

HIGHLIGHTS

- » Enabling SaaS through Virtualization
- » Engineering View: Oracle® on NFS
- » New Entry-Level Platform: FAS2000
- » Resources:
 - 10 NetBackup Questions Answered
 - Sneak Peek: New VMware Reports

Cutting Storage Use 50% Using A-SIS Deduplication

Jamie Ridge, Enterprise Systems engineer, NetApp

See how one company reclaimed half its backup and home directory storage overnight with virtually no infrastructure changes.

[More »](#)

BLOGGING WITH DAVE


Dave Hitz, NetApp Founder and EVP



"Benchmarks can make your head spin. Depending on how you measure, two systems go from a tie to EMC almost twice as slow to NetApp doing 25% more ops."

» [Dave's Blog](#)

DRILL DOWN

- » **New Entry-Level Platform: FAS2000**
Check out capacity and other specs for the newest NetApp hardware platform, plus a 3D interactive tour! 
- » **Using NetApp RAID-DP in Exchange Server 2007 Storage Designs**
Compare errors, write performance, data loss during reconstruction, and much more using RAID 10, RAID 5, and RAID-DP.
- » **"If you've never worked with NetApp, you'd be most surprised by..."**
Hundreds of NetApp users responded to the contest. Check out the winners.

ADMIN RESOURCES

- » **Disk-Based Backup for NetBackup™: Answers to Your Top 10 Questions**
NetBackup and SnapVault® integration, deduplication, NearStore® VTL, and more.
- » **SnapMirror® Technical Q&A**
A top engineer provides advice on troubleshooting, configurations, and more.

TIPS FROM THE TRENCHES

What's Hot with VMware and NetApp: Enhanced DR, Virtual Storage, NAS, VDI

Highlights from VMworld 2007

Check out the latest trends and stay cutting edge with early access to brand-new technical reports and demos involving:

- Best practices: VMware ESX 3.0 on NetApp
- Backup, recovery, and DR architectures
- VMware in NAS environments
- Cost-effective desktop virtualization


[More »](#)

Virtual Reality: Building an Architecture Capable of Evolving in Any Direction

Sammie Carter, Brian Bouterse, and Patrick Bowen
The Friday Institute, North Carolina State University

The challenge: Enable rapid application delivery to remote classrooms.

See why the team chose VMware on NAS and how they've achieved on-the-fly provisioning, multi-dimensional scalability (blades, storage, or connectivity), and the ability to rapidly deploy VMs or bare metal.



[More »](#)

ENGINEERING TALK

Oracle on NetApp Part 1: Using NFS for Simplicity/Manageability

Sanjay Gulabani, sr. Database Performance engineer, NetApp
Darrell Suggs, technical director and cache architect, NetApp

A behind-the-scenes look at how NetApp Engineering has driven standards, enhancements, and best practices for Oracle on NFS, plus updates on the impact of emerging standards like NFSv4 and Direct NFS.

[More »](#)

[FEEDBACK](#)



Jamie Ridge
Enterprise Systems Engineer, NetApp

Jamie has over 15 years of experience in Information Technology with Fortune 500 companies. He has been an Enterprise Systems engineer with NetApp for 3.5 years, working with some of NetApp's largest and most critical enterprise accounts. Prior to joining NetApp, he spent over five years at Sun Microsystems. Jamie holds a variety of certifications from Microsoft, Sun, and NetApp and has a Masters of Business Administration degree from the University of Houston.

A Case Study: Cutting Storage Use 50% Using A-SIS Deduplication

By Jamie Ridge

When it comes to storage, getting the most bang for your buck has never been more critical than it is now. I recently worked with a Fortune 500 company that was struggling to justify its storage spend. The company was spending about 80 cents on secondary storage (for backup and DR) for every dollar spent on primary storage. Because the IT budget was allocated on a per-project basis, business units across the company were painfully aware of how much of each project budget was being spent on secondary storage that-while critical for data protection-did not contribute directly to end user productivity or the bottom line.

Using NetApp Advanced Single Instance Storage (A-SIS) deduplication technology-a new feature from NetApp that is bundled in Data ONTAP® 7.2.2 at no additional cost-the customer was able to drop secondary storage spending to only about 25 cents per dollar spent on primary storage, an almost 70% reduction.

In this article I'm going to describe the customer's environment, explore the particular storage challenges the IT team faced, and discuss the A-SIS deduplication implementation and resulting benefits.

Storage and Disaster Recovery Environment

The customer's IT environment consists of two primary data centers located in different states plus a number of remote offices around the globe. Currently, the customer has about 1,000 terabytes of primary and secondary NetApp storage. About 60 to 70% of their annual storage spending goes for NetApp storage with the remainder split between other storage vendors.

The customer uses NetApp SnapVault® and SnapMirror® software to provide backup and disaster recovery between its two primary data centers. SnapVault is used to eliminate the need for nightly incremental tape backups, while SnapMirror is used to mirror the critical data-corporate databases, etc.- between sites for greater data protection.

Remote sites from all over the world trickle their backups across company networks to the primary data centers for disaster recovery. While some of the remote sites utilize NetApp Open System SnapVault (OSSV) for space-efficient backups of non-NetApp storage, a variety of other backup products are in use to support other remote sites depending on need. The largest component of this consists of

OnDemand TechTalk Webcast:

Cut Data Storage Use by As Much As 95%: Using Deduplication to Drive Down Storage Costs




Figure 1. Customer backup and disaster recovery environment.

RELATED INFORMATION

- [A-SIS Deployment Guide \(PDF\)](#)
- [Top 10 A-SIS Questions](#)
- [Deduplication Comes of Age](#)
- [Dave's Blog: How Deduplication Fits the NetApp Master Plan](#)

Engineering Perspective: Deduplication Comes of Age

At its heart, A-SIS deduplication relies on the time-honored computer science technique of reference counting.

When A-SIS deduplication is enabled on a volume, it computes a database of fingerprints for all of the in-use blocks in the volume (a process known as "gathering"). Once this initial setup is finished, the volume is ready for deduplication.

To avoid slowing down ordinary file operations, the search for duplicates is done as a separate batch process. As the file system gets updated during normal use, WAFL creates a log describing the changes to its data blocks. This log grows until one of the following occurs:

- The administrator issues a sis start command
- The next time specified in the sis config schedule occurs
- The changes to the log exceed a predetermined threshold

Learn more. [Read the article.](#)

Veritas™ NetBackup™ backing up to NetApp NearStore® secondary storage in the two main data centers. The data change rate ranges from 10-20%, and all incremental backups are retained on disk for 60 days.

Backup Type	Media	SW Used	Frequency	Data Type
Full (includes all remote offices)	Tape	NetBackup	Once a month	All
Incremental	Disk	NetBackup	Daily	Structured data
Incremental	Disk	SnapVault	Daily	Unstructured
Incremental Remote Office	Disk	OSSV/SV	Daily	Unstructured
Incremental Remote Office	Disk	NetBackup	Daily	Unstructured

Customer Challenge

The storage challenge in this environment results mainly from the growth of backup data on secondary storage coming from Veritas NetBackup and other conventional backup software. The IT team was faced with constant space issues to accommodate these backups on secondary storage, and business units were bristling at the expense of adding secondary storage.

Deduplication to Reclaim Disk Space

To improve storage efficiency, the customer began looking at deduplication solutions. As the name implies, deduplication eliminates duplicate copies of blocks or files. This can substantially cut storage costs by reclaiming existing disk space for reuse and enabling IT teams to postpone purchasing additional storage capacity.

After considering solutions from several vendors, the customer chose NetApp A-SIS deduplication.

A-SIS is an integral part of the NetApp WAFL® file system. As a result, deduplication has very low overhead and occurs "behind the scenes," regardless of which applications you run or how you access the data. A-SIS operates in a post-processing mode, so you can schedule deduplication to occur at times of low activity or when set thresholds are reached. (Read [Deduplication Comes of Age](#) for an engineering perspective.)




Figure 2. Before and after A-SIS Deduplication.

This customer is typically very conservative, and often spends several months testing and kicking the tires before proceeding to roll out a new technology. In this case, however, once the IT team realized that A-SIS could be implemented in its existing NetApp environment with virtually no risk or additional investment, the move to A-SIS was accomplished very quickly. A-SIS was enabled across the customer's NetApp environment (the product is free); then the customer tested it in the lab, found out it worked as expected, and turned it on in production. No changes were required to the existing backup environment.

