



Technical Report

NetApp All Flash FAS Solution for Nonpersistent Desktops with VMware Horizon View

Joe Scott and Chris Gebhardt, NetApp
May 2017 | TR-4539

Abstract

This document describes the solution components used in the 1,500-user VMware Horizon View 7.0 deployment on NetApp All Flash FAS reference architecture validation. It covers the hardware and software used in the validation, the configuration of the hardware and software, use cases that were tested, and performance results of the completed tests. It also contains sections focusing on the new platforms in the AFF product line, the NetApp AFF A-Series, and their relative performance to the previous generation. It includes workload profiles that allow customers to see the different IO profiles that make up each different test performed in this reference architecture.

TABLE OF CONTENTS

1	Executive Summary.....	6
2	Reference Architecture Objectives.....	7
3	Introduction.....	7
3.1	Document Overview.....	7
3.2	NetApp All Flash FAS	8
3.3	ONTAP FlashEssentials.....	8
3.4	NetApp ONTAP 9.....	8
3.5	Storage Efficiency	9
4	Solution Infrastructure.....	10
4.1	Hardware Infrastructure	11
4.2	Software Components	12
4.3	Login VSI	12
4.4	Hypervisor-Based VMware Linked Clone Desktops.....	13
5	Storage Design	14
5.1	Aggregate Layout.....	14
5.2	Volume Layout	15
5.3	NetApp Virtual Storage Console for VMware vSphere.....	15
6	Testing and Validation: VMware Horizon View Linked Clone Desktops.....	15
6.1	Test Results Overview	15
6.2	Test for Provisioning 1,500 VMware Horizon View Hypervisor-Based Linked Clones	16
6.3	Boot Storm Test Using VMware Horizon View Administrator.....	20
6.4	Boot Storm Test Using Horizon View During Storage Failover	24
6.5	Monday Morning Login and Steady State	27
6.6	Monday Morning Login and Steady State During Failover.....	32
6.7	Tuesday Morning Login and Steady State	37
6.8	Tuesday Morning Login and Steady State During Failover.....	42
6.9	Refresh 1,500 Desktop VMs	47
6.10	Recompose 1,500 Desktop VMs.....	52
7	Conclusion of Testing and Validation	56
8	Data Center Efficiency with AFF A-Series	56
	Appendix: Workload Profiles.....	58
	Boot.....	58

Monday Morning Login and Steady State Workload Profile	58
Tuesday Morning Login and Steady State Workload Profile	59
References.....	59
Acknowledgements	60
Version History	60

LIST OF TABLES

Table 1) Nonpersistent desktop test results.	7
Table 2) Hardware components of server categories.	11
Table 3) Solution software components.	12
Table 4) Virtual desktop configuration.	13
Table 5) Test results overview.....	16
Table 6) Results for hypervisor-based linked clone provisioning of 1,500 virtual desktops.	17
Table 7) Efficiency results for each FlexVol volume.	20
Table 8) Results for linked clone boot storm.	20
Table 9) Results for linked clone boot storm during storage failover.	24
Table 10) Results for Login VSI login and steady-state knowledge worker workload.....	27
Table 11) Results for Login VSI login and workload during storage failover.....	32
Table 12) Results for Login VSI login and workload.....	38
Table 13) Results for Login VSI login and workload.....	43
Table 14) Results for Login VSI login and workload during storage failover.....	47
Table 15) Results for a recompose operation.....	52
Table 16) Capacity density of A-Series systems.	57
Table 17) Boot workload profile.	58
Table 18) Monday morning login and steady state workload profile.	58
Table 19) Tuesday morning login and steady state workload profile.	59

LIST OF FIGURES

Figure 1) Visual representation of inline compression and data compaction.....	10
Figure 2) Advanced drive partitioning v2.	10
Figure 3) Solution infrastructure for nonpersistent desktops.	11
Figure 4) VMware OS optimization tool.....	14
Figure 5) SSD layout.	15
Figure 6) Volume layout for nonpersistent desktops.	15
Figure 7) Creating 150 VMs in one pool.	16
Figure 8) Throughput and IOPS for linked clone provisioning.	18
Figure 9) Storage controller CPU utilization for linked clone creation.....	19

Figure 10) Storage efficiency savings.	20
Figure 11) Throughput and IOPS for Horizon View boot storm.	21
Figure 12) Storage controller CPU utilization for Horizon View linked clone boot storm.	22
Figure 13) Read/write IOPS for hypervisor-based linked clone boot storm.	23
Figure 14) Read/write ratio for Horizon View linked clone boot storm.	24
Figure 15) Throughput, IOPS, and latency for linked clone boot storm during storage failover.	25
Figure 16) Storage controller CPU utilization for boot storm during storage failover.	26
Figure 17) Read/write IOPS for VMware Horizon linked clone boot storm during storage failover.	26
Figure 18) Read/write ratio for linked clone boot storm during storage failover.	27
Figure 19) VSImax results for Login VSI login and workload.	28
Figure 20) Scatterplot of login times.	28
Figure 21) Throughput, IOPS, and latency for Login VSI login and workload.	29
Figure 22) Storage controller CPU utilization for Login VSI login and workload.	30
Figure 23) Read/write IOPS for Login VSI login and workload.	31
Figure 24) Read/write ratio for Login VSI login and workload.	32
Figure 25) VSImax results for Login VSI login and workload during storage failover.	33
Figure 26) Scatterplot of login times during storage failover.	33
Figure 27) Throughput, IOPS, and latency for Login VSI login and workload during storage failover.	34
Figure 28) Storage controller CPU utilization for Login VSi login and workload during storage failover.	35
Figure 29) Read/write IOPS for Login VSI login and workload during storage failover.	36
Figure 30) Read/write ratio for Login VSI login and workload during storage failover.	37
Figure 31) VSImax results for Login VSI login and workload.	38
Figure 32) Scatterplot of Login VSI login times.	39
Figure 33) Throughput, IOPS, and latency for Login VSI login and workload.	39
Figure 34) Storage controller CPU utilization for Login VSI login and workload.	40
Figure 35) Read/write IOPS for Login VSI login and workload.	41
Figure 36) Read/write ratio for Login VSI login and workload.	42
Figure 37) VSImax results for Login VSI login and workload.	43
Figure 38) Scatterplot of Login VSI login times.	44
Figure 39) Throughput, IOPS, and latency for Login VSI login and workload.	44
Figure 40) Storage controller CPU utilization for Login VSI login and workload.	45
Figure 41) Read/write IOPS for Login and VSI login and workload.	46
Figure 42) Read/write ratio for Login VSI login and workload.	47
Figure 43) Throughput, IOPS, and latency during a refresh operation.	48
Figure 44) Storage controller CPU utilization during a refresh operation.	49
Figure 45) Read/write IOPS during a refresh operation.	50
Figure 46) Read/write ratio during a refresh operation.	51
Figure 47) Per-volume storage efficiency savings.	51
Figure 48) Throughput, IOPS, and latency during a recompose operation.	53
Figure 49) Storage controller CPU utilization during a recompose operation.	54

Figure 50) Read/write IOPS during a recompose operation.....	55
Figure 51) Read/write ratio during a recompose operation.....	55
Figure 52) Per-volume storage efficiency savings.....	56
Figure 53) User per platform density.....	57
Figure 50) Data center density for AFF8040 and A700s.....	57

1 Executive Summary

The decision to virtualize desktops affects multiple aspects of an IT organization, including infrastructure and storage requirements, application delivery, end-user devices, and technical support. In addition, correctly architecting, deploying, and managing a virtual desktop infrastructure (VDI) can be challenging because of the large number of solution components in the architecture. Therefore, it is critical to build the solution on industry-proven platforms such as NetApp® storage and industry-proven software solutions from VMware. VMware and NetApp provide leading desktop virtualization and storage solutions, respectively, for customers to successfully meet these challenges. A desktop virtualization solution provides workspace mobility, centralized management, consolidated and secure delivery of data, and device independence.

The criteria for determining the success of a virtual desktop implementation must include the end-user experience. The end-user experience must be as good as or better than any previous experience on a physical PC or virtual desktop.

As a workload, desktop virtualization is highly variable, and it includes cycles of heavy reads, heavy writes, and everything in between, along with varying block sizes for writes. Therefore, it should be no surprise that storage is often the leading culprit of unmet end-user performance expectations. The NetApp All Flash FAS (AFF) solution, easily handles highly variable desktop virtualization workloads to deliver an end-user experience that exceeds expectations.

Another criterion of project success is solution cost. Storage has often been the most expensive part of the virtual desktop solution, especially when storage efficiency and flash acceleration technologies were lacking. It was also common practice to forgo an assessment. Skipping this critical step meant that companies often overbought or undersized the storage infrastructure because information is the key to making sound architectural decisions that result in wise IT spending.

NetApp has many technologies that help customers reduce the storage cost of a virtual desktop solution. Technologies such as inline deduplication, inline compression, data compaction, advanced drive partitioning, and thin provisioning help reduce the total amount of storage required for VDI. Storage platforms that scale up and scale out with NetApp ONTAP® software deliver the right architecture to meet the customer's price and performance requirements. NetApp helps customers achieve their cost and performance goals while providing rich data management features.

With the NetApp All Flash FAS solution, customers might pay as little as US\$29 per desktop (AFF A200) for storage when deploying at scale. This figure includes the cost of NetApp hardware and software and three years of 24/7 premium support with four-hour parts replacement. A similarly low ratio of dollars per desktop can be achieved within the portfolio of All Flash FAS platforms.

With VMware and NetApp, companies can accelerate the virtual desktop end-user experience by using NetApp All Flash FAS storage for VMware Horizon View. NetApp All Flash FAS storage, powered by the AFF8000 and AFF A-Series systems, is the optimal platform for using high-performing solid-state drives (SSDs) without adding risk to desktop virtualization initiatives.

When a storage failure prevents users from working, that inactivity translates into lost revenue and productivity. That is why what used to be considered a tier 3 or 4 application is now critical to business operations. Having a storage system with a robust set of data management and availability features is key to keeping users working and lessens risk to the business. ONTAP has multiple built-in features, such as active-active high availability (HA) and nondisruptive operations, to help improve availability and to seamlessly move data in the storage cluster without affecting users.

NetApp also provides the ability to easily increase storage system capacity by simply adding disks or shelves. There is no need to purchase additional controllers to add users when additional capacity is required. When the platform requires expansion, additional nodes can be added in a scale-out fashion and managed within the same management framework and interface. Workloads can then be nondisruptively migrated or balanced to the new nodes in the cluster without users ever noticing.

2 Reference Architecture Objectives

The reference architecture described in this document is a 1,500-desktop design using NetApp AFF8040, VMware Horizon View 7.0, and VMware vSphere 6.0. The reference architecture validates a nonpersistent 1,500-desktop architecture created with VMware Horizon View 7.0 using hypervisor-based linked clones. The testing covered common administrative tasks including boot storms, login storms, and steady-state operations. These tests determined the time to complete, storage response, and storage utilization.

In all tests, end-user login time, guest response time, and maintenance activities performance were excellent. The NetApp All Flash FAS system performed well, averaging less than 50% controller utilization during most operations and use cases. All test categories demonstrated that, based on the 1,500-user workload, the AFF8040 system is capable of supporting 2,666 users while still being able to fail over in the event of a failure. At a density of 2,666 VMs on an AFF8040 system with the same I/O profile, storage for VDI can be as low as US\$39 per desktop. This figure includes the cost of hardware, software, and three years of 24/7 premium support with four-hour parts replacement. A similarly low ratio of dollars per desktop can be achieved within the broader portfolio of AFF platforms. Table 1 lists the results of testing for 1,500 nonpersistent desktops.

Table 1) Nonpersistent desktop test results.

Test	Time to Complete	Peak IOPS	Peak Throughput	Average Storage Latency
Provision 1,500-VM pool	~80 min	86,621	2038MBps	457us
Boot storm test (50 VMware Horizon View concurrent power-on operations)	~14 min	134,722	3103MBps	373us
Boot storm test during storage failover (50 VMware Horizon View concurrent power-on operations)	~13 min	109,975	2376MBps	7.83ms
Login VSI Monday morning login and workload	19.21 sec/VM	43,528	808MBps	465us
Login VSI Monday morning login and workload during failover	19.93 sec/VM	69,700	1126MBps	802us
Login VSI Tuesday morning login and workload	19.37 sec/VM	27,159	620MBps	450us
Login VSI Tuesday morning login and workload during failover	19.65 sec/VM	44,250	1106MBps	707us
Refresh existing VM pool	~25 min	111,789	2079MBps	723μs
Recompose VM existing VM pool	~149 min	98,404	2222MBps	465us

3 Introduction

This section provides an overview of the NetApp All Flash FAS solution for VMware Horizon View 7.0, explains the purpose of this document, and introduces Login VSI.

3.1 Document Overview

This document describes the solution components used in the 1,500-user VMware Horizon View 7.0 deployment on the NetApp All Flash FAS reference architecture validation. The document covers the

hardware and software used in the validation, the configuration of the hardware and software, use cases that were tested, and performance results of the completed tests.

The testing included the following criteria:

- Launch and boot storm tests of 1,500 desktops (with and without storage node failover) using VMware Horizon View linked clones
- Login VSI login and steady-state workload tests using Login VSI 4.1.4.2 (with and without storage node failover)

3.2 NetApp All Flash FAS

NetApp ONTAP has evolved to meet the changing needs of customers and help drive their success. ONTAP provides a rich set of data management features and clustering for scale-out, operational efficiency, storage efficiency, and nondisruptive operations to offer customers one of the most compelling value propositions in the industry. The IT landscape is undergoing a fundamental shift to IT as a service. This model requires a pool of computing, network, and storage resources to serve a range of applications and deliver a range of services. Innovations such as ONTAP are fueling this revolution.

The 1,500 desktops were hosted on a 2-node NetApp All Flash FAS active-active storage system running NetApp ONTAP 9 configured with 24 800GB SSDs.

3.3 ONTAP FlashEssentials

NetApp ONTAP FlashEssentials is the power behind the performance and efficiency of All Flash FAS. ONTAP is a well-known operating system, but it is not widely known that ONTAP, with its WAFL® (Write Anywhere File Layout) file system, is natively optimized for flash media.

ONTAP FlashEssentials encapsulates key features that optimize solid-state-drive (SSD) performance and endurance, including the following:

- NetApp data-reduction technologies, including inline compression, inline deduplication, and inline data compaction, can provide significant space savings. Savings can be further increased by using NetApp Snapshot® and NetApp FlexClone® technologies. Studies based on customer deployments have shown that total data-reduction technologies have enabled up to 933-times space savings.
- Inline data compaction provides continued innovation beyond compression and deduplication, further increasing storage efficiency.
- Coalesced writes to free blocks maximize performance and flash media longevity.
- Flash-specific read path optimizations enable consistently low latency.
- Parallelized processing handles more requests at once.
- Software-defined access to flash maximizes deployment flexibility.
- New advanced drive partitioning (ADPv2) increases storage efficiency and further increases usable capacity by almost 20%.
- Data fabric readiness enables live workload migration between flash and hard disk drive (HDD) tiers on the premises or to the cloud.
- Quality of service capability safeguards service-level objectives in multiworkload and multitenant environments.

3.4 NetApp ONTAP 9

NetApp ONTAP 9 is a major advance in the industry's leading enterprise data management software. ONTAP 9 combines new levels of simplicity and flexibility with powerful capabilities and efficiencies. Customers can integrate the best of next-generation and traditional technologies, incorporating flash, cloud, and software-defined architectures while building a foundation for their data fabric. Plus, new

customers and existing Data ONTAP 8.3 environments can quickly and easily use the rich data services delivered by ONTAP 9.

An essential feature for VDI deployed on shared enterprise storage is the ability to deliver consistent and dependable high performance. High performance must be coupled with nondisruptive operations, high availability, scalability, and storage efficiency. Customers can depend on ONTAP 9 and All Flash FAS to provide these essential elements.

Built on ONTAP unified scale-out architecture, All Flash FAS consistently meets or exceeds the high-performance demands of VDI. All Flash FAS also provides rich data management capabilities, such as Integrated Data Protection and nondisruptive upgrades and data migration. These features allow customers to eliminate performance silos and seamlessly integrate All Flash FAS into a shared infrastructure. ONTAP delivers enhanced inline deduplication and a completely new inline data-compaction capability that significantly reduces the amount of flash storage required, with no effect on system performance. It also provides industry-leading ecosystem integration with database applications that makes administration of databases and storage systems far more efficient than other flash storage solutions on the market.

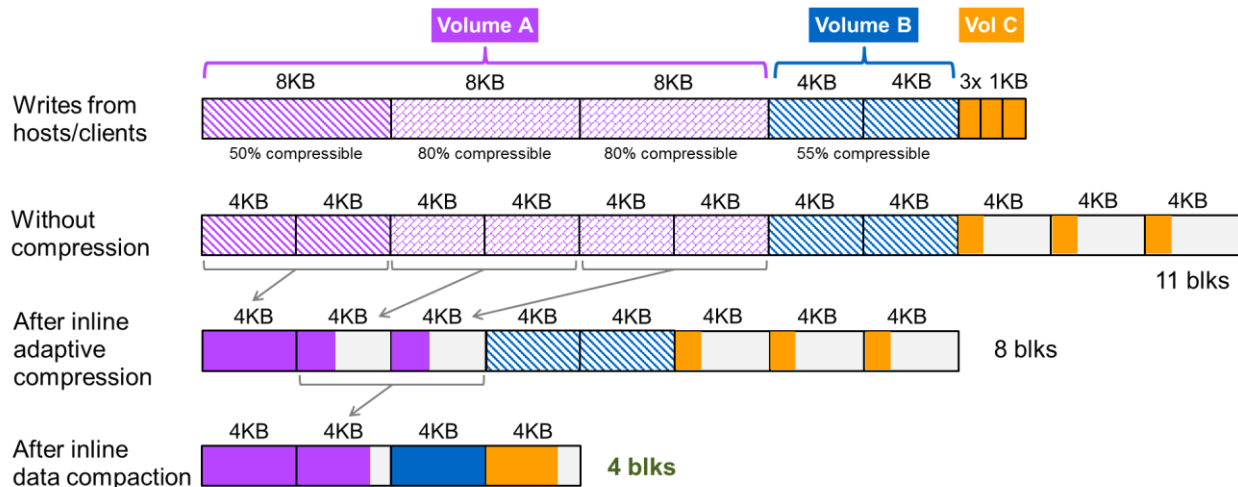
NetApp is a global enterprise scale-out storage and data management fabric provider, and ONTAP has been an industry-leading operating system since 2012. On-site ready but cloud connected, ONTAP is a complete solution that is future-proof in a rapidly changing technology environment.

3.5 Storage Efficiency

Simply stated, storage efficiency enables you to store the maximum amount of data within the smallest possible space at the lowest possible cost. The following NetApp storage efficiency technologies can help you realize maximum space savings:

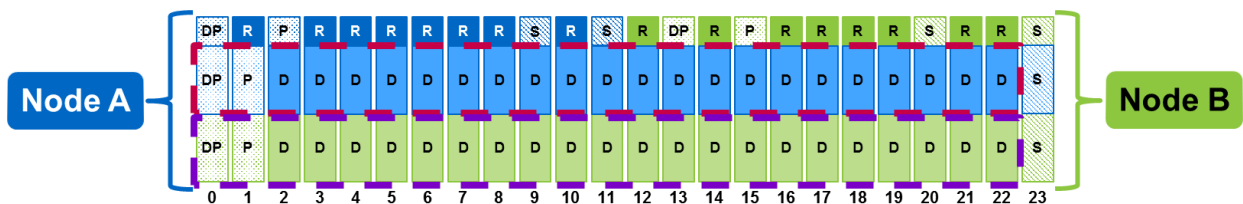
- **Inline compression.** Data compression reduces the disk space required, regardless of storage protocol, application, or storage tier. Inline compression also reduces the data that must be moved to SSDs, thereby reducing the wear on SSDs. Furthermore, the reduced amount of data being written can deliver an overall performance increase.
- **Inline deduplication.** Data deduplication cuts storage requirements by reducing redundancies in primary, backup, and archival data. Inline deduplication of zeros speeds up VM provisioning by 20% to 30%. Further improvements to inline deduplication in ONTAP 9 provide additional efficiency by extending elimination of duplicate data to blocks in memory and SSDs.
- **Inline data compaction.** New in ONTAP 9, NetApp inline data compaction provides significant storage savings by compressing and coalescing small I/O together into single-block writes. This change further reduces the disk space required and the associated wear on SSDs. Figure 1 demonstrates how compressed data that is smaller than 4K can be stored in a block with other data smaller than 4K.

Figure 1) Visual representation of inline compression and data compaction.



- **Snapshot technology.** NetApp Snapshot technology provides low-cost, instantaneous, point-in-time, space-efficient copies of the file system (volume) or LUN by preserving ONTAP architecture and WAFL consistency points without affecting performance. NetApp SnapManager® software for Oracle automates and simplifies Oracle database management with backup, recovery, restore, and cloning features with no downtime.
- **Thin provisioning.** Thin provisioning, implemented by NetApp at the NetApp FlexVol® volume and LUN level, defers storage purchases by keeping a common pool of free storage available to all applications.
- **Thin replication.** Thin replication is at the center of the NetApp data protection software portfolio, which includes NetApp SnapMirror® and NetApp SnapVault® software. SnapVault thin replication enables more frequent backups that use less storage capacity because no redundant data is moved or stored. SnapMirror thin replication protects business-critical data while minimizing storage capacity requirements.
- **RAID DP.** NetApp RAID DP® technology protects against double disk failure without sacrificing performance or adding disk-mirroring overhead.
- **FlexClone volumes.** FlexClone virtual cloning reduces the need for storage by enabling multiple, instant, space-efficient writable copies.
- **Advanced drive partitioning (v2).** As is shown in Figure 2, advanced SSD partitioning with the latest ONTAP 9 release further increases usable capacity by almost 20%.

Figure 2) Advanced drive partitioning v2.

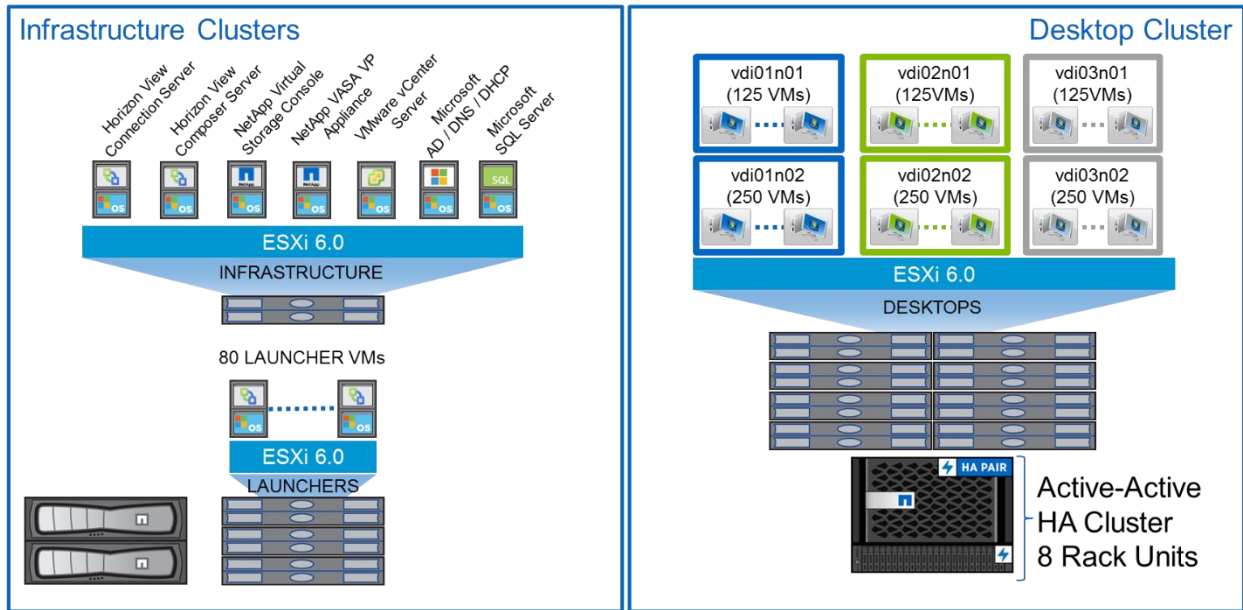


4 Solution Infrastructure

This section describes the software and hardware components of the solution. Figure 3 shows the solution infrastructure, which includes one Horizon View Connection Server and one Horizon View Composer Server.

