



Technical Report

# NetApp All Flash FAS Solution for Nonpersistent Desktops with Citrix PVS

Bhavik Desai, Chris Gebhardt, Brian Casper, Joe Scott, NetApp

September 2016 | TR-4518

## TABLE OF CONTENTS

<b>1</b>	<b>Executive Summary .....</b>	<b>5</b>
<b>2</b>	<b>Reference Architecture Objectives .....</b>	<b>6</b>
<b>3</b>	<b>Introduction .....</b>	<b>6</b>
3.1	Document Overview .....	6
3.2	NetApp All Flash FAS.....	7
3.3	ONTAP FlashEssentials .....	7
3.4	NetApp ONTAP 9.....	7
3.5	Storage Efficiency .....	8
3.6	Performance Capacity.....	9
3.7	Performance Capacity Terms.....	10
<b>4</b>	<b>Solution Infrastructure .....</b>	<b>11</b>
4.1	Hardware Infrastructure.....	12
4.2	Software Components.....	13
4.3	Login VSI.....	14
4.4	PVS Desktops .....	14
<b>5</b>	<b>Storage Design.....</b>	<b>15</b>
5.1	Aggregate Layout.....	15
5.2	Volume Layout .....	15
5.3	NetApp Virtual Storage Console for VMware vSphere .....	16
<b>6</b>	<b>Testing and Validation: Nonpersistent PVS Desktops .....</b>	<b>16</b>
6.1	Test Results Overview .....	16
6.2	Storage Efficiency .....	17
6.3	Boot Storm Test Using XenDesktop.....	17
6.4	Boot Storm Test Using XenDesktop During Storage Failover .....	20
6.5	Steady-State Login VSI Test .....	23
	<b>Conclusion.....</b>	<b>30</b>
	<b>Appendix: Workload Profiles .....</b>	<b>30</b>
	Boot .....	30
	Login and Workload.....	31

<b>References .....</b>	<b>31</b>
<b>Acknowledgements .....</b>	<b>32</b>
<b>Version History.....</b>	<b>32</b>

## LIST OF TABLES

Table 1) Nonpersistent desktop test results.....	6
Table 2) Hardware components of server categories. ....	12
Table 3) Solution software components.....	13
Table 4) Virtual desktop configuration. ....	14
Table 5) Test results overview. ....	17
Table 6) Efficiency results for each FlexVol volume. ....	17
Table 7) Results for PVS boot storm. ....	18
Table 8) Results for full-clone boot storm during storage failover. ....	20
Table 9) Results for Login VSI login and workload. ....	23
Table 10) Results for Login VSI login and workload during storage failover.....	26

## LIST OF FIGURES

Figure 1) Visual representation of inline compression and data compaction. ....	8
Figure 2) Advanced drive partitioning v2. ....	9
Figure 3) Performance capacity graphic. ....	10
Figure 4) Solution infrastructure for nonpersistent desktops.....	12
Figure 5) SSD layout. ....	15
Figure 6) Volume layout for nonpersistent desktops.....	16
Figure 7) Storage efficiency savings.....	17
Figure 8) Throughput and IOPS for PVS boot storm. ....	18
Figure 9) Storage controller CPU utilization for PVS boot storm.....	19
Figure 10) Read/write IOPS for PVS boot storm.....	19
Figure 11) Read/write ratio for PVS boot storm. ....	20
Figure 12) Throughput, IOPS, and latency for PVS boot storm during storage failover.....	21
Figure 13) Storage controller CPU utilization for boot storm during storage failover. ....	21
Figure 14) Read/write IOPS for full-clone boot storm during storage failover.....	22
Figure 15) Read/write ratio for full-clone boot storm during storage failover.....	22
Figure 16) VSImax results for Login VSI login and workload.....	23
Figure 17) Scatterplot of Login VSI login times.....	24
Figure 18) Throughput, IOPS, and latency for Login VSI login and workload.....	24
Figure 19) Storage controller CPU utilization for Login VSI login and workload. ....	25
Figure 20) Read/write IOPS for Login VSI login and workload. ....	25