A-SIS deduplication is now run once a week over the weekend to eliminate duplicate data that accumulates in backup volumes during the week.

Impressed with the results achieved on backup volumes, the IT team next looked at other types of data on secondary storage and enabled A-SIS on NearStore volumes storing user home directories. These directories contain the standard mix of Microsoft® Office documents, PDF files, image files, etc. NetApp is the first major vendor to support deduplication on this type of data.

(Most of the company's internal customers rely on high-performance primary storage for home directories; this is considered a best practice within the IT organization. However, since IT costs are tracked by project, some highly cost-sensitive internal customers have chosen to store home directories on less expensive NearStore secondary storage.)

Deduplication Impact Space Savings

After execution, just over 50% of the disk space in the NetBackup environment was recovered as reported by A-SIS. The reduction in storage use on home directory volumes was even more significant-about 67% of storage capacity was reclaimed, which is well above the typical space recovery of 50% for A-SIS enabled on home directories. (This suggests a higher-than-normal level of file duplication between

A-SIS Deployment and Implementation Guide

Want to learn even more about A-SIS deduplication? This comprehensive guide gives you everything you need to understand, plan, and deploy deduplication across your NetApp environment, including:

- A complete technology overview
- Supported configurations
- Setup and operation details
- Best practices for deployment
- Using A-SIS with cluster and DR configurations
- Troubleshooting

[Download the guide.](#)

87% Space Savings in 60 Days

The County of Sacramento has a backup environment that includes approximately 2.1TB of data. The agency saves daily incremental Snapshot copies for two weeks, and keeps weekly full backups for 60 days. Every day, approximately 5% of files and 0.5% of blocks change.

By deploying a NetApp storage solution that includes A-SIS deduplication technology, the agency saw results that included:



- ~ 87% space savings at 60 days
- Backup 16x faster
- Restores 6x faster: 5 min vs. 30 min

Listen to [Keith Scott, County of Sacramento IT Analyst](#), describe the [deployment](#).

How Data Deduplication Fits into the NetApp Master Plan

Following is an excerpt from a recent post on NetApp Founder Dave's Blog:

Buying less storage is the small picture. The big picture is that we want to help customers create a disk-based copy for all of their primary storage. ...

Interesting things start to happen when you create a disk-based copy of everything. Instead of doing searches on primary storage, which could hurt performance, why not search the

users.)

	Space Usage before Deduplication	Space Usage after Deduplication	Space Savings
NetBackup Vols	150TB	73.5TB	51%
Home Directories	50TB	16.5TB	67%

The reduction in the use of secondary storage yielded significant benefits. In the past, the company had bought one secondary storage system for every primary storage system. Since enabling A-SIS, the IT team has dropped that to one secondary storage system for every two primary storage systems, and cut secondary storage spending to about 25 cents for every dollar spent on primary storage. As a result, more of the IT budget has been allocated for primary storage purchases and the IT department has been able to pass on cost savings to company business units and end users.

No Operational Impact

Naturally, the company had some concerns that deduplication would affect storage performance. Home directory users have reported no impact whatsoever. The company remains cautious with regard to enabling A-SIS on storage used by more demanding applications and users, but is definitely considering it as they grow more comfortable with the technology.

IT Benefits

From the standpoint of the IT team, A-SIS requires virtually no ongoing management. Once it is enabled, it runs in the background without constant monitoring or tinkering. Because A-SIS helps keep free space available on secondary storage, the team now spends a lot less time dealing with disk space shortages and can focus more time on other value-added tasks. Overall, deduplication makes it easier for the team to get the job done within the constraints of a tight IT budget.

Expanding the Use of Deduplication

In the future, the IT team plans to enable deduplication on existing SnapVault volumes. NetApp SnapVault provides time-based deduplication that greatly reduces the amount of data that must be stored versus traditional methods. After an initial baseline copy, subsequent SnapVault backups transfer and retain only blocks that have changed since the last backup, eliminating the creation of multiple copies of unchanged files. Adding space-based deduplication using A-SIS will further reduce the disk space used by SnapVault by eliminating redundancy due to different users or applications storing copies of the same data, as was clearly the case in the company's home directory volumes given the space recovered by A-SIS.

Enabling A-SIS has been a big win for this company and it looks forward to further expanding use of the technology. Business units are much happier with the reduced allocations for storage spending, and the IT team is even happier with its storage infrastructure.

Beyond the Hype What You Should Know about Space Savings

When evaluating deduplication space savings, you should have two goals:

First, examine your backup data. It is reasonable to expect 5:1 to 20:1 space savings (over time) with a backup change rate of 2%. If you store more than 20 backup copies on disk or your change rate is less than 2%, your deduplication ratio will increase. Conversely, if you retain a lower number of backup copies or if your data change rate is higher than 2%, your overall deduplication ratio will be reduced.

Next, examine your non-backup data. Are there any opportunities to reduce duplicate data on those volumes? Generally speaking, if you can reduce this data by 1.25-1.75:1, deduplication would be economically feasible. Think of it as receiving a "storage rebate" by reducing the storage capacity of these volumes by 20-40%.

As costs today can easily exceed \$10,000 per TB, saving just a few terabytes across your enterprise can easily have a significant impact on your overall IT budget.

secondary copy? If the people running decision support systems want their own copy of a critical database, why not clone the secondary instead of paying for a whole new copy? Why not create lots of cloned copies for the test and development team preparing to upgrade to the next version of Oracle or SAP?

When you create a copy of everything, and add functionality like Snapshot copies and clones, what you end up with is a smart copy infrastructure that can completely change the way you think about data management.

This won't happen overnight. We understand that. But anything that helps people reduce the cost of creating copies helps us achieve our vision more quickly. In the short run, data deduplication helps customers save space and save money, but what's more important is that by reducing the cost of copies, it helps us achieve our master plan.

Read the [full Dave's Blog post on Deduplication](#)

This feature was a collaborative effort of members of the NetApp technical team who attended and presented at VMworld 2007, including (alphabetically) :

VMWORLD 2007

- Darrin Chapman, senior manager, DPR Product Engineering, NetApp
- Peter Learmonth, consulting systems engineer, NetApp
- Stephen Manley, technical director, Data Protection Group, NetApp
- Mike Slisinger, Professional Services consultant, NetApp
- Vaughn Stewart, consulting systems engineer, Infrastructure Virtualization
- Kim Weller, SAN consulting systems engineer, NetApp

What's Hot with VMware and NetApp: Enhanced DR, Virtual Storage, NAS, and VDI

Highlights from VMworld 2007

This month thousands of people are visiting San Francisco to attend VMworld. At VMworld, NetApp sponsored a variety of presentations and hands-on labs, plus shared a slew of demos and new technical papers the team spent the past few months developing.

Some of the hottest topics at the show and in the NetApp booth included:

1. **Comprehensive data protection.** NetApp technologies streamline backup, recovery, and disaster recovery across a wide range of different environments and help optimize network infrastructures through the use of unique intelligent data movement capabilities and space-saving features. Additionally, at VMworld NetApp unveiled a sneak preview of a NetApp and VMware co-developed failover utility that enables seamless disaster recovery for virtual machines.
2. **Storage efficiency and virtualization.** VMware environments can consume a substantial amount of physical storage capacity. NetApp advanced RAID, data deduplication, and thin provisioning with policy-based space management capabilities can help significantly reduce provisioned yet unused storage.
3. **Shared infrastructure and cost savings with NAS.** With Virtual Infrastructure 3, organizations can leverage ubiquitous low-cost Ethernet networks to present storage to ESX clusters. Network-attached storage (NAS) delivers simplicity and cost savings, and magnifies the benefits of storage virtualization and storage-based backup/recovery/replication tools.
4. **Cost-saving virtual desktop solutions.** VMware and NetApp have partnered on and jointly presented a virtual desktop infrastructure (VDI) solution that reduces the storage and time to deploy requirements by 90% relative to traditional array-based solutions. At VMworld, NetApp demonstrated the deployment of 100 virtual desktops in under five minutes!

Because you're a Tech OnTap member, we're pleased to bring you early access to all of the brand-new technical reports and many of the demos and presentations NetApp shared at VMworld. If you weren't able to attend VMworld in person, this is definitely the next best thing.

