFAs powered by FlexiRemap to meet data-access demands more efficiently than conventional SSD arrays or RAID systems. FlexiRemap is crucial for speeding up random-write access patterns, which are often a major cause of storage bottlenecks. FlexiRemap ensures that AccelStor storage devices meet the stringent performance demands of mission-critical applications.

2.2. Simplified Virtualization

AccelStor provides seamless storage solutions for virtualization environments, especially for the VMware platform. NeoSapphire Professional AFAs offload storage operations by using VAAI block implementation, providing an optimized user experience under VDI applications. AccelStor has developed a NeoSapphire management plug-in for the VMware vSphere Web Client with which hypervisors can perform NeoSapphire AFA volume management tasks. This plug-in is a very handy tool for virtualization users.

2.3. True High Availability

The NeoSapphire High Availability AFA is built on a shared-nothing architecture consolidated dual storage nodes that are completely independent and self-sufficient systems. Each node has its own set of resources, including CPU, memory, network and a dedicated storage pool. These nodes execute synchronous replication with each other via a high-speed, low-latency InfiniBand FDR interlink. This eliminates the risk of failure caused by shared components. This approach, specifically designed to provide no single point of failure and minimal service interruptions, is the way to true high availability.

3.Virtualization solution

VMware is the global leader in cloud infrastructure & digital workspace technology, it accelerates digital transformation for evolving IT environments. Two VMware products were used in the architecting of this solution.

- **VMware vSphere 6.7** is the virtual platform selected for virtualizing server and desktop infrastructures. vSphere provides a powerful, flexible, and secure foundation for business agility.
- **VMware Horizon 7.3** is used to deploy and manage the virtual desktops pool. It also provides end users access to the virtual desktops, applications, and online services through a single digital workspace.

3.1. VMware Horizon components

The essential Horizon components used in this paper are detailed in table 2.

Table 2: VMware Horizon components

vCenter Server	Orchestrates VMware Virtual Infrastructure by configuring, provisioning, and managing the data centers.
View Connection Server	Authenticates users through Microsoft Active Directory and brokers client connections.
View Administrator	Web-based administration platform for View Infrastructure components.
View Client	Software that is used to access View desktops.
View Agent	Provides communication between the View Connection server and virtual desktops.
View Composer	Service used to create pools of linked clone VM images to reduce storage and management requirements.
Client Devices	Personal computing devices used by end users to run the View Client.

3.2. View Composer linked clone

Linked clones (non-persistent desktop) is one out of the three options available in Horizon View Composer for desktop deployment. When persistent desktops are used, the administrator usually has to consider additional storage and other administrative requirements. Simplifying the overall solution, this reference architecture uses linked clones to streamline the storage capacity and management.

An advantage of linked clones is their fast creation, allowing users to set up distinct virtual machines for specific operations or tasks. Linked clones are therefore excellent for joint projects and teamwork, for example, quality assurance, stress testing, and debugging.

4. Infrastructure and test configuration

This section provides information about the infrastructure configured for testing the virtual desktop solution. It includes hardware components used such as storage, host, and networking, also detailed software used and applications under test. These components are illustrated in Figure 1.




Figure 1: Overall test environment

4.1. Hardware Infrastructure

4.1.1. Storage architecture

AcceleStor P710 was the selected model for architecting the solution presented in this paper. For an I/O pattern of 4K random writes, P710 can offer a sustained 700,000 IOPS while preserving latency under 1ms. This outstanding benchmark result uses the same demanding I/O pattern seen for virtual desktops, making easy for AcceleStor to design, configure and plan the best AFA solution for VDI in a worst-case scenario.

The storage layout was carefully designed to fulfill the best performance, availability, and capacity growth. Based on Horizon 7 sizing limits and recommendations, up to 500 VMs could be allocated in a single LUN, but consideration must be taken for IOPS and latency capacity, also for maintenance operations such as rebalancing, new desktops pools creations and desktops high availability. This reference uses 5 volumes to allocate the virtual desktops and an additional volume is created to place the replica virtual machine, to isolate read and write I/O commands and easily monitor storage performance. Each volume was created with 2TB capacity. This configuration is described in Figure 2.




Figure 2: Storage layout for desktops infrastructure

A different storage environment is used to store the data center (virtual) servers, the AccelStor NS3405 AFA is selected to for this job. This model provides enough IOPS capacity for the most important virtual machines of the whole infrastructure, the data center servers. Also, the Login VSI benchmark tool was deployed in this storage array, making sure that the test will run smoothly without compromising the desktops environment.

Two volumes are created, one for the data center servers and another for Login VSI servers, each volume with 2TB of storage capacity, as described in Figure 3.