Figure 21) Read/write ratio for Login VSI login and workload.....	26
Figure 22) VSImax results for Login VSI login and workload during storage failover.....	27
Figure 23) Scatterplot of login times during storage failover.....	27
Figure 24) Throughput, IOPS, and latency for Login VSI login and workload during storage failover.....	28
Figure 25) Storage controller CPU utilization for Login VSI login and workload during storage failover.....	28
Figure 26) Read/write IOPS for Login VSI login and workload during storage failover.....	29
Figure 27) Read/write ratio for Login VSI login and workload during storage failover.....	29

# 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 Citrix. Citrix 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 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 solution, with the AFF8000 platform, 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 (ONTAP) 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\$39 per desktop for storage when deploying at scale. This figure includes the cost of NetApp 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 portfolio of All Flash FAS platforms.

With Citrix and NetApp, companies can accelerate the virtual desktop end-user experience by using NetApp All Flash FAS storage for Citrix XenDesktop. NetApp All Flash FAS storage, powered by the AFF8000 system, 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 the users working and lessens the risk to the business. ONTAP has multiple built-in features to help improve availability, such as active-active high availability (HA) and nondisruptive operations, and to seamlessly move data in the storage cluster without user impact.

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 in order 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 the users ever noticing.

## 2 Reference Architecture Objectives

The reference architecture described in this document is a 1,500-desktop design using NetApp AFF8040, Citrix XenDesktop 7.6, Citrix Provisioning Services (PVS) 7.6, and VMware vSphere 6.0. It validates a nonpersistent 1,500-desktop architecture created with Citrix PVS. 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 All Flash FAS 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
Boot storm test (50 Citrix XenDesktop concurrent power-on operations)	~30 min	13,785	123MBps	0.24ms
Boot storm test during storage failover (50 Citrix XenDesktop concurrent power-on operations)	~30 min	14,700	51.9MBps	0.37ms
Login VSI test login and workload	22.9 sec/VM	26,276	392MBps	0.24ms
Login VSI test login and workload during failover	26.38 sec/VM	31,409	388MBps	0.43ms

## 3 Introduction

This section provides an overview of the NetApp All Flash FAS solution for Citrix XenDesktop and PVS, 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 Citrix XenDesktop and PVS 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 tests completed.

The testing included the following criteria:

- Launch and boot storm test of 1,500 desktops (with and without storage node failover) using Citrix XenDesktop
- Login VSI login and steady-state workload test using benchmark mode in 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 (ITaaS), a model that 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 two-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 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 933x 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 premises or to the cloud.
- Quality of service (QoS) capability safeguards service-level objectives in multiworkload and multitenant environments.

### 3.4 NetApp ONTAP 9

ONTAP 9 is a major advance in the industry's leading enterprise data management software. It 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. It 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 impact to 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 further reduces the disk space required and 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.



### 3.6 Performance Capacity

NetApp or partner system engineers should perform initial sizing using NetApp, OS, and application vendor best practices and typically use NetApp internally available tools to determine optimal solution sizing. After initial sizing, NetApp recommends basing all incremental performance sizing, monitoring, capacity planning, and workload placement on available performance capacity. This is a departure from NetApp's previous recommendation, which was to size workloads to use less than 50% CPU utilization. This previous recommendation had the benefit of being easy to make and understand but is far less nuanced and more prone to guesswork than our current recommendation to use performance capacity planning for sizing.


NetApp's best practice for sizing a SAN AFF business-processing (AFBP) environment is to use performance capacity to size each node to less than 50% of performance capacity on each controller in an HA pair. By sizing this way, you can maintain acceptable low latency in the event of a takeover. The cost of this approach is that you sacrifice a little of the steady-state top-level performance.

Before getting started, note that the discussion is based on the assumption of a transactional workload that uses IOPS and latency as principal metrics. With that said, let's define some terms and then finish the discussion with a real-world example of how OnCommand® Performance Monitor (OPM) 7 can be used to determine the total performance capacity and performance capacity available on a given set of controllers.

For more information about performance capacity, see [TR-4515: Best Practices for AFF Business-Processing SAN Workloads](#).

### 3.7 Performance Capacity Terms

Figure 3) Performance capacity graphic.



The optimal point identifies the total performance capacity a resource has before latency increases more quickly than IOPS do.