RELATED INFORMATION

- [NetApp/VMware Best Practices](#) (pdf, updated September '07)
- [5 Ways to Use NetApp Snapshot Copies with VMware VI3](#)
- [Love VMware. Hate Backups?](#)

NetApp Snapshot™ Copies and VMware

NetApp Snapshot technology is ideally suited for use with VMware. A recent Tech OnTap article explored five uses of Snapshot and other NetApp technologies derived from it in VMware environments:

- Near-instantaneous VM backup
- Fast and flexible VM recovery
- Accelerated data management through cloning
- Disaster recovery
- Application backup and management

Learn more. [Read Five Ways to Use NetApp Snapshot Copies in VMware Environments.](#)

NetApp Solutions for VMware Environments

In this Webcast, VMware vice president Brian Byun and NetApp vice

1. Comprehensive Data Protection

Backing up and protecting data has become a hot-button issue in the virtualization community and is the number one pain point in large VMware environments. Consolidating tens of physical servers onto a single server often creates a disproportionate ratio of data to physical server bandwidth and can lead to artificial consolidation limits and drive up the number of servers required. As environments grow, traditional backup applications based on a backup agent per machine become increasingly difficult to manage and full backups of each individual VM start to consume excessive space.

NetApp offers a variety of tools enabling instantaneous-and potentially fully automated-backup creation that enables you to copy data off-site much more efficiently and cost effectively than with tape. Backups are streamlined with NetApp solutions to only send and store changed data blocks. Once on NetApp disk, data is immediately usable and can be exposed to a variety of advanced NetApp features for data management, including deduplication technologies and cloning capabilities.

NetApp also offers multiple options for implementing disaster recovery in VMware environments. NetApp SnapMirror™ technology enables flexible and low-cost disaster recovery for existing NetApp users while NetApp ReplicatorX™ enables data replication and recovery across a broad range of platforms and storage systems.

Early access for Tech OnTap members:



- [Disk-to-Disk Backup for NetApp/VMware](#) (PDF - technical report)
A technical overview of how NetApp SnapVault™ technology can be implemented to increase the storage availability of VMware ESX environments. Includes a detailed plan for the set up, configuration, and functional testing of a disk-to-disk backup environment.
- [DR for VMware using NetApp SnapMirror and MetroCluster](#) (PDF -technical report)
A technical overview of the way NetApp disaster recovery technologies can be implemented to increase the availability of VMware ESX environments by providing protection at a local, campus, or regional level. Includes a detailed plan for the set up, configuration, and functional testing of a high-availability storage environment.
- [Simply Addressing the DR Needs of Virtual Infrastructures](#)
(VMworld presentation BC33 - PDF)
Reviews common deployments of SnapMirror with VMware Infrastructure 3, including using synchronous or asynchronous replication over IP or Fibre Channel connections and the flexibility of mixing FC and SATA disks in a replicated environment.

Plus, during the show NetApp shared a sneak peek into the future of disaster recovery for VMware environments. NetApp and VMware have co-developed a new failover utility that enables seamless automated discovery and failover for virtual machines.

Additional resources:

- [Streamlined DR for VMware ESX in Heterogeneous Environments](#)
(Tech OnTap article)
Real-world examples of the way IT teams have leveraged NetApp ReplicatorX to implement cost-effective disaster recovery across multi-vendor storage environments and migrate from ESX 2.1 to ESX 3.0 with only 25 minutes of downtime.
- [SnapMirror: Behind the Music](#) (Tech OnTap article)
One of the original NetApp SnapMirror engineers provides an overview of core product design and the way it is used in a cross-section of real-world deployments. Plus, discover the tricks that the professionals use to avoid or resolve common challenges.

2. Storage Efficiency and Virtualization

president Patrick Rogers discuss the advantages of server-to-storage virtualization and the way NetApp and VMware products complement each other.

Senior systems administrator Jason Guibert of Loyola Marymount University describes his real-world experience deploying and using VMware and NetApp. Key topics include HA, backup and recovery, and cloning.

In a follow-up Q&A, Webcast participants asked Jason about specific technical aspects of his deployment, including iSCSI, RDM implementation, cloning, and failover.

[Watch the Webcast.](#)

Organizations face various challenges when managing storage for their virtual infrastructure deployment. First, data redundancy is prevalent because VMware environments are typically deployed with templates: having 100 Windows™ 2003 servers means you have 100 copies of identical OS and patch data. Second, VMware offers multiple layers of storage virtualization, each of which results in excess capacity. Consider this: If virtual disks and their underlying VMware data stores are all 70% full, the actual storage utilization is only 49%! Ouch!

Storage virtualization—specifically, virtualization using unique NetApp functionality—can address these challenges, significantly reducing provisioned yet unused storage. NetApp A-SIS deduplication, thin provisioning at the data-store level, and FlexClone™ technology all allow storage admins to address more storage than is physically available and to significantly reduce storage costs.



Early access for Tech OnTap members:

- **Best Practices for Implementing VMware VI3 with NetApp FAS Systems**
(PDF - technical report)
Reviews deployment options involving VI3 with NetApp FAS systems and explores topics including NFS, FCP, iSCSI, VMFS, RDM, storage layouts, and more.
- **Getting Familiar with NetApp Storage Virtualization Features** (demo*)
An introduction to NetApp storage virtualization features including flexible volumes, thin-provisioning, and space management. Shows how users can rapidly provision and resize storage entities (FlexVols / LUNs / VMFS datastores) and thin-provision flexible volumes and LUNs. Finally, this demo addresses how space management policies eliminate risks associated with thin provisioning by allowing storage to respond to capacity growth.
- **Rapidly Provisioning Multi-Use Copies of DataStores with FlexClone** (demo*)
Shows how NetApp FlexClone technology enables the creation of thin-provisioned copies of datastores and the virtual machines they contain. This demo elaborates on how FlexClone technology is critical to data recovery of virtual disks or data and demonstrates how FlexClones streamline failover testing for VMware environments being replicated with SnapMirror.
- **Eliminate data redundancy with NetApp A-SIS** (demo*)
Shows how A-SIS deduplication reduces capacity requirements to house Virtual Machines. This demo illustrates the process of de-duplicating an existing datastore, reviews capacity savings after de-duplication, and explains how to schedule de-duplication. It also addresses the advantages that A-SIS brings when data is being replicated with SnapMirror.

Additional resources:

- **NetApp and VMware ESX Server 3.0: Building a Virtual Infrastructure from Server to Storage** (PDF - technical report)
Discusses virtual storage solutions that reduce cost, increase storage utilization and fault tolerance, and address the challenges of backing up and restoring virtual environments.
- **Benefits of Consolidation: From Server to Storage** (podcast)
Phil Brotherton, senior director of Enterprise Alliances and Solutions at NetApp, discusses how storage infrastructure consolidation can increase manageability and responsiveness of VMware environments while lowering TCO.
- **NetApp A-SIS Deduplication: Top 10 Customer Questions Answered**
(Tech OnTap article)
Get the answers to the most popular questions about A-SIS technology, licensing, configuration, performance, and support.

3. Shared Infrastructure and Cost Savings with NAS

With the incorporation of NFS into VI3, organizations can now leverage existing IP networks and NFS storage for VMware. When deployed using NetApp best practices,

NFS-based data stores achieve high levels of redundancy and performance.

Further, NFS delivers space efficiencies through automatic thin provisioning of VMs, removal of the file system layer (VMFS), and the ability to deploy more virtual machines per data store. Finally, NFS data stores enable more flexible recovery and streamlined data replication using storage-based tools.



Early access for Tech OnTap members:

- **Building Virtual Infrastructures with NAS** (VMworld presentation IP43)
This presentation deck demonstrates how implementing VMware VI3 with NAS simplifies storage provisioning, streamlines virtual machine deployment, and enables instantaneous backup and recovery. This presentation also includes a case study detailing a real-world VMware VI3 on NAS deployment.

Additional resources:

- **Virtual Reality: An Architecture Capable of Evolving in Any Direction**
(Tech OnTap article)
See how one IT team built a fully virtual software-as-a-service architecture. Includes a discussion on the various protocols the team evaluated, pros and cons, and why the team decided to implement VMware on NAS.
- **VMware over NFS** (Blog post)
NetApp consulting systems engineer and SAN specialist Nick Triantos recently blogged about his experience testing out NFS over VMware. His post shares observations about throughput, latency, thin provisioning, utilization, and more.

4. Cost-Effective Desktop Virtualization

VMware and NetApp have partnered on and jointly presented a VDI solution that reduces storage and time to deploy requirements by 90% relative to traditional array-based solutions.



