Technical Report

Quantifying the Value of Running VMware on NetApp

Vaughn Stewart, Dan Chilton, NetApp
July 2010 | TR-3856

WORKFLOW COMPARISON OF NETAPP UNIFIED STORAGE VS. TRADITIONAL STORAGE ARRAY

Virtualization is changing everything about the structure of information technology. Traditional storage arrays with historic architectures are less effective in unlocking the value in virtualized environments. NetApp® unified storage solutions are at the forefront of providing solutions that empower customers to realize all of the benefits of a virtualized infrastructure. NetApp solutions enable customers to quickly deploy solutions that are storage efficient and provide improved performance.

TABLE OF CONTENTS

1 INTRODUCTION

Virtualization is changing everything about the structure of information technology. NetApp unified storage solutions are at the forefront of providing solutions that empower and enable customers to realize all of the benefits of a virtualized infrastructure. With NetApp solutions, customers have the power to quickly deploy solutions that are storage efficient and provide improved performance.

To demonstrate the value of running a virtual infrastructure on NetApp's unified virtual storage architecture, we conducted a series of tests to compare NetApp features and performance to those of a competitive, traditional storage array. We concentrated on the tasks that administrators use on a daily basis to manage their virtual infrastructures. We believe that our tests show VMware® deployments on traditional storage arrays require choices among storage efficiency, high availability, and performance. With NetApp solutions, customers obtain space-efficient rapid deployment of virtual machines (VMs), high availability, and improved performance. In addition, efficiency in the workflow, design, and performance of NetApp VMware virtualized solutions provides value in reduced capital expenditures and operational expenditures.

This technical report details the tests we conducted, which support our conclusions.

KEY BENEFITS

The tests that we conducted reveal the gains of running a virtual infrastructure on NetApp's unified virtual storage architecture, as compared to a traditional storage array. Our tests show that implementing VMware on NetApp provides the following key benefits:

- **95% reduction in data store provisioning time.** The NetApp unified architecture required 5 minutes and 38 seconds to provision, as opposed to 1 hour and 45 minutes for the traditional storage array.
- **97% reduction in VM deployment provisioning time**. Our tests showed that using the provisioning and cloning capabilities of the Virtual Storage Console (VSC) to create 80 VMs took 16 minutes and 56 seconds. This includes the time to customize each of the 80 VMs. This contrasts with the traditional storage array 80 VM deployment, which took 10 hours and 5 minutes.
- **NetApp provides 99% gain in usable capacity because the solution requires 99% less space.** The traditional storage array required 1.01TB for 80 VMs, in contrast to the NetApp FlexClone® solution, which required only 11.26GB.
- **NetApp performance compared to traditional storage array.** NetApp performance, as compared to the traditional storage array, was 30% better than RAID 6, 25% better than RAID 5, and 10% better than RAID 10.
- **Effect of Flash Cache on performance.** Upon adding the Flash Cache to the NetApp solution, we found the performance to be 52% better than the peak traditional storage array performance using RAID 10.

2 TEST METHODOLOGY

We compared the NetApp solution to the traditional storage array by conducting a series of tests that address common administrative tasks routinely performed in virtualized environments. These tasks included provisioning shared storage as well as creating or cloning large numbers of VMs. At a high level, we conducted the following test cases for both the NetApp and traditional storage array configurations:

1. Measure the amount of time required to provision shared data stores to vSphere™ hosts.
2. Measure the amount of time required to provision 80 VMs.
3. Review the data protection options available to virtualized environments.
4. Evaluate the storage efficiency technologies and measure the actual storage savings provided.
5. Measure the performance using realistic workloads when combining storage efficiency and data protection technologies.

See Table 1 and Table 2 for a comparison of the storage technologies used for testing.

Table 1) Storage technologies used for testing.

Traditional Midtier Fibre Channel Storage Array	NetApp FAS3140 Midtier Unified Storage Array
• 2 x 8GB of storage controller cache • 152 x 146GB, 15K RPM Fibre Channel (FC) disk drives • Dual storage controllers	• 2 x 4GB of storage controller cache • 40 x 450GB, 15K RPM SAS disk drives • 2 x 256GB Flash Cache expansion modules

Table 2) Storage networking technologies used for testing.

Traditional Midtier Fibre Channel Storage Array	NetApp FAS3140 Midtier Unified Storage Array
• 4Gb Fibre Channel network • Brocade 200E FC switch	• GbE Ethernet network • Cisco 4948 Ethernet switch

Figure 1 shows the details of the NetApp and traditional storage array configurations used for the tests. In both cases, we configured a VMware vSphere data center as follows:

- 4 x IBM x3650 servers configured with 8 Intel® Xeon™ E5430 CPUs at 2.66GHz and 36GB memory running vSphere 4 update 1
- 80 VMs running Windows® Server 2003 and SP2; each VM configured with 1GB memory and 1 virtual CPU

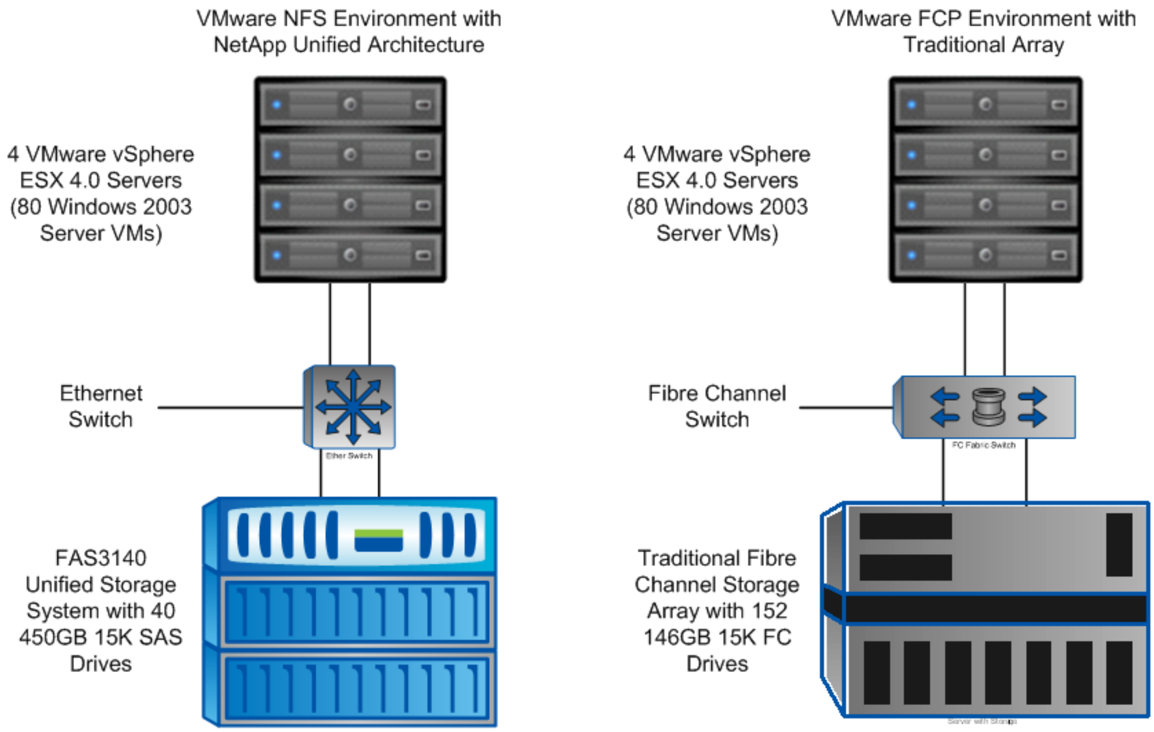

Figure 1) NetApp and traditional storage array test configurations.