Figure 3 shows the solution infrastructure for nonpersistent desktops.

Figure 3) Solution infrastructure for nonpersistent desktops.



4.1 Hardware Infrastructure

During solution testing, 22 servers were used to host the infrastructure and the desktop VMs. The desktops and infrastructure servers were hosted on discrete resources so that the workload to the NetApp All Flash FAS system could be measured precisely. It is both a NetApp and an industry best practice to separate the desktop VMs from the infrastructure VMs. That is because noisy neighbors or bully virtual desktops can affect the infrastructure and have a negative impact on all users, applications, and performance results.

Options to neutralize this problem include leveraging intelligent quality-of-service policies in ONTAP to eliminate noisy neighbor behavior and using intelligent sizing to account for infrastructure VMs. Options also include putting infrastructure VMs on an existing or separate NetApp FAS storage system. For this lab validation, we used a separate NetApp FAS storage system (not shown) to host the infrastructure and the launcher. In the real world, infrastructure VMs can be on the same All Flash FAS system that is hosting the virtual desktops. Table 2 lists the hardware specifications of each server category.

Table 2) Hardware components of server categories.

Hardware Components	Configuration
Infrastructure and Launcher Servers	
Server quantity	8 Fujitsu Primergy RX300 S8
CPU model	Intel Xeon CPU E5-2630 v2 at 2.60GHz (6-core)
Total number of cores per server	12 cores (24 with hyperthreading)
Memory per server	64GB–256GB
Desktop Servers	
Server quantity	14 Fujitsu Primergy RX2540 M1

Hardware Components	Configuration
CPU model	Intel Xeon CPU E5-2670 v2 at 2.30GHz (12-core)
Total number of cores per server	24 cores (48 with hyperthreading)
Memory per server	512GB
Storage	
NetApp controller	AFF8040A
Disk shelf	1 DS2246
Disk drives	24 800GB SSDs

4.2 Software Components

This section describes the purpose of each software product used to test the NetApp All Flash FAS system and provides configuration details. Table 3 lists the software components and identifies the version of each component.

Table 3) Solution software components.

Software	Version
NetApp FAS	
NetApp ONTAP	9.0 RC
NetApp VSC for VMware	6.2
VMware Software	
VMware ESXi	6.0.0, 2494585
VMware vCenter Server	6.0.0
VMware Horizon View Administrator	7.0.0, 3633490
VMware Horizon View Client	4.0.0, 3677815
VMware Horizon View Agent	7.0.0, 3634043
VMware vSphere PowerCLI	6.3, 3737840
Workload Generation Utility	
Login VSI Professional	Login VSI 4.1.4.2
Database Server	
Microsoft SQL Server	11.12.1 build 14008
Microsoft SQL Server Management Studio	2.6.0.5031

4.3 Login VSI



Login Virtual Session Indexer (Login VSI) is the industry-standard load-testing tool for testing 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 major hardware and software vendors and is recommended by both leading IT analysts and the technical community. Login VSI is vendor independent and works with standardized user workloads; therefore, conclusions based on Login VSI test data are objective, verifiable, and replicable.

SBC-oriented and VDI-oriented vendor organizations that are committed to enhancing the end-user experience in the most efficient way use Login VSI as an objective method of testing, benchmarking, and improving the performance and scalability of their solutions. VSImax provides proof (vendor independent, industry standard, and easy to understand) to innovative technology vendors of the power, the scalability, and the gains of their solutions.

Login VSI-based test results are published in [technical white papers](#) and are presented at conferences. Login VSI is used by end-user organizations, system integrators, hosting providers, and testing companies. It is also the standard tool used in all tests executed in the internationally acclaimed Project Virtual Reality Check.

For more information about Login VSI or for a free test license, refer to the [Login VSI](#) website.

4.4 Hypervisor-Based VMware Linked Clone Desktops

The Windows 10 base image was created with the virtual hardware and software listed in Table 4. The VM hardware and software were installed and configured according to [Login VSI documentation](#). View Composer cloned a base VM to each volume and then created linked clones from the base templates.

Table 4) Virtual desktop configuration.

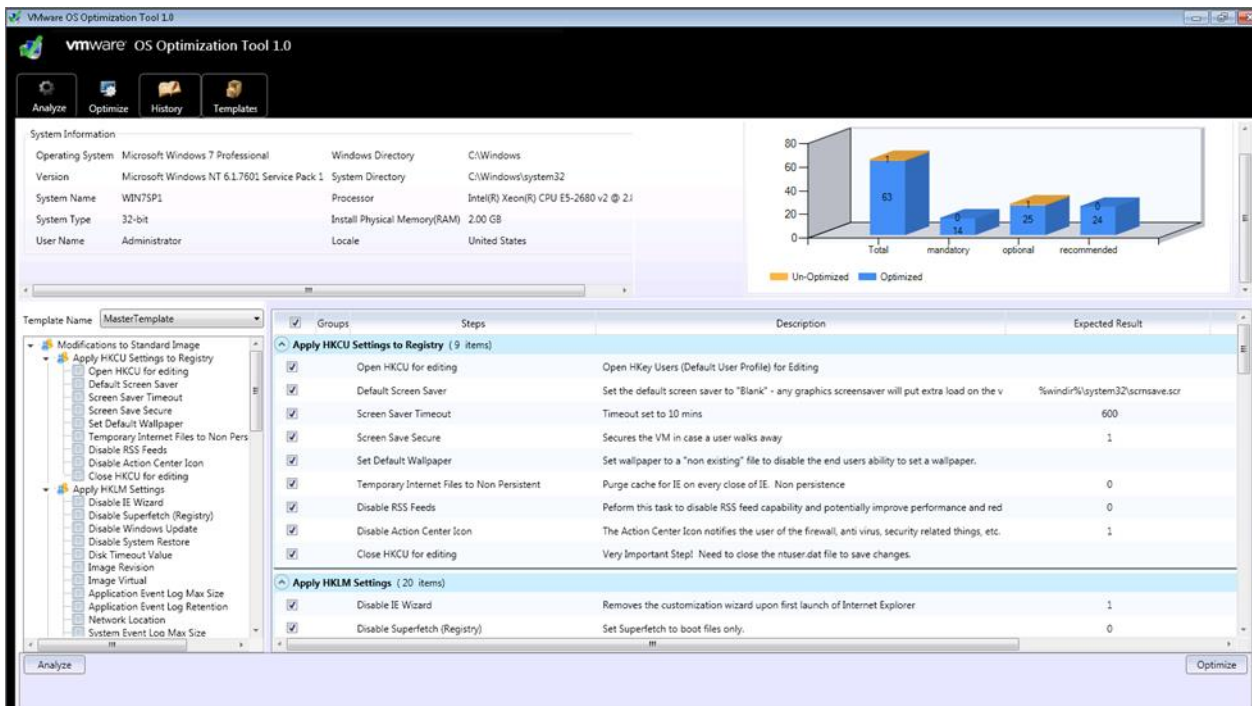
Desktop	Configuration
Desktop VM	
VM quantity	1,500
VM hardware version	11
vCPU	2 vCPUs
Memory	2GB
Hard disk size	32GB
Hard disk type	Thin provisioned
Desktop Software	
Guest OS	Windows 10 (64-bit)
VM hardware version	ESXi 6.0 and later (VM version 11)
VMware tools version	9.10.0.2476743
Microsoft Office	2016 version 16.0.4266.1001
Adobe Acrobat Reader	11.0.00
Adobe Flash Player	11.5.502.146
Java	7.0.130
Doro PDF	1.82
VMware Horizon View Agent	7.0.0, 3634043

Desktop	Configuration
Login VSI target software	4.1.4.2

Guest Optimization

Guest OS optimizations were applied to the template VMs used in this reference architecture. Figure 4 shows the VMware OS optimization tool that was used to perform the guest optimizations.

Figure 4) VMware OS optimization tool.



Although it might be possible to run desktops without guest optimizations, the impact of not optimizing must first be understood. Many recommended optimizations address services and features (such as hibernation, Windows updates, and system restores) that do not provide value in a virtual desktop environment. Running services and features that do not add value would decrease the overall density of the solution and increase cost. That is because these services and features would consume CPU, memory, and storage resources in relation to both capacity and I/O.

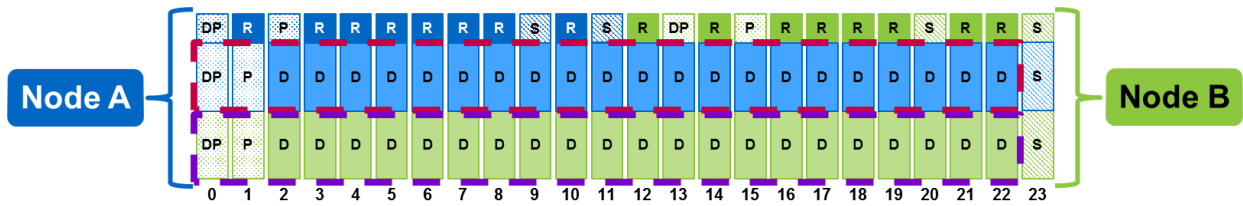
5 Storage Design

This section provides an overview of the storage design, aggregate layout, volume layout, and Virtual Storage Console (VSC).

5.1 Aggregate Layout

In this reference architecture, we used advanced drive partitioning (ADPv2) to partition the 24 800GB SSDs across the 2 nodes of an AFF8040 controller. This architecture is shown in Figure 5.

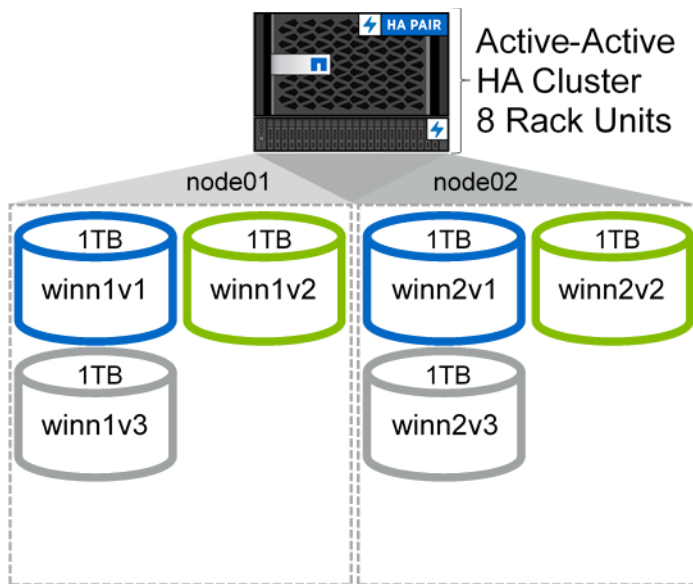
Figure 5) SSD layout.



5.2 Volume Layout

To adhere to NetApp best practices, all volumes were provisioned with the NetApp VSC. Figure 6 shows how the volumes were arranged.

Figure 6) Volume layout for nonpersistent desktops.



Note: A root volume for the SVM is present but not depicted.

5.3 NetApp Virtual Storage Console for VMware vSphere

NetApp VSC was used to provision the datastores in this reference architecture. NetApp VSC for VMware vSphere provides integrated, comprehensive, end-to-end virtual storage management for your VMware infrastructure, including discovery, health monitoring, capacity management, provisioning, cloning, backup, restore, and disaster recovery.

6 Testing and Validation: VMware Horizon View Linked Clone Desktops

This section describes the testing and validation of VMware Horizon View linked clone desktops.

6.1 Test Results Overview

Table 5 lists the high-level results that were achieved during the reference architecture testing.

Table 5) Test results overview.

Test	Time to Complete	Peak IOPS	Peak Throughput	Average Storage Latency
Provision 1,500-VM pool	~ 80 min	86,621	2038MBps	457 us
Boot storm test (50 VMware Horizon View concurrent power-on operations)	~14 min	134,722	3103MBps	373 us
Boot storm test during storage failover (50 VMware Horizon View concurrent power-on operations)	~13 min	109,975	2376MBps	7.83 ms
Login VSI Monday morning login and workload	19.21 sec/VM	43,528	808MBps	465 us
Login VSI Monday morning login and workload during failover	19.93 sec/VM	69,700	1126MBps	802 us
Login VSI Tuesday morning login and workload	19.37 sec/VM	27,159	620MBps	450 us
Login VSI Tuesday morning login and workload during failover	19.65 sec/VM	44,250	1106MBps	707 us
Refresh existing VM pool	~25 min	111,789	2079MBps	723 µs
Recompose VM existing VM pool	~149 min	98,404	2222MBps	465 us

Note: CPU and latency measurements are based on the average across both nodes of the cluster. IOPS and throughput are based on a combined total of each.

6.2 Test for Provisioning 1,500 VMware Horizon View Hypervisor-Based Linked Clones

This section describes test objectives and methodology and provides results from testing the provisioning of 1,500 VMware Horizon View hypervisor-based linked clones.

Test Objectives and Methodology

The objective of this test was to determine how long it would take to provision 1,500 VMware Horizon View virtual desktops.

To set up for the tests, 1,500 VMware Horizon View hypervisor-based linked clones were created using a Windows PowerShell script for simplicity and repeatability. Figure 7 shows one line of the script completely filled out to demonstrate what was done for one pool of 150 VMs. The script shown in Figure 6 contains the entire script that was used to create the pools.

Figure 7) Creating 150 VMs in one pool.

```
$vcserver = "vcsa6.vdi.rtp.netapp.com"
$domain = "vdi.rtp.netapp.com"
$username = "administrator"
$sleep = "300"

function CreateNonpersistentPool ($poolid, $dataStorePaths, $numvms, $templatePath)
{
    $persistence = "Persistent"
    $customizationSpec = "win10vdi"
    $overcommit = "Aggressive"
    $parentVMPPath = "/VDI/vm/win10linked"
```

```

$parentSnapshotPath = "/test2"
$vmFolderPath = "/VDI/vm"
$resourcePoolPath = "/VDI/host/Desktops/Resources"

Write-Host "Creating $numvms desktops named" $poolid "in datastores " $dataStorePaths
Get-ViewVC | Get-ComposerDomain -domain $domain -Username $username | Add-
AutomaticLinkedClonePool -Pool_id $poolid -displayName $poolid -namePrefix $poolid-{n:fixed=3}"
-ParentVMPath $parentVMPath -ParentSnapshotPath $parentSnapshotPath -vmFolderPath $vmFolderPath
-resourcePoolPath $resourcePoolPath -dataStoreSpecs "[Unbounded,OS,data]$(($dataStorePaths))" -
HeadroomCount $numvms -minimumCount $numvms -maximumCount $numvms -UseSeSparseDiskFormat $true -
SeSparseThreshold 0 -OrganizationalUnit "OU=Computers,OU=LoginVSI" -UseTempDisk $false -
PowerPolicy "AlwaysOn" -SuspendProvisioningOnError $false -CustomizationSpecName
$customizationSpec -Persistence $persistence
}

#Create pools below
CreateNonpersistentPool -poolid "pooln1v1" -dataStorePaths "/VDI/host/Desktops/win10n1v1" -numvms
260 -templatePath win10n1v1
CreateNonpersistentPool -poolid "pooln2v1" -dataStorePaths "/VDI/host/Desktops/win10n2v1" -numvms
260 -templatePath win10n2v1
CreateNonpersistentPool -poolid "pooln1v2" -dataStorePaths "/VDI/host/Desktops/win10n1v2" -numvms
260 -templatePath win10n1v2
CreateNonpersistentPool -poolid "pooln2v2" -dataStorePaths "/VDI/host/Desktops/win10n2v2" -numvms
260 -templatePath win10n2v2
CreateNonpersistentPool -poolid "pooln1v3" -dataStorePaths "/VDI/host/Desktops/win10n1v3" -numvms
260 -templatePath win10n1v3
CreateNonpersistentPool -poolid "pooln2v3" -dataStorePaths "/VDI/host/Desktops/win10n2v3" -numvms
260 -templatePath win10n2v3

sleep $sleep

#Entitle pools below
Get-Pool | Add-PoolEntitlement -sid (Get-User -IncludeUser $false -IncludeGroup $true -Name
"LoginVSI").sid

```

For this testing, we chose specific pool and provisioning settings that would stress the storage while providing the most granular reporting capabilities. NetApp does not advocate using or disabling these features because each might provide significant value in the correct use case. NetApp recommends that customers test these features to understand their impacts before deploying with these features enabled. These features include, but are not limited to, persona management, replica tiering, user data disks, and disposable file disks. Table 6 lists the provisioning data that was gathered.

Table 6) Results for hypervisor-based linked clone provisioning of 1,500 virtual desktops.

Time	Average Latency	Peak IOPS	Avg. IOPS	Peak Throughput	Avg. Throughput	Peak Storage CPU	Avg. Storage CPU
~145 min	457 us	86,621	41,495	2038MBps	764MBps	100%	46%

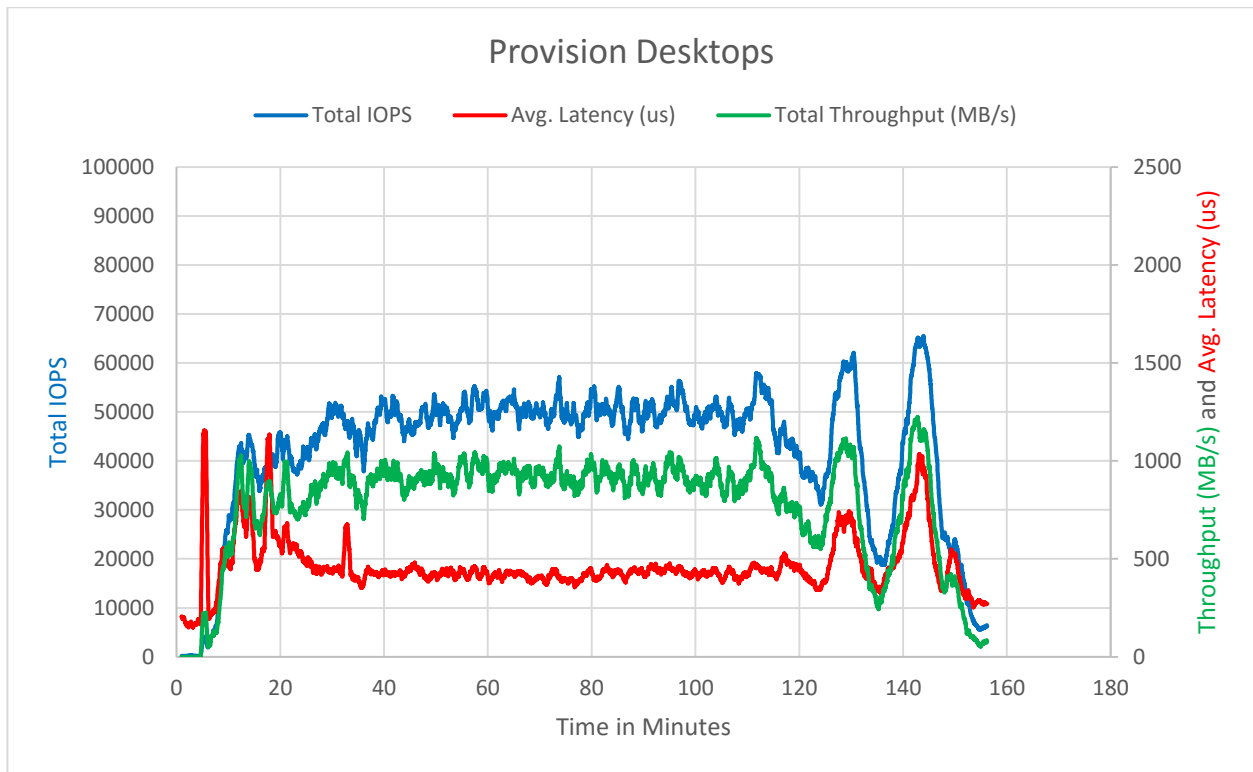
Note: All desktops had the status `Available` in VMware Horizon View.

Note: CPU and latency measurements are based on the average across both nodes of the cluster. IOPS and throughput are based on a combined total of each.

Throughput and IOPS

During the provisioning test, the storage controllers had a combined peak of 81,002 IOPS, 2693MBps throughput, and an average of 46% utilization per storage controller with an average latency of 457 us. Figure 8 shows the throughput and IOPS for linked clone creation.

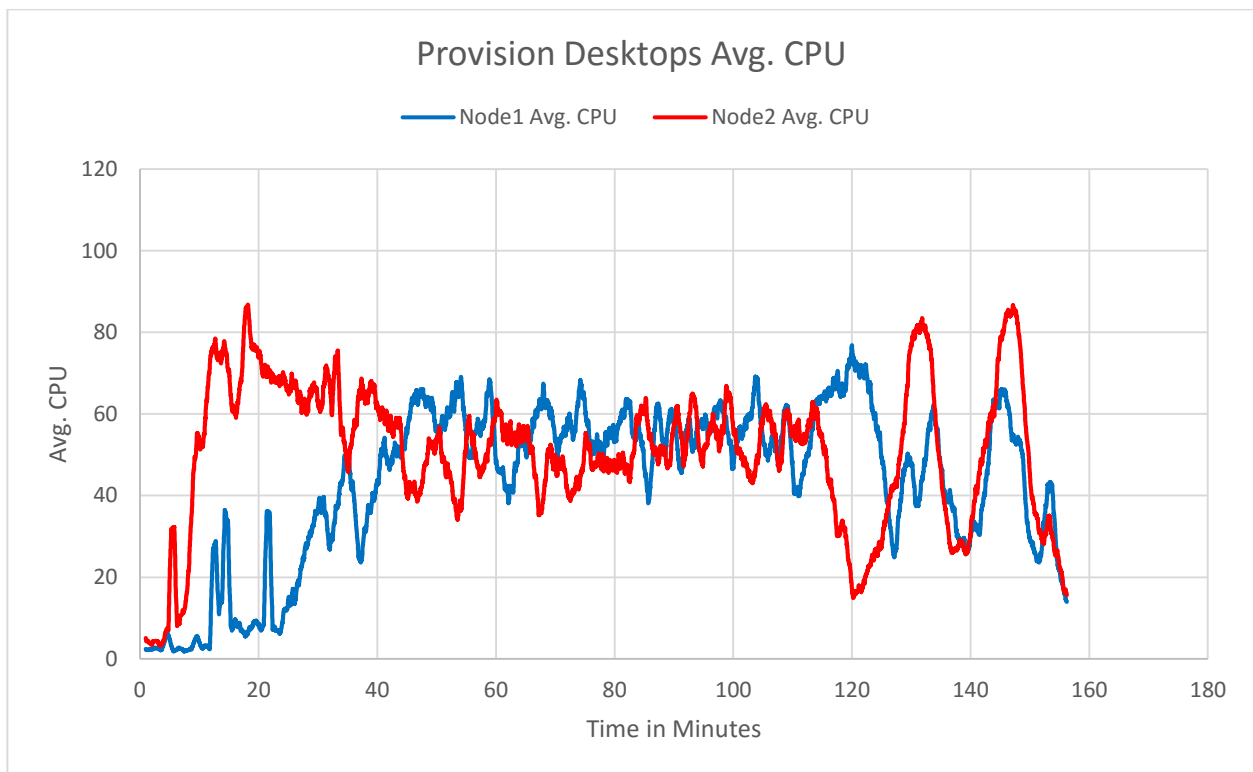
Figure 8) Throughput and IOPS for linked clone provisioning.



Storage Controller CPU Utilization

Figure 9 shows the storage controller CPU utilization across both nodes of the two-node NetApp cluster. The utilization average was 46% with a peak of 100%.

Figure 9) Storage controller CPU utilization for linked clone creation.