Early access for Tech OnTap members:

- **Deploying VDI with NetApp FAS Systems** (PDF - technical report)
Technical overview of the architecture and benefits of deploying VDI on NetApp FAS systems.
- **Optimizing Storage for Virtual Desktops** (VMworld presentation DV16)
A look at the way NetApp storage can drive down costs and management of a virtual desktop solution. Also investigates how A-SIS (Advanced Single Instance Storage) deduplication and FlexClone technology can help minimize storage requirements and ease patch management.
- **VMware VDI on NetApp** (demo*)
See the power of VDI on NetApp with the deployment of 100 virtual desktops in under five minutes!

** Demo may take a few seconds to load; please be patient.*

TECH ONTAP ARCHIVE



**Sammie Carter, Brian Bouterse,
and Patrick Bowen**

The Friday Institute, North Carolina State University

Sammie Carter (left) designed and built the Friday Institute's overall technology infrastructure. He currently manages IT Operations, serves as project manager for all IT-related projects, and specializes in security, systems integration and interoperability, virtualization, grid, and HPC. He works closely with Brian Bouterse (center), a networking and systems architect who has been designing and managing projects for several years and specializes in networking services, software delivery, Web 2.0, and SOA. The team's third core member, Patrick Bowen (right), has worked on the Friday Institute's infrastructure development since August 2005 and plans to enter NCSU as a freshman in the fall of 2007.

Virtual Reality: Building an Architecture Capable of Evolving in Any Direction

By Brian Bouterse, Patrick Bowen, and Sammie Carter

In June, Tech OnTap showcased the NetApp *Kilo-Client* Engineering lab, a 1500-node diskless server farm. This article details how an end user IT team is leveraging the virtual environment concept to deliver applications to remote sites.

The Friday Institute recently developed a new architecture involving server virtualization, advanced cloning technology, and end-to-end management to deliver software services on demand.

While our pilot project focused on the specific needs of K-12 education, this approach could be leveraged by any organization dealing with large numbers of remote sites and minimal remote staff and expertise. Overall, we believe this infrastructure is efficient enough to be adapted for almost any situation requiring the ability to flexibly provision compute, software, and storage resources.

In this article, we introduce the project and resulting architecture, then highlight five of the ways it provides flexibility unimaginable in traditional environments:

- **Multidimensional scalability**, including the ability to choose from various connectivity options (we piloted VMware on NFS but plan to look at iSCSI as 10GbE adoption increases)
- **On-the-fly provisioning and reprovisioning**, which enable VM creation in under two minutes and efficient management of tens or even hundreds of VMs
- **Synergy between server and storage virtualization** to combine rapid processing power allocation with highly efficient storage flexibility
- **Open management tools** capable of managing the entire environment with a single interface
- **Adaptability to tackle problems from new angles** and, for example, easily shift from delivering lots of applications per server to aggregating processing power across servers

Goal: Enable Software as a Service for K-12

Technologies such as virtualization, software as a service (SaaS), and service-oriented architecture (SOA) are transforming IT in virtually every sector—except education. The situation is particularly dire in K-12, where IT is typically extremely understaffed even when budgets are available.

The Friday Institute, a research arm of the North Carolina State University (NCSU)

RELATED INFORMATION

- [Friday Institute: Enabling Software as a Service Through Virtualization](#)
- [NetApp Engineering's 1500-Node Diskless Server Farm](#)
- [Analyst Report: A Practical Application for Grid Computing](#)
- [VMware in NFS Environments](#)

About the Friday Institute

The Friday Institute for Educational Innovation is part of North Carolina State University's College of Education. The mission of the Friday Institute is to advance education through innovation in teaching, learning, and leadership. We conduct research, create resources, advocate to improve teaching and learning, and provide services to educators and policymakers. Our work focuses on innovations that help prepare all students, from preschool through college, to live and work successfully in the 21st century.

The work described in this article is central to one of the five current major initiatives of the Friday Institute: Technology Infrastructures for K-12 Schools. Our other initiatives are Teaching and Learning in the 21st Century, Evaluation of Educational Innovations, Leadership Development, and Policies to Support Innovation. More information is available at www.fi.ncsu.edu.

A detailed 6-page technical paper,

College of Education, is dedicated to addressing this imbalance. In 2007, the Institute's IT team launched an initiative to create an end-to-end architecture capable of delivering software as a service.

Rather than bring more technology into the classroom, our team envisioned an architecture that would enable any instructor or student working from a classroom or online to request and connect to a particular operating system/application running on a remote server without operator intervention and with minimal management overhead.

On-demand computing for K-12 brings applications such as lab—simulation or graphing software to support 21st—century skill development. We believe 100Mb for the clients, and 1 GbE between the IBM blades and NetApp storage will ensure a rich user experience.

Specific goals at the outset included:

- Consolidate many users onto a single server blade to get the most from slim IT budgets
- Increase resource flexibility through the use of virtual machines
- Ensure wait time is less than 2 minutes... every time
- Minimize storage space requirements

The resulting architecture was the product of a successful collaboration between our team and industry leaders including IBM, Cisco, and NetApp. Without their support the project would be more proprietary (less open), would have taken longer to develop, and the solution would not be nearly as rich.

Solution: A SaaS Architecture Built on Virtualization

In our pilot setup, an IBM BladeCenter with five HS20 blades connects to a NetApp FAS940 for backend storage via a Cisco Catalyst 3560G Ethernet Switch. VMware Virtual Infrastructure 3 (VI3) provides server virtualization.

Funding has been received to extend this computing paradigm to support regional shared service alliances, serving both K-12 and higher education within North Carolina.

To manage the architecture we chose openQRM software from Qlusters, which allows us to manage all elements of the infrastructure from a single interface. Qlusters and the Friday Institute are working closely together to ensure openQRM provides a seamless, central management experience for system administrators.

Instead of storing the boot image for each virtual machine (VM) on a local disk, all VMs connect to their virtual disks via NFS over Ethernet.

This architecture, and specifically a key NetApp technology called FlexClone®, enables us to create real-time clones of existing disk—images to meet or exceed our wait time goal of under two minutes. When a resource request is received, a new virtual machine is rapidly provisioned using FlexClone to eliminate the time-consuming data copies that are part of the standard process. We built a number of "golden" volumes in advance, each containing the target operating system and any required applications on that operating system. The process of VM creation then becomes:

1. Use FlexClone to create a "golden" volume
2. Mount and register the cloned volume with VI3 via NFS
3. Register a virtual machine within VI3
4. Power on the virtual machine

Because FlexClone volumes can be created instantaneously, the wait time for VM provisioning is dramatically shortened. In demonstrations, we are able to have 30 Windows® Server 2003 and 20 Linux® FC6 virtual machines cloned, mounted, and booted in 103 seconds.




Figure 1. Conceptual architecture.




Figure 2. Physical architecture.

PowerPoint presentation, and video demonstration of the solution are available on the [Friday Institute Project Overview Site](#).

NetApp Kilo-Client: A Practical Application for Grid Computing

The Friday Institute team architecture takes the virtual environment concept and focuses on how this can be used for application delivery to remote sites. The team used the NetApp Engineering Kilo-Client lab as a model in its early design process.

The Kilo-Client environment includes:

- Use of standard components
- "Podule" architecture
- Automated provisioning
- Centralized monitoring and management
- "Greening" the grid

Learn more:

- [A Practical Application for Grid Computing](#) (PDF)
- [NetApp Engineering's 1500-Node Diskless Server Farm](#) (Tech OnTap article)

A detailed six—page technical paper, PowerPoint presentation, and video demonstration of the solution are available on the [Friday Institute Project Overview Site](#). The basic architecture for supporting VMs is quite similar to that of VMware's [Virtual Desktop Infrastructure](#). With VDI companies can host individual desktops inside virtual machines. Users access these desktops remotely from a PC or a thin client using via RDP. Since applications are managed centrally from the data center, organizations gain better control over their desktops. Installations, upgrades, patches, and backups can be done with more confidence without user intervention.

Result: A Fully Virtual Environment Provides Levels of Flexibility

Inconceivable in Traditional Environments

Our resulting framework is a virtual, open environment that enables us to rapidly adapt to technology changes and emerging needs. Hardware resources can be allocated and deallocated as needed, and the architecture can grow in any dimension—blades, storage, or connectivity. Plus, the software architecture supports rapid deployment of VMs, bare metal, or both, so the architecture can adapt dynamically to a wide range of user requests.