3 RAID CHOICES

In VMware environments, shared data stores are commonly terabytes in size and host many VMs, which generally are sized from 10GB to 50GB. Critical requirements for virtualized environments are high availability, performance, and storage efficiency.

Traditional storage arrays traditionally offer three parity protection schemes: RAID 5, RAID 6, and RAID 10. Each parity protection scheme offers benefits and limitations. When considering these options, administrators must make trade-offs based upon availability, performance, and storage efficiency requirements. When performance and storage efficiency are the key concerns with a tolerance for reduced data protection, RAID 5 is recommended. RAID 6 is recommended for environments that prioritize data protection and storage efficiency but are willing to sacrifice some degree of performance. When performance and data protection are the key requirements and the user can accept poor storage efficiency, RAID 10 is the recommended option. It is clear that customers using traditional storage arrays must choose a parity protection model based on which requirements are most important to them. These are difficult choices, particularly in a VMware environment, where the inability to survive a double disk failure can mean disaster for multiple applications.

NetApp unified data storage offers one parity protection scheme: RAID-DP®. RAID-DP fulfills the requirements for high availability, performance, and storage efficiency.

See Table 3 for a comparison of the different data protection schemes.

Table 3) Characteristics of different data protection schemes.

Data Protection Scheme	Characteristics
Traditional storage RAID 5	• Provides greater uptime than no data protection by surviving single disk failures • Performance is moderately affected by having to write data and parity across all disks, that is, a data write and one parity write per disk • Offers storage efficiencies and capex savings (RAID overhead is in a range of 12.5% to 20%)
Traditional storage RAID 6	• Offers optimal uptime by surviving double disk failures • Performance is significantly affected by having to write data and parity across all disks, that is, a data write and two parity writes per disk • Offers storage efficiencies and capex savings (RAID overhead is 12.5%)
Traditional storage RAID 10	• Offers optimal uptime by surviving double disk failures • Helps performance by writing only data to each disk, that is, data is actually written to two discrete sets of disks • Does not provide storage efficiencies and has a negative impact on capex (RAID overhead is 100%)
NetApp RAID-DP	• Offers optimal uptime by surviving double disk failures • Helps performance by centralized parity system and NVRAM • Offers storage efficiencies and capex savings (RAID overhead is 12.5%)

With NetApp unified systems, customers no longer need to choose between satisfying one requirement or another. Instead of having to choose between storage efficiency and data protection, data protection and performance, or storage efficiency and performance, customers can choose NetApp RAID-DP to meet all three requirements. RAID-DP provides high-performance design, storage efficiency, and excellent usable storage capacity, especially when compared to a RAID 10 configuration.

4 TEST CASE 1: MEASURE THE AMOUNT OF TIME REQUIRED TO PROVISION DATA STORES TO HYPERVISORS

For this test, we measured the elapsed time required to provision a shared data store on both the NetApp solution and traditional storage array, present that data store to the vSphere hosts, and validate that they can be accessed successfully. The initial storage provisioned was large enough eventually to provide 2TB of usable space. Initially, both the storage and vSphere hosts are powered on and connected to each other, but no storage objects are created on either the NetApp or traditional storage array.

We planned to use this same provisioned storage during the performance tests. Therefore, we made sure that the provisioning method included all of the available disk spindles on both the NetApp and traditional storage arrays to support workloads consisting of high numbers of random read and write operations using small request sizes. Also, the method used to configure the storage was recommended and supported by both NetApp and the traditional storage vendor we used in our tests.

Upon completion of the provisioning, the NetApp configuration had one 2TB Network File System (NFS) data store shared across all four ESX nodes. The traditional storage array configuration had a 1.9TB spanned Virtual Machine File System (VMFS) data store shared across all four ESX nodes.

We found that the time required to provision a data store to the vSphere hosts was reduced by 96% when the traditional storage array was replaced by the NetApp unified architecture. In addition, the traditional storage array did not offer plug-ins to the VMware vSphere server.

4.1 RESULTS WITH TRADITIONAL STORAGE ARRAY

According to the best practices published by the traditional storage array vendor, virtualized environments with more than a few VMs that use FC local unit numbers (LUNs) should use multiple LUNs striped together on the ESX hosts using a technique that stripes the individual LUNs into a single data store from the perspective of the ESX hosts. We followed this best practice during the test case.

To duplicate this configuration, we performed the following steps. The test required 1.75 hours to complete for RAID 6.

1. Using the standard management interface available with the traditional storage array, create a series of 19 RAID 6 groups, each with 6 data disks and 2 parity disks (6+ 2).
2. Using the standard management interface available with the traditional storage array, create a single 100GB LUN on each of the RAID groups and assign ownership of the LUNs to one of the storage controllers.
3. Using the standard management interface available with the traditional storage array, mask the LUNs using the WWPN names from the four vSphere hosts so that each of the vSphere hosts can see all of the LUNs.
4. On the first vSphere host, create a spanned VMFS data store of approximately 1.9TB that includes all 19 of the LUNs created in step 2.
5. Configure multipathing using VMware Native multipath. Note that the three remaining vSphere servers recognize the spanned VMFS as a shared data store.
6. After the space efficiency, deployment time, and performance tests are complete, destroy the configuration.
7. Repeat these steps for RAID 5 and RAID 10 configurations. (An additional 1.75 hours is needed for each configuration.)

4.2 RESULTS WITH NETAPP UNIFIED ARCHITECTURE

Using VSC to automate the provisioning and cloning capabilities of the NetApp environment greatly simplifies provisioning. This tool integrates with VMware vSphere and provides a simple GUI that allows you to complete all the required provisioning steps according to NetApp best practices (see TR-3749). Depending on the architecture chosen for the environment, provisioning might include the following:

- Provisioning FC, iSCSI, FCoE, and NFS data stores
- Masking LUNs
- Configuring NFS export security
- Configuring native vSphere multipath to make sure of high availability and load balancing

The provisioning and cloning capabilities of VSC make it possible for you to complete all of these provisioning tasks rapidly, without ever leaving the VMware vSphere management interface.

VSC requires that storage already be provisioned on the NetApp platform; therefore, we used the NetApp System Manager (NSM) application to create a single 11.6TB aggregate spanning all 40 of the disks, composed of two 16-disk RAID-DP groups and one 8-disk RAID-DP group. We then executed the following steps using the NetApp VSC tool to provision the storage for the vSphere hosts:

1. Install the NetApp VSC plug-in on the VMware vSphere server.
2. Select the NetApp VSC icon within VMware vSphere and add NetApp storage controller log-in credentials.
3. Within the vSphere client, select Home→Inventory→Datastores, then select the data center. Right-click and select NetApp Provision Datastore, select Storage controller, select NFS, select aggregate, and create a 2.0TB thin-provisioned data store with auto grow feature.
4. Review the configuration and create**.**

To create a 2.0TB data store rapidly, NetApp VSC completes all steps, including creating the 2.0TB FlexVol® volume and NFS exports, and setting optimal NFS and network settings.