The optimal point can be determined by either:

- Finding the “knee” in the performance curve, where an increase in utilization leads to more rapidly increasing latency. Generally, performance curves are fairly flat at lower utilizations, but latency increases as the number of IOPS increases. There is a point in the curve where the rate of increase in latency starts accelerating more rapidly than the increase in the number of IOPS being served.
- An alternative to the knee of the curve method is to target a specific latency value and draw a horizontal line at that latency value. The optimal point is the point where the IOPS curve intersects the latency threshold line you have just drawn.

Total performance capacity (or optimal point) = performance capacity used + performance capacity.

#### Performance Capacity Used

Performance capacity used can be defined as the amount of the useful capacity of the resource that has been consumed. As noted earlier, the remaining useful capacity is performance capacity available.

Performance capacity used = optimal point – performance capacity available.

#### Performance Capacity Available

Performance capacity available is derived by subtracting the performance capacity used from the total performance capacity (or simply performance capacity) of a node. The performance capacity is identified by the optimal point. From Figure 3, we see that:

total performance capacity (or optimal point) = performance capacity used + performance capacity available.

#### Operating Point

The operating point is the location on the curve where your resource is currently operating. This point illustrates the number of IOPS served at a given latency.

## Unsafe Zone

The unsafe zone in a performance context can be defined as the portion of the performance graph that is above and to the right of the optimal point. Performance in these areas has higher latencies, and little increases in IOPS yield larger increases in latency.

## Safe Zone

The safe zone is the area of the performance graph that is below and to the left of the optimal point. This is the area of the graph where you see the highest throughput relative to latencies and is the area within which you want to operate. In order to maintain consistent low-latency high performance, the operating point needs to stay inside the safe zone.

## Differences Between Total Performance Capacity and Utilization

Performance capacity used incorporates both utilization and latency. Performance capacity of up to 100% is acceptable; however, performance capacity can rise above 100%, at which point a little more output increases latency much more quickly. Utilization as a measure is fine for workloads that are not sensitive to latency such as batch workloads.