Advantage 1. Grow in Any Dimension: Blades, Storage, or Connectivity

All portions of our system are virtualized, including the servers, disk shelves, and storage appliances. This fact ensures scalability, as any virtualized resource can have physical resources added to its pool. For example, if a storage system requires better performance, one can scale vertically by adding in additional shelves; to support more concurrent connections, one could add additional storage systems (scaling horizontally). A similar analogy can be made with the processing hardware (blades).

VMware VI3 supports NFS, iSCSI, and Fibre Channel plus RDM (Raw Device Mapping), while NetApp storage provides native multiprotocol support for these and other protocols. This offered significant flexibility when it came to choosing a connection methodology and enables us to easily adapt the architecture in response to evolving requirements and technologies.

We considered all five connection methodologies described in the technical report [NetApp and VMware ESX Server 3.0 Storage Best Practices](#) (PDF - technical report) before settling on NFS. Of the five, we did not directly test iSCSI HBAs or Fibre Channel because these solutions were deemed to be too costly.

The following table summarizes our findings with each technology.

Technology	Media	SW Used
NFS	<ul style="list-style-type: none"> • Outperformed iSCSI • Simple setup/admin 	<ul style="list-style-type: none"> • SW iSCSI LUN is required for VMware swap file
iSCSI w/ SW Initiator	<ul style="list-style-type: none"> • Less expensive than FCP 	<ul style="list-style-type: none"> • Requires TOE for performance
RDM over iSCSI	<ul style="list-style-type: none"> • Doesn't require host OS for Xfer • VM talks directly w/ storage 	<ul style="list-style-type: none"> • Maintenance of mount points
iSCSI w/ HBA (not tested)	<ul style="list-style-type: none"> • Less expensive than FCP 	<ul style="list-style-type: none"> • HBAs add significant expense for large deployments
Fibre Channel (not tested)	<ul style="list-style-type: none"> • Fast and reliable 	<ul style="list-style-type: none"> • High cost and complexity

Ultimately, we settled on NFS to meet our current needs because it was easier to set up and manage. NFS fits well in the K-12 environment, as Ethernet is prevalent, and using standard network interface cards was very attractive from a cost perspective. From a storage standpoint, NFS eliminates a layer of storage virtualization and any wasted storage that may be incurred in such a configuration. With NFS we were able to realize higher storage utilization than we could achieve with a traditional VMFS configuration with FCP or iSCSI. Finally, in our testing we were pleased with the overall performance of NFS.

(More information about using NFS to support VMware environments is available in the Tech OnTap article "[What's Hot with VMware and NetApp: Enhanced DR, Virtual Storage, NAS, and VDI](#)")

In the longer term, we believe that iSCSI will be the best choice once the protocol matures and 10 Gigabit Ethernet becomes widely adopted. For both NFS and iSCSI, as speeds approach 10 Gb speeds, the computational overhead severely taxes the

CPU unless a TCP/IP offload engine (TOE) is being used. The NCSU Centaur Networking Lab provides a great overview of the problems experienced with 10GbE network as related to TOE cards in a recent paper, "[1Gbps and 10Gbps Ethernet Server Scalability](#)" (PDF). The switch to iSCSI avoids the complexity of managing the exports file, while simultaneously providing a more straightforward management paradigm via iSCSI's initiator and target framework.

Advantage 2. On-the-Fly Provisioning and Reprovisioning

Our key requirement was the ability to create new virtual machines in under two minutes. Regardless of the approach you use, you ultimately have to clone or copy the volume or volumes that a virtual machine will use. We looked at VMware native capabilities, but, since VMware has to work with whatever storage and server hardware is available, copying is naturally the default approach. Normal VMware provisioning takes two to three times longer than this process. We also explored offerings from various storage vendors, including Equallogic, IBM, and Apple, but all of them used bit-for-bit copying under the hood, giving NetApp and FlexClone a clear advantage.

Using FlexClone saves us both time and space. When you create a FlexClone volume on NetApp storage, it only consumes additional space as changes are made to the clone volume. This is why you can create a clone so quickly (in seconds). We did some measurements, and (in our environment) making a FlexClone volume only consumed about 0.07% of the space used by the original volume to create the clone as measured immediately after clone creation.

We can create dozens of clones without consuming a significant amount of disk space up front. Of course, as files are added or changed, space consumption increases, but each VM has many gigabytes of system files and software that never changes, so when you're dealing with tens or even hundreds of VMs the space savings are huge. No matter what you're doing inside a VM, FlexClone is consistently maximally efficient versus other approaches. This fits perfectly with our software-as-a-service model where you only pay for what you actually use.

A video demonstration is available on the [Friday Institute Project Overview Site](#).

Advantage 3. Synergy Between Server and Storage Virtualization

The core technologies we used in this architecture—FlexClone and VMware ESX—have tremendous synergy. The complementary relationship between the two makes it a simple and straightforward process to connect the components together. VMware provides an excellent mechanism to provision multiple virtual machines on shared hardware such as CPU, memory, networking, and so on. NetApp storage virtualization, which truly is unlike any other storage platform on the market today, allows us to provision dozens of virtual machines instantly and without consuming any storage outside of what is required to store the first VM.

Incidentally, of all the components of the architecture, NetApp was the easiest and the fastest to configure and the most reliable. We had no technical issues with the system crashing, and we also got great support from our local team. Finally, NetApp has a wealth of best practices, deployment guides, and other resources that we were able to draw on during the deployment.

Advantage 4. Single Open Management Platform

The tricky part of virtualization is management, especially when you're talking about managing potentially large numbers of virtual machines in an environment where you are constantly turning them over plus virtualized storage. Because our solution will likely be deployed in environments with limited IT resources, management becomes that much more important.

One trend we've noticed in the virtualization area lately is that tools are becoming available that manage a broader range of technologies. We chose Qlusters' openQRM software for exactly this reason. It not only manages virtual machines, it can also manage bare metal and has ties into NetApp storage as well. This gives us one solution that can manage the whole infrastructure.

Advantage 5. Adaptability to Tackle New Problems

While we're currently looking at deployments in environments where server virtualization is needed to more finely partition resources, there are other situations where a job requires the complete resources of one or multiple blades, and bare metal makes more sense.

By developing a fully virtualized environment, we've created an architecture that is flexible enough to attack problems from different angles. We have the ability to leverage the benefits of virtualizing everything to deliver many applications per server. However, we can also easily go the other direction to provision bare metal blades with

diskless booting to meet the high node-count requirements of HPC without significant architectural changes. Our hardware architecture, our software provisioning architecture, and our management platform all will support this.

Summary

We believe the implications of this framework are significant. By leveraging the unique capabilities of a variety of products plus the overwhelming generosity and collaboration from partners, including IBM, Cisco, NetApp, VMware, and Qlusters, we've created an architecture that offers flexibility and scalability no traditional environment can match.



Sanjay Gulabani, Senior Database Performance Engineer, NetApp (pictured)
Darrell Suggs, Technical Director and Cache Architect, NetApp

After spending eight years at Oracle Corporation, Sanjay joined NetApp with a focus on improving NFS and iSCSI client performance, writing best practices, and troubleshooting escalated issues. He works closely with Darrell, who focuses on performance engineering, has authored several key performance tools, and has helped guide NFS performance enhancements and the development of FlexShare™. Other database experts who provided invaluable contributions to this article include Steve Daniel, director of Database Platforms and Performance Technology and Vasu Subbiah, director of Product Management.

Oracle on NetApp Part I: Using NFS for Simplicity and Manageability

By Sanjay Gulabani, Darrell Suggs, Steve Daniel, and Vasu Subbiah


This is the first in a three-part series about Oracle® storage architectures. Next month's article explores how NetApp has optimized storage for Oracle in FC SAN environments, and November will feature a technical case study involving Oracle on NetApp using NAS, iSCSI, and FC SAN.

There is no simple "one size fits all" solution when it comes to storage system design for Oracle. Performance, availability, and of course budget requirements vary dramatically. Choosing a storage protocol may be one of the most critical decisions most IT teams make—although many may not think of it as a decision at all.

IT teams have traditionally deployed block-based protocols to support Oracle databases and applications, but a large number of companies today run Oracle on the Network File System (NFS). Many of these organizations—ranging from small IT shops to major Fortune 100 companies, including Oracle itself—use Network Appliance™ storage. NetApp has been a leader in the NFS market since 1992 and has pioneered and legitimized the use of NFS with Oracle.