The following screen shot shows the NetApp VSC provisioning tool summary screen. Note that VSC is a plug-in that is integrated with the VMware vSphere client, providing storage management capabilities to the VMware administrator.

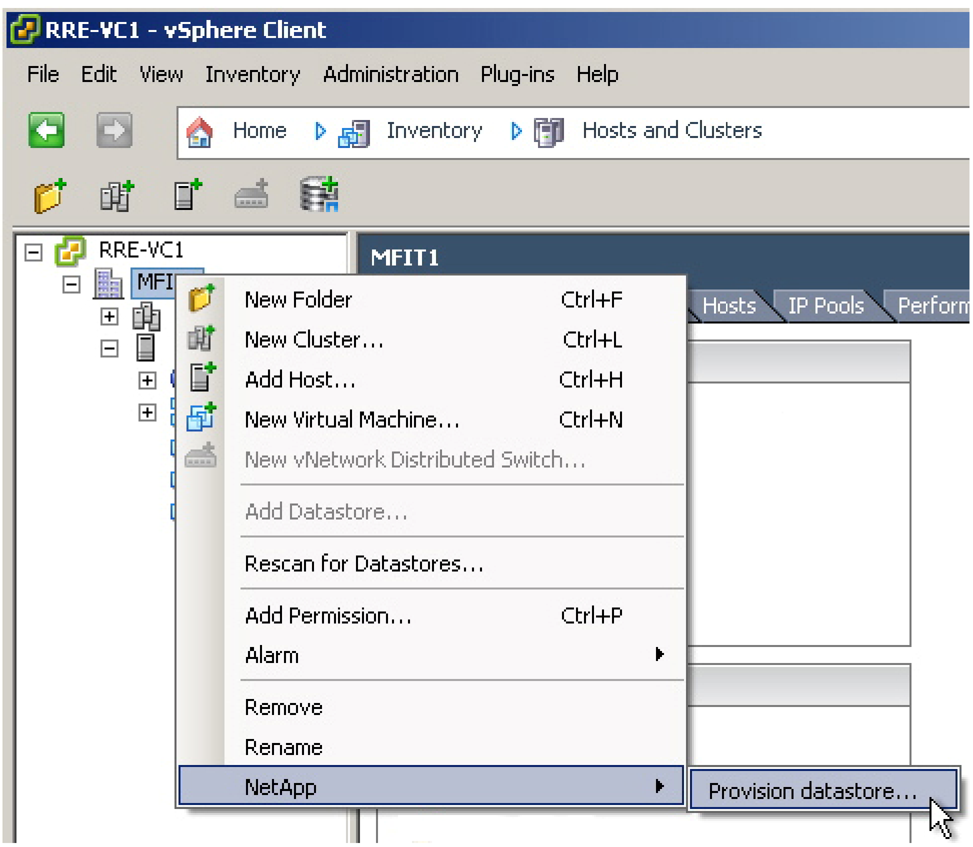

Figure 2 illustrates the dramatic reduction in time that is realized when using the provisioning and cloning capabilities of VSC, as compared to provisioning the traditional storage array. When we combined the time to create the aggregate and provision the storage, our testing showed that NetApp unified architecture required 5 minutes and 38 seconds to provision, as opposed to 1 hour and 45 minutes for the traditional storage array (95% less time).

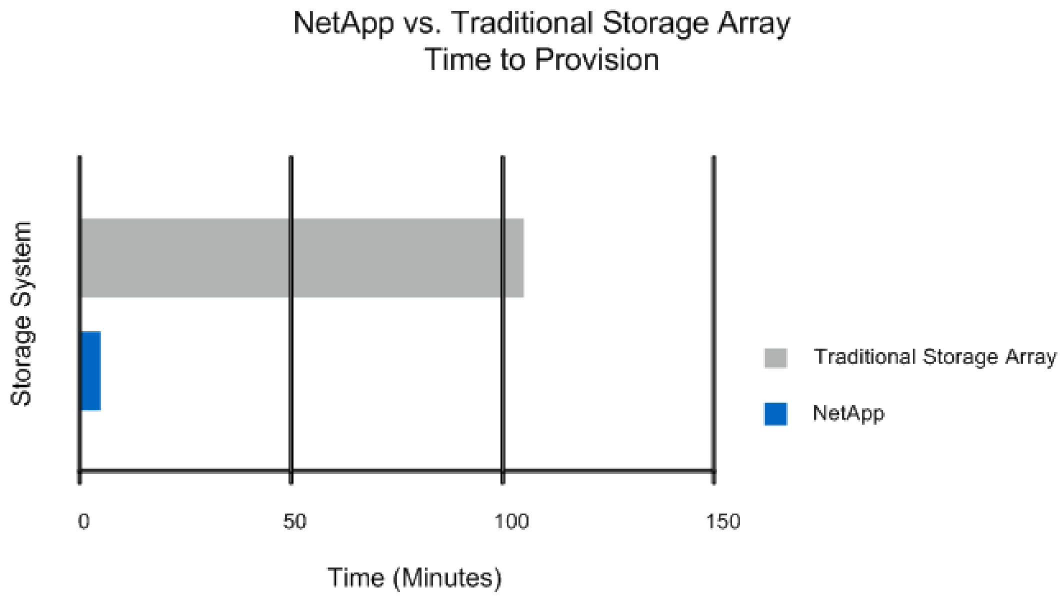

Figure 2) Provisioning test results.

5 TEST CASE 2: MEASURE THE AMOUNT OF TIME REQUIRED TO DEPLOY 80 VMS

For this test, we measured the elapsed time required to provision 80 VMs across the four vSphere hosts in the test configuration (20 VMs per host).We used the shared data store created on both the NetApp solution and traditional storage array from test case 1.

We performed the following steps to configure a golden master VM. We then used this golden master VM to clone 80 VMs.

1. Use the local storage of vSphere host 1 to create the golden master VM.
2. Install Windows Server 2003 Enterprise Edition and SP2 as the guest OS.
3. Create a 5GB VMDK for the Windows installation (c:\) and a second 10GB VMDK to store data files used by IOMeter during the performance testing (d:\).
4. Follow best practices from TR-3749 to align the guest OS to the VMDKs.
5. Install the IOMeter dynamo application and run a short IOMeter test to create a 10GB IOMeter data file on the 10GB VMDK.

After creating the golden master VM, we used the cloning technologies available to both the NetApp and traditional storage array in a VMware environment to create a total of 80 additional VMs using the shared data store created in test case 1. These 80 VMs required a total of 1.28TB of space.

We created a customization script that would be applied to all 80 cloned VMs. The customization process allowed us to provide a distinct identity to each VM.

5.1 DEPLOYING VMS ON TRADITIONAL STORAGE ARRAY

After investigating the available options for cloning VMs for the traditional storage array, we concluded that the best option was to use the standard VMware cloning technology. In using the standard VMware cloning technology, we followed these guidelines:

- Use the VMware vServer to initiate the manual VM clone process.
- Initiate the manual process 80 separate times, using the golden master VM as the source.
- Make sure that the 80 VMs are created evenly across the four vSphere hosts (20 VMs per host).

Because of the manual nature of the cloning process available to the traditional storage array, it took 10 hours and 5 minutes to create 80 VM clones and apply the customization script.