Customer Impact (Test Conclusions)

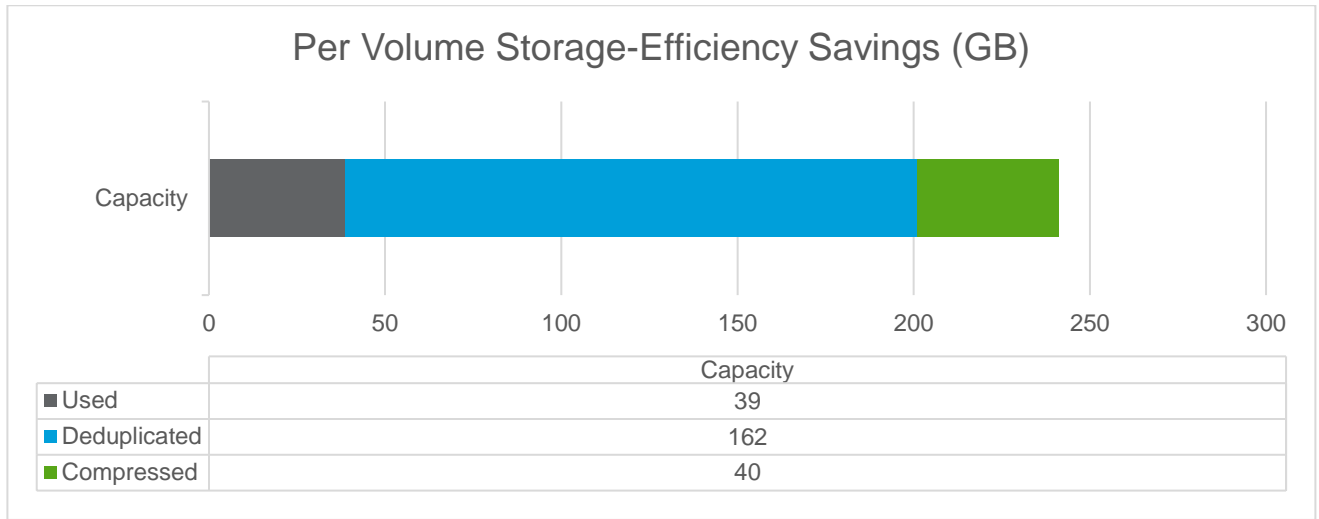
During the provisioning of 1,500 persistent desktops, the storage controller had enough performance capacity to perform a significantly greater number of concurrent provisioning operations. On average, the NetApp All Flash FAS system and systems from other all-flash vendors provision at the rate of approximately 5.8 seconds per desktop.

This test leverages hypervisor-based cloning, meaning all deduplication savings reported are from inline efficiencies only. It should be noted that the rate at which the storage acknowledges clone requests can exhaust the default throttling limits employed by the Horizon View provisioning process. So that provisioning completes in a smooth and successful manner, adjustments to the default throttling setting might be necessary. In our testing, the maximum concurrent provisioning and power-on settings were set to 10 simultaneous operations.

Storage Efficiency

During the tests, all space efficiency technologies were turned on. These technologies included inline compression, inline dedupe, inline zero elimination, and data compaction. On average after Horizon View VM creation and boot, a 6.2:1 efficiency ratio, or 83% storage efficiency, was observed at the FlexVol level. At the aggregate level, which accounts for physical blocks stored, we saw another 18% savings from compaction in addition to the 83% savings seen at the FlexVol level. Figure 10 shows the significant difference in storage efficiency savings.

Figure 10) Storage efficiency savings.



Because of the synthetic nature of the data used to perform these tests, these results are not typical of real-world savings. In addition, although thin provisioning was used for each volume, thin provisioning is not an inline storage-reduction technology and therefore was not reported on. Table 7 lists the efficiency results from the testing.

Table 7) Efficiency results for each FlexVol volume.

Capacity Used	Total Savings	Dedupe Savings	Compression Savings
39GB	203GB	162GB	40GB

6.3 Boot Storm Test Using VMware Horizon View Administrator

This section describes test objectives and methodology and provides results from boot storm testing.

Test Objectives and Methodology

The objective of this test was to determine how long it would take to boot 1,500 virtual desktops with Horizon View. This situation could occur, for example, after maintenance activities and server host failures.

This test was performed by enabling all desktop pools from within Horizon View and observing when the status of all VMs in VMware Horizon View changed to *Available*. Table 8 lists the boot storm data that was gathered.

Table 8) Results for linked clone boot storm.

Time	Average Latency	Peak IOPS	Avg. IOPS	Peak Throughput	Avg. Throughput	Peak Storage CPU	Avg. Storage CPU
~14 min	373 us	134,722	43,471	3,103	872MBps	99%	38%

Note: All desktops had the status *Available* in VMware Horizon View.

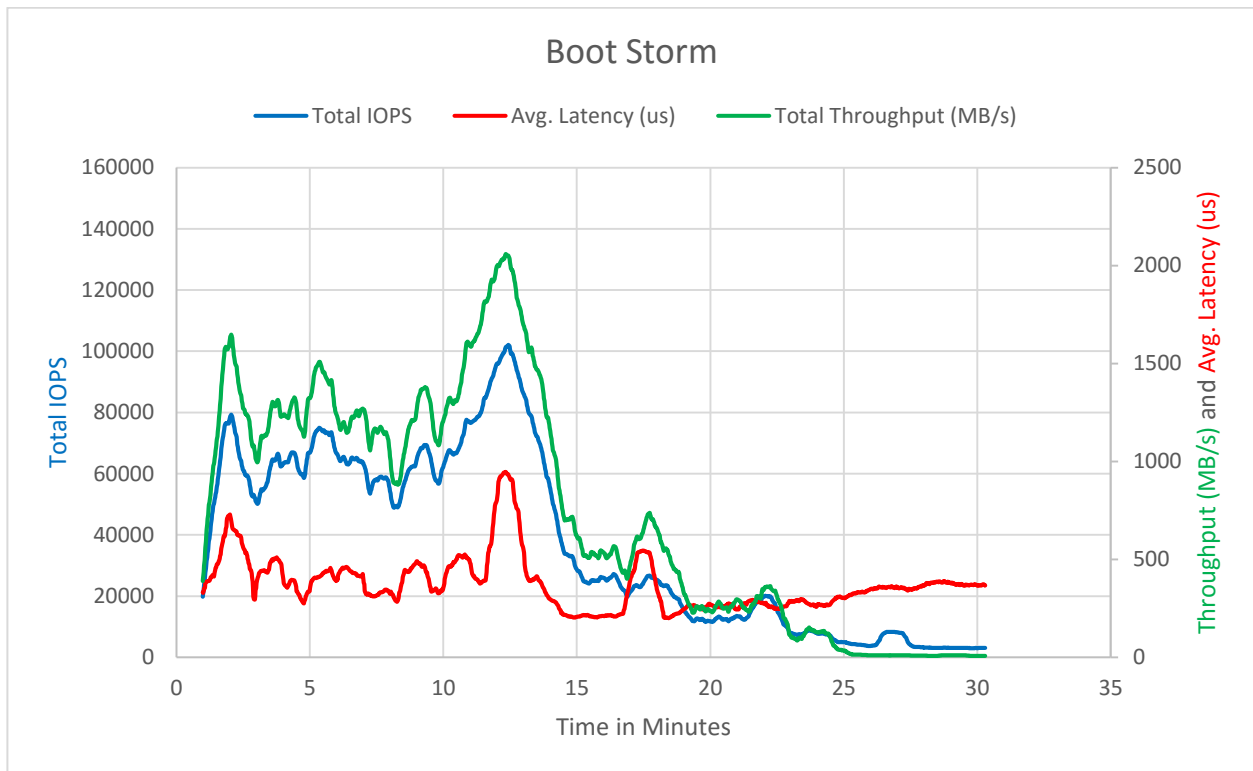
Note: The time taken to boot the VMs was a result of Horizon View registering the virtual machines.

Note: CPU and latency measurements are based on the average across both nodes of the cluster. IOPS and throughput are based on a combined total of each.

Throughput, IOPS, and Latency

Figure 11 shows testing results for throughput and IOPS for a Horizon View boot storm.

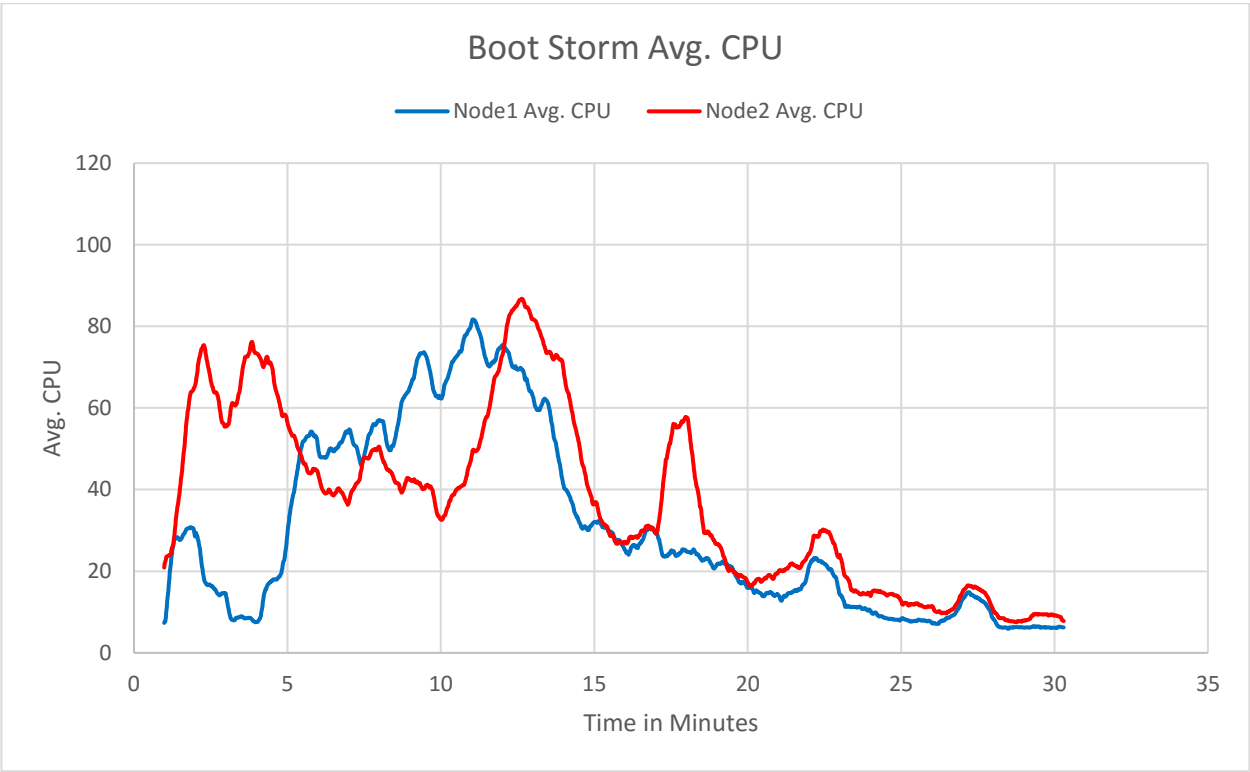
Figure 11) Throughput and IOPS for Horizon View boot storm.



Storage Controller CPU Utilization

Figure 12 shows testing results for storage controller CPU utilization for a Horizon View linked clone boot storm.

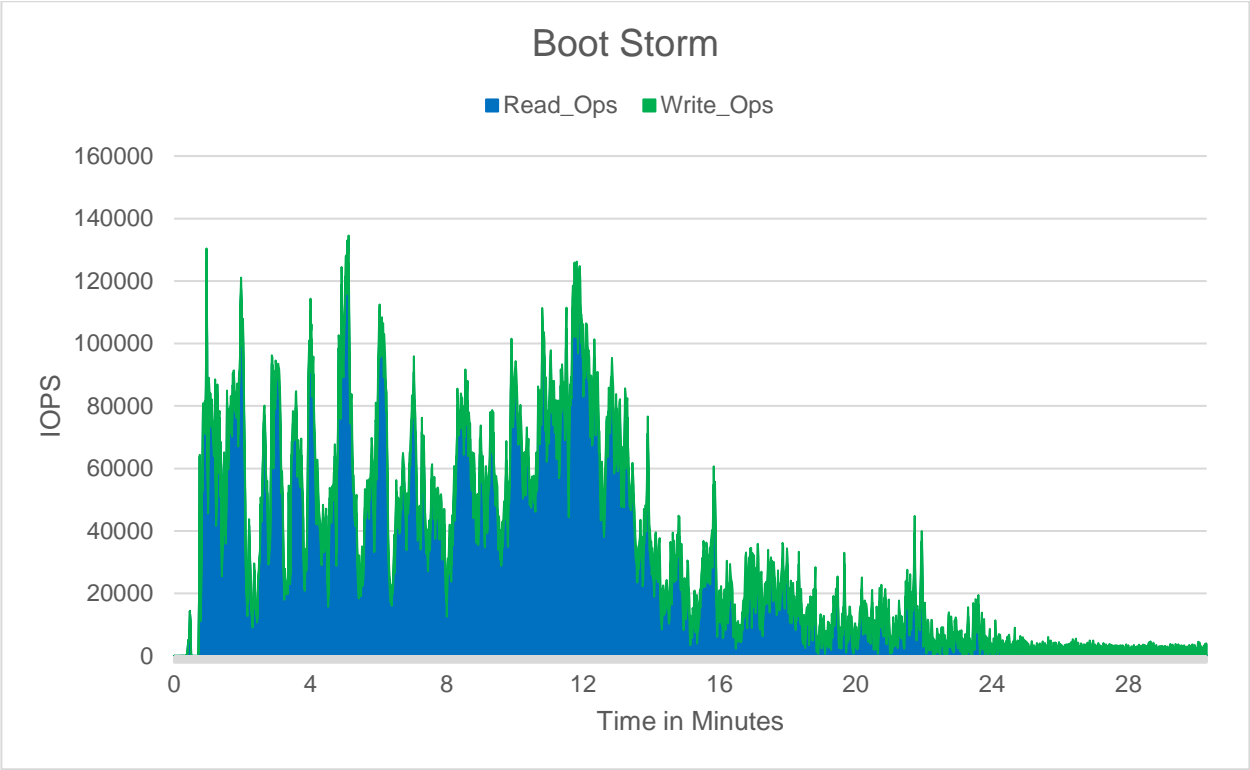
Figure 12) Storage controller CPU utilization for Horizon View linked clone boot storm.



Read/Write IOPS

Figure 13 shows testing results for read and write IOPS for a hypervisor-based linked clone boot storm.

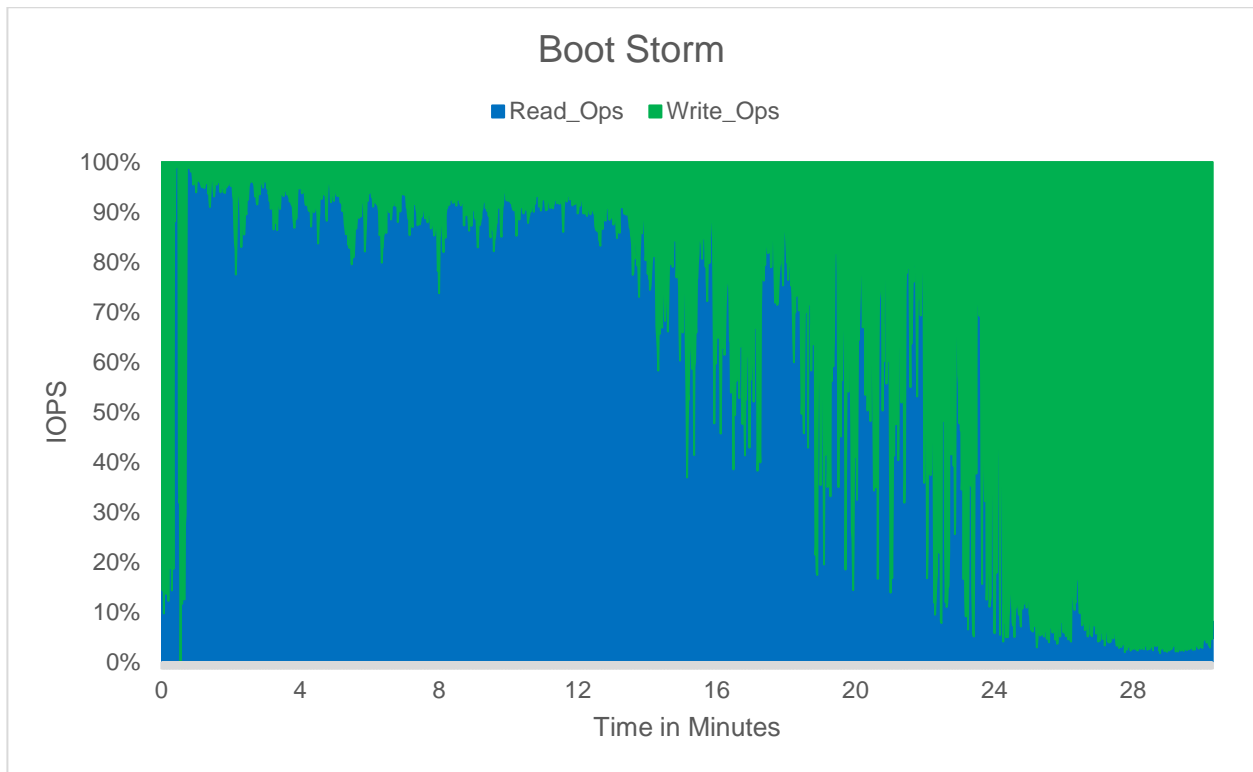
Figure 13) Read/write IOPS for hypervisor-based linked clone boot storm.



Read/Write Ratio

Figure 14 shows the read and write ratio for a Horizon View linked clone boot storm.

Figure 14) Read/write ratio for Horizon View linked clone boot storm.



Customer Impact (Test Conclusions)

Although the total time it took to boot the VMs was ~14 minutes, storage had plenty of performance left, as seen by the storage CPU utilization used on the controllers.

6.4 Boot Storm Test Using Horizon View During Storage Failover

This section describes test objectives and methodology and provides results from boot storm testing during storage controller failover.

Test Objectives and Methodology

The objective of this test was to determine how long it would take to boot 1,500 virtual desktops if the storage controller had a problem and was failed over. This test used the same methodologies and process that were used in section 6.3, "Boot Storm Test Using VMware Horizon View Administrator." Table 9 shows the data that was gathered for the boot storm during storage failover.

Table 9) Results for linked clone boot storm during storage failover.

Time	Average Latency	Peak IOPS	Avg. IOPS	Peak Throughput	Avg. Throughput	Peak Storage CPU	Avg. Storage CPU
~13 min	7.83 ms	109,975	70,795	2,376	1,313	100%	77%

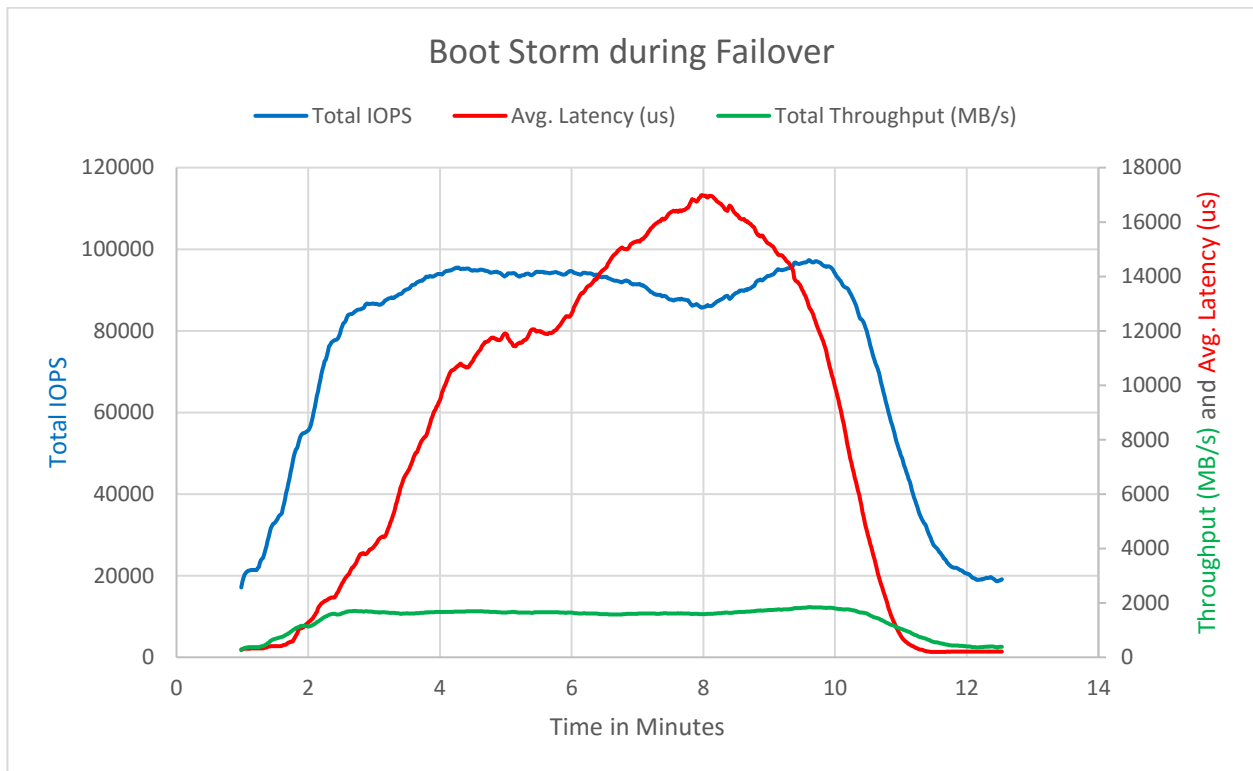
Note: All desktops had the status of `Available` in VMware Horizon View.

Note: CPU and latency measurements are based on the average across both nodes of the cluster. IOPS and throughput are based on a combined total of each.

Throughput, IOPS, and Latency

Figure 15 shows throughput, IOPS, and latency for a linked clone boot storm during storage failover.

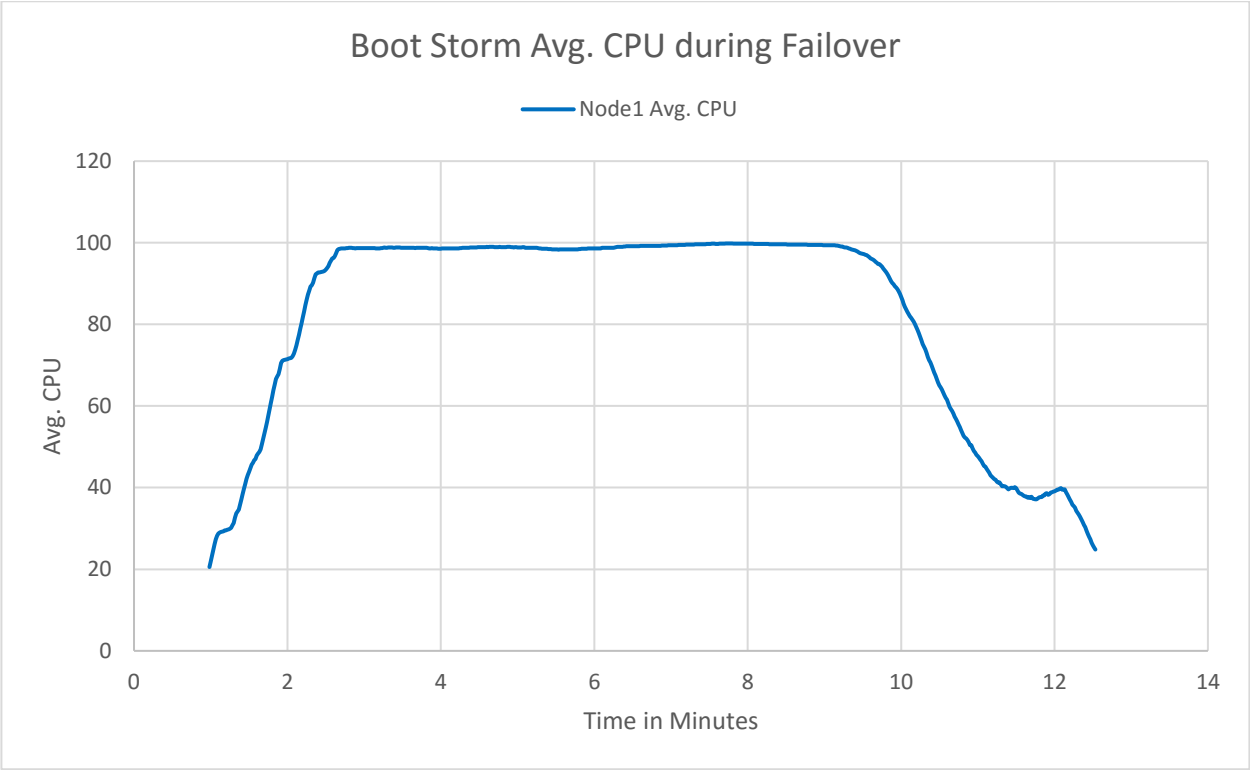
Figure 15) Throughput, IOPS, and latency for linked clone boot storm during storage failover.



Storage Controller CPU Utilization

Figure 16 shows storage controller CPU utilization for a linked clone boot storm during storage failover.

