

HIGHLIGHTS

- ▣ Exchange Archival Case Study
- ▣ Real-World Storage Performance
- ▣ Hospital Prescribes Oracle on NFS
- ▣ Resources:
 - Demystifying Dedupe Savings Ratios
 - High-Performance Workloads Guide

From Months to Minutes: Rapid Apps Provisioning Case Study

Rich Angalet, Manager, Sprint

Dale Elmer, Director of IT Operations Management, Sprint

Want to mount up to 1TB of storage in just a few seconds? How about deploying a standard three-server/database/Web environment in minutes with one only person? The Sprint team shows you how.

[More ▣](#)

BLOGGING WITH DAVE

Dave Hitz, NetApp Founder and EVP

"Controversy: NetApp Outperforms EMC in SAN Database Benchmark." One key takeaway from this result is that turning on a simple feature like snapshots can radically change performance. Don't let a bad experience with EMC's snapshots scare you away from NetApp's."

▣ Dave's Blog

DRILL DOWN

- ▣ High-Performance Random-Access Workloads Config and Tuning Guide
Read up on best practices NetApp has developed for improving the performance of its storage systems for demanding workloads.
- ▣ iSCSI Software Boot for Windows
Step-by-step guide on how to prepare, configure, and boot servers using NetApp.

ADMIN RESOURCES

- ▣ **Demystifying Dedupe Savings Ratios**
Confused about how to get the real dedupe savings ratios? Watch this video and learn how to get apples-to-apples comparison between vendors.
- ▣ **Accelerate Application Development**
Hear how Marvell Semiconductor Inc. reduced database cloning time by 83% to deploy applications faster. Plus, an online [technical Q&A transcript](#) from the Webcast.

TIPS FROM THE TRENCHES

Oracle on NFS: One Hospital's Prescription for Maximum Flexibility

Jess Carruthers, Project Manager

Architecting a storage environment for Oracle® can be a tricky proposition. While there is no simple "one size fits all" solution, this article takes a closer look at one busy medical center that chose NFS to meet its Oracle storage needs.



[More ▣](#)

Implementing Exchange Archival: Mailbox Management and Compliance

Shaun Mahoney, Consulting Systems Engineer, NetApp

E-mail archival often has two components, archival for mailbox management and journaling for regulatory compliance. Learn about the planning and implementation processes for a large financial company on how it was able to:

- Migrate most users to Exchange Server 2007
- Add archival capability to get the volume of e-mail under control and eliminate the need for PST files
- Add journaling to meet regulatory compliance


[More ▣](#)

ENGINEERING TALK

Real-World Storage Performance Benchmarking NetApp Versus EMC

Stephen Daniel, Director, Database Performance, NetApp

The NetApp FAS3040 head to head with the EMC CLARiiON CX3 Model 40 with snapshots enabled. Get the full download on what makes SPC-1 real life and results of the benchmark.



[More ▣](#)

FEEDBACK

TECH ONTAP ARCHIVE



Rich Angalet
Manager, Sprint

With over 25 years of IT experience, Angalet (left) has expertise in operations, hardware, operating systems, network, data centers, facilities, and automation. He is currently responsible for the implementation of Sprint's 4S initiative. Rich attended Rutgers University and enjoys motorcycles, classic cars, and the outdoors.

Dale Elmer
Director of IT Operations Management, Sprint

Dale Elmer started in IT in 1976 at Centel, a small utilities company that merged with Sprint in 1993. Dale held various IT positions within Sprint. Prior to his current role, Dale was the director of Quality Assurance for Sprint's 4S initiative.

Simplicity, Speed, Standards, and Stability: Provisioning Model for Rapid App Deployment

By Rich Angalet and Dale Elmer

At Sprint, new applications undergo testing by our Test Environment Operations team before we release them into production for use by internal customer groups. This follows a familiar development/test/production cycle.

But, as the organization and IT infrastructure grew in both size and complexity, we started experiencing significant slowdowns in our ability to deploy test and production environments, sometimes moving into months before the underlying infrastructure could be arranged.

With more and more solutions being developed and awaiting internal release to production, our team began looking for ways to support the more rapid innovation and application deployment cycles Sprint required.

This is when we began discussing the factors keeping us from achieving rapid deployment of our test environments. This is also when we birthed our initial "4S" services provisioning model. Having heard about the success of [Australia's Telstra](#) with its OmniPresence "Storage Everywhere" project, our Sprint team started to consider how a more flexible, server-and-storage-farm infrastructure could help us more quickly provision, tear down, and restore our test and production environments. This services provisioning model ended up surrounding the four S's we viewed as critical to achieving this goal:

- **Simplicity:** We knew the current deployment process—involving multiple teams and multiple IT layers—was way too complex. Our new infrastructure would have to take fewer steps and require less support from various groups to deploy test environments.
- **Speed:** We set out to achieve a zero-hour service level agreement for delivery of test environments to our customers. This was a high bar to achieve, since our typical environment delivery cycle could take weeks or even months.
- **Standards:** We had so many variations in our environment. It was unbelievable. By standardizing on the key infrastructure components to be used in the server, storage, database, middleware, and application environments, we thought it would simplify our efforts and deliver what we needed a lot faster. We wanted to get to the point where we could quickly build 50 to 100 hundred servers exactly alike, if we wanted.
- **Stability:** We wanted our deployments to be stable enough to be built up, then torn down or rebuilt just as quickly, with the ability to quickly reprovision the freed capacity for another test or production environment.