5.2 DEPLOYING VMS ON NETAPP UNIFIED ARCHITECTURE

The provisioning and cloning capabilities of the VSC, combined with NetApp FlexClone, allow users to quickly and efficiently create, deploy, and manage the lifecycle of VMs from an easy-to-use user interface integrated into VMware vCenter™. It is ideal for virtual server, desktop, and dynamic cloud environments. The tool can be used to:

- Clone entire data stores or individual VMs
- Create, resize, and delete data stores
- Apply quest customization specifications to VMs
- Power up VMs
- Run deduplication operations
- Monitor storage savings
- Redeploy VMs after patching
- Import VMs into Virtual Desktop Infrastructure connection brokers and management tools

Unlike the traditional storage array, deploying cloned VMs for the NetApp environment is empowered by the NetApp vCenter plug-in. This tool, which integrates with VMware vSphere, provides a simple GUI that rapidly deploys prededuplicated VMs that are space efficient, properly aligned, and ready to perform. We performed the following steps to create the VMs:

1. From the VMware vSphere client, select the ESX host where the golden master VM is stored, and then select the VMs tab and right-click the golden master VM.
2. Select NetApp→Create Rapid Clones.
3. When the Rapid Clone Wizard opens, select the NetApp storage controller in which the data store is provisioned.
4. Select the data center that contains the four ESX nodes. This makes sure that the VMs are distributed evenly across all four ESX nodes.
5. Select thin-provisioned format.
6. Enter the number of clones, VM clone name, and starting number.
7. Select the option to apply customization and click Next.
8. Select the data store and click Apply.

Our tests showed that using the provisioning and cloning capabilities of the VSC to create 80 VMs took 16 minutes and 56 seconds. This includes the time to customize each of the 80 VMs. This contrasts with the traditional storage array 80 VM deployment, which took 10 hours and 5 minutes.

Figure 3 shows that the NetApp unified storage solution required 97% less time to create the 80 VMs, as compared to the traditional storage array using standard VMware cloning technology. Deployment time on traditional storage arrays is limited by historic storage architecture. The NetApp solution uses FlexClone technology to create prededuplicated VMs, providing immediate storage efficiency without requiring back-end processes to reclaim the space.

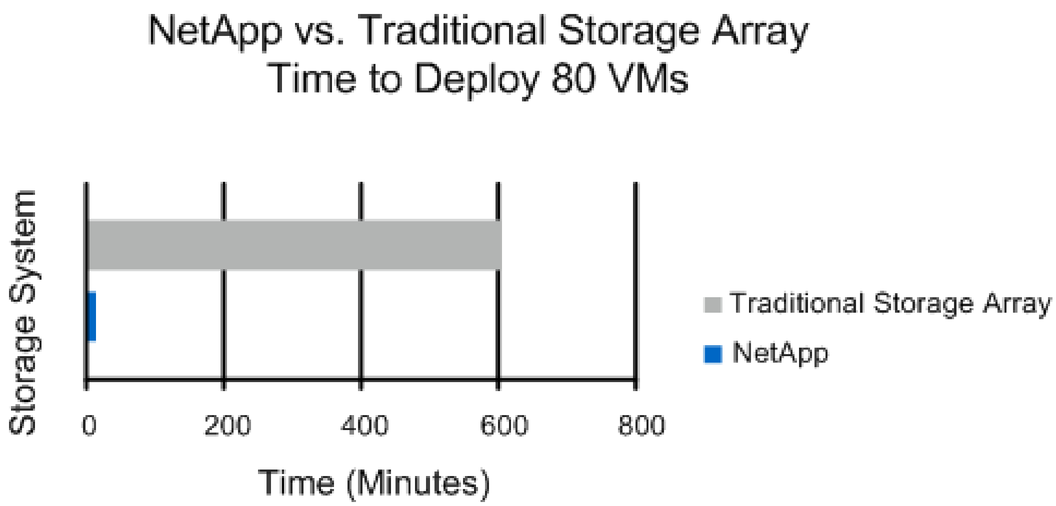

Figure 3) Deployment test results.

6 TEST CASE 3: STORAGE EFFICIENCIES PROVIDED BY STORAGE SAVINGS TECHNOLOGIES

For this test, we measured the storage efficiency of each solution by recording the space consumption reported by the shared data store after creating 80 VMs (20 VMs per host). The golden master VM was created with a 5GB VMDK for the Windows operating system and a 10GB VMDK containing a 10GB IOMeter test file. The golden master VM was created using VMware thin virtual disks. The VMware vSphere client reported the space used as 12.9GB for the golden master VM. The 80 customized Windows clones were deployed from the VMware thin golden master VM.

Before discussing the results, it is useful to review some of the storage efficiency technologies provided by VMware and storage arrays.

6.1 VMWARE THIN-PROVISIONED VIRTUAL DISKS

In this test, we used the VMware thin-provisioned virtual disk format. This format reduced the total storage consumed within the shared data store for both storage arrays by allocating storage on demand. Traditionally, virtual disks preallocate storage equal to the maximum size of their virtual disks. By enabling this feature, our test bed was able to reduce total storage consumption on both storage arrays by 19.4% (consuming 12.9GBs for each 16GB VMDK).

Note that the thin format is an option for the Fibre Channel-connected data store but it is the default format for NetApp NFS data stores.

The use of thin-provisioned disks in this testing did not introduce any latency into the I/O performance tests run for this technical report. Storage allocation for the IOMeter test file was included in the VM image, so storage for the IOMeter test file was preallocated.

6.2 STORAGE EFFICIENCY CAPABILITIES IN STORAGE ARRAYS

Traditional storage arrays require thick-provisioned LUNs and volumes that reserve the full amount of space provisioned. This is because LUNs are tied to the physical hardware disk sectors of the underlying RAID group. Thin LUN provisioning was added as a feature to storage arrays to provide more flexible, efficient space consumption. Storage array thin LUNs allow for LUNs to be provisioned from a pool of underlying disks that can be much larger than a traditional RAID group size. Thin LUNs can be provisioned initially with a small amount of space and can be configured to grow on demand, taking additional space from the pool when required. This feature adds flexibility but might or might not be recommended for performance-sensitive environments.

NetApp storage arrays provide storage savings at the time of VM creation and throughout the VM lifecycle. These technologies are FlexClone and deduplication, respectively. At the core, these technologies are identical and leverage NetApp's pointer-based virtualization layer WAFL®. The enablement of FlexClone and deduplication is the same technology used in NetApp Snapshot™ copy technology, SnapMirror®, SnapVault®, Flash Cache, and so on.

The benefit of a single architecture to reduce storage at the time of provisioning and throughout the VM lifecycle makes sure of optimal storage efficiencies.

Depending on the storage array capabilities, NetApp storage technologies can be combined to realize increased storage efficiencies. Using powerful NetApp FlexClone technology, the space consumption for multiple VMs is initially equivalent to the space consumed by the single master VM from which they were cloned. As data is written to the VMs, the additional space consumed is the delta between the space in the master VM and the clones. NetApp recommends every VM data set be deduplicated. The cloning process is a forward provisioning model that equates to reduplicating full VM images.

Deduplication, FlexClone, and FlexCache® work in tandem to provide transparent storage cache sharing (TSCS). TSCS is very similar to VMware transparent page sharing found in ESX/ESXi. TSCS allows for a single block-level object to be stored in the array's cache and made available to multiple external references. These references can be in many forms, including VM images, user data sets, and high-performance enterprise applications.