Figure 3: Storage layout for server infrastructure

4.1.2. Networking

For data storage connectivity, 10GbE SPF+ over iSCSI was used with two data ports of the storage array. Following VMware and AccelStor best practices, ESXi hosts are configured to guarantee the best storage performance with the P710. Some of these configurations include:

- Jumbo Frames are enabled in all relevant elements of the infrastructure.
- iSCSI configuration set login Timeout to 30
- iSCSI configuration set DelayedAck to disable
- Port binding and round-robin configured.
- Round robin IOPS limit adjusted to 1.
- VAAI is enabled by default.
- One of the storage path is used for vMotion

Two network switches are configured to interconnect storage and ESXi host. One 1GBps switch for the management and virtual machine network and, other of 10GBps for the data storage network.




Figure 4: Network switches

4.1.3. Host configuration

The data center infrastructure was set up within two Intel-based servers. One server to store the data center virtual servers and the other server to store Login VSI data server and Launchers. Server specs are shown in Table 3.

Table 3: Hardware components of server infrastructure

Hypervisor:	VMware ESXi, 6.7.0, 8169922	 x2 host
Model:	GS-R12P8E	
Processor Type:	Intel(R) Xeon(R) CPU E5-2450 v2 @ 2.50GHz	
Logical Processors:	32	
CPU capacity:	40 GHz	
Memory capacity:	192 GB	

The desktop infrastructure and test configuration were held within a Supermicro blade chassis, which consists of 8 nodes (SYS-F618R2-R72+). A total of 352 GHz of CPU capacity and 2TB of memory capacity is used among all the 8 nodes. The node configuration is shown in Table 4.

Table 4: Hardware components of desktops infrastructure

Hypervisor:	VMware ESXi, 6.7.0, 8169922	 x8 nodes
Model:	SYS-F618R2-R72+	
Processor Type:	Intel(R) Xeon(R) CPU E5-2450 v2 @ 2.50GHz	
Logical Processors:	40	
CPU capacity:	44 GHz	
Memory capacity:	256 GB	

4.2. Software components

This section describes the software used for the storage arrays, data center servers and Login VSI infrastructure used to test the virtual desktops. All virtual servers used in this infrastructure are built under Windows Server 2016 datacenter edition (64-bit). Table 5 describes the software version for each component.

Table 5: Software components

NeoSapphire AFA	System Image Version	
P710	2.2.0 (Build 3713)	
NS3405	2.1.0 (Build 3656P3)	
VMware software	OS	Version
ESXi	NA	6.7.0, 8169922
vCenter Server	NA	6.7.0, 8217866
Horizon View Administrator	Windows Server 2016 datacenter (64-bit)	7.3.2 build-7161118
Horizon View Composer	Windows Server 2016 datacenter (64-bit)	7.3.2 build-7161120
Horizon View Agent	NA	7.3.2 build-7161471
Horizon View Client	NA	4.7.0 build-7395453
Login VSI	OS	Version
Login VSI dataserver	Windows Server 2016 datacenter (64-bit)	4.1.32.1 - Pro
Launchers x20	Windows Server 2016 datacenter (64-bit)	NA
Data center	OS	Version
Domain Controller	Windows Server 2016 datacenter (64-bit)	NA
SQL database	Windows Server 2016 datacenter (64-bit)	2017.0140.17218.00
Microsoft SQL Server Management Studio	NA	v17.7 (14.0.17254.0)

4.3. Virtual desktop

The virtual desktops pool is created based on Linked Clone technology (Described in section 3.2). The parent VM is configured to test “knowledge” and “power” users, so it is constructed with 2 vCPU and 3.5 GB of memory (all memory is reserved by the ESXi hosts). Windows 10 version 1709, is the operating system used. It is important to notice and mention the Windows version used to compare other configurations, due to CPU, memory, and disk I/O resources consumption is different and has been increasing with the latest versions.

The Parent VM was configured following best practices from Windows, VMware and Login VSI to provide the best user experience and performance results. The VM was tuned up using VMware OS Optimization Tool (from VMware labs) with the “Login VSI like a pro” template provided by Login VSI. Additional configurations were made in the Windows operating system to boost OS login and improve software and applications response times.

Table 6 detailed the hardware and software configuration used for the parent and all linked clone virtual desktops.

Table 6: Virtual Machine (desktop) configuration

Hardware	vCPU	2
	Memory	3.5 GB (100% reserved)
	HDD	60 GB
OS	Edition	Windows 10 Pro
	System type	64-bit
	Version	1709
Software	MS office 2016, 32-bits: Word, Excel, PowerPoint and Outlook	
	Login VSI applications: Adobe Acrobat, Doro PDF writer, FreeMind, Photo Viewer.	
	Others: View agent, Internet Explorer.	