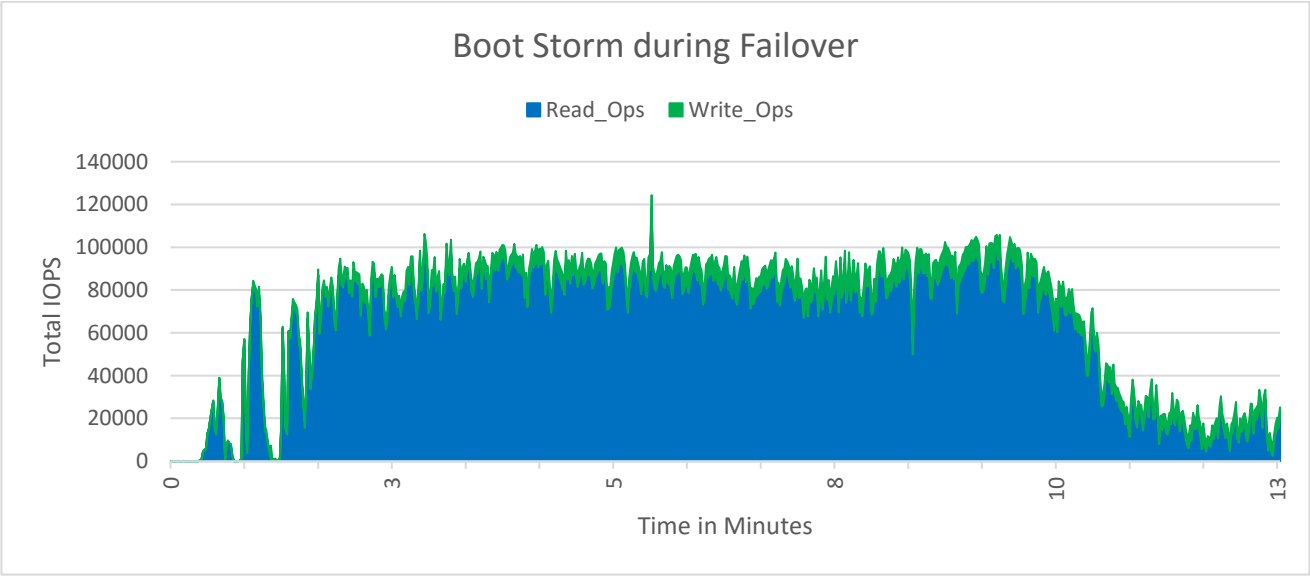
Figure 16) Storage controller CPU utilization for boot storm during storage failover.



Read/Write IOPS

Figure 17 shows read and write IOPS for a linked clone boot storm during storage failover.

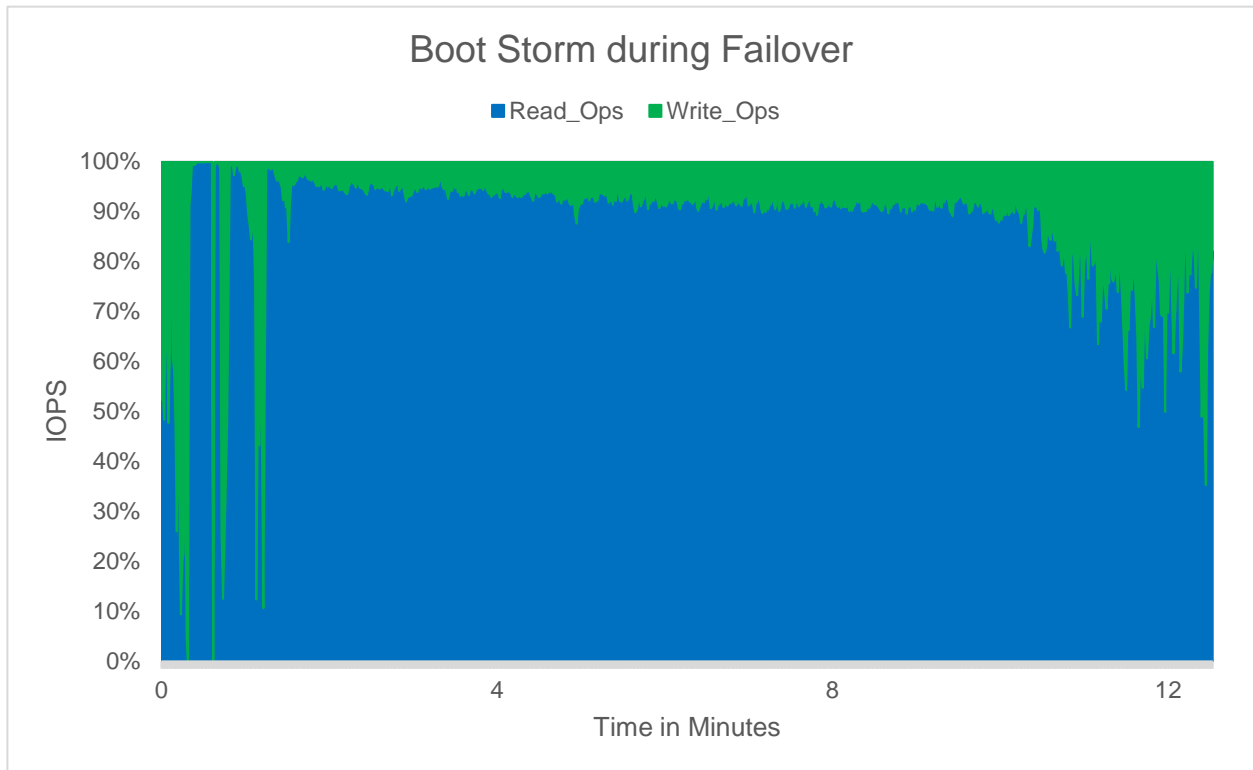
Figure 17) Read/write IOPS for VMware Horizon linked clone boot storm during storage failover.



Read/Write Ratio

Figure 18 shows read and write ratios for a linked clone boot storm during storage failover.

Figure 18) Read/write ratio for linked clone boot storm during storage failover.



Customer Impact (Test Conclusions)

Even during failover, the majority of time used for powering on the VMs was spent registering with Horizon View. Storage was not pushed to its limit even during the failover, as seen from the CPU utilization.

6.5 Monday Morning Login and Steady State

In this scenario, 1,500 users logged in after the VMs had already been logged into once, the profile had been created, and the desktops had been rebooted. During this type of login, user and profile data, application binaries, and libraries had to be read from a disk because they were not already contained in the VM memory. Table 10 lists the results for Monday morning login.

Table 10) Results for Login VSI login and steady-state knowledge worker workload.

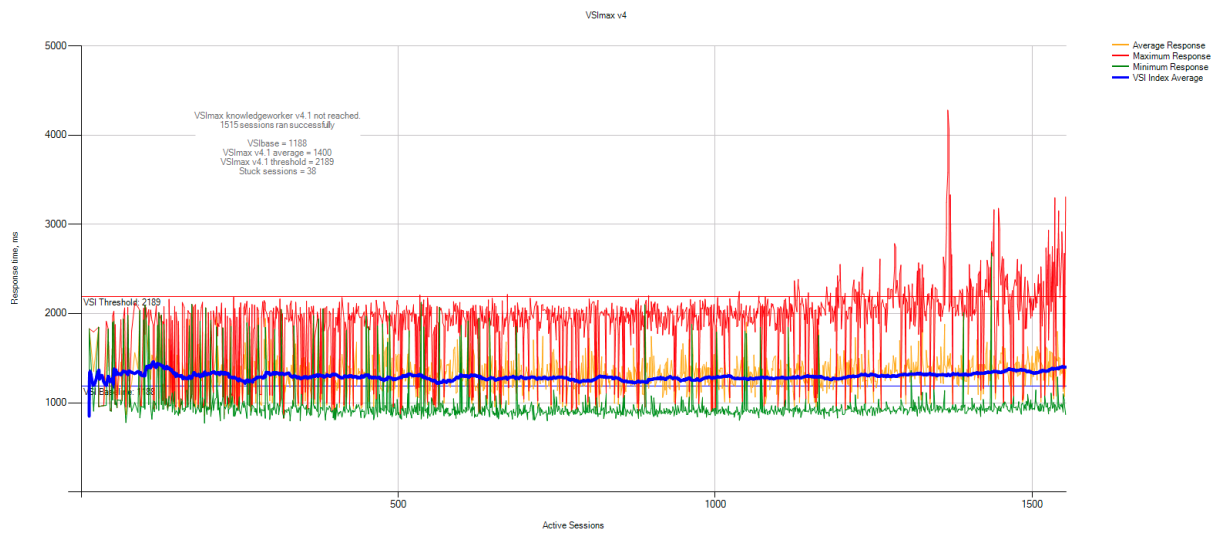
Avg. Login Time	Average Latency	Peak IOPS	Avg. IOPS	Peak Throughput	Avg. Throughput	Peak Storage CPU	Avg. Storage CPU
19.21 s	465 us	43,528	14,889	808MBps	233MBps	96%	31%

Note: CPU and latency measurements are based on the average across both nodes of the cluster. IOPS and throughput are based on a combined total of each.

Login VSI VSImax Results

Because the Login VSI VSImax v4.1 limit was not reached, more VMs could be deployed on this infrastructure. Figure 19 shows the VSImax results for Login VSI login and workload.

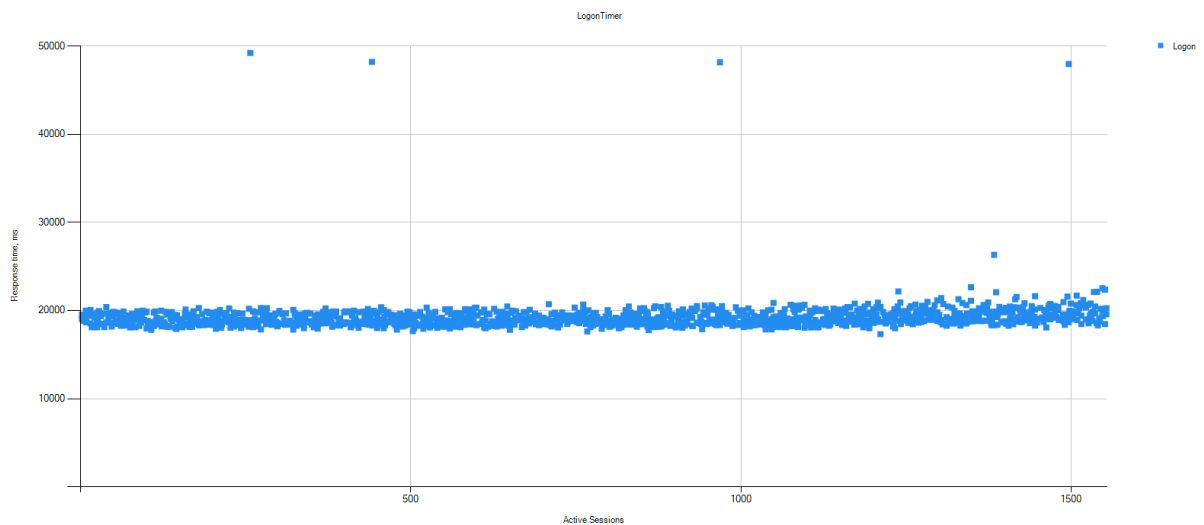
Figure 19) VSImax results for Login VSI login and workload.



Desktop Login Time

Average desktop login time was 19.21 seconds, which is considered a good login time. Figure 20 shows a scatterplot of the login times.

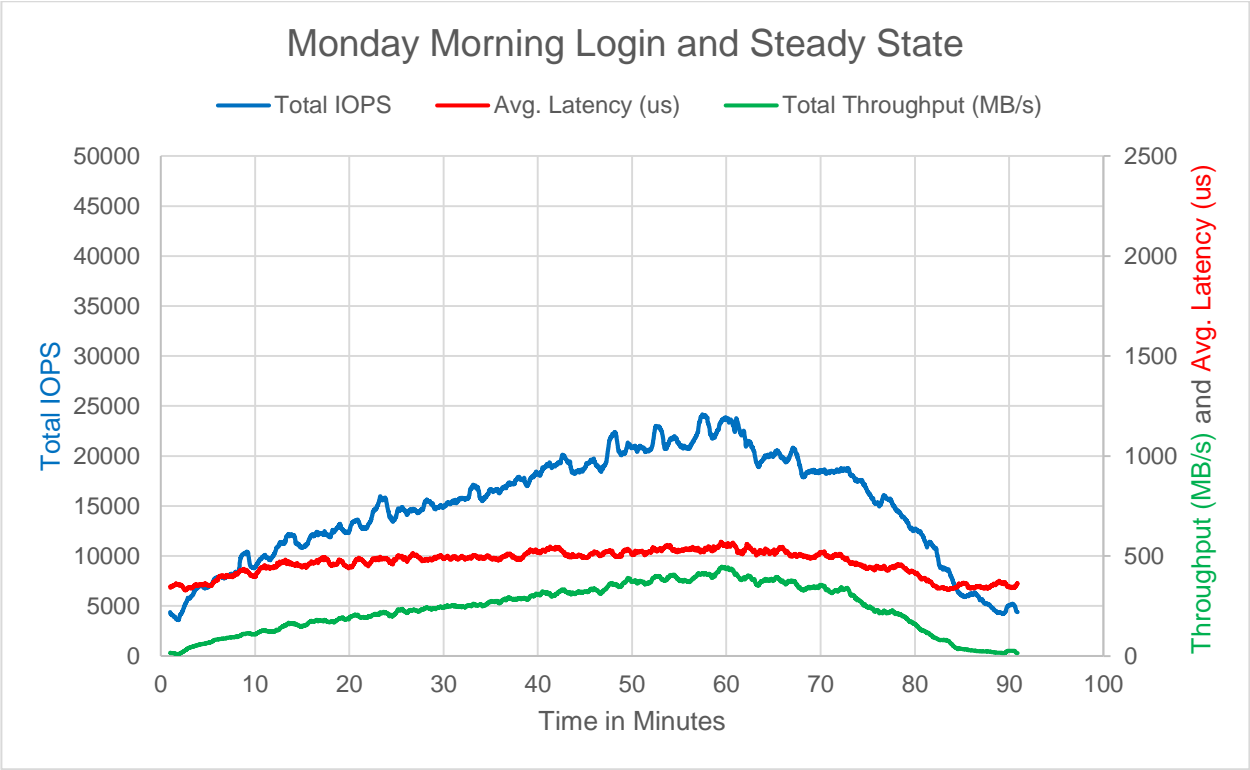
Figure 20) Scatterplot of login times.



Throughput, IOPS, and Latency

Figure 21 shows throughput, IOPS, and latency for Login VSI login and workload.

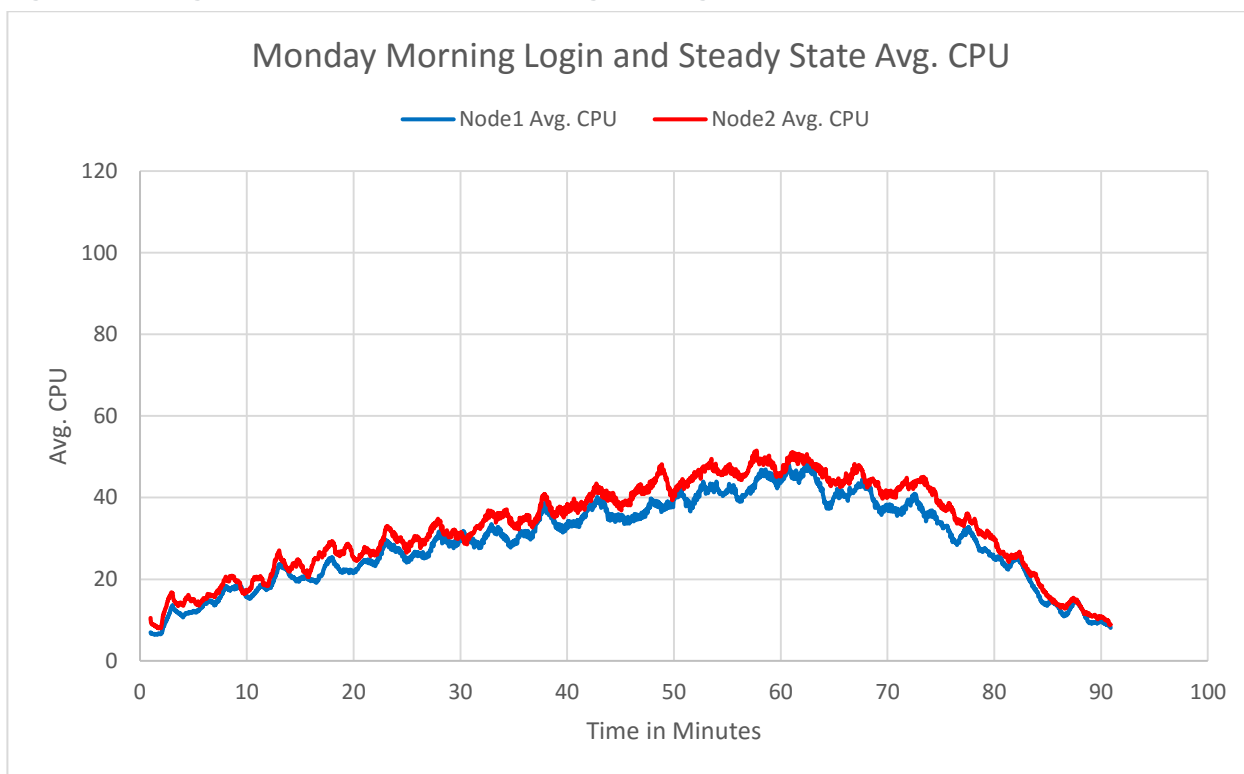
Figure 21) Throughput, IOPS, and latency for Login VSI login and workload.



Storage Controller CPU Utilization

Figure 22 shows results for storage controller CPU utilization for Login VSI login and workload during normal operation.

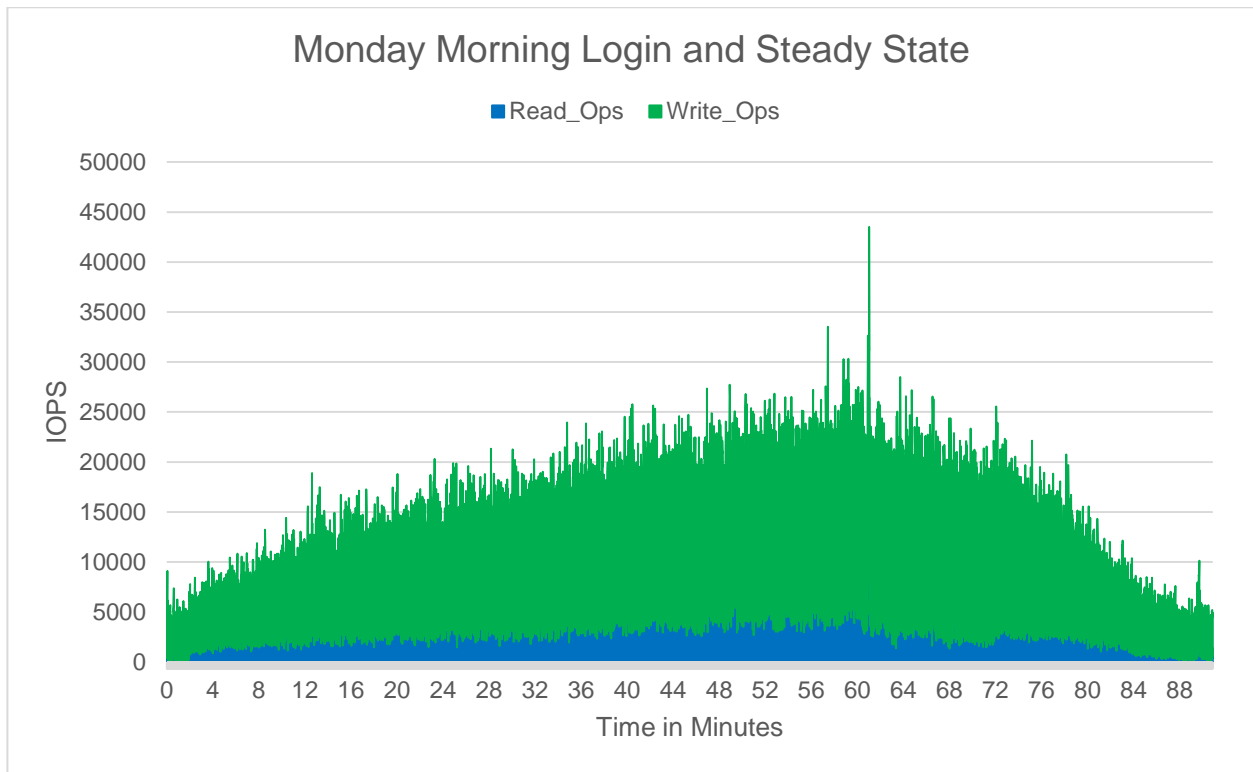
Figure 22) Storage controller CPU utilization for Login VSI login and workload.



Read/Write IOPS

Figure 23 shows results for read and write IOPS for Login VSI login and workload.

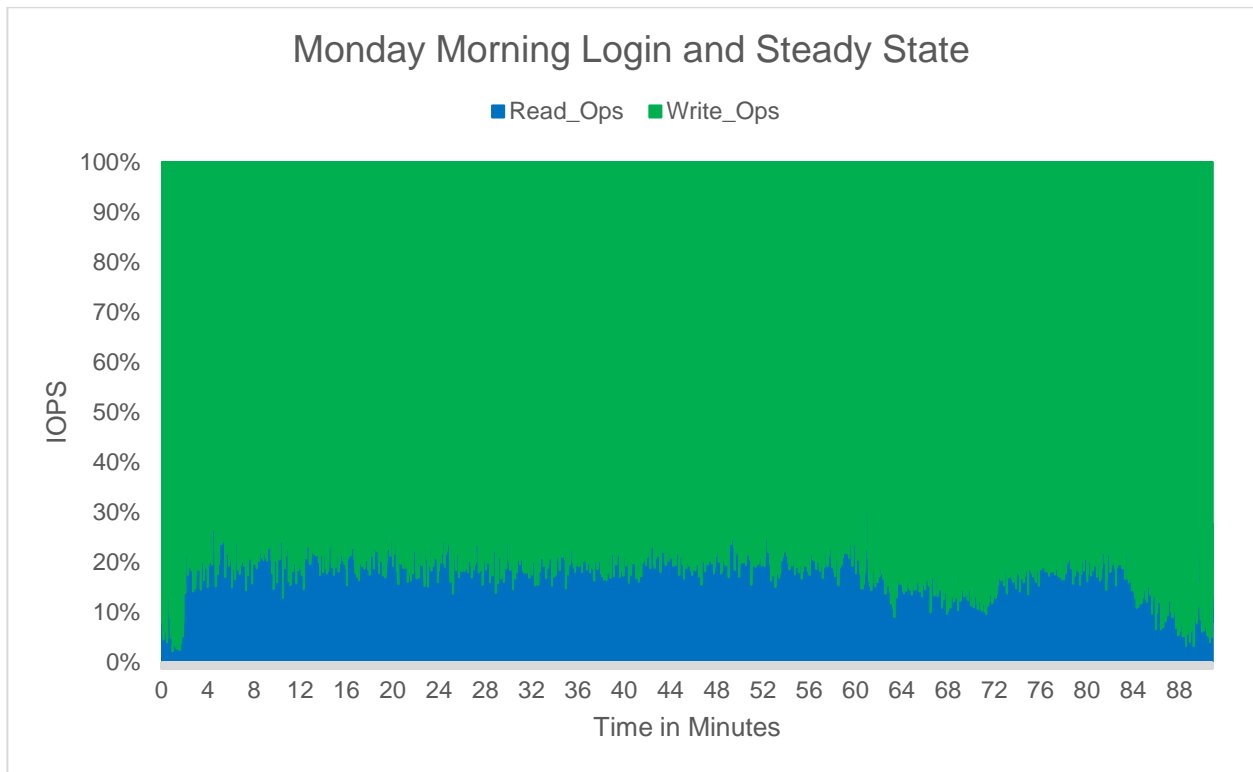
Figure 23) Read/write IOPS for Login VSI login and workload.



Read/Write Ratio

Figure 24 shows the read-to-write ratio for Login VSI login and workload.

Figure 24) Read/write ratio for Login VSI login and workload.



Customer Impact (Test Conclusions)

During the Login VSI test, the storage controller performed very well. The CPU utilization averaged less than 50%, latencies were under 1ms, and desktop performance was excellent. These results suggest that for this type of workload it might be possible to double the storage controller workload to 3,000 users total (1,500 per node) with excellent end-user performance and with the ability to tolerate a storage failover. This is further illustrated in the following section with results from a storage failover event.