Together, these technologies provide storage efficiency and performance acceleration. Figure 4 illustrates the value of NetApp FlexClone and deduplication technologies.

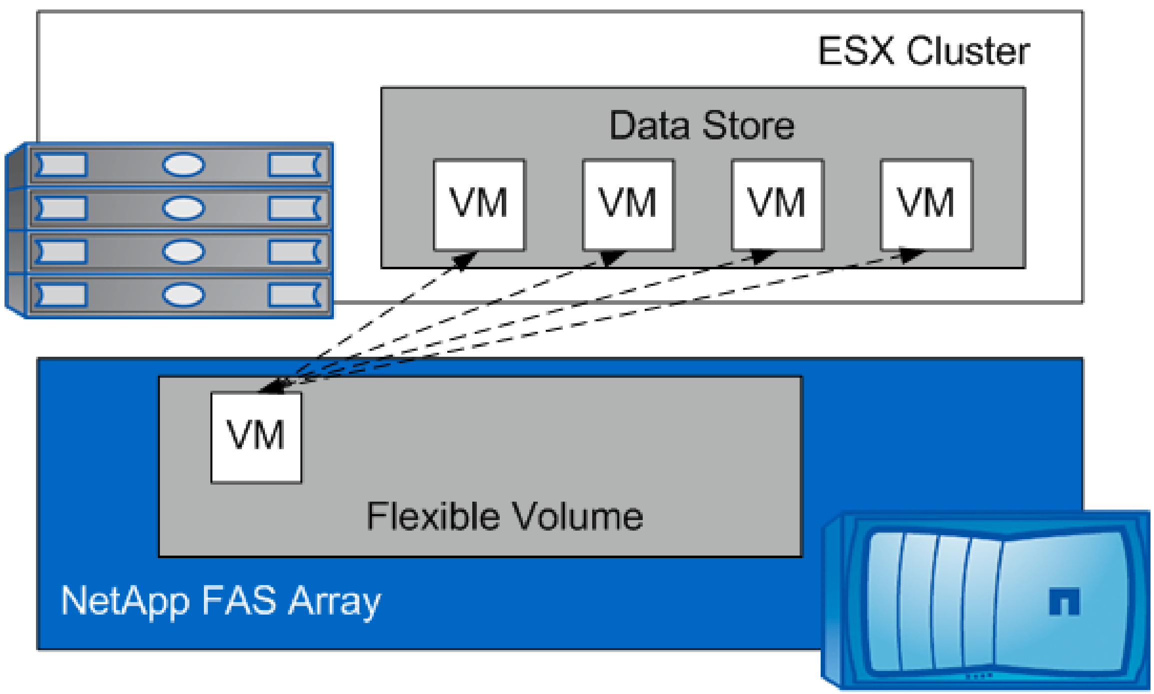

Figure 4) Value of NetApp FlexClone and deduplication.

6.3 TEST 4: RESULTS WITH TRADITIONAL STORAGE ARRAY

Spanned VMFS data stores are usually recommended for environments that need more space than can be provided by a single LUN VMFS. We used spanned VMFS in our traditional storage array design. We used the following RAID group sizes for the traditional storage array, using 7+1 RAID groups for RAID 5, 6+2 RAID groups for RAID 6, and 4+4 RAID groups for RAID 10. If we had used the maximum LUN size available, we would have used only two or three RAID groups, thus limiting the number of back-end spindles and affecting performance. In order to compare fairly with the NetApp solution that used 40 disks, we instead created a spanned VMFS that incorporated 19 LUNs, thus allowing a large number of disk spindles to be used for the mixed random workload. This configuration includes multiple storage array LUNs concatenated to provide a single data store. The purpose of this design is to increase the total I/O queue available to the data store. In this test, the VMs were deployed on a shared VMFS-spanned data store composed of 19 100GB thick-provisioned LUNs. Although thin-provisioned LUNs were an option for the traditional storage array, using thin LUNs is not a recommended best practice for performance-sensitive environments. Deduplication was not a feature offered by the traditional storage array. The total space consumed by the 80 VMs was 1.28TB. The net storage efficiency savings for the traditional storage array was 19%, all provided by the use of VMware thin virtual disk technology.

The following screen shot of the VMware vSphere client shows the properties of the spanned VMFS data store. VMware cloning technology was used to deploy 80 VMs. The capacity of the data store is 1.85TB, and the used space is 1.01TB.

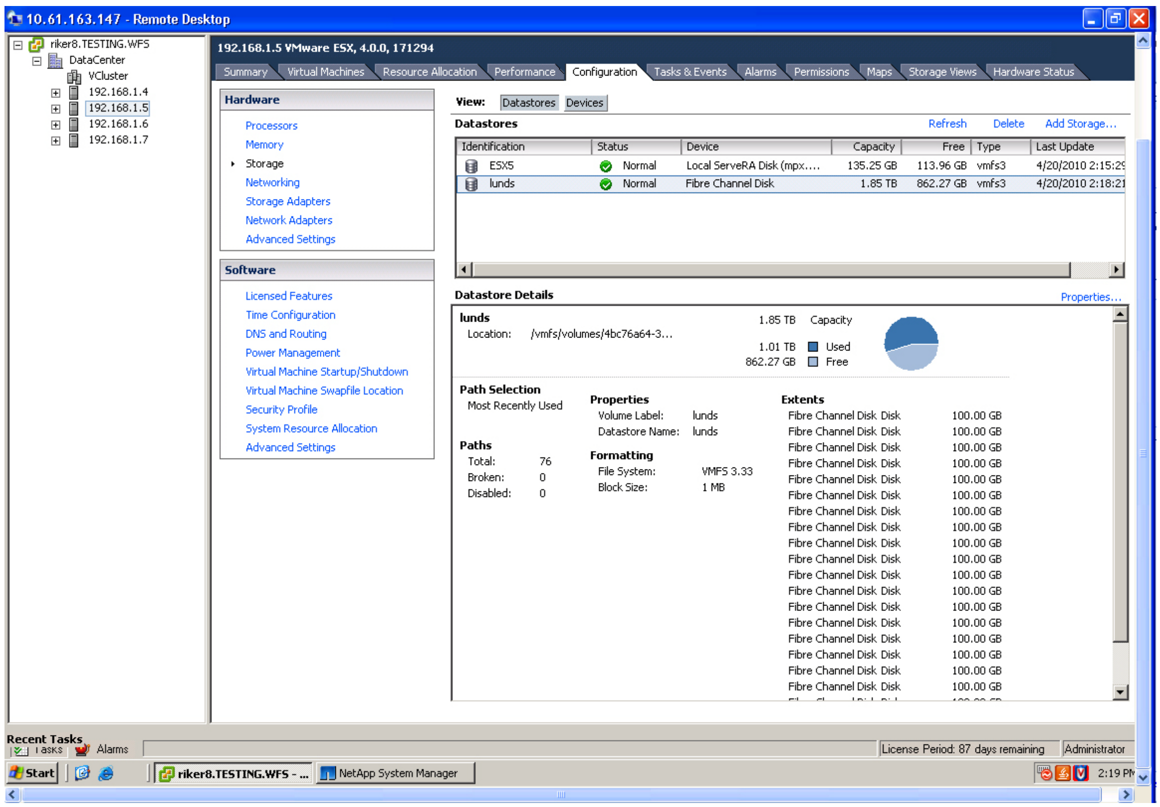

6.4 TEST 4: RESULTS WITH NETAPP UNIFIED ARCHITECTURE

We used the provisioning and cloning capabilities of NetApp VSC to deploy 80 VM clones from the VMware thin golden master VMDK. NetApp NFS uses thin VMDKs by default. This setting is controlled by the array and can be overridden by the virtual infrastructure administrator. VSC used NetApp FlexClone technology to create 80 prededuplicated VMs, thus providing immediate space savings. The space consumed by the 80 VMs was 11.26GB total.