RELATED INFORMATION

- [Telstra Delivers Storage Everywhere](#)
- [Friday Institute: Enabling Software as a Service Through Virtualization](#)

About Sprint Nextel

Sprint Nextel offers a comprehensive range of wireless and wireline communications services bringing the freedom of mobility to consumers, businesses, and government users. Sprint Nextel is widely recognized for developing, engineering, and deploying innovative technologies, including two robust wireless networks serving approximately 54 million customers at the end of 2007, industry-leading mobile data services, instant national and international push-to-talk capabilities, and a global Tier 1 Internet backbone.

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

Many types of organizations involve large numbers of geographically distributed locations. This can leave IT teams faced with supporting numerous remote sites—but with minimal remote staff and expertise.

To solve this dilemma, The Friday Institute developed an end-to-end architecture involving server virtualization, advanced cloning technology with NetApp FlexClone®, and end-to-end management to

Focusing on these four areas, we were able to develop a fast service provisioning model that now allows us to use just a few commands to rapidly roll out a server and install Oracle®, WebSphere, and any other application components. This model includes the automatic creation of storage volumes and the ability to allocate storage so rapidly that we've seen the system mount as much as 1TB of storage to a host in just a few seconds.

Project results:

- Ability to provision 1TB to a host in a few seconds
- Ability to automatically discover and apply protection policies to newly provisioned data sets
- Reduction of provisioning time of database, application, and storage on a server down to just 15 minutes

Frustration and Delays: The Impetus for Change

When we received a customer request to set up a new test environment, the process used to involve coordinating with a minimum of six or more teams. The server component might be provided by Facilities; the [storage] might be handled by our Systems Administration organization. Then, another group would install the middleware or database layer. Someone else would be responsible for the application side, or the installation of [BEA] WebLogic or [IBM] WebSphere, if those were required. For the most part, every piece of software was handled by other groups individually.

Since our deployment of environments depended so heavily on the availability of every team to do the work, it slowed the process down. By the time we had organized everyone's schedule, the time from request to deployment typically spanned from weeks to months.

We began to wonder how much faster we could deploy test environments if we minimized the reliance on other teams for execution of their component. What if, instead, we could just automate their portion of the implementation, based on their group's own preagreed policies and processes? If we could successfully separate out the "execution" side for any environment's deployment from the "policy" side associated with each group, what efficiencies could we gain?

To help test this premise, we decided to perform an informal pilot last year with the support of an executive sponsor. Using a core team of just four to five people, we were able to create a prototype process that was still largely manual. What the prototype allowed us to do, however, was prove our ability to dramatically reduce delivery time of a test or production environment to customers. In our pilot, we were able to go from 80 hours of calendar time down to just four hours and one person. Doing this allowed us to separate the actual time and labor it took to perform tasks, which were relatively small, from the group coordination and scheduling functions.

One time-saving deployment strategy used during the pilot was to replace the organization's traditional middleware and custom software installation processes with an installed image of the middleware application, created with NetApp Snapshot™ software operating on one of our NetApp FAS storage systems. Having previously been stored and catalogued in its installed form, the Snapshot copy could then be rapidly deployed to a new testing environment with a quick mount of the Snapshot copy using NFS to a designated target server.

With our initial theory proven, we embarked on a larger project leading to the development of our current 4S services provisioning model.

Moving to a Shared Services Delivery Platform

We set out to develop and test a new compute farm that would allow us to quickly provision and scale both server and storage resources to meet the needs of any host. After analyzing where we could gain the biggest win in deployment, we decided to focus first on automating the deployment of test environments consisting of Sun™ Solaris™ servers and an Oracle Database or BEA WebLogic or IBM WebSphere or Sun iPlanet. Applications with one or more of these components represented roughly 45% of the total test environments being deployed. These would be our first test case for our emerging services provisioning model.

We developed a server and storage farm, both running over a TCP/IP network and being managed by an infrastructure management software layer. The infrastructure management layer would be used to help automate and speed the delivery of resources from within the farm. A few key functions the infrastructure management

deliver software services on demand. The resulting framework is a virtual, open environment that offers flexibility unimaginable in traditional environments.

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

VSimplified Backup and Replication Management with NetApp Protection Manager

Protection Manager is an intuitive, policy-based management application for NetApp disk-based backup and replication technologies, including SnapMirror, SnapVault, and Open Systems SnapVault. This tool enables administrators to apply predefined policies to their data, thereby eliminating ambiguity and the potential for error inherent in manual management.

A new demo shows how Protection Manager solves three common management challenges:

- How do you make sure that everything is protected when data is distributed everywhere?
- How do you scale your data protection environment without spending all day on tedious manual tasks?
- How do you rapidly roll out global changes across all sites and systems?


[Watch the demo.](#)
[Read Three Backup and Replication Management Challenges Solved.](#)

layer needed to perform included:

- Provisioning
- Storage and backup management
- OS deployments using Jumpstart

These functions, and the farm's conceptual design, are shown in Figure 1. We anticipated each line in the figure, going from point A to point B, might involve significant custom development and integration work before we could turn the delivery of test environments and related services into a more automated, push-button process. When it got to the storage and backup management component, however, we were surprised at how little integration was required with NetApp Protection Manager, the application we chose to manage the process.