6.6 Monday Morning Login and Steady State During Failover

In this scenario, 1,500 users logged in after the VMs had already been logged into once, the profile had been created, and the desktops had been rebooted, but during a storage failover event. During this type of login, user and profile data, application binaries, and libraries had to be read from a disk because they were not already contained in the VM memory. Table 11 lists the results for Monday morning login and workload during storage failover.

Table 11) Results for Login VSI login and workload during storage failover.

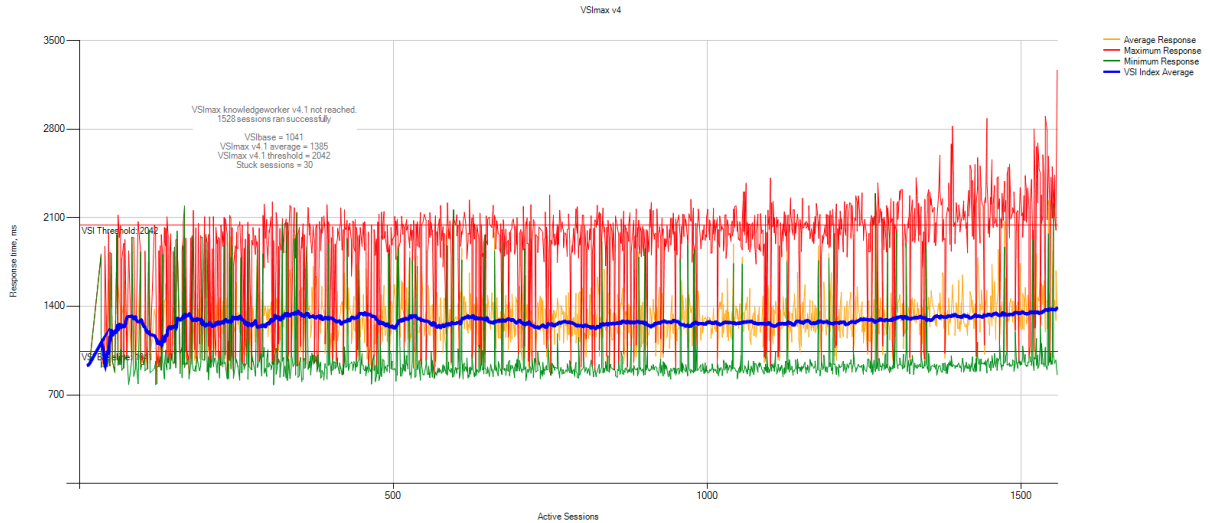
Avg. Login Time	Average Latency	Peak IOPS	Avg. IOPS	Peak Throughput	Avg. Throughput	Peak Storage CPU	Avg. Storage CPU
19.93 s	802 us	69,700	13,937	1126MBps	213MBps	100%	51%

Note: CPU and latency measurements are based on the average across both nodes of the cluster. IOPS and throughput are based on a combined total of each.

Login VSI VSImax Results

Because the Login VSI VSImax v4.1 limit was not reached, more VMs could be deployed on this infrastructure. Figure 25 shows the VSImax results for Login VSI login and workload during storage failover.

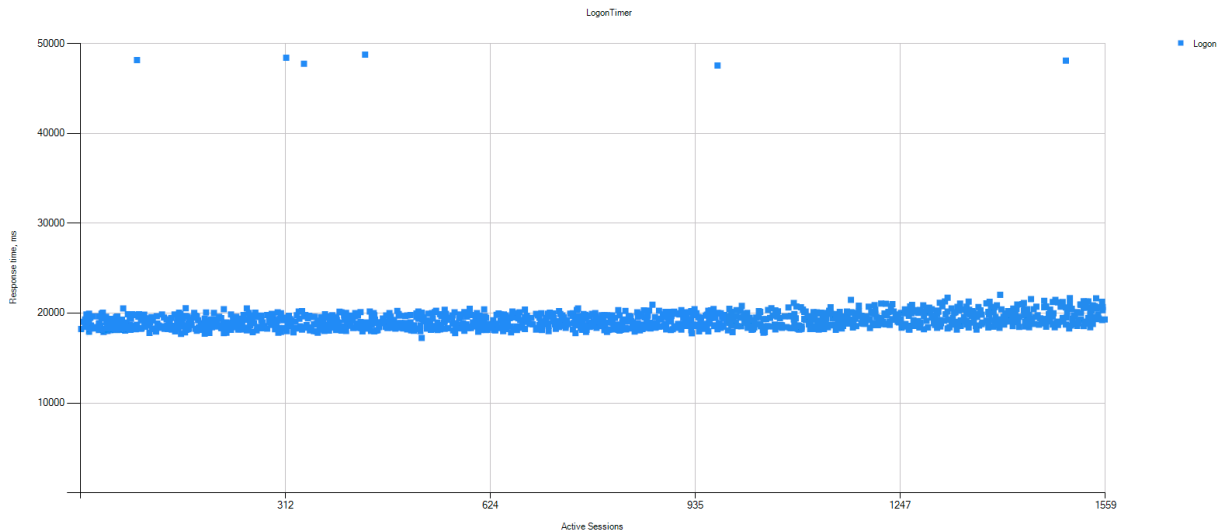
Figure 25) VSImax results for Login VSI login and workload during storage failover.



Desktop Login Time

Average desktop login time was 19.93 seconds, which is considered a good login time, especially during a failover situation and considering the amount of data each user had to download. Figure 26 shows a scatterplot of the login times during storage failover.

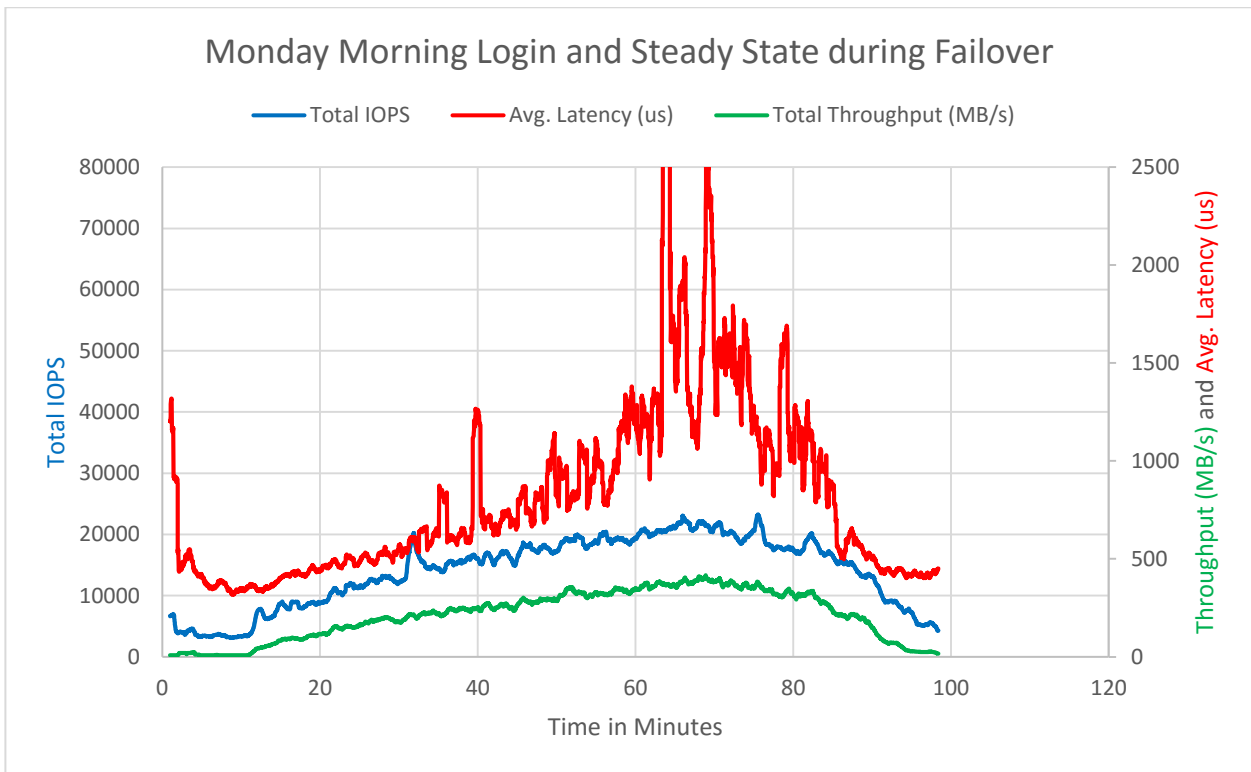
Figure 26) Scatterplot of login times during storage failover.



Throughput, IOPS, and Latency

Figure 27 shows throughput, IOPS, and latency for Login VSI login and workload during storage failover.

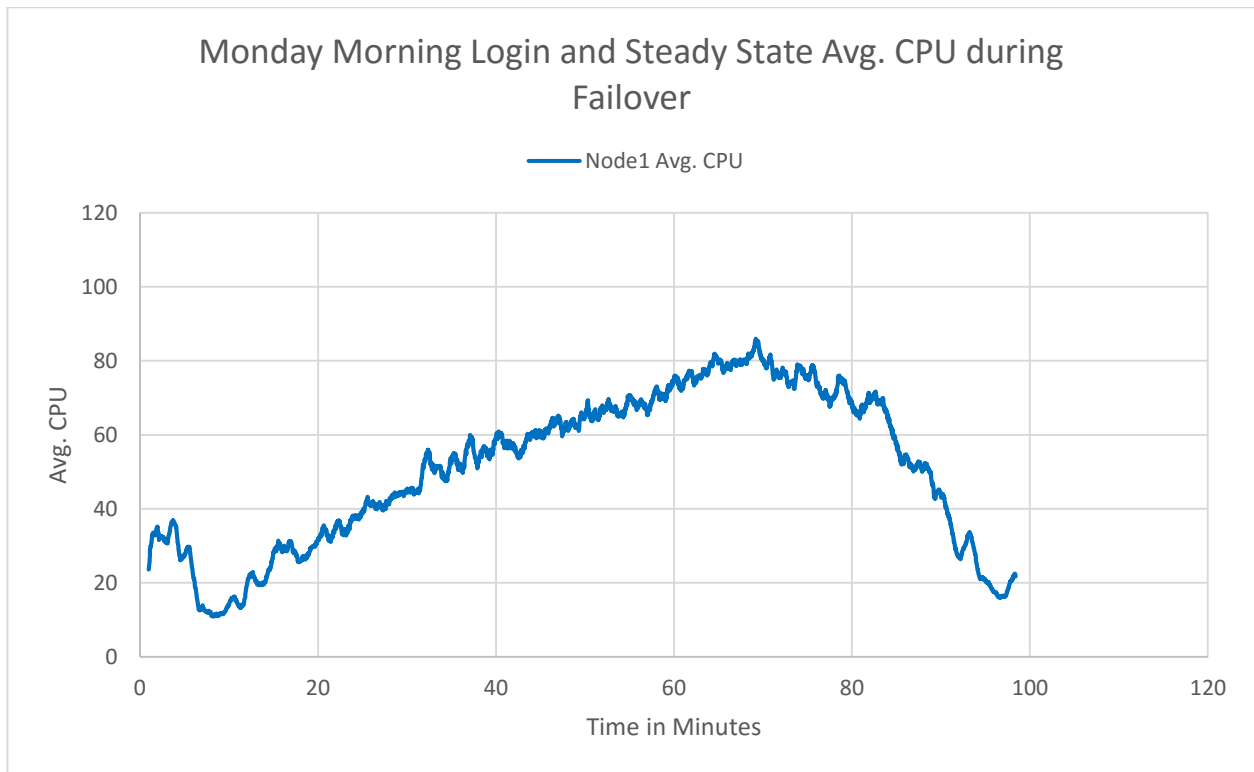
Figure 27) Throughput, IOPS, and latency for Login VSI login and workload during storage failover.



Storage Controller CPU Utilization

Figure 28 shows results for storage controller CPU utilization for Login VSI login and workload during storage failover testing.

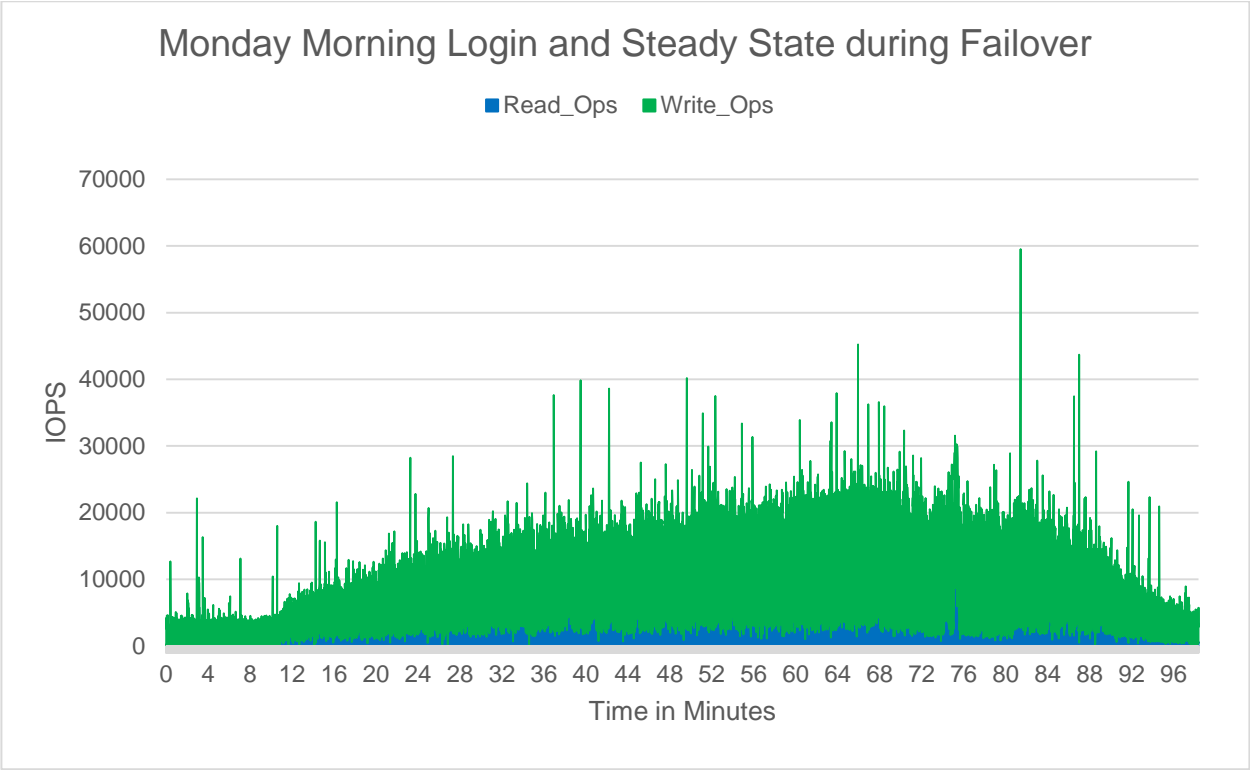
Figure 28) Storage controller CPU utilization for Login VSI login and workload during storage failover.



Read/Write IOPS

Figure 29 shows results for read and write IOPS for Login VSI login and workload during storage failover testing.

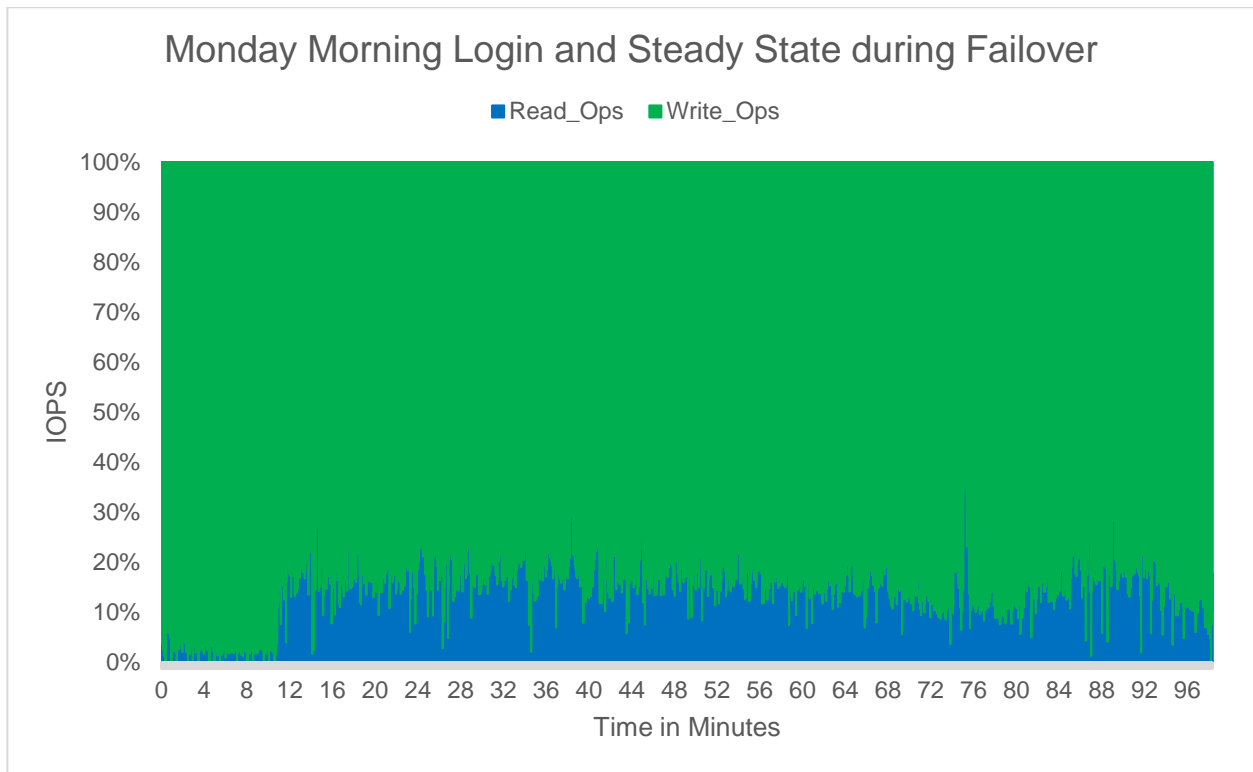
Figure 29) Read/write IOPS for Login VSI login and workload during storage failover.



Read/Write Ratio

Figure 30 shows the read-to-write ratio for Login VSI login and workload during storage failover testing.

Figure 30) Read/write ratio for Login VSI login and workload during storage failover.



Customer Impact (Test Conclusions)

During the Login VSI test, the storage controller performed very well. The CPU utilization averaged approximately 50%, latencies were stable, and desktop performance was excellent. These results suggest that for this type of workload it might be possible to double the storage controller workload to 3,000 users total (1,500 per node) with excellent end-user performance and with the ability to tolerate a storage failover.

6.7 Tuesday Morning Login and Steady State

This section describes test objectives and methodology and provides results from Tuesday Morning Login and Steady State testing.

Test Objectives and Methodology

The objective of this test was to run a Login VSI knowledge worker workload to determine how the storage controller performed and evaluate the end-user experience. This Login VSI workload first had the users log in to their desktops and begin working. The login phase occurred over a 60-minute period. The test was run for an additional 15 minutes, before which the first logoff command was sent. The Tuesday morning test was conducted immediately following a Monday morning test, simulating login profiles already in cache and application data already primed.

Login VSI Login and Workload Test

The desktops were configured so that every time a user connects, he or she gets the same desktop each time. In this scenario, 1,500 users logged in and executed the LoginVSI Knowledge Worker profile. Table 12 shows the results.

Table 12) Results for Login VSI login and workload.

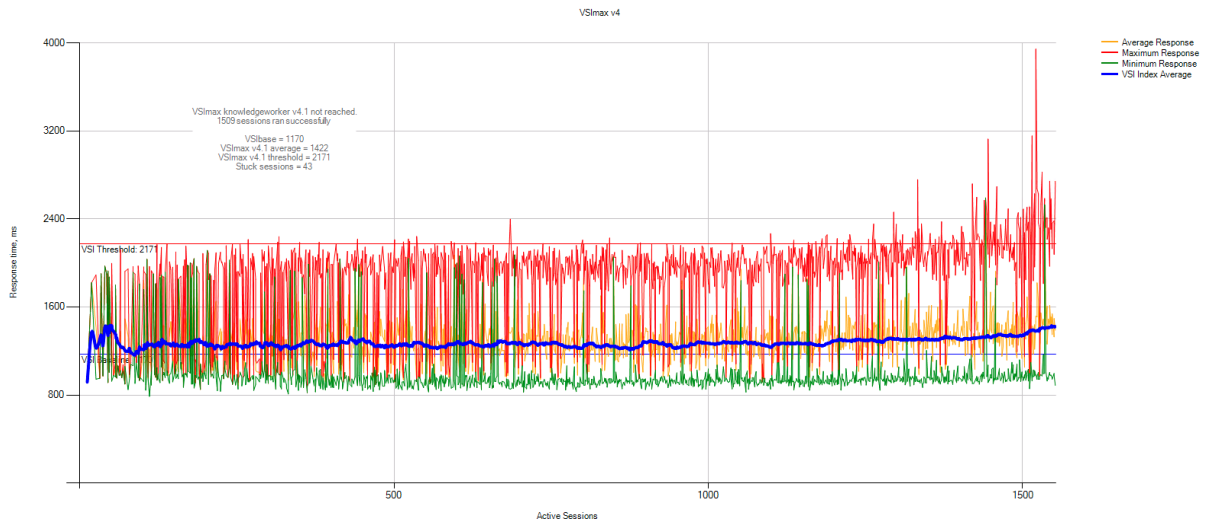
Avg. Login Time	Average Latency	Peak IOPS	Avg. IOPS	Peak Throughput	Avg. Throughput	Peak Storage CPU	Avg. Storage CPU
19.73 sec/VM	450 us	27,159	12,236	620MBps	191MBps	94%	28%

Note: CPU and latency measurements are based on the average across both nodes of the cluster. IOPS and throughput are based on a combined total of each.

Login VSI VSImax Results

Because the Login VSI VSImax was not reached, more VMs could be deployed on this infrastructure. Figure 31 shows VSImax results for Login VSI login and workload testing.

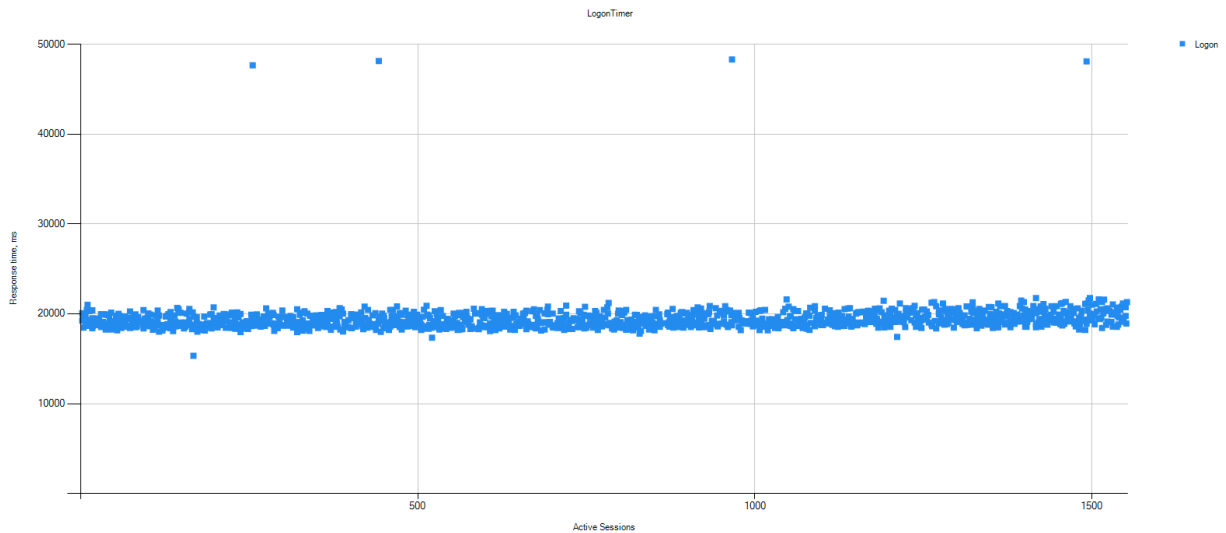
Figure 31) VSImax results for Login VSI login and workload.