5. Test methodology

This section defines the test objectives and describes the benchmark and monitoring tools used to collect significant data during the testing process.

5.1. Test objectives

The general objective is to determine the desktop capacity that can be deployed in this test environment using the AcceleStor P710 AFA while keeping excellent user experience indicators for Knowledge and Power Worker profiles.

Specifics objectives:

- Use Login VSI benchmark tool to saturate the entire environment and understand the behavior during the concurrency of the steady-state desktops.
- Determine the performance impact on the storage array during peak I/O activity such as provisioning, boot and login storms.
- Compare performance resources used by a Knowledge and Power Worker against the user-experience score reported by the Login VSI report.
- Evaluate the ideal maximum desktop density for AcceleStor NeoSapphire models using IOPS and latency key findings.

5.2. Login VSI

Login VSI (www.loginvsi.com) is the industry standard load-testing tool for measuring the performance and scalability of centralized Windows desktop environments, such as server-based computing (SBC) and VDI. Login VSI is used for testing and benchmarking by all the major hardware and software vendors and is recommended by both leading IT analysts and the technical community. Login VSI is 100% vendor independent and works with standardized user workloads and statistics—making all conclusions based on Login VSI test data objective, verifiable, and repeatable.

Login VSI-based test results are used and published in multiple [technical white papers](#) and presented at various IT-related conferences by our vendor customers. The product Login VSI is also widely used by end-user organizations, system integrators, hosting providers, and testing companies. It is also the standard testing tool used in all tests executed in the internationally acclaimed research project VDI Like a Pro (formerly known as Project Virtual Reality Check).

5.2.1. Workloads

Login VSI 4.1 uses different types of user workload that can be configured for testing scenarios. These

workloads feature a timer mechanism, which is executed multiple times within each workload and determines the performance of the system during testing. This reference architecture uses only Knowledge and Power Worker to prove the high storage performance of NeoSapphire AFA series and demonstrate how low-latency values can be kept no matter the workload.

Detailed information (from Login VSI) about the Knowledge and Power workloads is presented in table 7.

Table 7: Knowledge and Power Worker configuration.

Worker	Knowledge	Power
Definition	The Knowledge Worker is designed for 2(v)CPU environments. This is a well-balanced intensive workload that stresses the system smoothly, resulting in higher CPU, RAM and IO usage.	The Power Worker is designed for 2(v)CPU+ environments. This is a very intensive workload that puts maximum stress on the system, resulting in very high CPU, RAM and IO usage.
Used applications	• Adobe Reader • Freemind/Java • Internet Explorer • MS Excel • MS Outlook • MS PowerPoint • MS Word • Photo Viewer	• Adobe Reader • Freemind/Java • Internet Explorer • MS Excel • MS Outlook • MS PowerPoint • MS Word • Photo Viewer • Simulated application installs
Additional about the applications	It also uses native Windows apps (Notepad and 7-Zip) to complete the print and zip actions used by the workload meta functions.	For many of the applications in this workload, larger files and higher resolution media are used. It also uses native Windows apps (Notepad and 7-Zip) to complete the print and zip actions used by the workload meta functions.

The Power Worker workload uses the Login VSI Pro Library, this is an extra content that contains extra data files which are used by Login VSI Power Worker workload. The idea is to have a “larger” pool of “random” files (think of caching) that will be processed during a workload. The complete library is 9.3GB in size.

5.2.2. Workload resources usage

Every workload consumes different system resources depending on the number of applications opened in the machine. Table 8 contains the resource usage per workload, which provides crucial information for advanced users. These metrics are measured from a single user/VM, with all metrics based on the Knowledge Worker workload.

Note: These values will be different depending on the hardware used, virtual machine optimizations and user concurrency. The information provided below is only provided to help understand how they behave relative to one another.

Table 8: Login VSI 4.1 workloads resources usage

Workload Name	Apps Open	CPU Usage	Disk Reads	Disk Writes	IOPS	Memory	vCPU
Task	2-7	70%	79%	77%	6	1GB	1vCPU
Office	5-8	82%	90%	101%	8,1	1,5GB	1vCPU
Knowledge	5-9	100%	100%	100%	8,5	1,5GB	2vCPU
Power	8-12	119%	133%	123%	10,8	2GB	2vCPU+

5.3. VDI working-day

Besides the test of login and steady-state phases possible by Login VSI, additional data is collected for common operations for the “VDI working day”. Operations such as provisioning, booting, powering-off and, pool refresh are manually tested for the 500 VMs used in this solution. Also, pool refresh, one of the key linked clone maintenance operations is reported.