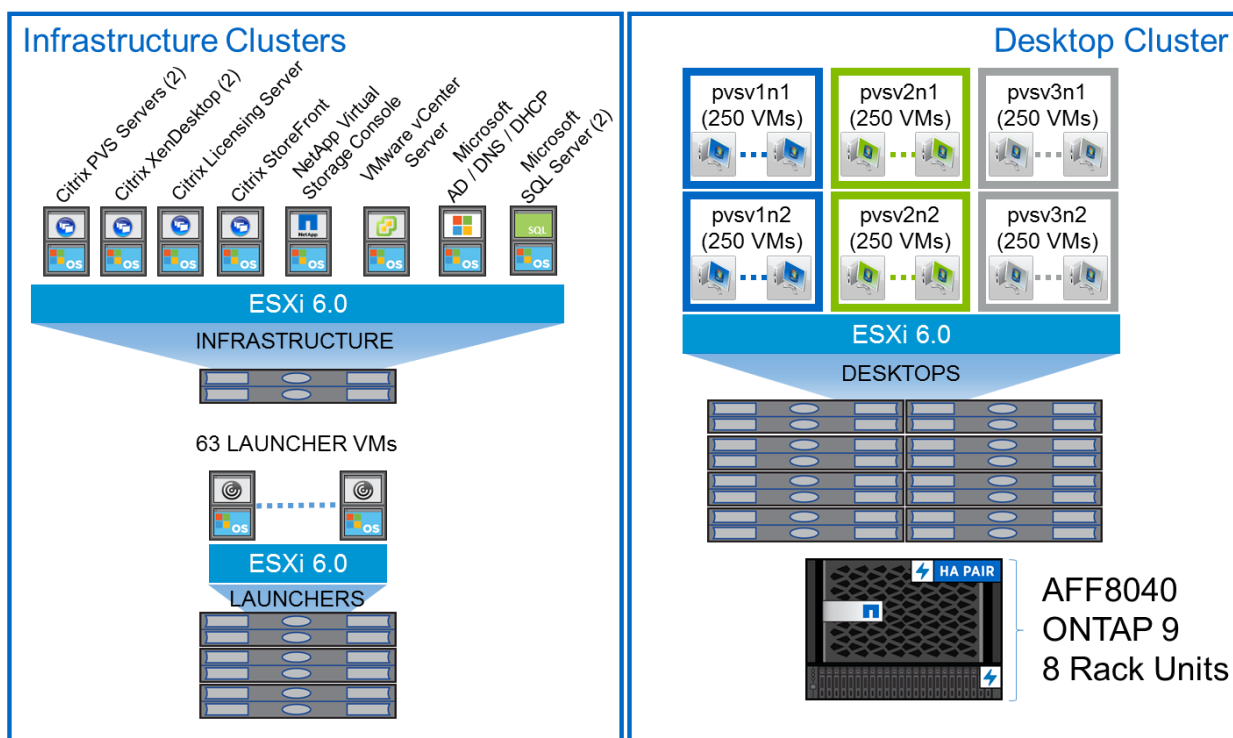
## 4 Solution Infrastructure

This section describes the software and hardware components of the solution. Figure 4 shows the solution infrastructure, which includes two Citrix XenDesktop VMs, two Citrix StoreFront VMs, one Citrix licensing server, and two Citrix provisioning servers.

All write cache volumes hosting write cache files and vDisks for 1,500 users were hosted on both storage controllers in the HA pair. Scale performance tests, such as boot, login, and steady-state tests, were performed to help validate that this All Flash FAS array is capable of hosting 1,500 users with the standard configuration of two storage controllers.

Figure 4 shows the solution infrastructure for nonpersistent desktops.

Figure 4) 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 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, using intelligent sizing to account for infrastructure VMs, or putting infrastructure VMs on an existing or separate NetApp All Flash FAS storage system. For this lab validation, we used a separate NetApp All Flash 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
<b>Infrastructure and Launcher Servers</b>	
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
<b>Desktop Servers</b>	

Hardware Components	Configuration
Server quantity	14 Fujitsu Primergy RX2540 M1
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
<b>Storage</b>	
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
<b>NetApp All Flash FAS</b>	
NetApp ONTAP	9.0 RC
NetApp VSC for VMware	6.2
<b>VMware Software</b>	
VMware ESXi	6.0.0, 2494585
VMware vCenter Server (VCSA)	6.0.0
<b>Citrix Software</b>	
Citrix XenDesktop	7.6
Citrix Provisioning Server	7.6
Citrix Licensing Server	11.12.1 build 14008
Citrix StoreFront Server	2.6.0.5031
Citrix XenDesktop SQL Server	11.0.3128 (64-bit)
<b>Workload Generation Utility</b>	
Login VSI Professional	Login VSI 4.1.4.2
<b>Database Server</b>	
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 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 the proof (vendor independent, industry standard, and easy to understand) to innovative technology vendors to demonstrate the power, the scalability, and the gains of their solutions.

Login VSI-based test results are published in [technical white papers](#) and 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 PVS Desktops

The PVS VM master device 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](#). The PVS imaging wizard was then used to create a base vDisk image from the master target device. For PVS VMs, a PVS write cache of 6GB was used when creating them through the XenDesktop Setup wizard.

Table 4) Virtual desktop configuration.

Desktop	Configuration
<b>Desktop VM</b>	
VM quantity	1,500
VM hardware version	11
vCPU	2 vCPUs
Memory	2GB
Hard disk size	32GB
Hard disk type	Thin provisioned
<b>Desktop Software</b>	
Guest OS	Windows 7 (64-bit)
VM hardware version	ESXi 6.0 and later (VM version 11)
VMware tools version	10.0.0.3000743
Microsoft Office	2013 version 15.0.4569.1506
Microsoft .NET Framework	4.5.1
Adobe Acrobat Reader	11.0.00

Desktop	Configuration
Adobe Flash Player	11.5.502.146
Java	7.0.130
Doro PDF	1.82
Citrix XenDesktop Agent	7.6
Citrix Receiver	14.1.200.13
Login VSI target software	4.1.4.2

## Guest Optimization

Guest OS optimizations were applied to the template VMs used in this reference architecture. Also, as part of creating the vDisk, the PVS Optimization tool was used to optimize the target with all values selected.


## 5 Storage Design

This section provides an overview of the storage design, the aggregate and volume layout, and VSC.