For the server farm, we deployed Sun Solaris servers, starting at 50 servers with an ultimate goal to move to about 100. For storage, we centralized on 100TB of NetApp FAS3000 series storage and 100TB of NetApp NearStore near-line storage as our primary and secondary storage. Systems are connected using an IP SAN that is architected to scale out to every corner of our data center. The architecture itself is built to scale to support thousands of hosts.



For our larger infrastructure management and provisioning component, we chose IBM Tivoli Provisioning Manager to help us manage, catalog, and automate provision prior server/operating system environments for reuse. To help us more quickly manage, back up, and provision the data sets associated with each test environment build, we evaluated a few backup and recovery applications before settling on the use of NetApp Protection Manager, in conjunction with other NetApp data protection tools like NetApp Snapshot, SnapMirror®, and SnapVault®.

Determining the best storage and backup management solution turned into one of the harder aspects of the model to implement.

Data Protection Tools Put to New Use: Fast Provisioning and Release of Storage Resources

When we began to architect the 4S model, we knew we wanted a very highly available and resilient backup capability so that we could avoid potential delays in the delivery of services. We also knew we wanted to do more management of the data sets than just backing up the data in the event of a local or broader system failure.

We needed a solution that would allow us to provision the storage and perform the subsequent teardown or release of that service so that the underlying storage assets could be reused. At the same point, we wanted to "checkpoint" the test environment so that we could stop it, yet resume it again at some point in the future, thereby freeing up our server capacity in the meantime for better utilization.

Host-Centric Compared to Storage-Centric View

When we compared NetApp Protection Manager against other host-based backup applications, we liked that it offered a storage-centric instead of a host-based view of backup data. If we wanted to provision storage services for a new test environment, mount the volume(s) to a host, then dismount them and remove them, we felt a host-based backup approach would cause some of the storage to be "orphaned" and without an associated host. This was a key reason we wanted to move the host out of the picture when it came to protecting and managing the data sets.

After testing Protection Manager, we liked the fact that it maintained a storage-focused oversight of the underlying physical storage, volumes, Snapshot copies, and data sets. More importantly, it had an important feature not available in the other solutions we looked at: the ability to autodiscover storage components and underlying storage volumes within our NetApp FAS and NearStore storage systems. This was huge for us, as it meant we would not have to build months' worth of custom script in order to allow the system to discover and report on or manage the current state of various data sets.

We also liked the fact that Protection Manager allowed us to group data sets or volumes with common protection requirements, then apply a predefined backup/restore policy to it. These turned into another type of provisioning policy within our model. The approach also fit with our original vision of reducing the various groups' involvement in *executing* the test environment builds, while still providing them


the oversight of *policy* surrounding what the builds should contain.

We set up our shared storage farm to be generic and simple with protection policies to back up our software clones, separate protection policies to back up user data, and separate protection policies for root volumes and storage system files. This process is outlined in Figure 2.

Examples of Rapid Storage Provisioning (and Reprovisioning)

NetApp Protection Manager worked well out of the gate. After just a few days spent implementing this aspect of the model and providing just two pages of written instructions, we were able to easily deploy new storage and test environments.

Now, to deploy a standard three-tier server environment (including three servers, each representing a database system, application, and Web environments and even terabytes of storage), we've reduced the provisioning time down to just 15 minutes and one person who now just needs to type in a few commands. That was a process that previously used to take hours or weeks to complete.




In another scenario, we had already provisioned resources for one test environment based on the customer's initial requirement for an Oracle10g™ database. After the environment had been provisioned, however, we were made aware of a new requirements for the environment to run on Oracle9i™ instead. This might have been a nightmare with our old provisioning style. With our new services model, we were able to use just enter a few commands to disconnect resources from Oracle10g and reconnect them again in Oracle9i.

Chasing Efficiency First, Cost Savings to Follow

One of the interesting things about this project was the fact that we didn't enter into it initially with cost savings uppermost in our minds. Sometimes, too much focus on cost savings can actually water down a project. Instead, we were provided the leadership to really tackle the issues correctly by focusing on the business issues first. Our number-one business issue was not about cost savings. It was about how we could speed up the deployment of applications and environments into production. We feel we accomplished that with our 4S model and are now beginning work on its wider acceptance throughout other facets of the organization. We are starting to see the benefits of this approach with subsequent cost savings in the form of better storage utilization and faster time to production.

The important thing for other groups to recognize is that this type of services delivery infrastructure doesn't remove their responsibilities and ownership of their piece of the infrastructure. Instead, it repackages their pieces in a way that allows them to be more proactive in setting and refining policy decisions and standards as we move forward. Also, by proving the merits of the project first with the support of a few core members and an executive sponsor, we were able to demonstrate a big win earlier that we hope will help the organization more quickly adapt to the benefits such an architecture now offers to other areas as well.