Important notes:

- The VDI-day operations could be sped up by modifying the Advanced Setting for the vCenter server from the Horizon View administrator. This test reference uses the default setting, to simplify the results analysis and make it easy to compare with other AccelStor AFA products and vendors solutions.
- These Advanced Settings are described in Table 9. It indicates, for example, that only 8 Linked Clones will be provisioning by View Composer at the same time, if this value is set higher, virtual machine creation will be boosted with the high performance of NeoSapphire AFA.

Table 9: vCenter Advanced Settings

Operation	Value
Max concurrent vCenter provisioning operations	20
Max concurrent power operations	50
Max concurrent View Composer maintenance operations	12
Max concurrent View Composer provisioning operations	8

5.4. Monitoring tools

The Analyzer tool by Login VSI is used after the test is completed to evaluate the performance of the system and to calculate the VSI_{max}. Other tools are used in this test for deeper analysis of the storage performance and gather the most important findings for this solution.

5.4.1. VMware ESXTOP

VMware ESXTOP displays information about the state of the physical servers running VMware ESXi. It lists CPU utilization for each physical processor, memory utilization, network bandwidth, disk performance, and storage adapters available for each ESXi host. ESXTOP is used in batch mode to capture this relevant information of the 8 nodes cluster used to host the virtual desktops.

5.4.2. vSphere client performance monitor

Provides real-time analysis during the test of all the servers in the cluster managed by vCenter. The total CPU and memory consumed were observed to define the number of servers needed for the cluster and make sure that no resource bottlenecks were caused. Also, "Recent Tasks" from vSphere client was used to measure times to complete different test operations.

5.4.3. NeoSapphire performance monitor

The NeoSapphire performance monitor provides essential information about the array performance, storage capacity, and storage computing resources.

6. Test validation and results

This section describes the results reported by Login VSI analyzer. Login VSI represents the test during the login time and steady state response time of the VMs for the whole system. To provide a full comprehension of the infrastructure, provisioning, booting and other VDI operations are also tested and reported.

Considerations and notes for test validation

- The results are presented in the order that should be considered for a VDI POC for easy understanding of the test environment and following a step-by-step VDI solution deployment.
- During booting, provisioning, login VSI test and other tested workloads, the storage controller CPU usage of P710 was always under 4% for the storage I/O operations, and used memory was under 16%.
- Storage latency reported in this paper is based on *Average guest /millisecond* from in the ESXi hosts, that is the latency observed directly from the datastores allocating the VMs. This latency is the sum of the *Average Driver MilliSec/Command* (latency of the storage arrays) and the *Average Kernel MilliSec/Command* (latency of the VMkernel). Notice that several vendors report the latency directly from their back-end storage array controllers, avoiding queuing overheads from network fabrics.
- All data from the hosts is reported as the average or total of the ESXi cluster, composed of 8 computing nodes. In general, the key metrics to analyze are the front-end CPU utilization, IOPS, and latency of the whole infrastructure.
- As is to be expected for VDI environments, write IOPS predominated during the steady-state test, as high as 89% for a Power Worker. On the other hand, reads IOPS predominated during booting, hitting 55%.
- To avoid slow response times in VDI environments and VM failures, CPU utilization of the ESXi infrastructure shouldn't be higher than 90% during any of the commons VDI operations. Only during booting time, peaks CPU usage hits 90%.
- Login VSI figures were scaled and modified to easy illustrate the relevant information only.
- Knowledge Worker test was performed enabling "benchmark mode". This test mode creates a default environment in which users can compare test results with each other. For a Power Worker workload, this mode is not available, but the equivalent configuration was used.
- For the extended steady-state test, in the Login VSI test wizard the option "*How long should the test keep running after the last user has been launched*" was configured to 30 minutes so that the new total test time was 78 minutes total.

6.1. Provisioning 500 desktops

Storage array performance

Storage array performance during the provisioning of the desktop pool. These numbers represent back-end IOPS and throughput. The storage array data reports:

- Total IOPS, a peak of 191K with an average of 38K IOPS.
- Total Throughput, a peak of 755MBps with an average of 141MBps.

Figure 5 shows back-end IOPS and throughput performance during the desktop pool provisioning.




Figure 5: Array performance when provisioning 500 desktops

Figure 6 shows inbound and outbound network traffic of the two data storage ports used during provisioning.




Figure 6: Array network traffic when provisioning 500 desktops

Hypervisor performance

Hypervisor performance of the tested environment during desktop provisioning. CPU utilization and storage adapter performance data are gathered between 30 minute time windows for the 8 ESXi nodes in the cluster, which was the average time to complete.