### 5.1 Aggregate Layout

In this reference architecture, we used advanced drive partitioning to partition the 24 800GB SSDs across the two 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 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.

**Note:** When using the NetApp VSC to set the NFS host optimizations, it sets the NFS.MaxQueueDepth to 64 for all versions of ESX. The default setting for vSphere 6 is 128, and the default setting in 5.x is 4,294,967,295. NetApp recommends using 64 as the setting when using a single-processor FAS. When using All Flash FAS, NetApp recommends the NFS.MaxQueueDepth be set to at least 128 if not higher to eliminate any queuing bottlenecks. The exception to this recommendation is when a customer has both a FAS2500 system as well as an All Flash FAS system connected to the same ESX cluster. If the setting were to remain at 64, the customer would not get the maximum performance from the All Flash FAS system. If the customer sets the NFS.MaxQueueDepth to a higher value, there is a risk of datastore disconnects on the FAS2500 system. Careful consideration should be given before making any changes to NFS.MaxQueueDepth.

```
Get-VMHost | Sort Name | Get-AdvancedSetting -Name NFS.MaxQueueDepth | Set-AdvancedSetting -Value 256
-Confirm:$false
```

## 6 Testing and Validation: Nonpersistent PVS Desktops

This section describes the testing and validation of nonpersistent Citrix PVS 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
Boot storm test (50 Citrix XenDesktop concurrent power-on operations)	~30 min	13,785	123MBps	0.24ms
Boot storm test during storage failover (50 Citrix XenDesktop concurrent power-on operations)	~30 min	14,700	51.9MBps	0.37ms
Login VSI test login and workload	22.9 sec/VM	26,276	392MBps	0.24ms
Login VSI test login and workload during failover	26.38 sec/VM	31,409	388MBps	0.43ms

**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 Storage Efficiency

During the tests, all space efficiency technologies were turned on. This included inline compression, inline deduplication, inline zero elimination, and compaction. On average after PVS VM creation and boot, a 3.3:1 efficiency ratio, or 70% storage efficiency, was observed at the FlexVol level. At the aggregate level, which accounts for physical blocks stored, we saw an additional 6% savings due to compaction in addition to the 70% seen at the FlexVol level. Figure 7 shows the significant difference in storage efficiency savings.

Figure 7) 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 a storage-reduction technology and therefore was not reported on. Table 6 lists the efficiency results from the testing.

Table 6) Efficiency results for each FlexVol volume.

Capacity Used	Total Savings	Dedupe Savings	Compression Savings
22.9GB	52.4GB	38.7GB	13.7GB

## 6.3 Boot Storm Test Using XenDesktop

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 XenDesktop, which might occur, for example, after maintenance activities and server host failures.

This test was performed by powering on all 1,500 VMs from within XenDesktop and observing when the status of all VMs in Citrix XenDesktop changed to `Registered`.

Table 7 lists the boot storm data that was gathered.

Table 7) Results for PVS boot storm.

Time	Average Latency	Peak IOPS	Avg. IOPS	Peak Throughput	Avg. Throughput	Peak Storage CPU	Avg. Storage CPU
~30 min	0.24ms	13,785	5,088.9	143MBps	39.62MBps	33%	9%

**Note:** All desktops had the status `Registered` in Citrix XenDesktop.


**Note:** The time taken to boot the VMs was a result of XenDesktop and PVS 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 8 shows testing results for throughput and IOPS for a PVS boot storm.


Figure 8) Throughput and IOPS for PVS boot storm.



## Storage Controller CPU Utilization

Figure 9 shows testing results for storage controller CPU utilization for a PVS boot storm.


Figure 9) Storage controller CPU utilization for PVS boot storm.



## Read/Write IOPS

Figure 10 shows testing results for read and write IOPS for a PVS boot storm.


Figure 10) Read/write IOPS for PVS boot storm.



## Read/Write Ratio

Figure 11 shows the read and write ratio for a PVS boot storm.

Figure 11) Read/write ratio for PVS boot storm.