TECH ONTAP ARCHIVE



Jess Carruthers
Project Manager

For the past nine years Jess Carruthers has managed storage systems from NetApp and competitors supporting Oracle® ERP, financial databases, clinical databases, document, and PACS imaging environments. In 2005 as a beta site for Data ONTAP® 7G, Jess and his team consolidated 28TB of Oracle and CIFS storage across nine NetApp systems to a FAS960c and a NearStore® system running SnapVault®. Utilization jumped from 50% to 76%. The team is currently expanding its environment to include three FAS clusters, three NearStore systems, and a FAS3070 at a remote DR site.

Customer Case Study Oracle on NFS for Maximum Flexibility

By Jess Carruthers

Architecting a storage environment for Oracle can be a tricky proposition. While there is no simple “one size fits all” solution, there is a wealth of knowledge available for anyone who wants to deploy Oracle with NetApp storage, and much of that information derives from direct customer experience. This article takes a closer look at one NetApp customer that chose NFS to meet its Oracle storage needs.

I work for a regional medical center with a total staff of nearly 15,000 people spread across three hospitals, a rehabilitation center, primary care clinics, extended-care centers, a research institute, and a hospice center. Medical staff includes 3,100 physicians representing more than 91 medical and surgical specialties.

We use Oracle as the back end for our clinical care and financial applications, so Oracle performance and availability are critical. A tight IT budget coupled with a growing demand for storage—for Oracle as well as other applications—has created significant challenges.

NetApp storage in conjunction with NFS has helped us meet our growing storage needs without breaking the budget. With our NetApp infrastructure, we can react to change more quickly than we can with other storage solutions. In this article I'm going to describe how we configure Oracle to run over NFS, discuss our data protection and disaster recovery strategy, and explain the advantages we get from this approach.

An Evolving Oracle Environment

Like many busy Oracle shops, ours is far from static. We're constantly evolving, changing, and adding to meet business demands. We run E-Business Suite as well as a diversity of applications for patient tracking, billing, and archiving. We have database instances running on Oracle Databases 8i, 9i, and 10g.

Two years ago, all our Oracle servers were monolithic HP and Sun™ servers with host-based HA for availability. Today, we're transitioning to Oracle Real Application Clusters (RAC) and slowly replacing or supplementing our big iron servers with clusters of cost-effective Linux® servers.

Today, we have a total of about 24 database instances, and we're adding about 1TB of new storage a month for Oracle and other purposes. This is where NFS comes in. We're adding storage and changing our storage configurations all the time, and NFS and NetApp make the whole process much simpler than if we were using SAN storage.

Why Oracle on NFS?

The main reason we prefer NFS for Oracle is that it allows us to adapt more easily. We can quickly expand or contract an NFS file system to meet changing needs. I

RELATED INFORMATION

- [Oracle on NetApp: NFS](#)
- [Oracle on NetApp: SAN](#)
- [Why Protocol Should Be Irrelevant](#)
- [SnapManager for Oracle](#)

Choosing a Storage Protocol for Oracle

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 a range of implementation options and tools for deploying and managing Oracle environments using NFS, Fibre Channel, or iSCSI, alone or in combination. Find out more specifics from previous Tech OnTap articles:

- [Oracle on NFS](#)
- [Oracle on Fibre Channel](#)
- [Why Protocol Should Be Irrelevant](#)

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 Sanjay Gulabani and NetApp technical director Mike Eisler. In addition, NetApp founder Dave Hitz recently blogged about the origins of direct NFS.

might take a 1TB aggregate and give each of eight applications 100GB while holding 200GB in reserve. We can then expand (or contract) any of those NFS volumes in a matter of minutes without disruption to ongoing operations. We can make changes whenever we need to, and it's no big deal.

Like most IT operations, we also have storage from other vendors. With our SAN hardware, expanding a volume is a lot more work. Expansion takes at least eight hours, and there's no way to automatically shrink a volume. With SAN storage systems, you find yourself overprovisioning to avoid getting caught short and putting off operations that are so easy on NetApp.

Another benefit of NAS over SAN is less administrative overhead. TCP/IP cards just work right out of the box. We don't have to spend a lot of time updating drivers and firmware as we would with HBA adapters. Our DBAs really prefer the NFS environment because it gives them much more autonomy. You can grow Oracle table spaces and data files, create or restore Snapshot™ copies of NetApp volumes without involving a storage administrator, and get the data you need up on any host (or multiple hosts) with a simple NFS mount—a particular benefit with Oracle RAC.

Today we have a total of 12 NetApp storage systems with raw capacity of 190TB. In addition to storage for our database/application needs, we recently added a NearStore R200 to support our radiology and mammography image archive (PACS), and we've deployed a NetApp FAS3070 for off-site vaulting with SnapVault, to support our objective of eliminating tape from our environment.

Configuring NFS for Oracle

In the beginning, everyone was afraid that performance would be an issue, but if you build your environment correctly, performance should not be a problem. When we configure a NAS environment for Oracle, we basically apply the same rules you would apply to a SAN. We create a private network and use redundant switches to create a redundant fabric. Since Fibre Channel SAN runs at 2Gb per second or faster, we meet or exceed that performance level by aggregating several Gigabit Ethernet connections together to give us bandwidth of anywhere from 2Gb to 6Gb per second depending on the application. Essentially, we're creating a dedicated SAN; it just runs a different protocol.

To get the best possible performance, some tuning is required for the TCP/IP stack and for NFS. Fortunately, NetApp has some great resources that tell you exactly what to do. (Find out more from a [recent Tech OnTap article](#).)