The following screen shot of the VMware vSphere client shows the properties of the NFS data store. The provisioning and cloning capabilities of VSC were used to deploy 80 VMs. The capacity of the data store is 1.95TB, and the used space is 11.26GB.

The following screen shot of the VMware vSphere client shows the properties of the NFS data store. The provisioning and cloning capabilities of VSC were used to deploy 80 VMs. The capacity of the data store is 1.95TB, and the used space is 11.26GB.

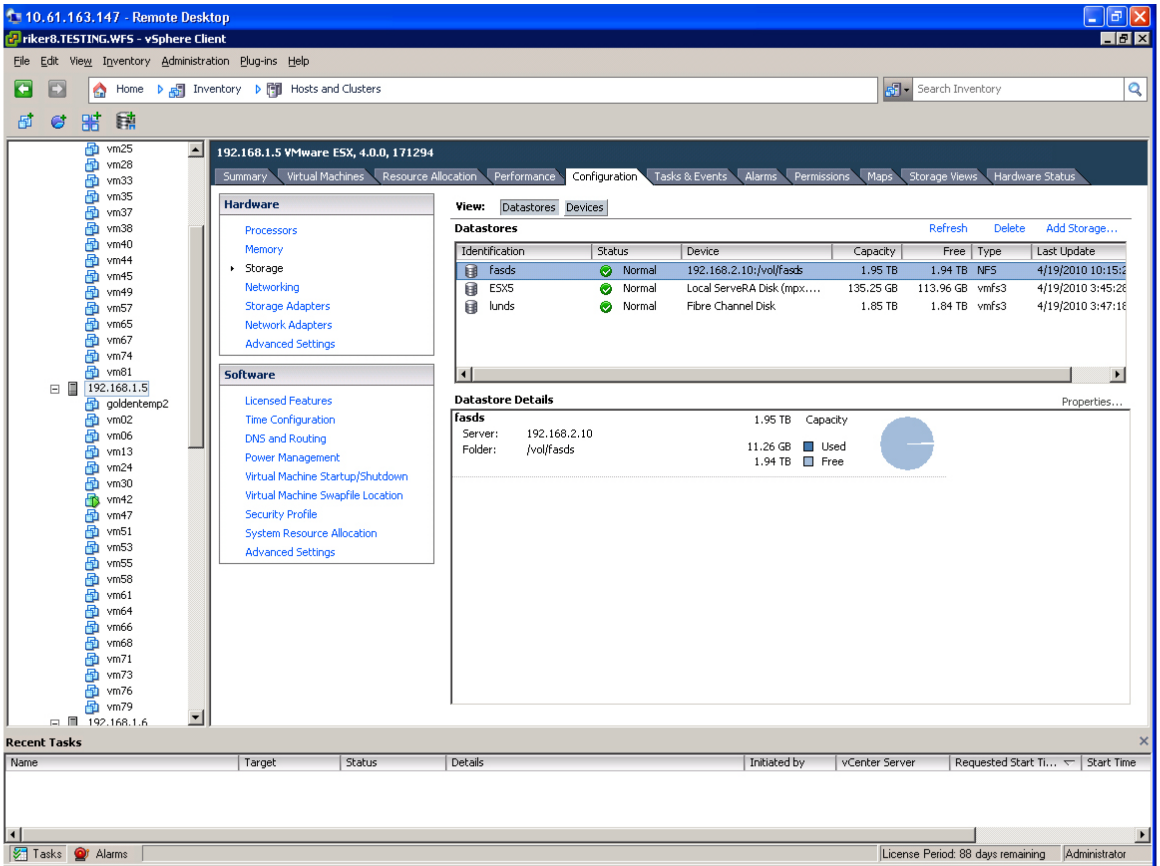

6.5 TEST 4: SUMMARY

The NetApp unified architecture allowed for a multiplication of storage savings when VMware thin provisioning was combined with NetApp thin provisioning and FlexClone. NetApp space savings are transparent and show in the data store and on the storage volume. Figure 5 illustrates how the combination of storage savings allowed for 99% space savings with the NetApp solution, as compared to the traditional storage array.

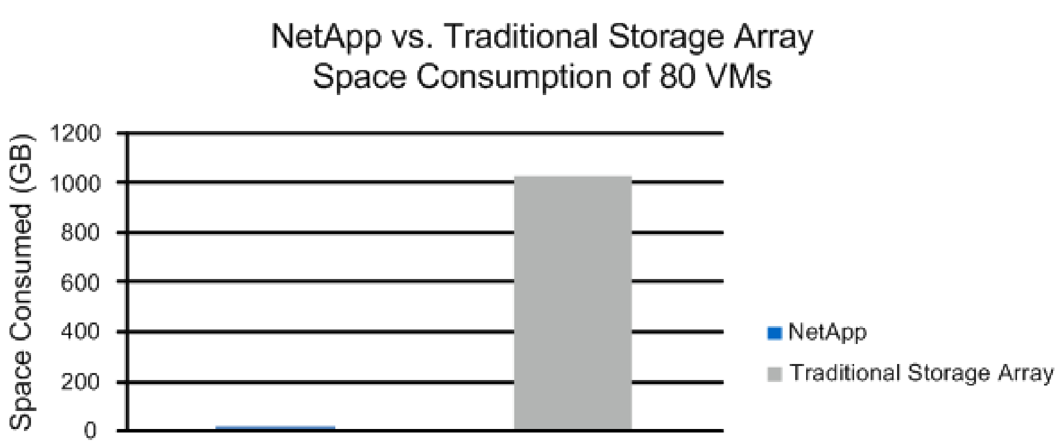

Figure 5) Storage efficiency test results.

The storage capacity required to run a virtual infrastructure was reduced 99% by replacing traditional storage arrays with the NetApp unified architecture. The storage administrator has the option to allow the virtual infrastructure administrator to manage storage savings by way of the NetApp vCenter plug-in. Traditional storage arrays realize the benefits of VMware thin provisioning, but they are unable to realize the benefits of FlexClone across the entire data store as provided by the NetApp unified architecture.

7 TEST 4: BENCHMARK PERFORMANCE COMPARISON WHEN COMBINING STORAGE EFFICIENCIES AND DATA PROTECTION TECHNOLOGIES

For the performance tests, we used the configuration that was created in tests 1, 2, and 3. The golden master VM from which all 80 VMs were cloned had the IOMeter dynamo application installed and a 10GB IOMeter data file created in the second VMDK, giving them all the capability to generate load in a performance test.