Desktop Login Time

Average desktop login time was 19.73 seconds. Figure 32 shows a scatterplot of the Login VSI times.

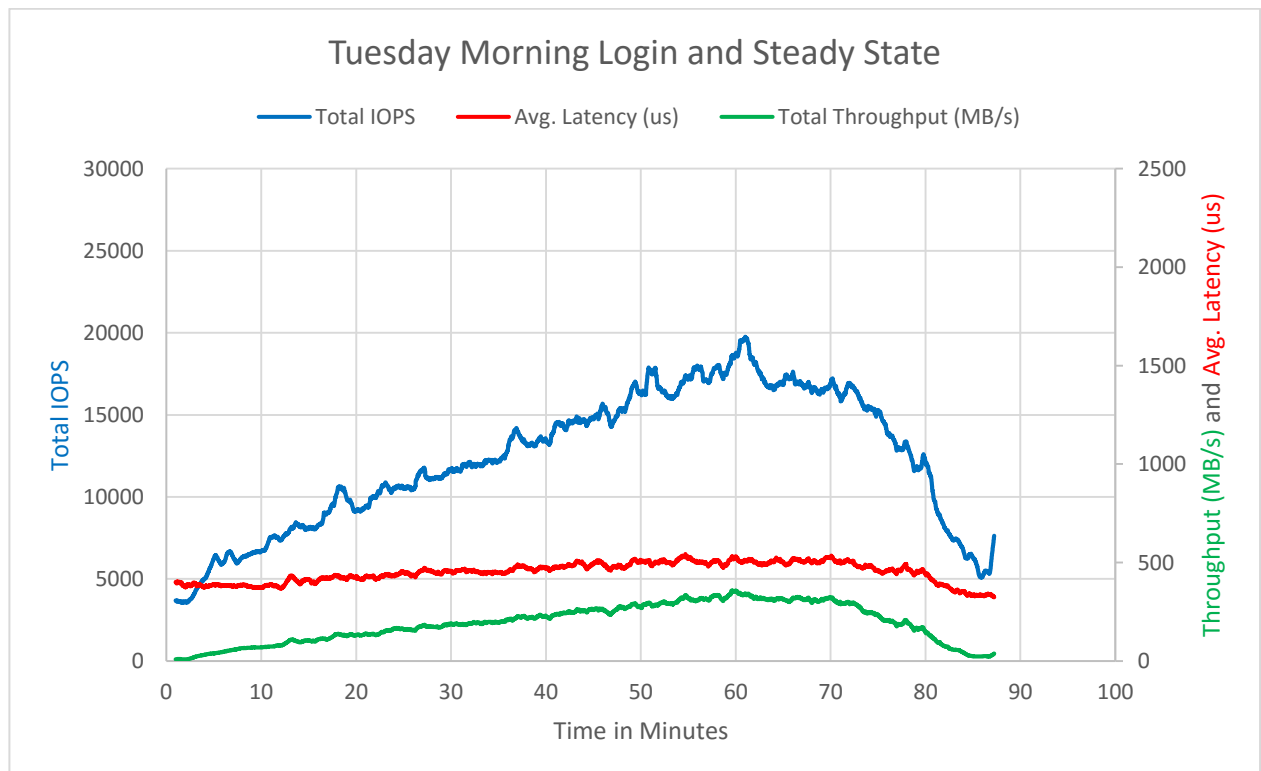
Figure 32) Scatterplot of Login VSI login times.



Throughput, IOPS, and Latency

Figure 33 shows results for throughput, IOPS, and latency for Login VSI login and workload testing.

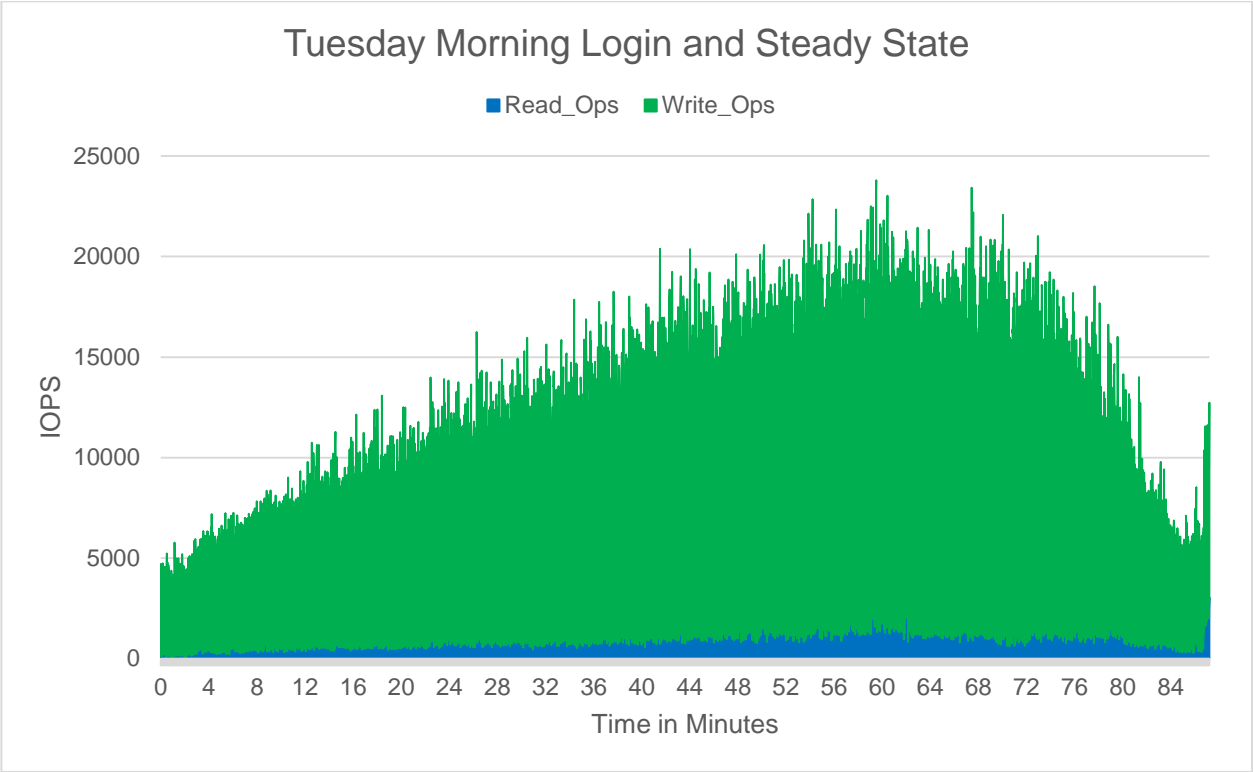
Figure 33) Throughput, IOPS, and latency for Login VSI login and workload.



Storage Controller CPU Utilization

Figure 34 shows results for storage controller CPU utilization for Login VSI login and workload testing.

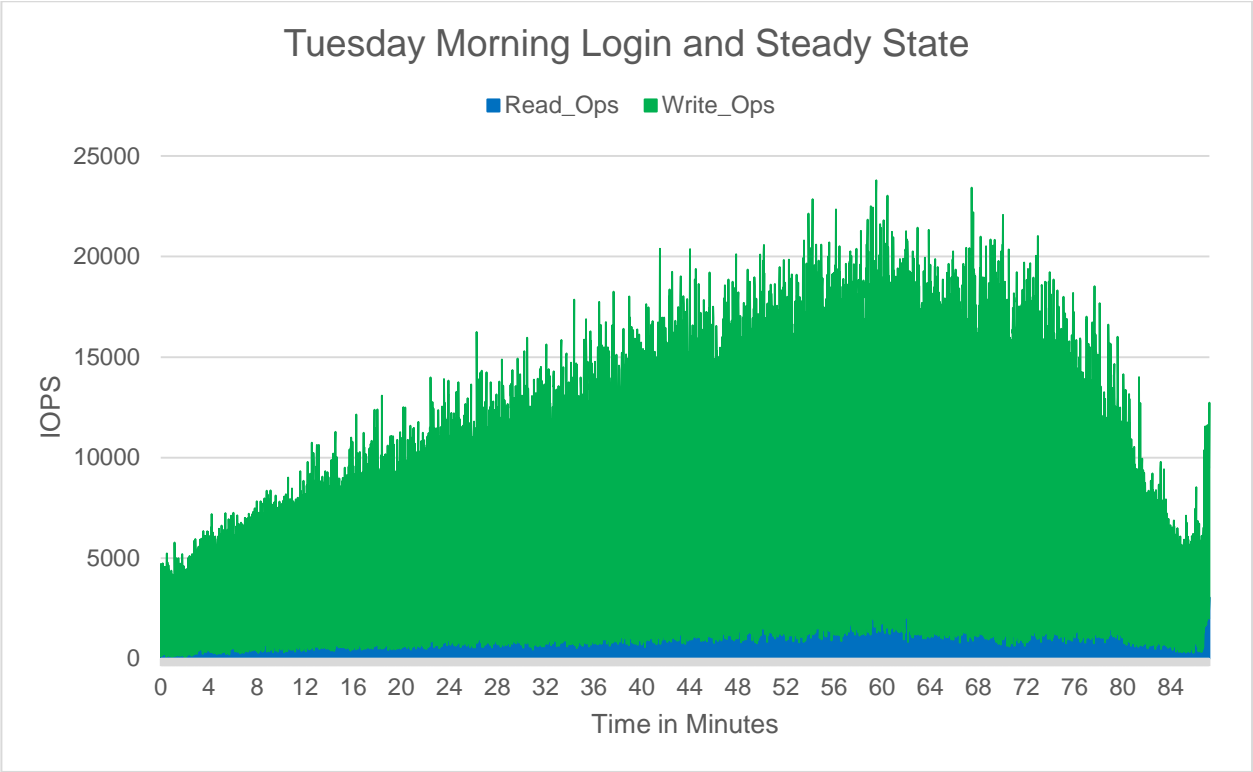
Figure 34) Storage controller CPU utilization for Login VSI login and workload.



Read/Write IOPS

Figure 35 shows results for read and write IOPS for Login VSI login and workload testing.

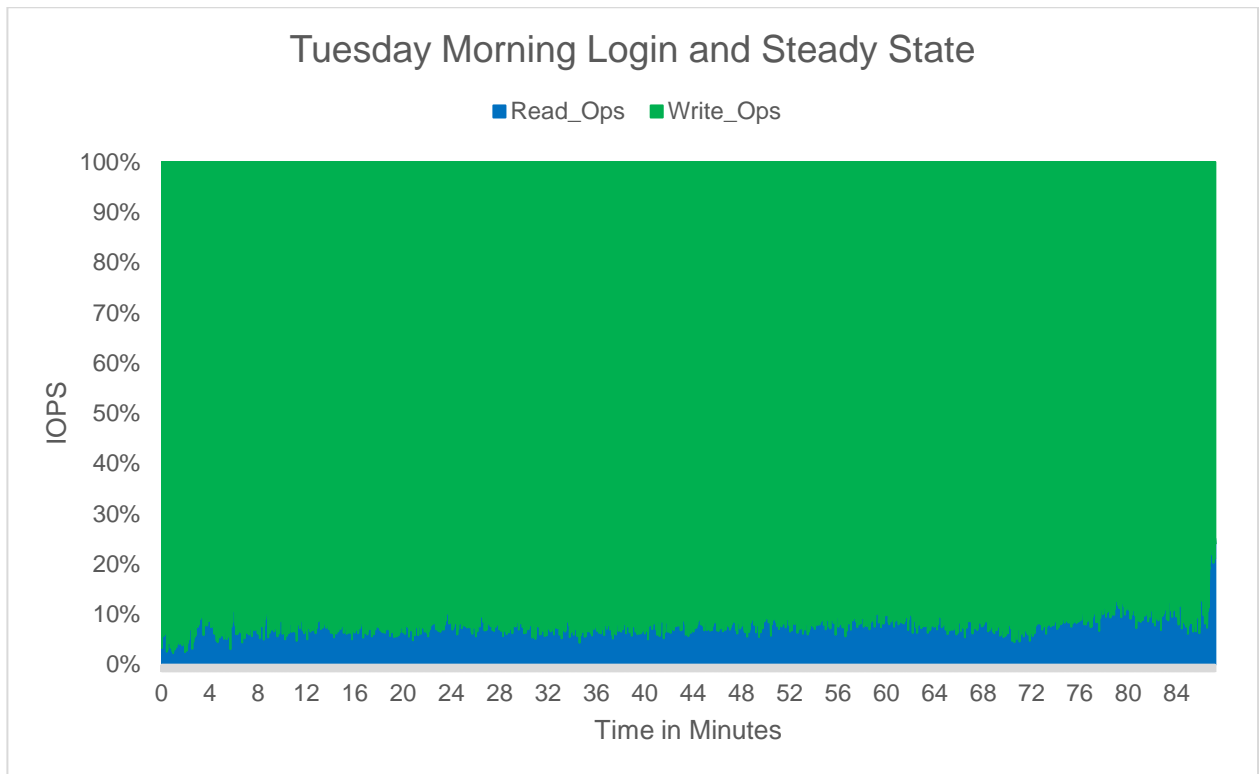
Figure 35) Read/write IOPS for Login VSI login and workload.



Read/Write Ratio

Figure 36 shows the read-to-write ratio for Login VSI login and workload testing.

Figure 36) Read/write ratio for Login VSI login and workload.



Customer Impact (Test Conclusions)

During the test, the storage controller performed very well. The CPU utilization was not high during this test, average latencies were under 1ms, and desktop performance was excellent. These results suggest that it might be possible to double the storage controller workload to 3,000 users or more and still maintain excellent end-user performance, provided that there is sufficient host infrastructure. The Login VSI test during storage failover described in the following section reinforces that point.

6.8 Tuesday Morning Login and Steady State During Failover

This section describes test objectives and methodology and provides results from Tuesday Morning Login and Steady State testing during a storage controller failover event.

Test Objectives and Methodology

The objective of this test was to run a Login VSI knowledge worker workload to determine how the storage controller performed and evaluate the end-user experience. This Login VSI workload first had the users log in to their desktops and begin working. The login phase occurred over a 60-minute period. The test was run for an additional 15 minutes, before which the first logoff command was sent. The Tuesday morning test was conducted immediately following a Monday morning test, simulating login profiles already in cache and application data already primed.

Login VSI Login and Workload Test

The desktops were configured such that every time a user connects, he or she gets the same desktop. In this scenario, 1,500 users logged in and executed the LoginVSI Knowledge Worker profile. After user and profile data was created, application binaries and libraries had to be read from disk for doing the test. Table 13 shows the results.

Table 13) Results for Login VSI login and workload.

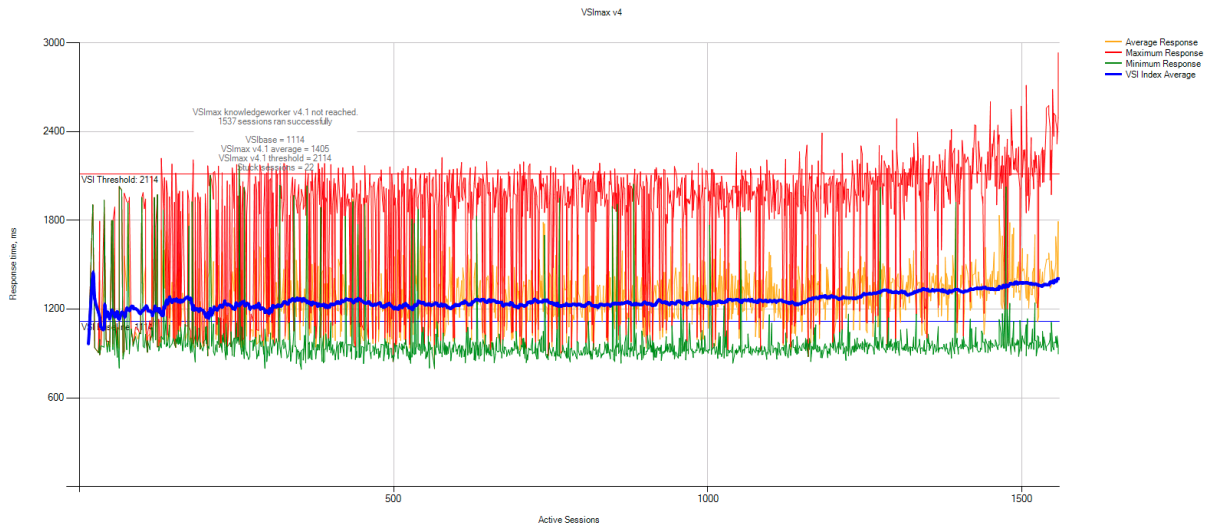
Avg. Login Time	Average Latency	Peak IOPS	Avg. IOPS	Peak Throughput	Avg. Throughput	Peak Storage CPU	Avg. Storage CPU
19.65 sec/VM	707 us	44,250	12,364	1106MBps	194MBps	100%	48%

Note: CPU and latency measurements are based on the average across both nodes of the cluster. IOPS and throughput are based on a combined total of each.

Login VSI VSImax Results

Because the Login VSI VSImax was not reached, more VMs could be deployed on this infrastructure. Figure 37 shows VSImax results for Login VSI login and workload testing.

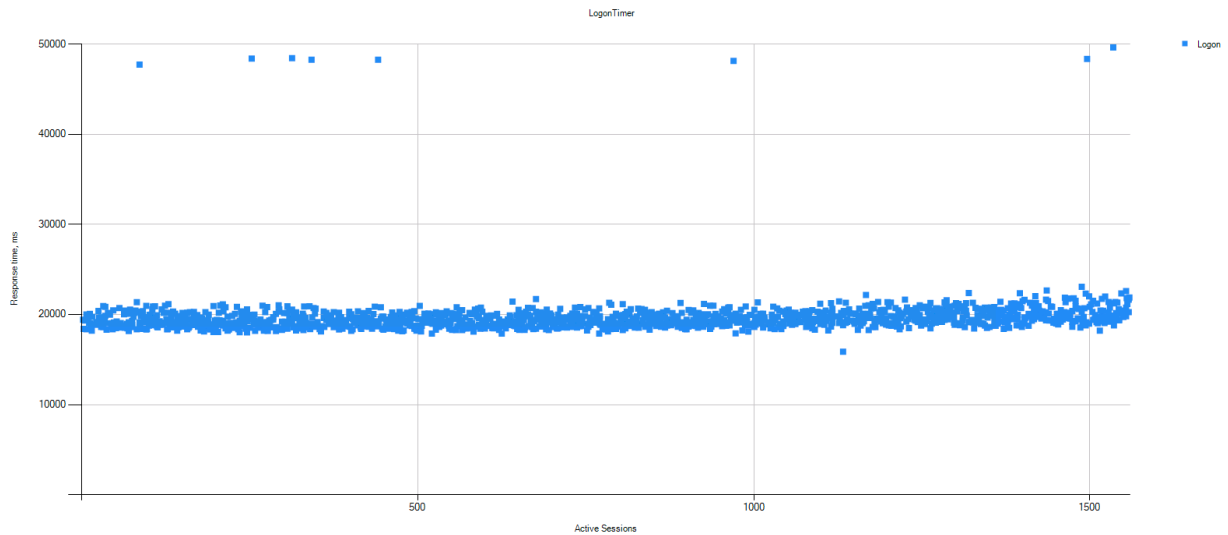
Figure 37) VSImax results for Login VSI login and workload.



Desktop Login Time

Average desktop login time was 19.65 seconds. Figure 38 shows a scatterplot of the Login VSI times.

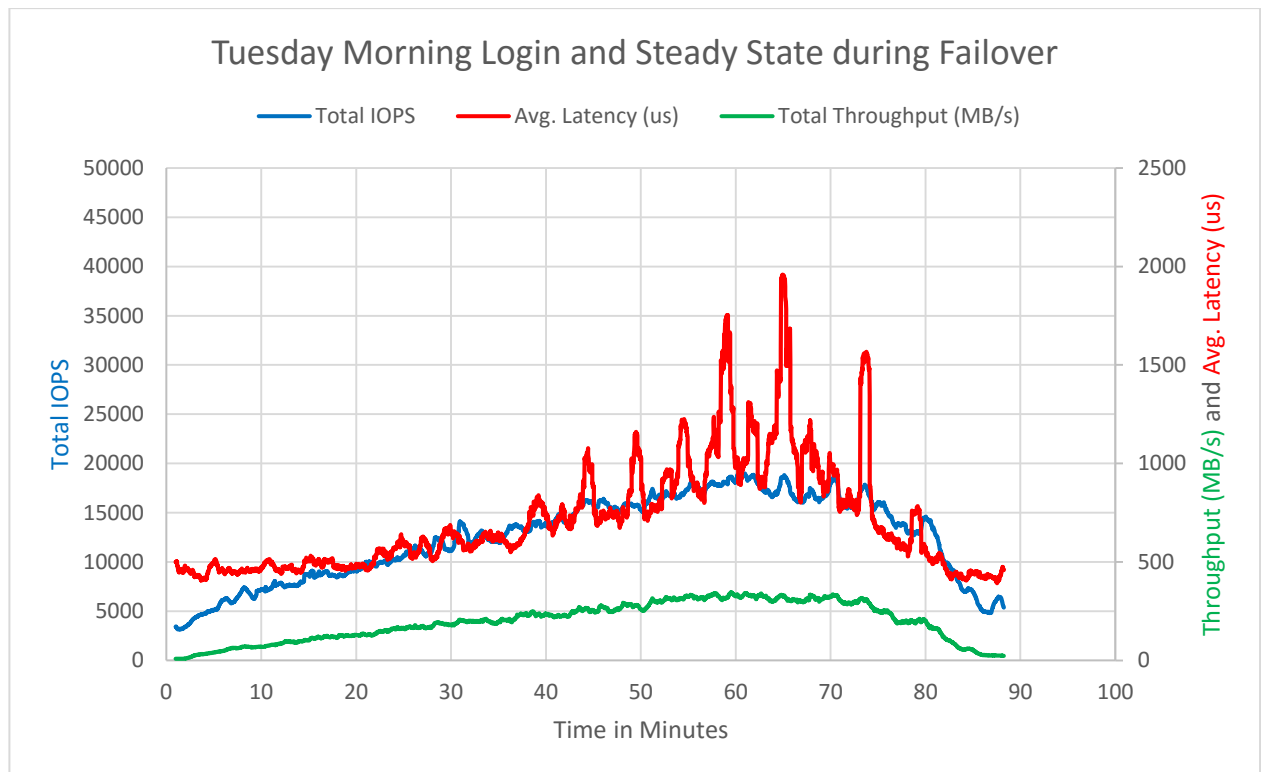
Figure 38) Scatterplot of Login VSI login times.



Throughput, IOPS, and Latency

Figure 39 shows results for throughput, IOPS, and latency for Login VSI login and workload testing.

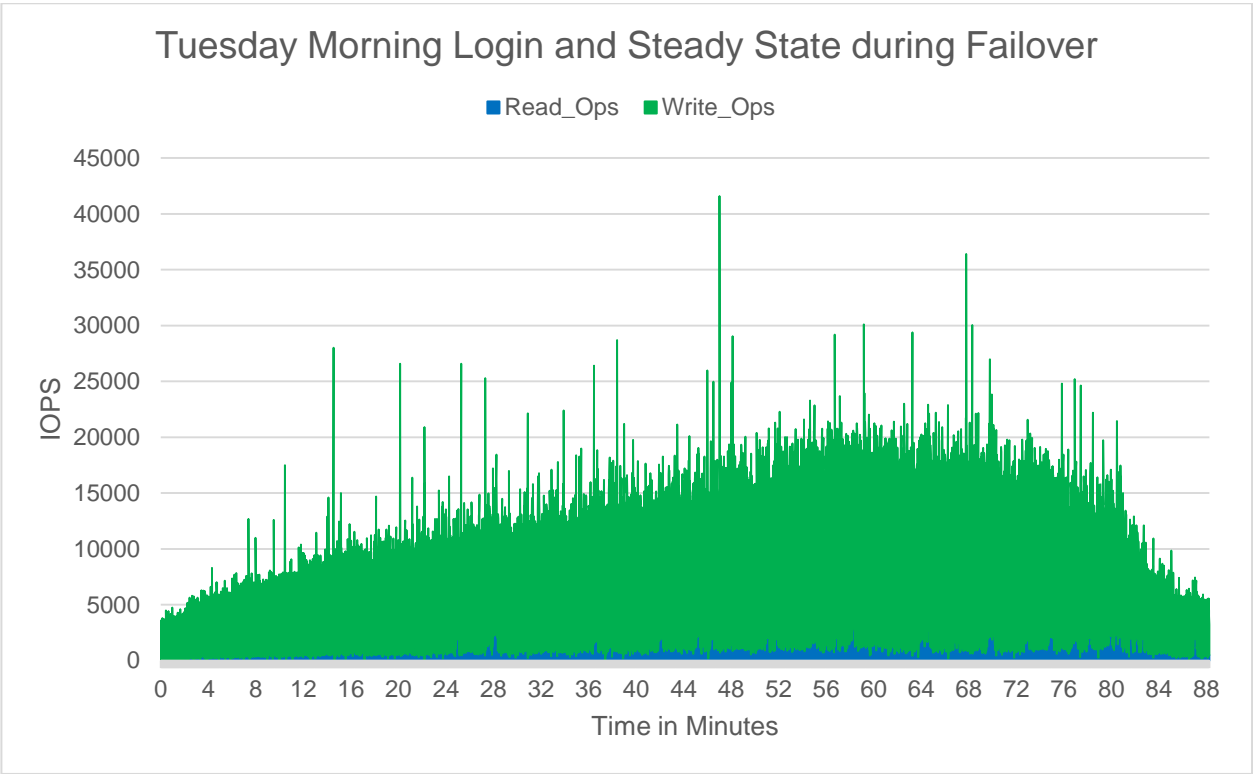
Figure 39) Throughput, IOPS, and latency for Login VSI login and workload.