Data Protection and Disaster Recovery

No matter if you choose NAS or SAN for Oracle, data protection and DR are going to be critical considerations. For our Oracle environment, we have an established recovery point objective (RPO) of five minutes and a recovery time objective (RTO) of four hours for any given database. We achieve these goals using a combination of NetApp SnapVault and NetApp SnapMirror®, providing both data protection and DR through a completely tapeless solution.


Every night at 1:00 a.m. we put all our Oracle Databases into hot backup mode and create a Snapshot copy of each one using a customized script. This takes about 10 minutes. Then starting at 1:15 a.m., SnapVault runs and vaults all those changes to our DR facility. We throttle the performance to 20MB per second so that we don't impact production applications. This gives us the equivalent of a full backup of every database every night with off-site storage. We maintain 20 days' worth of these nightly backups online. For database applications, we don't need to go back years. We just need to protect ourselves against possible application errors that would require us to roll back.

All our archive logs are stored in separate volumes, and we use NetApp SnapMirror to sync those volumes to the DR site every five minutes. The combination of the two allows us to meet our five-minute RPO and four-hour RTO.

To recover, we copy the SnapVault backup to a read/write volume and then replay the archive logs that have been stored using SnapMirror. We use SnapVault instead of SnapMirror to protect Oracle data volumes and department share data to reduce storage costs. Our primary volumes are on Fibre Channel disk, and our SnapVault

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](#)



volumes are on SATA disk. We keep five days of Snapshot copies on primary storage and 20 days of Snapshot copies at the DR site. We are also in the process of deploying NetApp deduplication on our vaulted data to further save on space.

Keeping Ahead of the Curve

My organization is very proactive when it comes to efficiency. We have business operations analysts on site to look at workflow and how things should interact. They do process redesign before we deploy an application to make sure we are making the most efficient use of new software and hardware.

Like everyone in IT, we face a constant struggle to stay ahead of demands and keep up with growth. We have three DBAs and three admins dedicated to UNIX® and Linux servers and storage. Because our NetApp storage systems configured for NAS are so easy to manage, it takes less than one full-time equivalent administrator to manage it all, freeing up resources for other work. Over the past few years, our total NetApp storage has grown by 80TB, but we haven't had to add additional staff to manage it. NetApp makes it possible for us to get more done in less time with less staff.

TECH ONTAP ARCHIVE



Shaun Mahoney
Consulting Systems Engineer, NetApp

Based in New York City, Shaun Mahoney serves as a business solutions architect focusing on archive and compliance solutions. Before joining NetApp, Shaun spent more than 10 years working in the financial services community. During the time he was a NetApp customer, Shaun helped architect and develop the NetApp LockVault™ product.

Implementing Exchange Archival: A Case Study

By Shaun Mahoney

When someone refers to e-mail archival they might be talking about:

- Archival for mailbox management or
- Journaling for regulatory compliance

Although the same application can be used to provide both, the two functions generally have completely separate environments with different hardware and software configurations and unrelated end goals. The fact that these two functions are so different often leads to confusion in discussions of e-mail archival, making it difficult for companies to figure out exactly what they need.

I recently worked with a large financial services company that is doing a major refresh of its messaging environment, including:

- Migrating most users to Exchange Server 2007
- Adding archival capability to get the volume of e-mail under control and eliminate the need for PST files
- Adding journaling to meet regulatory compliance

In this article, I'm going to look at this company as a case study to illustrate the differences between archival and journaling. Since journaling was the part I was most closely involved with, I'll focus particular attention on that, including the compliance planning process.

Customer Challenges and Goals Migration

The company needed to reduce the cost and complexity of its messaging solutions, improve availability, and increase efficiency. Like many companies in the finance area, this one had grown through acquisition and was using multiple messaging systems as a result. The Exchange migration project will address this problem and get the entire company running on a single up-to-date solution.

Archiving

At the same time, the company wanted to make its messaging environment more resilient and efficient, reduce backup costs, and avoid the need for PST files. Using e-mail archiving to move messages out of the active Exchange repository addresses these issues. Because the primary Exchange repository decreases in size, backup is simplified, and PST files are no longer needed to retain older e-mail.

OnDemand TechTalk Webcast:

E-Mail Archival
Solutions to Manage
Increased Data Growth

E-Mail Archive and
Compliance: Eliminate
the Need for Separate
Storage Silos Without
Sacrificing Performance

RELATED INFORMATION

- [Tips for Archive and Compliance Planning](#)
- [Impact of Exchange 2007](#)
- [NetApp: Viable for Exchange?](#)

Tips for Archive and Compliance Planning

What would you do if someone slid a 5¼" disk across the table and asked you to open up the WordStar document on it? This may seem like an extreme case, but when you're planning on retaining information for years there's a lot more to think about than just squirreling away the data.

In a recent Tech OnTap article, NetApp systems engineer Mike Riley provided six tips for improving your archive and compliance planning:

1. Avoid storage silos
2. Avoid proprietary data formats and technologies
3. Don't consciously limit your options
4. Plan for performance
5. Utilize storage virtualization
6. Think carefully about encryption

[Read the article](#) to find out more.