Following are the high-level objectives for the performance tests:

- Test both the NetApp unified storage array and the traditional storage array using the different RAID and space-efficiency technologies available on both systems. This means using RAID-DP and thin-provisioned data stores for the NetApp system and RAID 5, RAID 6, and RAID 10, along with thin VMDKs configured from a spanned VMFS partition, for the traditional storage array.
- Use IOMeter to generate a real-world workload to measure the performance. Our workload consisted of a 100% random mix of 75% reads and 25% writes using a 4KB request size, with two I/Os outstanding for each VM. This is the same workload we have used in previous joint NetApp and VMware performance tests. To see the details pertaining to these performance tests, use the following links to access TR-3808 and TR-3697:

  http://media.netapp.com/documents/tr-3808.pdf
  www.netapp.com/us/library/technical-reports/tr-3697.html

7.1 TEST 4: PERFORMANCE RESULTS FOR NETAPP VS. TRADITIONAL STORAGE ARRAY

When comparing the performance among traditional storage arrays, we used the throughput in input/output operations per second (IOPS) generated by the NetApp configuration as the baseline. In Figure 6, this baseline has a value of 100. The performance of the traditional storage array is presented as a percentage relative to the observed performance using the NetApp configuration.

The closer the percentage for the traditional storage array is to 100%, the closer the actual performance is to that using the NetApp unified storage system. For example, a value of 95% for the traditional storage array indicates that the performance in IOPS generated is 95% as good as the performance of the NetApp configuration. See Figure 6 for the results of the actual performance tests.

We found that the traditional storage array configured with RAID 10, RAID 5, and RAID 6 delivered approximately 90%, 75%, and 70% of the performance of the NetApp unified storage system, respectively. The results demonstrated that the NetApp unified architecture approach provides excellent performance in addition to providing storage efficiency, rapid deployment, and strong data protection. When compared to a traditional storage FCP array using RAID 5, RAID 6, or RAID 10, NetApp unified solutions are the clear leader.

Traditional storage array performance using RAID 10 surpasses the performance of RAID 5 and RAID 6 protection; however, those performance gains are offset by the dramatically decreased storage efficiency, which leads to higher capex and opex. NetApp unified storage architecture performance over NFS is competitive with the traditional storage array with a very demanding I/O load and real-world block size. The peak performance is achieved with storage efficiency and data protection, as illustrated in Figure 6.

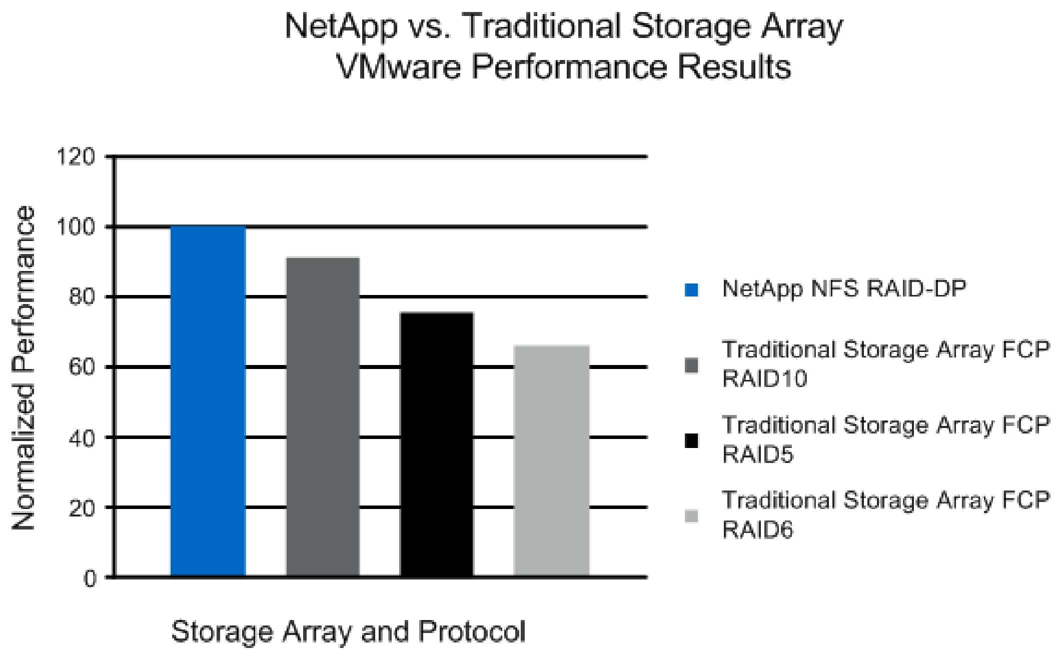

Figure 6) Performance test results.

7.2 VALUE OF NETAPP FLASH CACHE

The strong performance of the NetApp unified architecture is increased by implementing the Flash Cache. NetApp Flash Cache is an intelligent read cache that is unique to Data ONTAP®. While traditional storage arrays have implemented flash drives to act as storage disks for LUNs, there are inherent limitations with this design. Data sets must be architected in rigid LUN containers and assigned to the solid state storage. This requires storage administrators to perform the impossible task of consistently assigning and maintaining the hot data LUNs in solid state disk. Hot data is often a moving target, making it a challenge to migrate LUNs on and off the solid state disk drives continually as they become hot. The NetApp unified architecture overcomes this limitation by using solid state technology as a Flash Cache. Memory and storage-efficient hot data are migrated automatically and seamlessly at the block level without requiring storage administrator intervention. The performance increase with this technology is evident, as seen in the test results described in the following section.

In addition, Flash Cache introduces TSCS, which increases performance with storage area network (SAN) and network-attached storage (NAS), serving VM binaries and OS images, user data such as home directories and Outlook OST files, Exchange databases and DAGs, database clones in Oracle® test and dev, and so on.

7.3 TEST 4: RESULTS WITH NETAPP FLASH CACHE ADDED TO THE NETAPP UNIFIED ARCHITECTURE

For this round of performance tests, we installed a single 256GB Flash Cache card in each of the NetApp storage controllers and applied the license. Then we cloned 80 VMs from the golden master VM. The space savings described in test case 3 remained in place. We repeated the previously described IOMeter test.

Figure 7 provides the results of the performance tests using the Flash Cache. In comparing the performance between the NetApp unified storage controller and the traditional storage array, we used the throughput in IOPS generated by the NetApp FAS3140 configuration as the baseline. This value is represented on the graph using a value of 100. The performance of the traditional storage array and FAS3140 performance with the Flash Cache are then represented as a percentage relative to the observed performance using the NetApp configuration.

The closer that the percentage noted for the traditional storage array is to 100%, the closer the actual performance is to that observed using the NetApp unified storage system.

In this case, we found that the traditional storage array configured with RAID 10, RAID 5, and RAID 6 delivered approximately 65%, 55%, and 50% of the performance of the NetApp unified storage system, respectively. As illustrated in Figure 7, when we added the Flash Cache to the NetApp system, the results were 140% of the base NetApp results. These results demonstrate the value added by the NetApp Flash Cache for a VMware workload.