This article describes how NetApp engineering has worked to make NFS a robust storage option for Oracle. Key investments include:

- Enhancing NFS implementations for database
- Partnering to ensure high availability
- Driving industry standards including NFSv4
- Guiding emerging technologies like Oracle Direct NFS
- Developing integrated, Oracle-specific tools
- Defining Oracle on NFS best practices



We'll begin by examining why it makes sense to design an Oracle storage architecture based on NFS in the first place.

Why NetApp NFS for Oracle?

Oracle on NetApp NFS provides users with an easy to manage and flexible infrastructure capable of running a large number of database instances. Ethernet-based NFS facilitates simple provisioning, management, and sharing plus the flexibility to rapidly move and re-purpose database and storage resources. NFS implementations are also typically cost-effective as they do not require special Fibre Channel SAN expertise or training and allow IT teams to leverage standard Ethernet knowledge.

Oracle on NetApp NFS can be deployed for use in virtually any type of Oracle environment ranging from test/dev to 24x7 productions OLTP (including ERP and

RELATED INFORMATION

- [Best Practices for Oracle](#) (PDF)
- [Optimizing Oracle on NFS](#) (PDF)
- [Webcast: Data Center Solutions for Oracle](#)
- [Eisler's NFS Blog](#)
- [Dave's Blog: Oracle Optimizes Its Database for NFS](#)

NetApp Blogs Involving NFS and Databases

Updates, presentations, and observations on many of the topics discussed in this article are available from technical blogs maintained by author Sanjay Gulabani and NetApp technical director Mike Eisler. In addition, NetApp founder Dave Hitz recently blogged about the origins of Direct NFS.

All three blogs include an opportunity to post comments and submit questions:

- [Gulabani's Databases on NetApp Storage Blog](#)
- [Eisler's NFS Blog](#)
- [Dave's Blog: Oracle Optimizes Its Database for NFS](#)

Empowering DBAs with SnapManager for Oracle

For Oracle DBAs, backup and recovery is a complex task that may require frequent negotiations with both system and storage

custom applications) environments. NetApp customers have found that—when properly configured and tuned—Oracle on NetApp NFS delivers 80-90% of the transactional performance of Fibre Channel using identical hardware and test conditions, and NFS performance is typically comparable to iSCSI. NFS can achieve the same absolute performance levels as FCP through the addition of hardware resources.

The observed differential with FCP results primarily from two factors:

- FCP protocol processing (or iSCSI when an iSCSI HBA is in use) are offloaded from system CPUs to the HBA. Since all NFS processing is done on system CPUs, NFS requires slightly more CPU resources to do the same job.
- Fibre Channel networks currently have a per network bandwidth advantage (4Gbits/sec is common) versus IP networks (1Gbit/sec is common). As 10Gbits/sec Ethernet becomes more prevalent, this advantage will go away.

NetApp has done a careful evaluation of NFS performance versus FCP and other protocols on a variety of environments, including Oracle 10g™ on [Linux®](#), [Solaris™](#), [HP-UX](#), and [AIX](#), as well as [Oracle 10g RAC on Linux](#). You can refer to these reports for more information relative to your particular requirements.

Protocol Comparison	
NFS	FC SAN
<ul style="list-style-type: none"> • Simpler provisioning/management • Less expensive hardware support • Easy sharing to support RAC/grid 	<ul style="list-style-type: none"> • Highest performance per CPU • Utilizes host file system or ASM • Leverages FC investments

Enhancing NFS for Database Workloads

NFS version 3 (NFSv3, the version that is in wide use today) consists of server software—for instance, the software that runs on NetApp storage—and client software running on database servers or other hosts that require access to network storage. Achieving good performance and ensuring proper behavior require that both sides of the connection, client and server, are correctly implemented.

To ensure robustness on the server side, NetApp pairs its time-tested server code with a hardware platform that delivers proven availability and reliability through features such as active-active controller configuration and double-parity RAID.

The NFS client can be implemented as part of any of a number of operating systems, including Solaris, HP-UX, AIX, and Linux. NetApp engineering has worked closely with all of the major implementers of NFS to verify that client implementations are functionally correct and perform well for database workloads. NetApp has a team of interoperability engineers who have [carefully studied the requirements for a storage interconnect to recover correctly from error conditions](#) and have worked to make sure that NFS implementations consistently meet those requirements.

NetApp engineering's work with the engineering teams of various vendors has led to the wide implementation of the following NFS enhancements:

- Direct I/O is the ability to do NFS I/O directly from an application over a network to storage, bypassing the NFS client buffer cache. Direct I/O avoids redundant caching on the server, gives Oracle control over client I/O behavior, and supports data-sharing semantics required by Oracle RAC.
- Asynchronous I/O allows a database to manage many concurrent I/O transactions without processes blocking on individual transactions. This results in greater throughput by reducing process counts and overhead.
- Elimination of kernel bottlenecks that throttle client I/O enables support for large numbers of concurrent I/O operations.
- Additional features support Oracle RAC on NFS.
- Quorum and I/O fencing are supported, in conjunction with high-availability partners.

Partnering for High-Availability

Clustering is important to ensure high availability for critical database servers. NetApp has partnered with Oracle, Symantec, and Sun to ensure that NFS is supported for use with Oracle RAC, Veritas™ Cluster Server (VCS), and Sun™ Cluster, respectively.

NetApp and Sun have worked together to develop a package for Sun Cluster that

administrators.

SnapManager for Oracle integrates closely with the Oracle Database, allowing DBAs to manage their own data while offloading data management work from Oracle servers to NetApp storage. From a central management system, a DBA can monitor the state of all Oracle backups across multiple databases, even if they are running on different servers and operating systems.

In a recent article, NetApp database architect Alvin Richards took an in-depth look at the inner workings of SnapManager for Oracle, including:

- Snapshot copies
- Full and partial restores
- Cloning
- ASM integration

[Read the article](#) to find out more. You can also watch a [demo](#) of SnapManager for Oracle in action.

A Case Study: Architecting Storage for Testing/Development

Ensuring the effectiveness of your development and testing efforts is critical, but the simple act of provisioning dev/test environments can consume vast amounts of storage and become a development bottleneck.

The company profiled in this case study first moved to NFS to simplify sharing and data management. Ultimately, they settled on a solution using NetApp FlexClone, allowing them to create space-efficient clones of production data sets in almost zero time.

[Read the case study](#) to get the full details.

allows clustered servers to access shared storage via NFS rather than a shared Fibre Channel or SCSI connection to disk. The NetApp software includes support for I/O fencing and quorum to ensure proper behavior in the event of a server failure. The NetApp software package for Sun Cluster supports both Sun SPARC and x86 platform running Sun Solaris.

The VCS agent for NFS supports Oracle running on a broad range of operating systems (Linux, AIX, Solaris, etc.) and is a good option for customers who need a high-availability solution for single-instance failover.

Both the Sun package and the VCS agent for NetApp NFS are available on the [NOW™ \(NetApp on the Web\) site](#).

Driving Industry Standardization

NetApp chief scientist Steve Kleiman was one of the inventors of NFS, and over the years NetApp technical experts have taken leadership roles in helping define industry-wide standards. For example:

- Brian Pawlowski (NetApp chief technical officer) coauthored the NFSv3 specification and is cochair of the NFS version 4 (NFSv4) Working Group of the Internet Engineering Task Force (IETF).
- Garth Goodson (NetApp engineer) is co-author and editor of the pNFS portion of NFSv4.1.
- Mike Eisler (NetApp senior technical director) is the primary author of RFC 2203, which enhances the security of NFS, and recently coauthored RFC 3530, the NFSv4 specification. He is currently an editor for the NFSv4.1 specification. Prior to joining NetApp, Mike led the SEAM project at Sun that produced the first NFS implementation that used Kerberos V5 authentication.
- Tom Talpey (NetApp technical director) coauthored several of the NFS over RDMA specifications at IETF.
- Dave Noveck (NetApp senior engineer) is also a coeditor of the NFSv4.1 specification and a coauthor of RFC 3530.

NFSv4 is designed as an even more robust implementation of NFS and includes features that are beneficial in Oracle environments. For instance, the NFSv4 redirection feature allows a storage grid and a compute grid to mutually optimize I/O paths. A storage system can tell a compute server which storage system can best service particular requests to facilitate grid-based scale-out.