Support Your Compliance Technology with the Right Processes

The latest and greatest technology (from any vendor) is not enough to ensure you are in compliance. There are important processes that must be

The goals for the archival project were:

- Reduce costs
- Save disk space
- Limit impact to end users (ideally zero impact)

Compliance

For the financial company, compliance with SEC and other regulations is a major concern. The company had previously implemented a basic compliance strategy where it was journaling messages, but ultimately realized that it wasn't journaling in a manner that satisfied regulatory requirements. Fixing this problem was the driving force behind this part of the project. NetApp was brought in specifically to help address compliance needs, which is how I got involved.

The primary goals for the compliance project were:

- Reduce risk
- Ensure reliability
- Keep the solution as simple as possible

Establishing Compliance Objectives

When it comes to compliance, the most difficult thing is often determining what to save and for how long. If you choose too short a duration, you risk being out of compliance. If you choose too long a duration, in contrast, you may tie up resources, locking down data that you no longer need to retain. (In a properly implemented compliance system, once data is locked down, it can't be deleted under any circumstances until the retention period expires.)

When NetApp first sat down with the compliance team for this company, the team had identified the users for whom journaling was required. There were several business units that had different requirements for retention based on their business function and various regulatory requirements. However, they weren't sure how many different retention time frames they would need. They were originally considering five separate retention programs: three, five, six, seven, and eight years. This would have been very complex to manage, and they didn't have a good methodology to identify which users belonged in which program. It would also have made the project more costly.

Our first step was to determine what regulations they had to meet, what the time frames actually were, and to sit down with them to simplify the retention plan. Based on those discussions, we were able to get it down to two categories: three years and eight years—much more manageable.

Next I began digging into the specifics of the eight-year requirement. This requirement was based on a regional regulation that ranged from four to eight years depending on the region. Only the three-year requirement mandated that things be set for the specified time period. The other regulation mandated that the data be kept for the specified time period. This allowed us to set everything for only three years (further simplifying the journaling environment) and then change the record policy to review the records before the three years were up and extend the retention of information as necessary to satisfy the other regulations. The eight-year regulation is currently under review to be shortened to four years across all regions, so this step may ultimately save the company from unnecessarily locking down data for a period of four years.

Project Implementation

Migration and Archival

Both the implementation and ongoing management for the Exchange and the archival projects were outsourced to a third party. As part of the request for proposal (RFP) for each project, the company specified that NetApp be used for all back-end storage. The company already had a small amount of NetApp storage in place and clearly saw the benefit of using NetApp hardware and associated software tools, including SnapManager® for Exchange.

Storage for each project is provided using clustered NetApp FAS3070 storage systems. The company has one FAS3070 in each data center, and a stretched Microsoft® cluster spans the two data centers. The Exchange environment uses high-performance Fibre Channel disk and the iSCSI protocol for block access.

The archive environment uses its own clustered NetApp storage with a mix of Fibre Channel and SATA drives. Symantec® Enterprise Vault™ software is used to implement archiving. Enterprise Vault uses Microsoft SQL Server™ as its underlying database engine. All SQL Server data and indexes are on Fibre Channel storage

part of any compliance solution. Overlook them, and the auditors may never bother to take a look at your cool technology.

Here are some key processes you should know about and implement:

- All physical access and maintenance of the media or system must be audited and recorded and must have preapproval by a control source.
- You should have a security log recording any configuration changes made to your compliance solution.
- You need to be able to provide an auditable topology and workflow report of how data gets to your compliant storage.
- There must be an audit trail for outage remediation. In other words, should you experience an outage, what steps would be taken in order to bring the system back into compliance, and how did you know when you had reached that point?
- There must be a method to verify and audit the quality and accuracy of the data over time.


accessed using iSCSI, while the archive data itself is stored on CIFS shares on SATA disk for greater economy.

Journaling

Because it was for compliance, the journaling project was not outsourced. The company wanted to maintain direct ownership and control to ensure that it fully understood and minimized all potential risks. A preexisting dedicated team owns the project with responsibility for the software, servers, and storage.

The physical infrastructure for the journaling environment is very similar in principle to those for the Exchange and archival environments. A clustered NetApp FAS3070 storage system was deployed in each of the company's two data centers. Under normal operation, each data center provides journaling for roughly half the company's messaging. NetApp SnapMirror® software is used to mirror all journaled data between the two data centers so that two copies are maintained at all times for disaster recovery.

All Microsoft SQL Server data is stored on Fibre Channel disk and accessed using iSCSI. NetApp SnapManager for SQL Server is installed to simplify SQL Server data management. Index data is stored on CIFS shares on Fibre Channel storage. Journal data resides on CIFS shares on SATA storage. To ensure compliance, the volumes used for journal data are NetApp SnapLock® Compliance volumes. NetApp SnapLock Compliance provides WORM capabilities with retention support for data stored by Symantec Enterprise Vault.



Clustered Microsoft SQL and Symantec Enterprise Vault servers have been configured to span the two data centers. Each location has an active SQL and Enterprise Vault server with passive counterparts located in the other location for immediate failover if necessary.

Journaling is supported by a feature of Microsoft Exchange. Exchange maintains a journal mailbox. Every message that goes in or out of the environment gets copied to that mailbox. Enterprise Vault makes a copy of each message in the journal mailbox on its compliance storage.