Storage Controller CPU Utilization

Figure 40 shows results for storage controller CPU utilization for Login VSI login and workload testing.

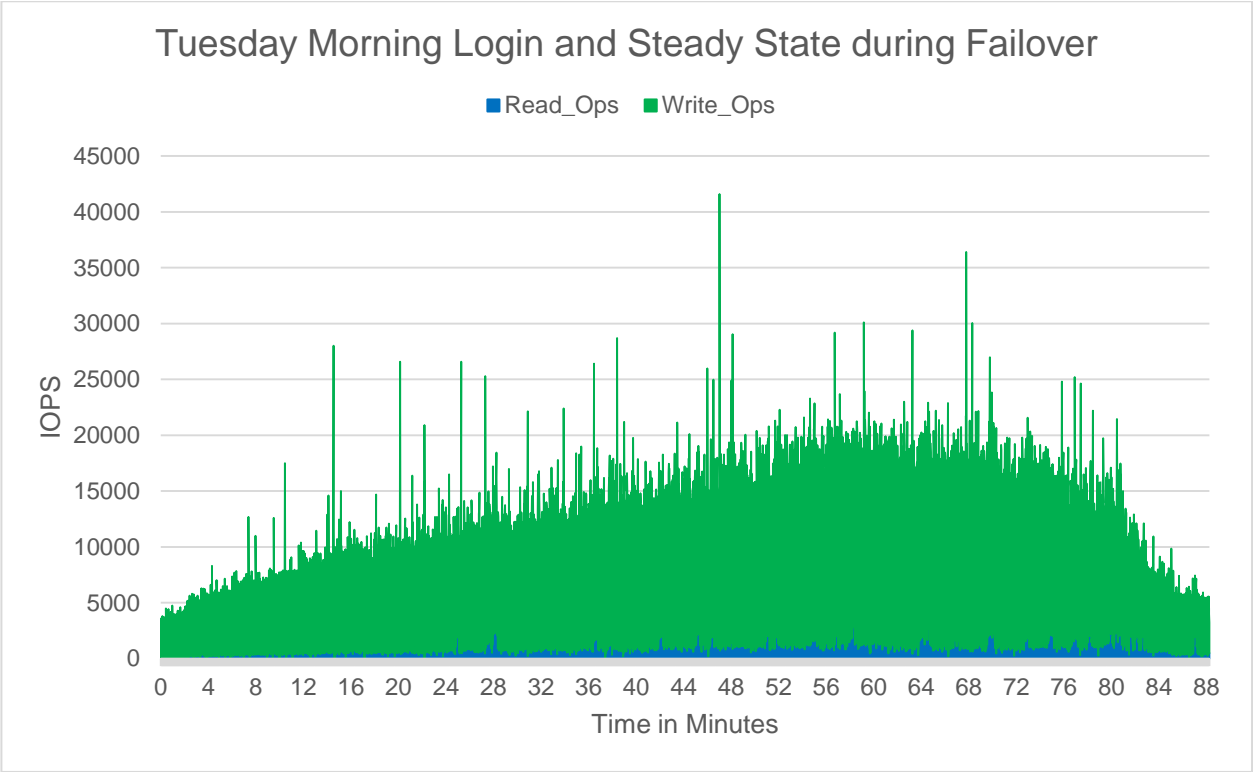
Figure 40) Storage controller CPU utilization for Login VSI login and workload.



Read/Write IOPS

Figure 41 shows results for read and write IOPS for Login VSI login and workload testing.

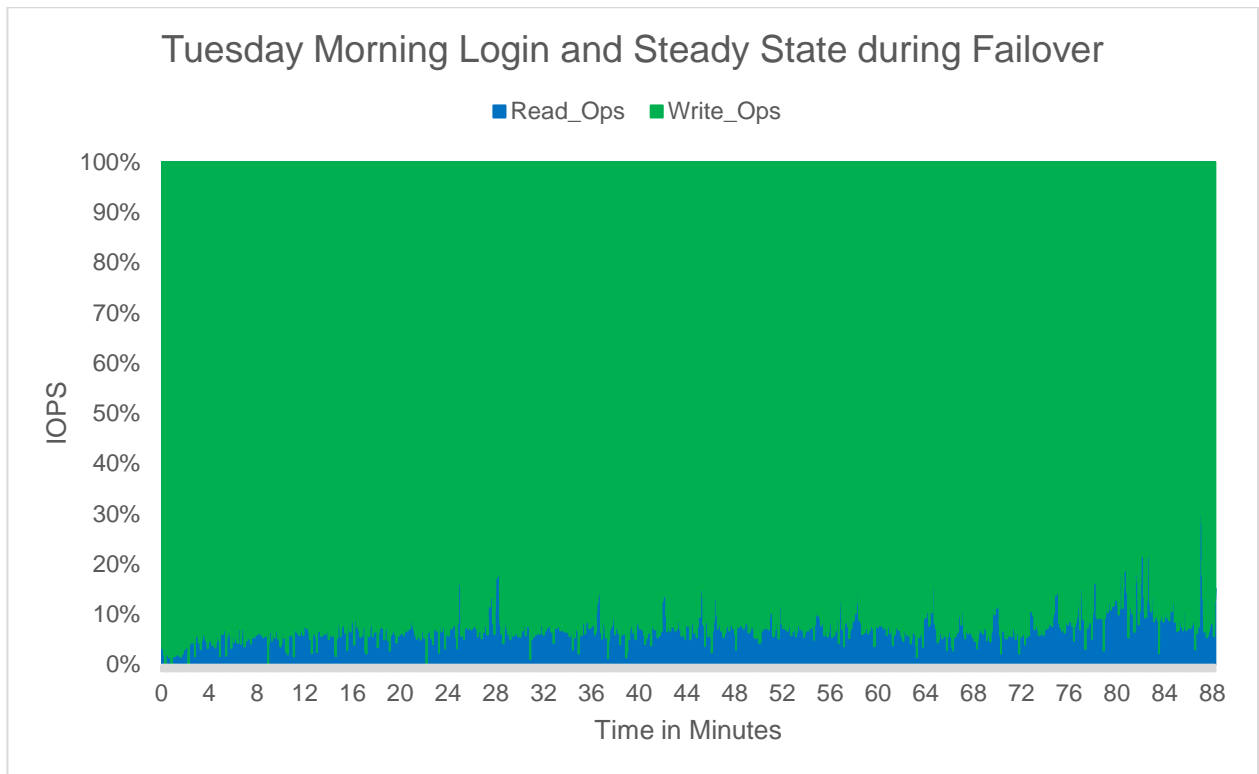
Figure 41) Read/write IOPS for Login and VSI login and workload.



Read/Write Ratio

Figure 42 shows the read-to-write ratio for Login VSI login and workload testing.

Figure 42) Read/write ratio for Login VSI login and workload.



Customer Impact (Test Conclusions)

During the test, the storage controller performed very well. The CPU utilization was acceptable during this test, average latencies were under 1ms, and desktop performance was excellent. These results suggest that it might be possible to double the storage controller workload to 3,000 users or more and still maintain excellent end-user performance, provided there is sufficient host infrastructure.

6.9 Refresh 1,500 Desktop VMs

In this scenario, 1,500 VMs are refreshed that had already been logged into, the profile had been created, and the desktops had been rebooted. Table 14 lists the results for the refresh operation workload profile.

The time to complete shows the point at which all refresh composer operations had completed and all VMs had reported as `Available` in Horizon View Administrator, with no pending View Composer operations.

Table 14) Results for Login VSI login and workload during storage failover.

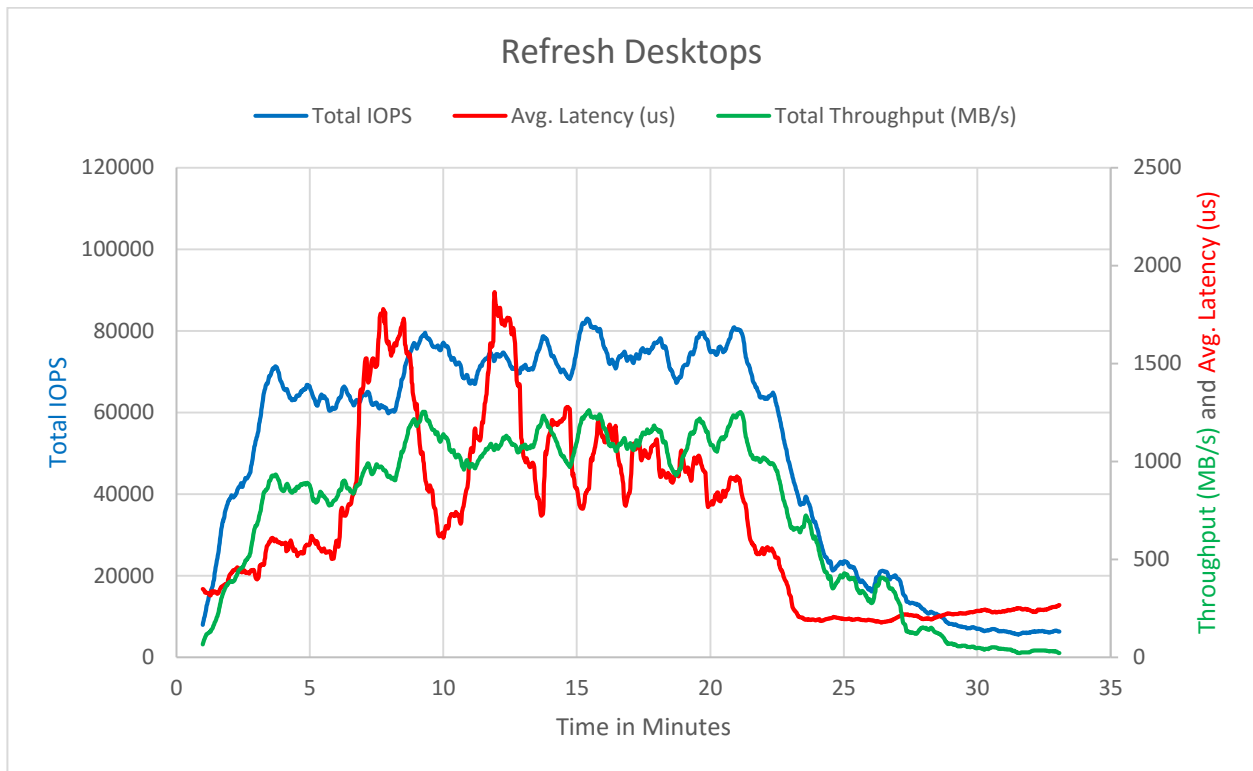
Time to Complete	Average Latency	Peak IOPS	Avg. IOPS	Peak Throughput	Avg. Throughput	Peak Storage CPU	Avg. Storage CPU
~25 min	723 us	111,789	53,920	2079MBps	786MBps	100%	60%

Note: CPU and latency measurements are based on the average across both nodes of the cluster. IOPS and throughput are based on a combined total of each.

Throughput, IOPS, and Latency

Figure 43 shows results for throughput, IOPS, and latency during a refresh operation.

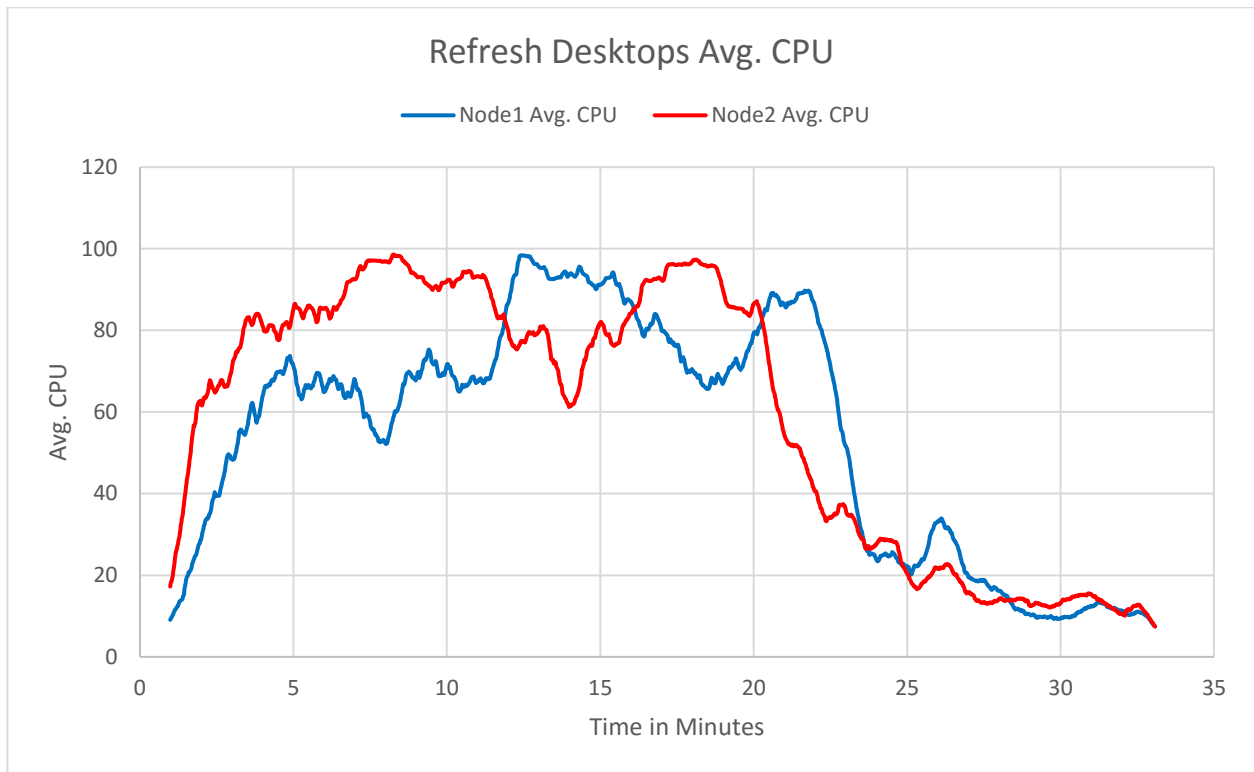
Figure 43) Throughput, IOPS, and latency during a refresh operation.



Storage Controller CPU Utilization

Figure 44 shows results for storage controller CPU utilization for Login VSI login and workload during storage failover testing.

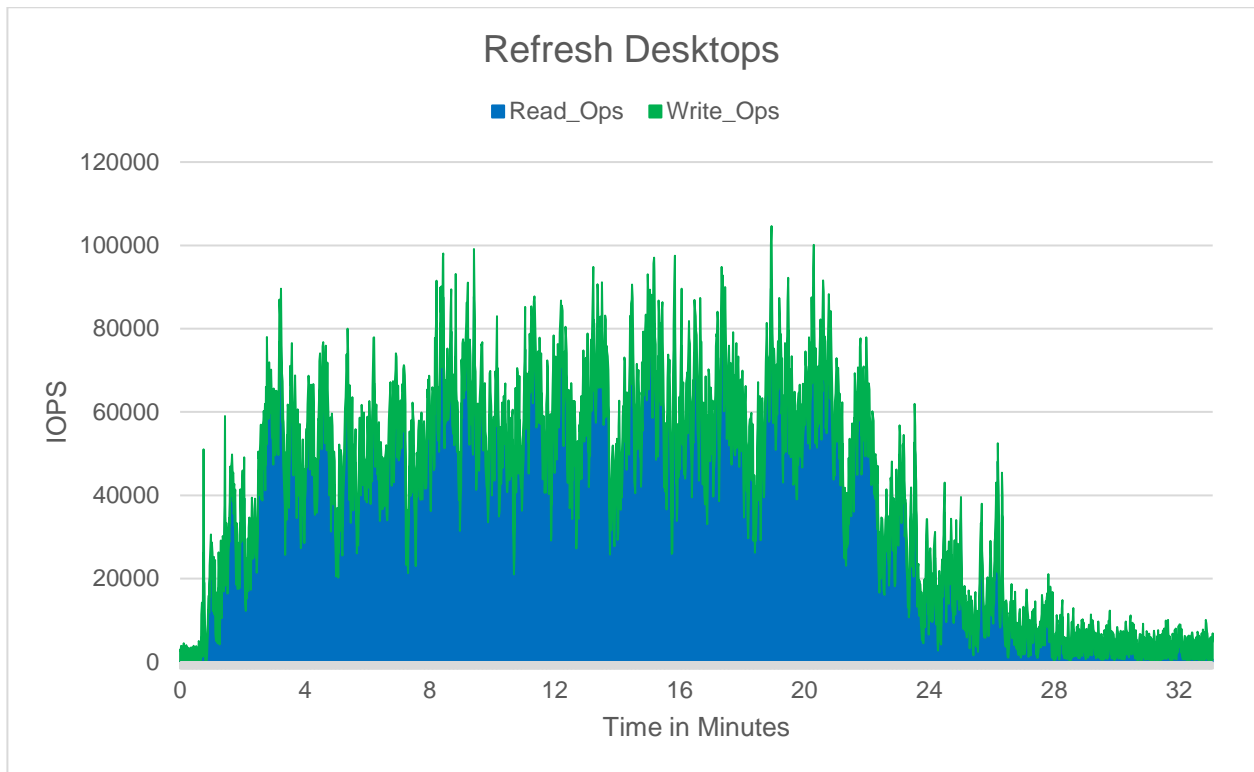
Figure 44) Storage controller CPU utilization during a refresh operation.



Read/Write IOPS

Figure 45 shows results for read and write IOPS during a refresh operation.

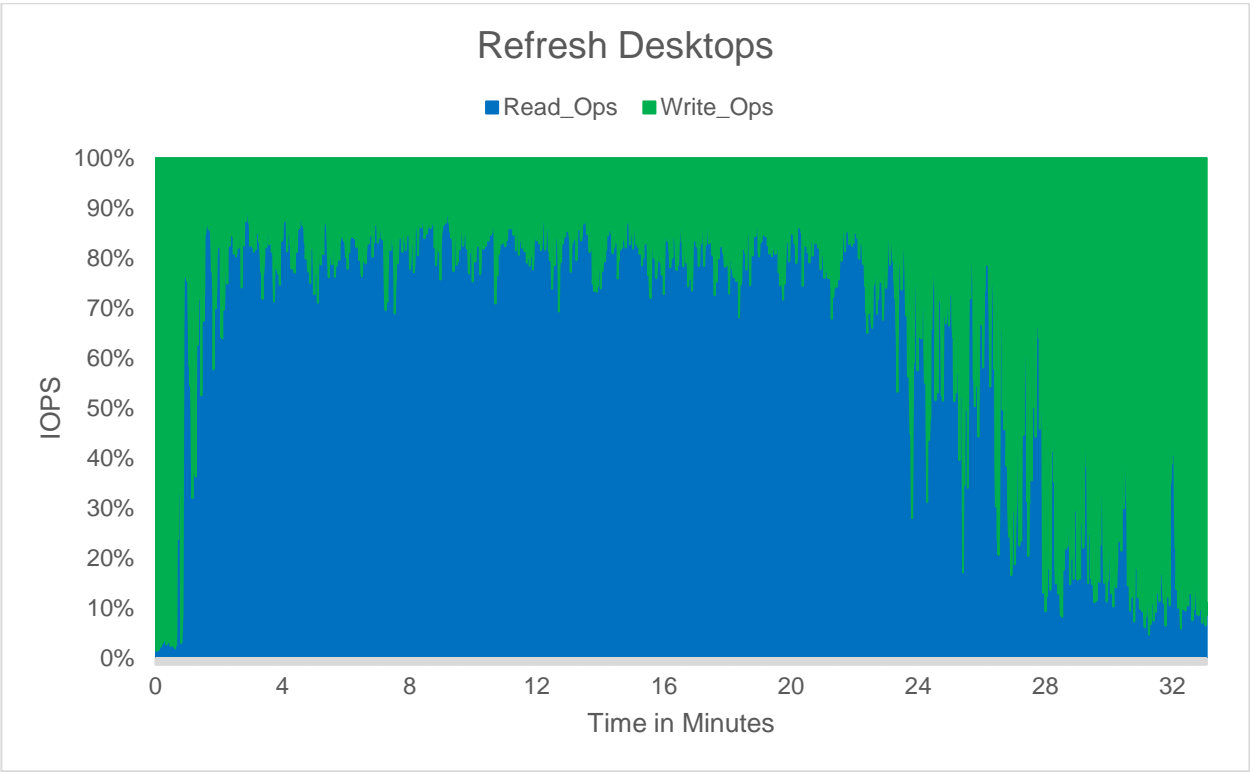
Figure 45) Read/write IOPS during a refresh operation.



Read/Write Ratio

Figure 46 shows the read-to-write ratio during a refresh operation.

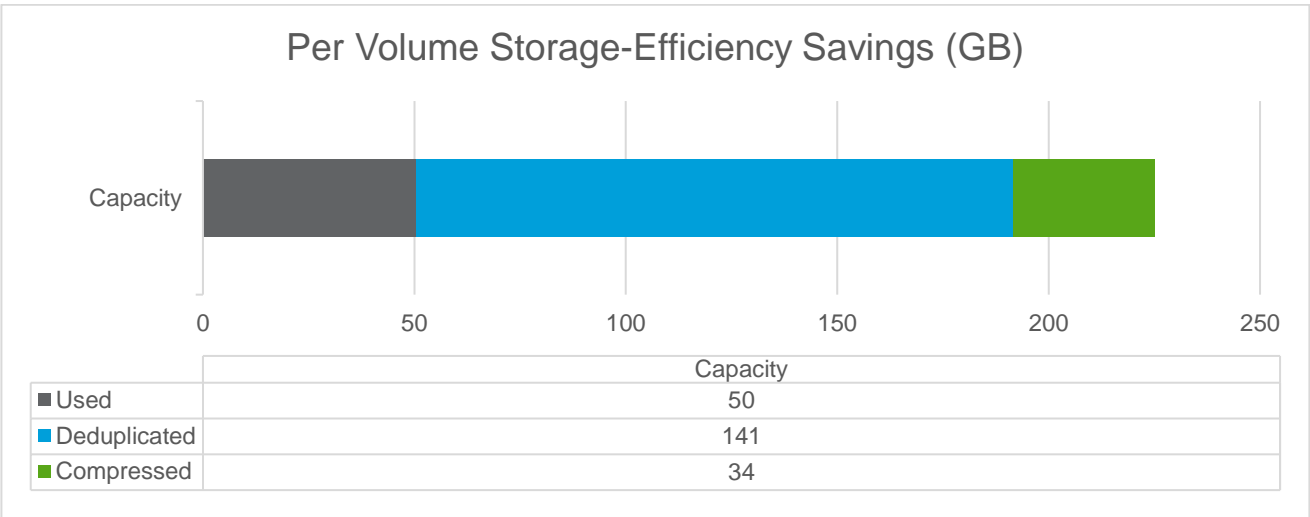
Figure 46) Read/write ratio during a refresh operation.



Storage Efficiency

After the refresh operation completed, we measured the space utilization on the volume. We had an efficiency ratio of 4.47:1. Figure 47 shows the storage savings.

Figure 47) Per-volume storage efficiency savings.