NFSv4 is currently gaining acceptance, and we expect to see a broader support from database applications for NFSv4 in the future. Current information about NFSv4 developments is available on [Eisler's NFS Blog](#).

Optimizing Database Access via Oracle Direct NFS

Direct NFS is an NFS client created by Oracle that is shipping as part of Oracle 11g™ (announced in July, 2007). Instead of relying on the operating system, Oracle's Direct NFS client generates NFS requests directly from the database.

In a recent blog post, NetApp founder Dave Hitz explained that [Oracle optimized its code specifically for NFS](#) based on the company's experience using NFS to run applications on tens of thousands of Linux servers accessing many petabytes of NetApp storage. NetApp worked closely with Oracle on Direct NFS development, providing technical consulting as well as functional and performance testing resources.

This Oracle client implementation delivers several unique advantages:

- Oracle can use the same NFS client, regardless of the platform it is running on, for more consistent and predictable performance.
- Because the client spans platforms, bugs can be identified and fixed faster.
- This will be the first NFS client validated for Oracle on Windows®.
- No mount option tuning is necessary. Existing clients require specific mount options for optimal Oracle performance, and the options differ for each operating system.
- It is easy to distribute NFS workloads across up to 16 network links. This has been difficult with traditional NFS clients.

Developing Integrated Storage Management Tools for Oracle

Oracle DBAs like running Oracle on NFS because it essentially gives them the ability to provision their own storage. Because Oracle can create files or make them bigger, the

DBA can simply create or grow tablespaces and other files as necessary, without having to contact a storage administrator to add LUNs or grow underlying LUNs every time. NFS also makes it easy to move storage or applications from one server to another; you don't have to do any complicated rezoning of switches, etc. Because NFS is a shared file system, by design it works with Oracle RAC.

Snapshot and FlexClone

Several of the capabilities of NetApp storage contribute to the ease of Oracle data management in NFS environments. NetApp Snapshot™ copies are space-efficient, point-in-time copies that can be created in a matter of seconds. To create an Oracle backup, you only need to put the database in hot backup mode for a few seconds while a Snapshot copy is being created. The impact to operations is usually negligible, so most DBAs schedule multiple Snapshot copies throughout the day as a safety net. You can retain up to 255 Snapshot copies of each storage volume and restore individual files from previous Snapshot copies or revert to a previous copy if a problem occurs.

Typically, IT shops using Oracle keep a number of Snapshot copies on primary storage to meet immediate needs. Snapshot copies can also be copied to secondary storage using NetApp SnapVault®, mirrored to a remote disaster recovery site using SnapMirror®, backed up to tape, or any combination of these options—whatever is necessary to meet data protection requirements. Because these operations occur on the storage system, there is no impact to running Oracle applications.

NetApp FlexClone® is another option that has proven to be highly useful in Oracle and other database environments. FlexClone allows you to create a writeable clone of any data volume. New disk space is consumed only as the original volume and the clone deviate from each other. This means that you can easily clone production data for test and development, data mining, or other purposes without the typical 2X storage requirement. As with Snapshot, you can create up to 255 FlexClone volumes of a single volume. (A recent Tech OnTap case study described the use of [FlexClone to streamline Oracle application development and test](#).)

NetApp has also created a number of value-added software tools specifically for Oracle environments.

SnapManager for Oracle

SnapManager® for Oracle is a management tool that simplifies the management of Oracle backup and recovery. It works across all storage protocols and integrates closely with Oracle ASM and RMAN. SnapManager for Oracle makes it easy to schedule the creation of consistent Snapshot copies and FlexClone volumes, enabling operations that previously required complicated scripts to be accomplished with a few clicks or scheduled for regular execution. (See sidebar for details.)

SnapValidator for Oracle

In rare fault conditions, data can get corrupted on the data path between server and storage. Oracle keeps a checksum on each block so that this corruption is detected, but in many cases it may be months before data is actually reread and checked, resulting in the need to recover from a very old backup. To address this issue, SnapValidator® for Oracle complies with the Oracle hardware assisted resilient data (HARD) initiative to verify the checksum and block offset every time a block is written. When a problem is detected the write is failed, forcing the server to repeat the write so that no restore is required.

This feature works particularly well with NFS. Traditional host file systems cannot support HARD validation due primarily to the mixture of database blocks and file system metadata in writes to the storage system. Under Oracle 9i™, HARD validation is practical only for raw disk and NFS; Oracle 10g extended that capability to Oracle ASM.

SnapLock Integration with Oracle

Oracle includes partitioning and transportable tablespace capabilities that make it possible to relocate subsets of your data to different storage. SnapLock® integration makes it possible to relocate content and messages from Oracle to write once, read many (WORM) volumes. This ensures that the data cannot be modified, helping to meet the requirements of Sarbanes-Oxley (SOX), HIPAA, and other regulatory guidelines. This feature is particularly advantageous with NFS because of the finer level of granularity that NFS provides; you can choose to use SnapLock on a particular file or sets of files. In a SAN environment, you would have to use SnapLock on an entire LUN.

Defining Oracle on NFS Best Practices

Although deploying Oracle on NFS is fairly straightforward, NetApp has formalized a

variety of best practices to optimize performance. The NetApp Technical Report library includes a number of implementation guides and best practices, including a 34-page guide to [Best Practices for Oracle](#) (PDF). NetApp engineers have also teamed with peers at leading OS vendors to document best practices for specific environments. A recent [joint paper by Sun and NetApp](#) (PDF), for example, provides proven guidelines for optimizing Oracle on NFS in a Solaris environment, including Oracle tuning options, appropriate mount options for the Solaris NFS client, optimal database layout on NetApp storage, and NFS tuning on NFS storage.

Summary: Choosing the Right Protocol for your Environment

If your primary objective is to create a simple, easily managed infrastructure that is capable of flexibly supporting a large number of Oracle database instances, NetApp NFS may be the right choice for you. NetApp has a proven track record of Oracle deployments on NFS, and the combination of NetApp storage features such as Snapshot and FlexClone with Oracle-specific software like SnapManager for Oracle, SnapValidator, and SnapLock can make the lives of DBAs and storage administrators much simpler.

Ultimately, of course, your choice of protocol depends on a wide variety of factors that include existing infrastructure, processes, and skill sets, in addition to the relative capabilities of each technology. NetApp offers excellent implementation options and tools for deploying and managing Oracle environments for NFS, Fibre Channel, and iSCSI, alone or in combination, so that customers deploying Oracle databases on NetApp storage realize best-of-breed capabilities on the protocols that fit their requirements.

For more information, visit NetApp at Oracle OpenWorld in San Francisco, November 11 -15, 2007. Stop by booth 1602 or attend one of four NetApp special sessions, including the following:

Ease Oracle Database Management using Advanced NFS Features in Oracle Database 11g

Oracle Database 11g provides Direct NFS, which simplifies storage management and enhances performance, scalability, and reliability. This session discusses how Oracle Database 11g leverages NFS to improve storage performance and reliability, regardless of operating environment and without the need for expert tuning. It also covers the attributes of an ideal Oracle storage system, including considerations to successfully implement NFS for Oracle.

For full details, check out the [Oracle OpenWorld Overview](#).

NetApp and Veritas™ NetBackup™ Top 10 Customer Questions Answered

Veritas NetBackup users are moving to disk-based backup solutions for a variety of reasons, including improved reliability, faster backups, and improved recovery. NetApp has partnered closely with Symantec to develop two advanced data protection solutions tailored for use with large NetBackup environments:

- **SnapVault™ for NetBackup** allows NetBackup customers to add NetApp NearStore™ and SnapVault software for a space-efficient method of protecting file and database data, all managed by NetBackup. This tool provides on-the-fly deduplication for a highly storage-efficient solution that leverages NetBackup management, monitoring, and reporting. Based on a joint development project between Symantec and NetApp, SnapVault for NetBackup is the first product to implement the Open Storage API and has been widely deployed.
- **NearStore Virtual Tape Library** (NearStore VTL) provides the highest possible sustained performance for optimizing tape backups and restores. NearStore VTL is an enterprise-class system that delivers superior speed, reliability, and integration for existing NetBackup environments. The combination of efficient VTL tape creation, fast restores from Shadow Tape, and hardware-accelerated compression dramatically improves the performance of traditional tape systems while reducing the overall costs of backup and restores.

Both options can be implemented with little or no impact on existing NetBackup environments.