Enterprise Vault also creates index entries based on the message subject, who it's from and to, basic message content, and so on and stores this information in a searchable index that supports e-discovery. The SQL database contains pointers between the data stored on disk and the metadata in the index. A master directory database contains pointers from all users to all messages and relevant index data. As you would expect, Enterprise Vault makes it possible to set a retention period for all messages in the journaling environment and can automatically delete those messages once the retention period has expired.

Solution Impact

At the time of this writing, both the Exchange and archive projects are installed, and users are moving to the new solutions. Compliance will go online after user migration is completed. It's not possible to determine the full impact of each project yet, but there are some clear benefits already. For instance, the advantage of having a single messaging environment across the company should be obvious.

We can also evaluate the archive and journaling projects against the stated objectives above. The goals of the archive project were reducing costs, saving disk space, and minimizing user impact. Because the archive is essentially transparent to end users, the user impact will be negligible. Costs are reduced by cutting the amount of expensive primary storage needed for Exchange and simplifying ongoing Exchange data management, particularly backup.

Costs are further reduced through the use of NetApp unified storage systems. The archive configuration uses both iSCSI and CIFS. Because NetApp storage supports multiple protocols, it is possible to provide all back-end storage with a single type of storage system. Solutions requiring separate storage for each protocol would be significantly more expensive. A single-vendor solution is naturally less complex than a multivendor one as well, so complexity is reduced, making management simpler.

The journaling project obtains similar benefits from the use of NetApp storage. The company also is able to take advantage of NetApp technologies such as space-

efficient Snapshot™ copies and FlexClone® volumes to further reduce space consumption and streamline data management and disaster recovery testing for SQL databases and index metadata. Standardizing on one storage platform gives the company less to buy (reducing costs), requires less space in the data center, and ultimately gives the company a greener solution that requires less power and cooling while performing significantly faster than competitors.

If we evaluate the journaling environment against the overall objectives of reducing risk, ensuring reliability, and keeping things simple, the plan discussed earlier meets the primary objective of reducing risk. The plan is also as simple as we can make it, and that simplicity should contribute to the reliability of the final journaling implementation. The hardware and software implementation for the journaling project is not only simple, but also as robust as we can make it, offering full failover for disaster recovery to ensure reliability.

Conclusion

The company has been very happy with the results achieved from these projects so far. It recognized from the outset that NetApp didn't just offer great products, but also had the expertise to partner with the company and help make the compliance project a success by creating a simple, robust compliance plan that meets regulatory requirements without locking down more data than necessary for longer than necessary.



Stephen Daniel

Director of Database Platforms and Performance Technology, NetApp

Steve Daniel has been at NetApp for seven years, focusing exclusively on database I/O and the impact to storage systems. Today, he focuses on database performance and reliability on both NAS and SAN. Before joining NetApp, Steve honed his database skills during 12 years at Data General.

Real-World Storage Performance Benchmarking NetApp versus EMC CLARiiON

By Stephen Daniels

Determining the performance that SAN storage will deliver in the real world—running real applications—can be a challenge. Many vendors have only discussed informal or ad hoc benchmark results, making comparisons between storage systems difficult. What's more, benchmarks frequently use impractical system configurations (the smallest available disks, huge numbers of spindles), and they don't assess the performance impact of features that may be essential to your operations.

NetApp recently set out to run a benchmark comparison that would address some of these limitations with:

- A standard workload representative of real-world applications
- Configurations based on vendor best practices
- The use of Snapshot™ copies during the benchmark

For this testing we chose a standard benchmark, SPC-1, to assess the performance of a NetApp FAS3040 system versus that of an EMC CLARiiON CX3-40. In this article I'm going to describe the benchmark, the system configurations, and the results achieved. For more detailed information you can also review the published, fully audited results on [the Storage Performance Council Web site](#).

Highlights of Results NetApp Versus EMC CLARiiON

- Better performance with Snapshot enabled (97% versus 36% of baseline)
- Higher performance with fewer disks

NetApp SPC-1 Firsts

- Highest raw storage utilization
- Lowest cost per tested GB
- First ever result with RAID 6
- Only active result not using RAID 1/0

The SPC-1 Benchmark

The Storage Performance Council (SPC) is a vendor-neutral standards body with representatives from a wide range of storage vendors. The SPC has so far released two benchmarks:

- **SPC-1** generates a workload with characteristics of typical business applications such as database applications and serving e-mail with random I/O, queries, and updates and therefore seemed highly appropriate for our purposes.
- **SPC-2** is designed to simulate applications with large-scale, sequential data movement, so we didn't consider it a good representative workload for this study.

I consider SPC-1 to be the best benchmark available to model the way that databases stress storage systems. When SPC developed the benchmark, it studied how a variety of applications accessed storage and then modeled the workload based on those

RELATED INFORMATION

• [Storage Performance Council](#)

Complete SPC-1 Results

This article is just a summary of the testing we performed. For more information go to the [SPC benchmark results page](#) or click the following links to go directly to the full disclosure reports (pdf):

- [NetApp FAS3040](#)
- [FAS 3040 with Snapshot](#)
- [EMC CLARiiON CX3 Model 40](#)
- [CX3 Model 40 with SnapView](#)

measurements. The mix of operations is representative of a broad class of applications—roughly half of all commercial applications.