Customer Impact (Test Conclusions)

During the refresh test, the storage controller performed very well. Because refresh is considered a maintenance operation, the refresh test could be further throttled to have a lower impact on storage. This test was conducted as quickly as possible, refreshing all desktops in all pools at once. The default concurrent composer operations value (12) was used for the test. The space utilization after the operation was still excellent, despite the idle desktop traffic in the background.

6.10 Recompose 1,500 Desktop VMs

In this scenario, 1,500 VMs were recomposed that had already been logged into, the profile had been created, and the desktops had been rebooted. Table 15 lists the results for the refresh operation workload profile. The refreshed image uses a new snapshot, including a 1.1GB security patch for Windows 10.

The time to complete shows the point at which all recompose composer operations completed and all VMs reported as `Available` in Horizon View Administrator with no composer operations pending.

Table 15) Results for a recompose operation.

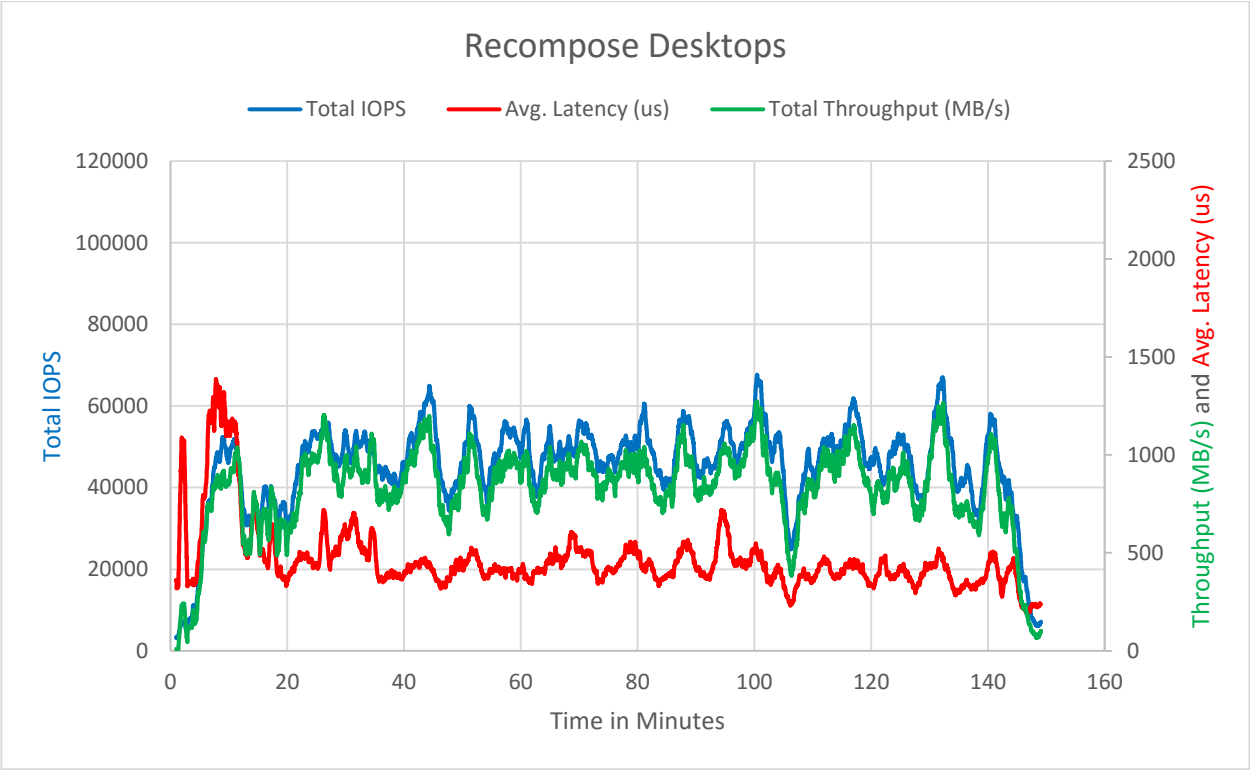
Time to Complete	Average Latency	Peak IOPS	Avg. IOPS	Peak Throughput	Avg. Throughput	Peak Storage CPU	Avg. Storage CPU
~149 min	465 us	98,404	44,842	2222MBps	826MBps	100%	51%

Note: CPU and latency measurements are based on the average across both nodes of the cluster. IOPS and throughput are based on a combined total of each.

Throughput, IOPS, and Latency

Figure 48 shows the throughput, IOPS, and latency during a recompose operation.

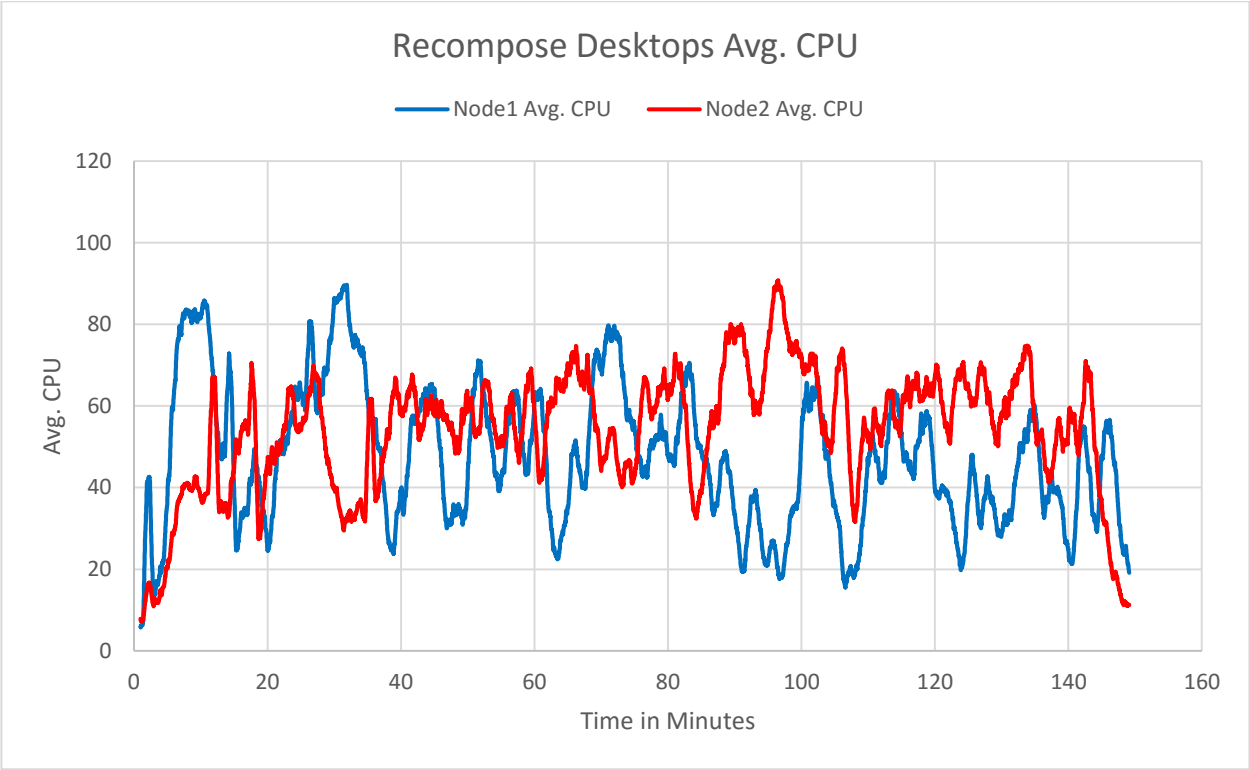
Figure 48) Throughput, IOPS, and latency during a recompose operation.



Storage Controller CPU Utilization

Figure 49 shows results for storage controller CPU utilization during a recompose operation.

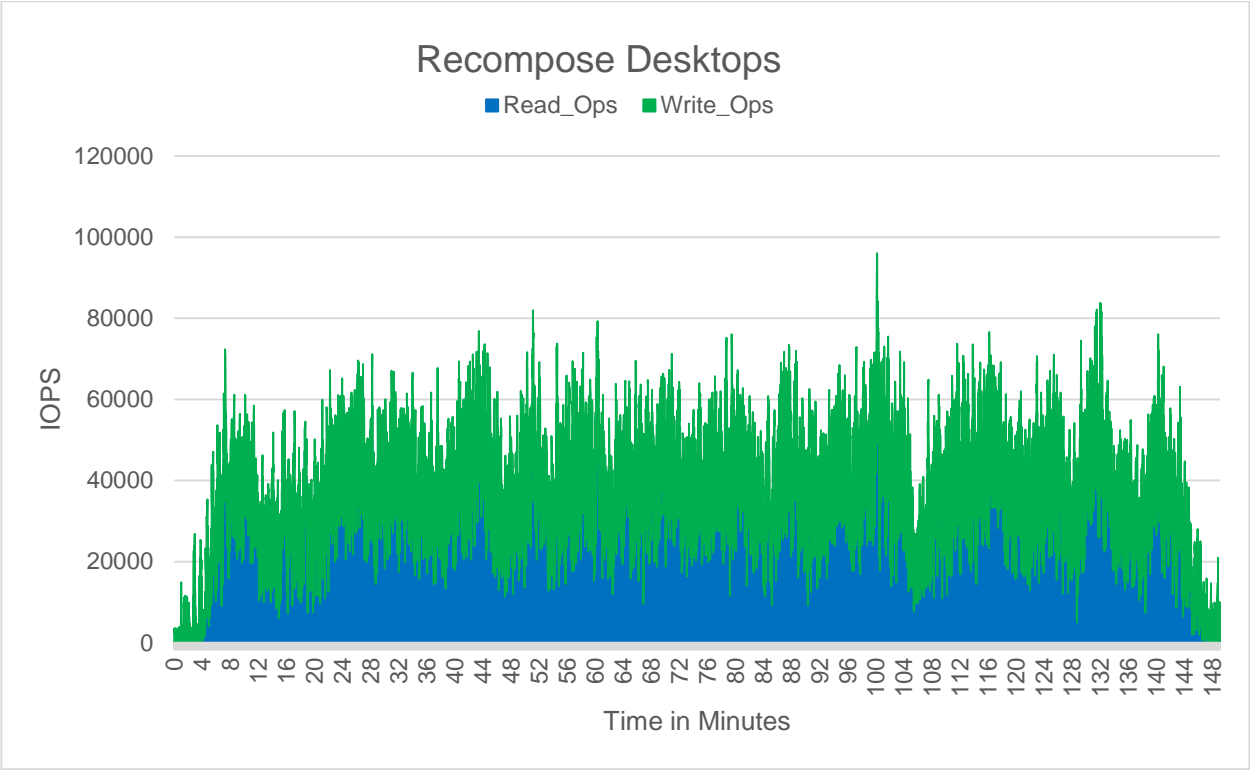
Figure 49) Storage controller CPU utilization during a recompose operation.



Read/Write IOPS

Figure 50 shows results for read and write IOPS during a recompose operation.

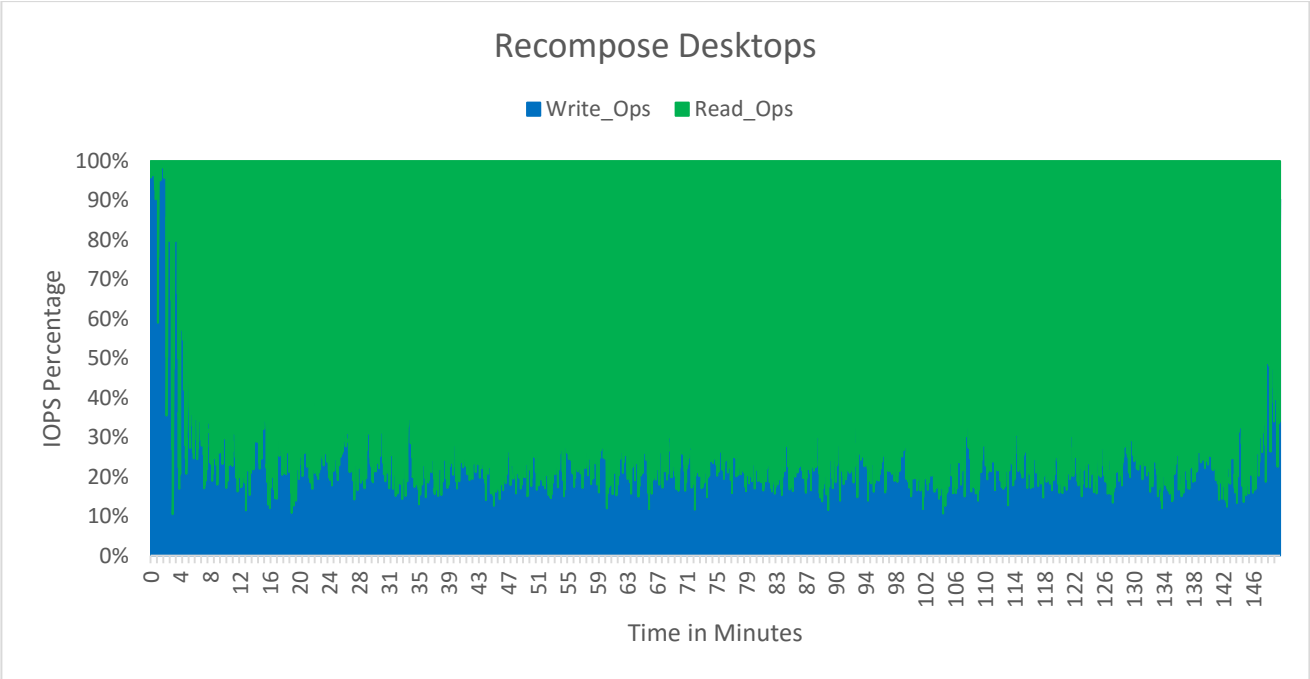
Figure 50) Read/write IOPS during a recompose operation.



Read/Write Ratio

Figure 51 shows the read-to-write ratio during a recompose operation.

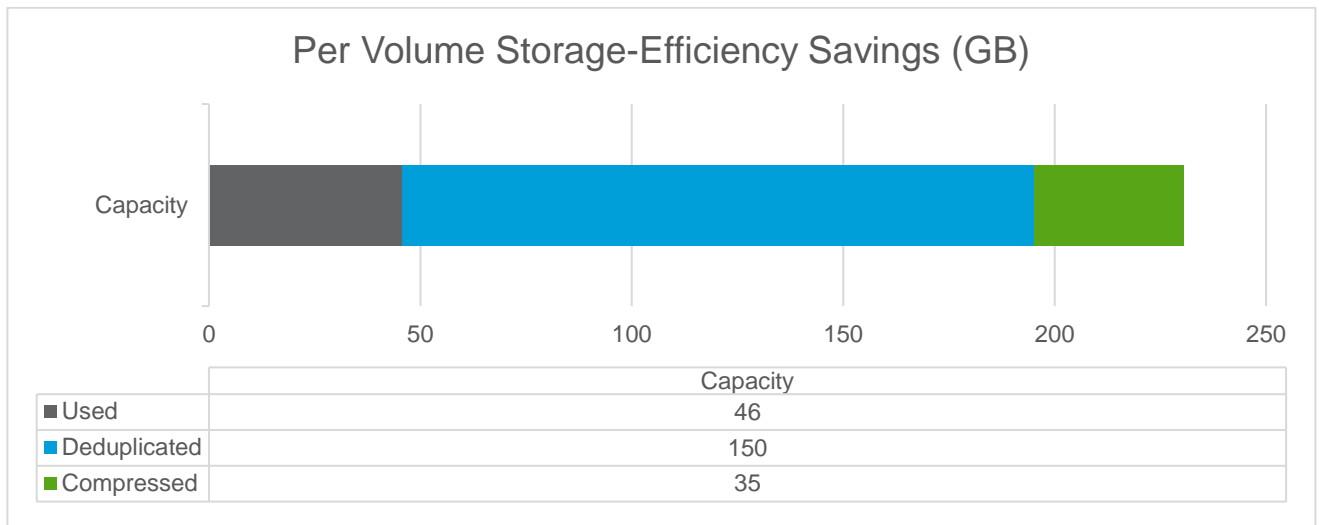
Figure 51) Read/write ratio during a recompose operation.



Space Efficiency

Following the recompose operation, space utilization on storage had an efficiency ratio of 5.05:1. Space utilization per FlexVol volume increased by 7GB, while the underlying image size increased 1.1GB. This increase means that the space per VM increase was ~30MB. Figure 52 shows the storage efficiency savings.

Figure 52) Per-volume storage efficiency savings.



Customer Impact (Test Conclusions)

During the recompose operation, the storage controller performed very well. The test was run unthrottled except by the default number of concurrent composer operations (12) in Horizon View Administrator. The test could be further throttled by adding a delay between recomposing each pool or lowering the number of concurrent operations to limit impact. The latency remained within optimal range during the test, meaning that user operations would not be affected during off-peak hours.

7 Conclusion of Testing and Validation

In all tests, the end-user login time and guest response time performance was excellent. The NetApp All Flash FAS system performed very well for the variety of tests. The system reached peak IOPS of 69,700 during Login VSI testing and averaged 50% CPU utilization on the storage node on which the tests ran. All test categories demonstrated that with the 1,500-user workload tested in this solution, the AFF8040 storage system could do significantly more while still being able to fail over in the event of a failure.

The following key findings were observed during the reference architecture testing:

- The NetApp All Flash FAS solution was able to easily meet all IOPS requirements of the 1,500-user workload (boot, login, steady state, logout) at an ultralow latency of approximately 1ms, delivering an excellent end-user experience. The storage configuration can easily support more users.
- During all login and workload scenarios, the Login VSI VSImax was not reached.

8 Data Center Efficiency with AFF A-Series

In February 2017, NetApp released its latest AFF platform, the A700s. The AFF A700s is a 4U All Flash array that can support 24 internal SSDs and can support up to 192 additional drives (8 external shelves).

Due to the dramatic change in both performance and capacity, we conducted performance tests on both the AFF8000 and AFF A-Series systems.

Capacity Density

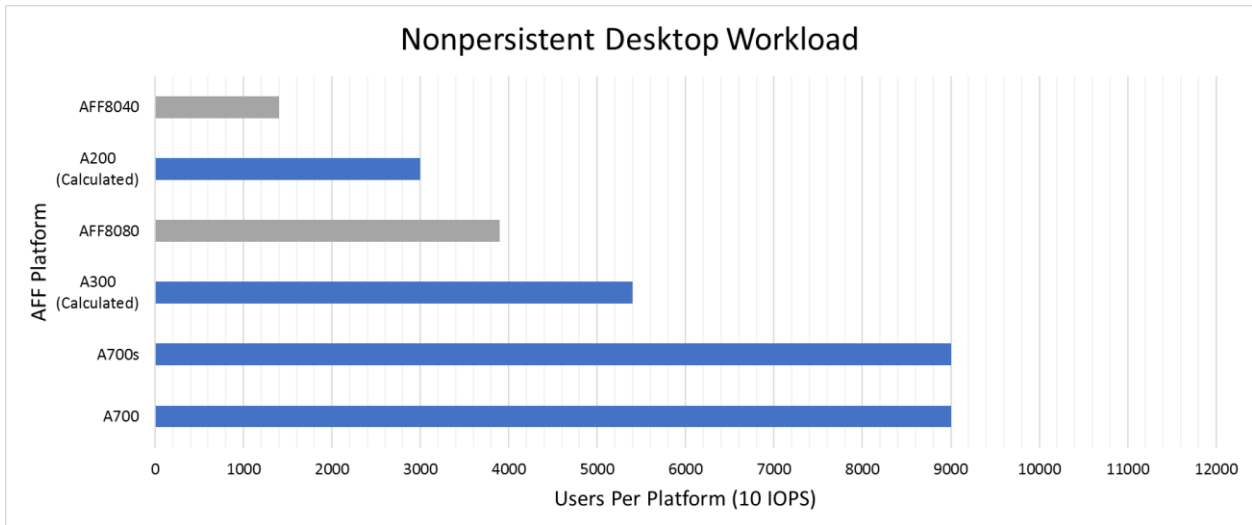
Table 16) Capacity density of A-Series systems.

Platform	SSDs (15.3TB)	Data Center Footprint	Usable Capacity	Effective Capacity (4:1)
A700s	24 internal drives	4U	287TB	1.15PB
A700s plus 8 shelves	24 internal drives 192 external drives	20U	2.6PB	10.3PB

Performance Density

Using the Tuesday Morning Workload Profile in the Appendix, we found that the A700/A700s could perform nearly 6 times the IOPS of the AFF8040. This number demonstrates that over 9,000 desktops can be supported in a 4U form factor. While we used 960GB SSDs for our testing, the 15.3TB SSDs are available and provide an effective capacity of 1.15PB.

Figure 53) User per platform density.



Note: The data was generated using the Tuesday Morning Workload Profiles with VDBench provided in the Appendix.

Figure 54) Data center density for AFF8040 and A700s.



The A700/A700s provide lower latencies and all of the feature-rich data management capabilities of ONTAP in a hyper dense form factor.

Appendix: Workload Profiles

In this section we break down the different I/O profiles for each of the different workloads in this paper. These I/O profiles were seen by [vscsiStats](#) at the ESXi layer. This information demonstrates that each workload is different, and each workload affects the storage controller differently.

These workload profiles can be used across different platforms as a standard workload for proof of concept. Many vendors just use a single block size to replicate what VDI is. As you can see from the following profiles, there are many different I/O sizes and read-write ratios during each of the test cases. Using tools such as Login VSI allows customers to generate loads that are more representative than a single I/O size. However, if an existing VDI deployment is in use, vscsiStats can be captured and replayed with vdbench to match the storage workload closely.

What this methodology does not capture is the real data and how factors such as deduplication, compression, and compaction affect the workload and storage efficiency. Nor does this capture the end-to-end server and network impact of the workloads. This is where tools such as Login VSI can bring reality to a given workload.

Boot

Here is a workload profile that matches what was observed from ESXi.

Table 17) Boot workload profile.

Operation Size	Workload Percentage	Read Percentage	Seek Percentage
512B	6	98	100
1KB	1	97	100
2KB	1	89	100
4KB	18	76	100
8KB	6	60	100
16KB	11	48	100
32KB	14	99	50
48KB	10	99	.1
>=64KB	33	100	.1

Monday Morning Login and Steady State Workload Profile

Here is a workload profile that matches what was observed from ESXi. Notice that the read percentages are higher than that of the Tuesday workload below. This difference is due to the user logging in for the first time after a reboot and the profile being read as well as the applications being loaded.

Table 18) Monday morning login and steady state workload profile.

Operation Size	Workload Percentage	Read Percentage	Seek Percentage
512B	2	84	100
1KB	1	63	100
2KB	4	52	100

4KB	39	36	100
8KB	10	40	100
16KB	11	45	100
32KB	21	95	50
48KB	1	73	.1
64KB	3	69	.1
128KB	3	61	.1
256KB	2	82	.1
512KB	3	39	.1

Tuesday Morning Login and Steady State Workload Profile

Here is a workload profile that matches what was observed from ESXi. The Tuesday workload occurs when a user has already logged in once prior to doing any work. The profiles and applications are already in memory.

Table 19) Tuesday morning login and steady state workload profile.

Operation Size	Workload Percentage	Read Percentage	Seek Percentage
512B	1	30	100
1KB	1	11	100
2KB	3	2	100
4KB	51	17	100
8KB	9	5	100
16KB	5	20	100
32KB	5	47	50
48KB	1	6	.1
64KB	19	95	.1
128KB	2	2	.1
256KB	1	4	.1
512KB	2	1	.1

References

The following references were used in this technical report:

- NetApp All Flash FAS Overview Data ONTAP 8.3.2
<http://www.netapp.com/us/media/tr-4505.pdf>

- VMware vSphere 6 on NetApp Clustered Data ONTAP
<http://www.netapp.com/us/media/tr-4333.pdf>
- FlexPod Datacenter with NetApp All Flash FAS and VMware Horizon
<http://www.netapp.com/us/media/nva-1110-fp-design.pdf>
- NetApp All Flash FAS Datasheet
<http://www.netapp.com/us/media/ds-3582.pdf>
- NetApp All Flash FAS Solution for VMware Horizon 6 and vSphere Virtual Volumes
<http://blogs.netapp.com/br/media/tr-4428.pdf>
- VMware Horizon 7 Documentation
https://www.vmware.com/support/pubs/view_pubs.html

Acknowledgements

The authors thank Eric Lacell and Scott Lane for their contributions to this document.

Version History

Version	Date	Document Version History
Version 1.0	September 2016	Initial release
Version 1.1	May 2017	Added data center efficiency section to include A700 testing and added workload profiles.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2017 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.