A recent TechTalk [Webcast](#) highlighted how two IT organizations deployed disk-based backup solutions using NetBackup, one with NearStore VTL and the other using SnapVault for NetBackup. A related chat session generated a variety of additional questions (see [chat transcript](#)).

Since Tech OnTap members have expressed interest in NetBackup, we've compiled answers to 10 commonly asked questions:

1. [How do NetApp solutions reduce the cost of D2D backup?](#)
2. [Which versions of NetBackup are compatible with NetApp technologies?](#)
3. [Which applications and databases can SnapVault for NetBackup protect?](#)
4. [We use SnapVault for backup to a NetApp NearStore. Can NetBackup manage this process?](#)
5. [Can SnapMirror™ protect SnapVault for NetBackup destination volumes?](#)
6. [Does SnapVault for NetBackup offer data deduplication?](#)
7. [Does NetApp A-SIS deduplication work with NetBackup?](#)
8. [What advantages does NearStore VTL provide compared with backing up to a NetBackup disk storage unit \(DSU\)?](#)
9. [Can NetBackup communicate via NDMP to a NearStore VTL?](#)
10. [Does the NearStore VTL support data deduplication?](#)

1. How do NetApp solutions reduce the cost of D2D backup?

RELATED INFORMATION

- [SnapVault for NetBackup](#)
- [NearStore VTL](#)
- [Demo: NetBackup Storage Optimization](#)
- [Optimizing Data Protection](#) (pdf)

Whiteboarding Session: NetBackup Integration with NetApp Disk

Engineers from NetApp and Symantec have worked hard to improve disk-based backup by integrating NetBackup management capabilities with NetApp SnapVault software and NearStore secondary storage.

In this whiteboarding session, two technical leads take an in-depth look at the resulting technology.

Topics include:


- Customer demands driving integration
- Technical integration points
- The benefits of the joint solution versus tape and other disk-based backup solutions
- How data reduction using single instance storage works

[Watch the session.](#)

If you use NetApp secondary storage, SnapVault software is typically the most efficient way to perform D2D backups. SnapVault leverages the unique NetApp Snapshot™ technology to replicate only blocks that have changed since the last backup operation. This is much more efficient than traditional backup methods, which usually copy an entire file even when only a small number of blocks were modified.

SnapVault for NetBackup brings these advantages to file and application data being backed up by NetBackup 6.0 or greater. It reduces backup storage requirements by using "on-the-fly" deduplication, and easily integrates into existing NetBackup configurations.

If the data you want to back up isn't on NetApp storage but instead resides on EMC, Hitachi, HP, IBM, or any other storage platform, NearStore VTL provides the speed and reliability of disk-based backup technologies without requiring any changes to your existing software or backup process. In addition to reducing storage requirements through the use of advanced compression technology, NearStore VTL reduces costs by simplifying backup management and increasing the ROI of the existing tape infrastructure.



2. Which versions of NetBackup are compatible with NetApp VTL and SnapVault technologies?

NearStore VTL works with any version of NetBackup. SnapVault for NetBackup works with NetBackup 6.x. or greater.

3. Which applications and databases can SnapVault for NetBackup protect?

SnapVault for NetBackup currently supports several leading databases, including Oracle™, Microsoft™ SQL Server™, Sybase, and Lotus Domino Server. SnapVault for NetBackup support for Microsoft Exchange Server is planned for 2008.

Supported NetBackup clients can also be installed within VMware virtual machines when used with SnapVault for NetBackup. Please refer to the NetBackup-supported platforms matrix found at support.veritas.com for full details about VMware support.

4. We use SnapVault to vault data from NetApp primary storage to a Netapp NearStore. Can NetBackup manage this backup process?

Yes. Starting with NetBackup 6.0, SnapVault backups can be managed using the NetBackup interface. With NetBackup 6.0 you can manage and monitor the SnapVault replication and catalog the files that are backed up using SnapVault.

5. Can SnapMirror™ protect SnapVault for NetBackup destination volumes?

Yes. In fact, data replicated with SnapMirror does not lose its space efficiency during the replication process. SnapMirror support is available in two ways: as part of NetBackup versions 6.5 and 6.0 MP5, or separately from Symantec for use with NetBackup 6 MP4. SnapMirror provides a highly storage-efficient method for protecting NetBackup data residing within SnapVault destination volumes for cost-effective disaster recovery.

6. Does SnapVault for NetBackup offer data deduplication?

Yes. SnapVault for NetBackup incorporates time-based deduplication, which can reduce backup storage requirements by up to 90% (sometimes even more!). This form of deduplication is done "on the fly" and requires no temporary storage. As customers run the same NetBackup jobs over, they naturally back up many of the same files, which may have changed very little. However, after a file is backed up once using SnapVault for NetBackup, only the blocks that changed since the previous backup are saved, eliminating the storage waste that comes from storing the same data over and over again in successive backups. This is made possible through close integration with NetBackup and understanding of the NetBackup data format.

7. Does NetApp A-SIS deduplication work with NetBackup?

A-SIS can deduplicate existing NetBackup images on a NearStore secondary storage system. However, it does not convert them into SnapVault images. SnapVault for NetBackup is required to convert existing images to a combined SnapVault/NetBackup format. Support for A-SIS deduplication is planned for a future release of SnapVault for NetBackup in order to further increase storage efficiencies.

8. What advantages does NearStore VTL provide compared with backing up to a NetBackup disk storage unit (DSU)?

NearStore VTL provides significant advantages over basic DSUs in the areas of management, scalability, and performance. It appears exactly like a tape library to NetBackup, enabling it to be easily managed and shared without changing processes and without requiring backup administrators to become disk specialists. Because it is optimized specifically for tape I/O, NearStore VTL typically delivers 5 to 10 times higher aggregate throughput than a corresponding disk system, enabling it to meet the most challenging data center backup windows.

9. Can NetBackup communicate via NDMP to a NearStore VTL?

Yes. NetBackup can stream backups to a NearStore VTL in any format supported with physical tape, including NDMP. You can typically improve NDMP performance by up to 30% by writing to a VTL instead of tape.

10. Does the NearStore VTL support data deduplication?

NetApp has announced plans to offer an upgrade that will enable NearStore VTL data deduplication in a future software release. Unlike other VTL systems, NearStore VTL deduplication technology has been designed to meet data center performance requirements by supporting both in-line and post-process methods of removing redundant data.

TECH ONTAP ARCHIVE

"People who've never worked with NetApp storage would be most surprised by..."

With characteristic candor and creativity, hundreds of NetApp users completed this sentence. Contest entrants reviewed the top 20 semi-finalists, and their votes about what non-NetApp users don't realize determined our winners.

In general, more than a third of you emphasized NetApp ease of use and simplicity, 13% highlighted unique NetApp technology like WAFL® and Data ONTAP®, and 7% referenced performance and reliability.

"...the level of dedication within the entire NetApp team to offer or recommend the solution that is truly in the best interest of the customer, even if it is the less profitable route."

- Gillian Grant (Canada)
Grand Prize Winner of a NetApp Swiss Army Watch



Tech OnTap T-shirt Winners:



"...Replacement disks arriving before you even knew you had a failure."

- Tech OnTap member (U.K.)

"...how much more they could do with their time and , most of all, the nice peace of mind they can achieve."

- Garda Vasile, IT manager (Romania)

"...the ability to completely scale from 20MB to 1,000TB in one single system with every connectivity protocol."

- Tech OnTap member (U.S.A.)

"...the fact that NetApp actually offers a strong enterprise Fibre Channel and SAN solution, and not only NAS."

- Tim McGue, sr. technical team lead (U.S.A.)



"...how much more they get to see their children now that hours of maintenance are a thing of the past!"

- Tech OnTap member (U.K.)

"...a product line that works as advertised, doesn't require advanced degree in quantum mechanics to deploy."

- Michael Royer, director, Enterprise Systems, Ticketmaster (U.S.A.)

"...the simplicity of NetApp technology. This company is not a hardware company, it is an intelligent software company."

- Joanne Major (Canada)



"...the glowing red ball in the sky as they drive home at regular time (as opposed to spending time in the data center)."

- Rajeev K. Karamchedu, Computer Systems Manager (U.S.A.)



"...the fact that NetApp actually offers a strong enterprise Fibre Channel and SAN solution, and not only NAS."

- Tech OnTap member (Switzerland)

"...the much lower divorce rate among NetApp storage administrators than those who use other storage systems."

*- Reinoud Reynders,
IT manager (Belgium)*