## Customer Impact (Test Conclusions)

Although the total time it took to boot the VMs was approximately 30 minutes, the majority of that time period was the VM registering with the PVS server and then with XenDesktop. Storage had plenty of performance left, as seen by the storage CPU utilization used on the controllers. The bottleneck in this test was the PVS servers, which had CPU utilization at 100%. More PVS servers could have been added to distribute the load and improve the boot and login time.

## 6.4 Boot Storm Test Using XenDesktop 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 XenDesktop.”

Table 8 shows the data that was gathered for the boot storm during storage failover.

Table 8) Results for full-clone boot storm during storage failover.

Time	Average Latency	Peak IOPS	Avg. IOPS	Peak Throughput	Avg. Throughput	Peak Storage CPU	Avg. Storage CPU
~30 min	0.38ms	14,700	4,223.95	51MBps	18.28MBps	40%	13%


**Note:** All desktops had the status of `Registered` in Citrix XenDesktop.

**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 12 shows throughput, IOPS, and latency for a full-clone boot storm during storage failover.


Figure 12) Throughput, IOPS, and latency for PVS boot storm during storage failover.



## Storage Controller CPU Utilization

Figure 13 shows storage controller CPU utilization for a full-clone boot storm during storage failover.


Figure 13) Storage controller CPU utilization for boot storm during storage failover.



## Read/Write IOPS

Figure 14 shows read and write IOPS for a full-clone boot storm during storage failover.


Figure 14) Read/write IOPS for full-clone boot storm during storage failover.



## Read/Write Ratio

Figure 15 shows read and write ratios for a full-clone boot storm during storage failover.

Figure 15) Read/write ratio for full-clone boot storm during storage failover.



## Customer Impact (Test Conclusions)

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

## 6.5 Steady-State Login VSI Test

This section describes test objectives and methodology and provides results from steady-state Login VSI testing.

### Test Objectives and Methodology

The objective of this test was to run a Login VSI knowledge worker workload in benchmark mode 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 48-minute period. This value cannot be changed in benchmark mode. The test was run for an additional 10 minutes, before which the first logoff command was sent.

### Login VSI Login and Workload Test

The desktops were configured such that every time a user connects, that user gets a fresh/random desktop. This makes the Login VSI test an initial login test, which does profile creation because the user is logging on to the desktop for the first time. In this scenario, 1,500 users logged in and downloaded the Login VSI content package containing user data to be used by Login VSI during the test. After user and profile data was created, application binaries and libraries had to be read from a disk for doing the test. Table 9 shows the results.

Table 9) 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
22.9 sec/VM	0.25ms	26,276	13,024	392MBps	202MBps	87%	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 16 shows VSImax results for Login VSI login and workload testing.


Figure 16) VSImax results for Login VSI login and workload.



## Desktop Login Time

Average desktop login time was 22.9 seconds, which is considered a good login time for a first login when users download a lot of data. Figure 17 shows a scatterplot of the Login VSI times.


Figure 17) Scatterplot of Login VSI login times.



## Throughput, IOPS, and Latency

Figure 18 shows results for throughput, IOPS, and latency for Login VSI login and workload testing.

Figure 18) Throughput, IOPS, and latency for Login VSI login and workload.




## Storage Controller CPU Utilization

Figure 19 shows results for storage controller CPU utilization for Login VSI login and workload testing.




Figure 19) Storage controller CPU utilization for Login VSI login and workload.



## Read/Write IOPS

Figure 20 shows results for read and write IOPS for Login VSI login and workload testing.


Figure 20) Read/write IOPS for Login VSI login and workload.



## Read/Write Ratio

Figure 21 shows the read to write ratio for Login VSI login and workload testing.

Figure 21) 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 there is sufficient host infrastructure. The Login VSI test during storage failover described in the following section reinforces that point.

## Login VSI Login and Workload During Storage Failover Test

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 10 lists the results for Monday morning login and workload during storage failover.

Table 10) 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
26.3 sec/VM	0.43ms	31,409	12,372	388MBps	193.3MBps	97%	50%

**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 22 shows the VSImax results for Login VSI login and workload during storage failover.


Figure 22) VSImax results for Login VSI login and workload during storage failover.



## Desktop Login Time

Average desktop login time was 26.3 seconds, which is considered a good login time, especially considering the amount of data each user had to download during a failover situation. Figure 23 shows a scatterplot of the login times during storage failover.