Figure 7 shows the ESXi cluster CPU usage, with peaks of 20% and an average of 12% usage during the desktop pool provisioning time, indicating an excellent CPU usage performance from the infrastructure.




Figure 7: Hypervisor CPU utilization when provisioning 500 desktops

Figure 8 shows read and write IOPS from the ESXi cluster during provisioning. A high peak of read operations are seen during the replica machine creation, when the cloning of the linked clones begins, reads and writes shows parallel performance. Data collected reports:

- For writes, a peak of 15K with averages of 4.8K IOPS
- For reads, a peak of 34.2K with averages of 3.8K IOPS




Figure 8: Read and write IOPS when provisioning 500 desktops

Figure 9 shows the total throughput of the ESXi cluster during provisioning. Data collected reports average values of 98MBps writes and 146MBps reads.




Figure 9: Throughput when provisioning 500 desktops

Figure 10 shows average latency across the storage adapters of the ESXi cluster during provisioning. Data collected reports a 1.17ms peak latency and an average of 0.59ms.




Figure 10: Average latency when provisioning 500 desktops

6.2. Booting 500 desktops

Storage array performance

Storage array performance during the boot time. These numbers represent back-end IOPS and throughput. The storage array data reports:

- Total IOPS, a peak of 200K with an average of 65K IOPS.
- Total Throughput, a peak of 788MBps with an average of 166MBps.

Figure 11 shows back-end IOPS and throughput performance during the desktop pool booting.




Figure 11: Array performance when booting 500 desktops

Figure 12 shows inbound and outbound network traffic of the two data storage ports used during booting.




Figure 12: Array network traffic when booting 500 desktops

Hypervisor performance

Hypervisor performance for the tested environment during boot time. CPU utilization and storage adapter performance data are gathered between an 8 minute windows for the 8 ESXi nodes in the cluster, but the average time to complete is only 2 minutes.

Figure 13 shows the ESXi cluster CPU usage, with peaks of 90% and an average of 62% usage during the 2 minute boot time. Peak CPU usage indicates that this is an optimal architected scenario for 500 desktops concurrency.




Figure 13: Hypervisor CPU utilization when booting 500 desktops

Figure 14 shows read and write IOPS for the ESXi cluster during the boot time. Data collected reports:

- For writes, a peak of 22K with averages of 9.4K IOPS.
- For reads, a peak of 15.8K with averages of 11.7K IOPS.




Figure 14: Read and write IOPS when booting 500 desktops

Figure 15 shows the total throughput for the ESXi cluster during booting. Data collected reports average values of 60MBps writes and 448MBps reads.




Figure 15: Throughput when booting 500 desktops

Figure 16 shows the average latency across the ESXi cluster during booting operations. The data reports a 1.4ms peak latency with an average of 0.68ms.




Figure 16: Average latency when booting 500 desktops

6.3. Login VSI test

Results of Login VSI for 500 Knowledge and Power Worker workloads are described in this section, the benchmark tool is used to represent the login and steady-state of the tested environment.

In the Login VSI report, there are two values that are important to note: **VSIbase** and **VSImax**:

- **VSIbase**: A score reflecting the response time of specific operations performed in the desktop workload when there is little or no stress on the system. A low baseline indicates a better user experience, resulting in applications responding faster in the environment.
- **VSImax**: The maximum number of desktop sessions attainable on the host before experiencing degradation in host and desktop performance.

Login VSI measures the total response time of several specific user operations being performed within a desktop workload in a scripted loop. The baseline is the measurement of the response time of specific operations performed in the desktop workload, which is measured in milliseconds (ms).

6.3.1. Knowledge Worker results

Table 10 details the results generated by the Login VSI Analyzer, reported for the Knowledge Worker test.

Table 10: Knowledge Worker test overview

Testname	500users_Office 2016_knowledgeworker
VSImax v4	500 Sessions & Baseline 1087 ms
Benchmark mode	Enabled
VSI Threshold reached?	NO
VSIbaseline average response time (ms)	1087
VSImax average response time threshold (ms)	2088
VSImax threshold was reached at sessions	WAS NOT REACHED
VSI response time threshold headroom	1014
Sessions not responding	0
Corrected VSImax is	500
Total Sessions configured	500
Total Sessions successfully launched	500
Total Timeframe of test in seconds	2880
Average session launch interval in seconds	5.76
Amount of active launchers during test	20
Average session capacity per launcher	25

Figure 17 shows VSImax v4 results for a Knowledge Worker workload where the Login VSI VSImax was not reached. The active sessions ratio shows that there is still enough performance for more desktops to be added to the test environment before hitting the system saturation.