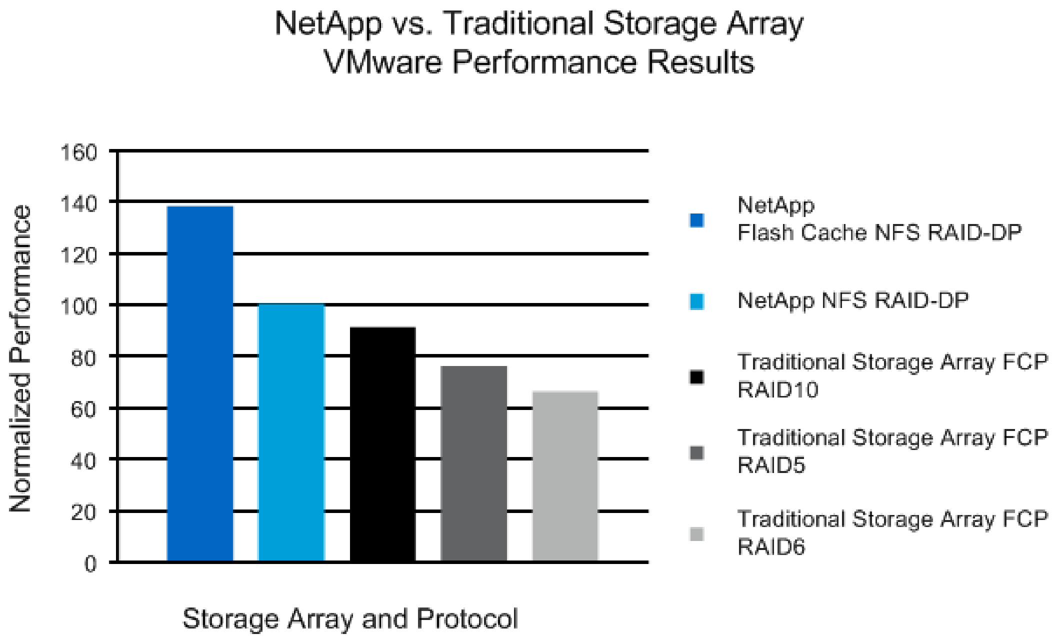

Figure 7) Performance test results with Flash Cache.

8 SUMMARY

The goal of this technical report is to demonstrate the gains of running a virtual infrastructure on NetApp's unified virtual storage architecture, as compared to a traditional storage array. We found in our tests that implementing VMware on NetApp provides the following key benefits:

- **95% reduction in data store provisioning time.** The NetApp unified architecture required 5 minutes and 38 seconds to provision, as opposed to 1 hour and 45 minutes for the traditional storage array.
- **97% reduction in VM deployment provisioning time**. Our tests showed that using the provisioning and cloning capabilities of the VSC to create 80 VMs took 16 minutes and 56 seconds. This includes the time to customize each of the 80 VMs. This contrasts with the traditional storage array 80-VM deployment, which took 10 hours and 5 minutes.
- **NetApp provides 99% gain in usable capacity because the solution requires 99% less space.** The traditional storage array required 1.01TB for 80 VMs, in contrast to the NetApp FlexClone solution, which required only 11.26GB.
- **NetApp performance compared to traditional storage array.** NetApp performance, as compared to the traditional storage array, was 30% better than RAID 6, 25% better than RAID 5, and 10% better than RAID 10.
- **Effect of Flash Cache on performance.** Upon adding the Flash Cache to the NetApp solution, we found the performance to be 52% better than the peak traditional storage array performance using RAID 10.

Figure 8 illustrates the comparative efficiency and performance of the NetApp unified storage solution to the traditional storage array.

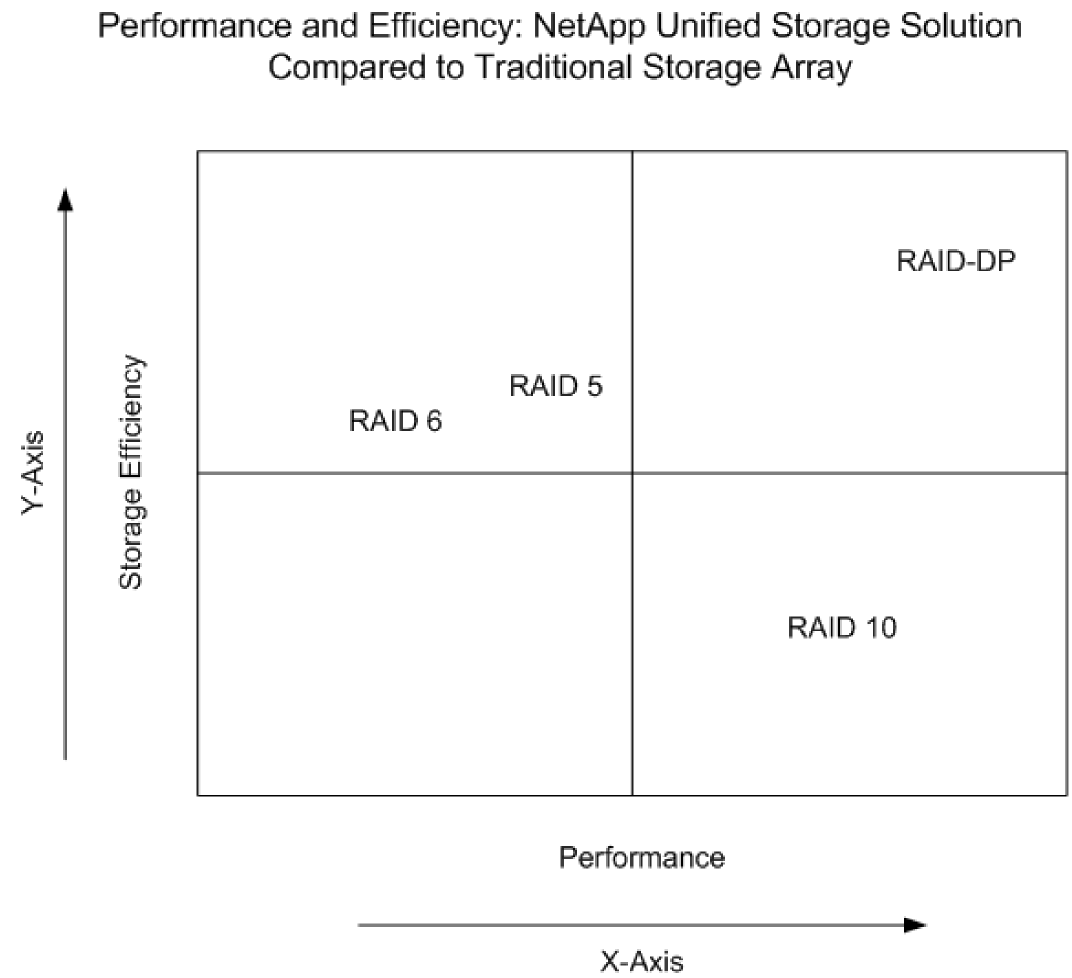

- RAID-5 (4+1) = 20% Parity Overhead
- RAID-6 (6+2) = 25% Parity Overhead
- RAID-10 (4+4) = 100% Parity Overhead
- RAID-DP (14+2) = 12.5% Parity Overhead

Figure 8) NetApp unified storage solution compared to the traditional storage array.

We believe this technical report fairly represents the ability of NetApp solutions to provide simplicity, agility, and high performance. The gains we documented should provide capex and opex savings by reducing hardware purchase and lowering ongoing power costs. We understand that you might not run VMware on the traditional storage platform against which we tested, so we invite you to let us prove these results in your data center; we can also do so by virtualizing your existing arrays with our V-Series.