Figure 23) Scatterplot of login times during storage failover.



## Throughput, IOPS, and Latency

Figure 24 shows throughput, IOPS, and latency for Login VSI login and workload during storage failover.


Figure 24) Throughput, IOPS, and latency for Login VSI login and workload during storage failover.



## Storage Controller CPU Utilization

Figure 25 shows results for storage controller CPU utilization for Login VSI login and workload during storage failover testing.


Figure 25) Storage controller CPU utilization for Login VSI login and workload during storage failover.



## Read/Write IOPS

Figure 26 shows results for read and write IOPS for Login VSI login and workload during storage failover testing.


Figure 26) Read/write IOPS for Login VSI login and workload during storage failover.



## Read/Write Ratio

Figure 27 shows the read to write ratio for Login VSI login and workload during storage failover testing.

Figure 27) Read/write ratio for Login VSI login and workload during storage failover.



## Customer Impact (Test Conclusions)

During the Login VSI test during storage failover, the storage controller performed very well. The CPU utilization averaged at 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.

## Conclusion

In all tests the end-user login time and guest response time performance were excellent. The NetApp All Flash FAS system performed very well for the variety of tests, reaching peak IOPS of 31,409 during Login VSI test while averaging 50% CPU utilization on the storage node on which the tests were running. 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 very easily meet all IOPS requirements of the 1,500-user workload (boot, login, steady state, logout) at an ultralow latency of approximately 1ms, delivering 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.

## 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 (<https://communities.vmware.com/docs/DOC-10095>) 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 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 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 the ESXi layer.

Block Size (KB)	% Read	% Random	Workload Percentage
512	18%	100%	6%
1,024	12%	100%	2%
2,048	6%	100%	1%
4,096	2%	10%	64%
8,192	1%	10%	10%

Block Size (KB)	% Read	% Random	Workload Percentage
16,384	1%	10%	9%
32,768	2%	0.1%	6%
>=65,536	1%	0.1%	2%

## Login and Workload

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

Block Size (KB)	% Read	% Random	Workload Percentage
512	4%	100%	4%
1,024	1%	100%	2%
2,048	2%	100%	1%
4,096	2%	25%	44%
8,192	4%	25%	4%
16,384	5%	25%	3%
32,768	3%	25%	41%
>=65,536	1%	0.1%	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>
- TR-4138: Citrix XenDesktop on NetApp Storage Solution Guide  
<http://www.netapp.com/as/media/tr-4138.pdf>
- TR-4342: NetApp All Flash FAS Solution for Persistent and Nonpersistent Desktops with Citrix XenDesktop and XenApp  
<http://www.netapp.com/us/media/tr-4342.pdf>
- NetApp All Flash FAS Datasheet  
<http://www.netapp.com/us/media/ds-3582.pdf>

- Citrix XenDesktop 7.6 Blueprint  
[http://www.citrix.com/content/dam/citrix/en\\_us/documents/products-solutions/xendesktop-deployment-blueprint.pdf?accessmode=direct](http://www.citrix.com/content/dam/citrix/en_us/documents/products-solutions/xendesktop-deployment-blueprint.pdf?accessmode=direct)
- Citrix Product Document on XenApp 7.6 and XenDesktop 7.6  
<http://docs.citrix.com/en-us/xenapp-and-xendesktop/7-6.html>
- FlexPod Datacenter with Citrix XenDesktop/XenApp 7.7 and VMware vSphere 6.0 for 5000 Seats  
[http://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/cisco\\_ucs\\_xd77esxi60u1\\_flexpod.html](http://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/cisco_ucs_xd77esxi60u1_flexpod.html)

## Acknowledgements

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

## Version History

Version	Date	Document Version History
Version 1.0	June 2016	Original document
Version 1.1	September 2016	Added NFS.MaxQueueDepth recommendations, performance capacity and AFBP, 8-hour test determining performance capacity and 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 © 1994–2016 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, Go Further, Faster, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Fitness, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANshare, SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, SnapCopy, Snap Creator, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, SolidFire, StorageGRID, Tech OnTap, Unbound Cloud, vFiler, WAFL, and other names are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

TR-4518-0916