Figure 17: Knowledge Worker VSImax overview

Figure 18 shows the ESXi cluster CPU utilization during the Knowledge Worker test. Peak CPU usage when all the sessions are active at the same time reaches 71%.




Figure 18: Average CPU usage of ESXi cluster for a Knowledge Worker test

Figure 19 shows read and write IOPS during Login VSI Knowledge Worker test. Data collected reports:

- For writes, a peak of 3.4K with averages of 2K IOPS.
- For reads, a peak of 1.5K with averages of 292 IOPS.




Figure 19: Read and write IOPS for a Knowledge Worker test

Figure 20 shows read and write IOPS during Login VSI Knowledge Worker test. Data collected reports:

- Peak latency of 0.52ms and average of 0.43ms.
- Peak total IOPS of 4K and averages of 2.4K.




Figure 20: Average latency and total IOPS for a Knowledge Worker test

6.3.2. Power Worker results

Table 11 details the results generated by the Login VSI Analyzer, reported for the Power Worker test.

Table 11: Power Worker test overview

Testname	500users_Office 2016_powerworker
VSImax v4	498 Sessions & Baseline 1080 ms
Benchmark mode	Disabled
VSI Threshold reached?	NO
VSIbaseline average response time (ms)	1080
VSImax average response time threshold (ms)	2081
VSImax threshold was reached at sessions	WAS NOT REACHED
VSI response time threshold headroom	1031
Sessions not responding	2
Corrected VSImax is	498
Total Sessions configured	500
Total Sessions successfully launched	500
Total Timeframe of test in seconds	2880
Average session launch interval in seconds	5.76
Amount of active launchers during test	20
Average session capacity per launcher	25

Figure 21 shows VSImax v4 results for a Power Worker workload. Same as the Knowledge Worker test, Login VSI VSImax was not reached. The active sessions ratio shows that there is still enough performance for more desktops to be added to the test environment.




Figure 21: Power Worker VSImax overview

Figure 22 shows CPU utilization during the Power Worker test. Peak CPU usage when all the sessions are active at the same time reaches 82%.




Figure 22: Average CPU usage of ESXi cluster for a Power Worker test

Figure 23 shows read and write IOPS during Login VSI Power Worker test. Data collected reports:

- For writes, a peak of 4.4K with averages of 2.2K IOPS.
- For reads, a peak of 1.4K with averages of 289 IOPS.




Figure 23: Read and write IOPS for a Power Worker test

Figure 24 shows read and write IOPS during Login VSI Power Worker test. Data collected reports:

- Peak latency of 0.7ms and average of 0.46ms.
- Peak total IOPS of 5.5K and averages of 2.7K.




Figure 24: Average latency and total IOPS for a Power Worker test

6.3.3. Login VSI extended steady-state

The average host and storage resources for a 30 minutes steady-state are reported in this section.

Steady-state is still the most representative workload of VDI environments, even when extreme cautions must be taken during login and boot storms. By default, the Login VSI test is constructed within a 2 min steady-state window, which starts when the last session is activated. For a complete understanding of VDI during the steady-state, the Knowledge and Power Worker workloads are configured with an extended steady-state time, so the average values of the infrastructure could be analyzed and compared.

Another test was done within an 8 hours window for a Power Worker only. This extra extended steady-state test simulates the real approach to the VDI-working-day. Due to the length of the test, is not possible to gather the raw data, only the key values are reported, but similar values are found during the 30 minutes test.

Table 12 shows the key steady-state data from the ESXi infrastructure datastores of an 8 hours test. Note that this data is used as a reference only.

Table 12: 8 hours steady-state performance for a Power Worker.

8 hours Steady-state	Avg. CPU usage	Avg. Write IOPS	Avg. Read IOPS	Avg. Total IOPS	Avg Storage Latency
Power Worker	80%	2,900	387	3,400	0.2ms

Knowledge Worker: 30 minutes steady-state

Figure 25 shows the average CPU utilization during the Knowledge Worker test. Average CPU usage during the 30 minutes of the steady-state is 65%.




Figure 25: Average CPU usage of ESXi cluster for a steady-state Knowledge Worker

Figure 26 shows the average read and writes IOPS during the Knowledge Worker test. Data collected reports average writes of 2.6K and average reads of 296 IOPS.




Figure 26: Read and write IOPS for a steady-state Knowledge Worker

Figure 27 shows the average latency and total IOPS during the Knowledge Worker test. The data reports average latency of 0.48ms and an average of 3.2K IOPS. This results also show that the average IOPS consumed per a Knowledge Worker desktop is 6.4.




Figure 27: Average latency and total IOPS for steady-state Knowledge Worker

Power Worker: 30 minutes steady-state

Figure 28 shows the average CPU utilization during the Power Worker test. Average CPU usage during the 30 minutes of the steady-state is 78%.




Figure 28: Average CPU usage of ESXi cluster for a steady-state Power Worker

Figure 29 shows the average read and writes IOPS during the Power Worker test. Data collected reports average writes of 3.2K and average reads of 377 IOPS.




Figure 29: Read and write IOPS for a steady-state Power Worker

Figure 30 shows the average latency and total IOPS during the Power Worker test. Data collected reports average latency of 0.56ms and average of 3.9K IOPS. This results also show that the average IOPS consumed per a Power Worker desktop is 7.8.




Figure 30: Average latency and total IOPS for steady-state Power Worker

6.3.4. Workloads comparison

Using Login VSI Analyzer is also possible to compare different test results. Knowledge and Power Worker test are compared to see the system saturation impact of the two different tested workloads.

Figure 31 compares Knowledge and Power Worker workloads. The Login VSI analyzer reports:

- For a Knowledge Worker, a baseline of 1,080ms and a threshold of 2081.
- For a Power Worker, a baseline of 1,087 and a threshold of 2088




Figure 31: VSImax 4.1 comparison between Knowledge and Power Workers

Steady-state results of the previous section were also compared using ESXTOP raw data from the ESXi infrastructure. This data is summarized in table 13.

Table 13: 30 minutes steady-state performance comparison between Knowledge and Power Worker.

Steady-state workloads	Avg. ESXi CPU usage	Avg. Write IOPS	Avg. Read IOPS	Avg. Total IOPS	Avg Storage Latency
Knowledge	65%	2,600	296	3,200	0.48ms
Power	78%	3,200	377	3,900	0.56ms
% difference	17%	19%	21%	18%	14%

From table 12, it can be observed that there is a consumption difference of computing and storage resources from a Knowledge to a Power Worker workload. The growth difference is about 18% for ESXi cluster CPU and IOPS and 14% for storage latency (seen from the ESXi hosts storage adapters). After comparing this data with the Login VSI values of figure 31, it is concluded that even though the resources of the whole infrastructure increase for different workloads, the user experience (applications response time) will be preserved across the 500 desktops.

6.4. Power-off and pool refresh

To complete the common operations of a VDI production environment, desktop pool power-off and refresh were tested, and the most important values are reported in table 14.

Table 14: Desktop pool average values for power-off and pool refresh operations.

Pool of 500 Linked Clones	Time to complete	Peak CPU usage	Peak IOPS	Peak Throougput	Peak Latency
Power off	~7 min	4.2%	1.4K	22MBps	0.42ms
Refresh	~14 min	0.6%	527	34MBps	0.14ms

Low-performance impact is seen during the test of these two operations from the ESXi infrastructure (cluster CPU, IOPS, and latency). It is important to note that the desktop pool refresh only takes 14 minutes to complete. This operation, (that is important for Linked Clones maintenance) is recommended to run at least once or twice per month, to avoid the linked clone (snapshot) to grow as large as its attached replica. This maintenance can be easily implemented during weekends or night times without affecting the VDI-day cycle at all.

6.5. Storage capacity

Throughout the different performance tests executed for the VDI workloads, storage capacity was also evaluated. Linked clones, besides making the overall system infrastructure easy to manage, also makes the capacity easy to design. Due to reduction techniques used by NeoSapphire AFA like deduplication, even with the use of full-clone desktops, the capacity-efficiency save is proven to be over 98%.

Figure 32 shows that after booting all the virtual desktops, only 420GB consumed is seen from the NeoSapphire dashboard, this represents 4% of the total physical capacity, which does not include the thin-provisioned storage. The average capacity consumed by a single linked-clone is only 0.84GB.




Figure 32: Capacity used after booting from NeoSapphire dashboard.

7. Conclusion

Through the tests performed in this paper, we can confirm that VMware Horizon, along with the AcceleStor NeoSapphire P710 can deliver more than sufficient IOPS, while keeping sub-millisecond latency over the entire VDI-day common operations. The 500 Power Worker desktop test averages 3,900 IOPS with 0.56ms latency during the Login VSI steady-state test, this is remarkable, considering that results are observed from the hosts and the desktops themselves.

One of the key highlights of the test is found when comparing Knowledge and Power Worker results. It is noted that while computing and storage resources are slightly increased, the end-user login time and response time reported from Login VSI are preserved. Also, for both tested scenarios the Login VSI VSI_{max} was not reached, showing that there is still extra density room to fit more desktop to the tested system environment.

8. AccelStor sizing criteria for VDI

For an accurate VDI storage selection, the storage array must be able to handle all the I/O patterns that occur throughout the VDI production day and storage capacity requirements. Using the test validation and conclusions reported in this paper, booting and steady state operations are used to calculate desktop density for AccelStor NeoSapphire models.

Table 15 compares the performance resource consumption during boot and the Power Worker steady state. The key metric to evaluate is the sum of write and read IOPS. An 83% IOPS (x6 times) increase is seen during boot periods, when compared to steady state.

Table 15: Booting and steady-state performance comparison.

Pool of 500 Linked Clones	Avg. write IOPS	Avg. read IOPS	R+W IOPS	Avg. Latency
Booting	9.4K	11.7K	21.1K	0.68ms
30 minute steady-state	3.2K	377	3.6K	0.56ms
% difference	67%	97%	83%	18%

Using these key results, an average of 7.2 IOPS per desktop is observed for a Power Worker during the steady-state and 42 IOPS for booting. Also, during the steady-state test an average of 50% 4KB block size is reported by VMware vscsiStats tool. Other common block sizes are 32KB, 16KB and 65K, representing an extra average of 30% of the transfer size distribution.

Using tools such as Login VSI, allows VDI environments to generate workloads that are more representative than a single I/O size (4KB). This is optimal for testing POC environments, but large (production) environments might be difficult to replicate. Fortunately, it is possible to estimate desktop density by oversizing critical storage resources, configuring similar I/O patterns and, later use these profiles with IOMeter to run a comparable test. IOMeter is a well-known storage testing software tool that simulates application input/output to main storage on hard disk, SSD or storage network devices, but because of its use of synthetic workloads, caution must be considered.

AccelStor NeoSapphire is tested under the “worst-case scenario” that VDI environments could be stressed for the steady state. Its official benchmark report that with an I/O pattern of 4KB random writes, IOPS performance ranges from 360K to 700K for different NeoSapphire models.

Capacity results are also considered for the storage sizing. Even though less than 1GB per desktop is consumed through the test (due to the use of linked clones), capacity overheads are added based on best practice for desktop growth-ratio.

When sizing for desktop density and capacity with AccelStor NeoSapphire AFA models, the following conditions are assumed per desktop:

- Power Worker applications performance.
- Windows 10 pro version 1703 with MS office 2016.
- 2 vCPU, 3.5GB of memory and 60GB HDD.
- Average of 30 IOPS per desktop.
- 6GB per VM linked clone.

These conditions make easy to select the right NeoSapphire model based on the desktop density. When full-desktop density is used for NeoSapphire models, average storage cost won't exceed US\$30. This is an impressive value considering the use of a Power Worker and the conditions listed above.

AccelStor NeoSapphire AFA desktop density

Model	NS3405	P710	H510	H710
Desktop Density	Up to 1,000	Up to 2,000	Up to 4,000	Up to 20,000
Usable Capacity	5 TB	11 TB	22 TB	110 TB
Effective Capacity	25 TB	55 TB	110 TB	550 TB
Performance	360K IOPS	700K IOPS	360K IOPS	600K IOPS

Forget about undesired boot storms and the write-intensive random VDI IOPS, no matter the workload put on an Accelstor NeoSapphire AFA, FlexiRemap will ease performance troubles without compromising an excellent user experience, while providing the best data protection, integrity and SSD lifespan, all these benefits along the best VDI storage value per user-desktop.

Appendix

About Login VSI

Login VSI provides performance insights 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. The world's leading 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 www.loginvsi.com

Login VSI accepts no responsibility regarding this publication in any way and cannot be held accountable for any damages following from, or related to, any information contained within this publication, or any conclusions that may be drawn from it.

References

Please see the below references for more information on AccelStor and other products used in this technical report:

- AccelStor All-Flash Array VMWare ESXi 6.0 iSCSI Multipath Configuration Guide
https://www.accelstor.com/wp-content/uploads/2018/05/AccelStor-All-Flash-Array-VMWare-ESXi-6.0-iSCSI-Multipath-Configuration-Guide_v1.1.pdf
- FlexiRemap Technology
https://www.accelstor.com/category/flexiremap_technology/
- VMware vSphere
<https://docs.vmware.com/en/VMware-vSphere/index.html>
- VMware Horizon 7 Documentation
<https://docs.vmware.com/en/VMware-Horizon-7/index.html>
Login VSI Installation Guide
- <https://www.loginvsi.com/documents/documentation/Login-VSI-Installation-Guide.pdf>

Version history

Date	Version	History Description
August 2018	Version 1.0	Initial release