As you are probably aware, the most commonly used set of database benchmarks are the TPC benchmarks, but those are designed to test the database server, and therefore the measured results depend more on how the server is configured than they do on the storage.

SPC-1 has the advantage that it just measures the performance of the storage system. It doesn't depend on how you configure the server and database. The workload it sends to storage, however, is a great proxy for the kind of load that today's business applications generate. In addition, SPC has strict rules for how the benchmark must be run, so you can compare results between vendors; all results are audited by a certified, independent auditor. All published results include lengthy disclosures including full pricing of the benchmarked system.

System Configurations

As I said in the introduction, we wanted to test system configurations similar to those that customers might be likely to deploy. In each case, we wanted to build a Taurus rather than a NASCAR, so we chose two midrange systems, the NetApp FAS3040 and the EMC CLARiiON CX3 Model 40, and configured them appropriately. We used the exact same server, identically configured, to run the benchmark against each storage system in all cases.

At this point you might be thinking, "If you're benchmarking a system from another vendor, how can I trust the results?" SPC actually takes this into account in its rules. In order for us to publish a benchmark of another vendor's gear, the SPC required us to certify that we made a good faith effort to demonstrate the true maximum performance of the equipment. To do that, we started with EMC's documents on how to tune the system for best performance, and we followed them closely. Then we spent a couple of months tuning and adjusting, trying to improve the performance. Everything that goes into the final result is documented in the full disclosure reports for each test.

Specifically, EMC docs were very clear that the best performance could only be obtained from the CLARiiON system using RAID 1/0 (mirroring and striping), so we configured the system that way rather than with RAID 5 or RAID 6 (double-parity RAID). The FAS3040 system was configured to use RAID-DP™, NetApp's double-parity RAID 6 implementation, which is the default on NetApp systems. On both systems we used 82% of the usable storage for the benchmark. This is a tunable option of the benchmark. The CLARiiON system was configured with 155 disks, while the NetApp system had 140. Both systems used 146GB 15,000 RPM 4Gb per second disks.

Other than that, we also made an effort to tune the memory management on each system for best results. EMC allows you to tune the read versus write cache. We tried all the variables until we got the best results. We found, for example, that by turning off the CLARiiON write cache for portions of the workload that did very little writing we left more write cache for the write-intensive portion of the benchmark. This improved performance significantly.

On the NetApp system, we found we could improve performance by changing the memory management policy to reflect the fact that most SPC-1 data is not referenced repeatedly. This policy change can be implemented with the following **priority** settings with Data ONTAP® 7.3:

```
priority on
priority set enabled_components=cache
priority set volume <volume-name> cache=reuse
```

The net effect of these commands is to tell the memory system to reuse memory for newer items more aggressively than it would normally. (The `enabled_components` subcommand is new in Data ONTAP 7.3. If you are using Data ONTAP 7.2 you can skip that command.)

A couple of the things we tuned are still being refined, so they are enabled by the **setflag** command. In future versions of Data ONTAP either these flags will become options or they will disappear as the system becomes self-tuning for these features.

```
priv set diag
setflag wafl_downgrade_target 0
setflag wafl_optimize_write_once 0
```

The "downgrade_target" command changes the priority of a process within Data ONTAP that handles incoming SCSI requests. This process is used by both FC SAN and iSCSI. If your system is not also running NAS workloads, then this priority shift improves response time.

We're explicitly calling out these settings because, based on our testing, we think they will yield performance benefits for online business application workloads. If you are interested, you can read more about them in a [recent NetApp technical report](#).

Testing the Performance Impact of Managing Snapshot Copies

Most benchmarks ignore the performance impact of commonly used features such as snapshots. In the early days of NetApp, few other vendors offered this capability, but today, almost all storage vendors offer some form of snapshot technology. NetApp customers use NetApp Snapshot technology throughout the day to create backups or checkpoints from which applications can be restarted, and we assume that the SnapView™ capability provided by EMC is used similarly.

When a feature becomes a regular part of daily operation across a wide base of users, it makes sense to test the performance impact of that feature. So, in addition to testing the maximum performance of each system, we also ran the same tests again while periodic snapshots were being created. On the NetApp system, we set up a schedule that created a Snapshot copy every 15 minutes and retained the most recent three Snapshot copies. This schedule ensured that during a five-hour-long benchmark run we both created and deleted a reasonable number of Snapshot copies.

On the CLARiiON, we reduced the snapshot load. During the three-hour-long sustainability portion of the SPC-1 benchmark, we took one snapshot an hour into the test. An hour later we deleted it and took another. Just before the test finished we deleted the second snapshot.

Benchmark Results

The results of our testing are summarized in Figure 1. The baseline results show the maximum performance without snapshots. As you can see, the results are similar. The NetApp system delivers about 31,000 SPC-1 I/O operations per second (IOPS), while the CLARiiON delivers a maximum of about 25,000 SPC-1 IOPS. The NetApp system used for these tests has a list price of \$421,730.49, while the EMC system has a list price of \$517,851.02. This corresponds to \$13.61 per IOP for NetApp versus \$20.72 per IOP for the EMC configuration.




Figure 1) SPC-1 performance of NetApp versus EMC with and without snapshots.

An interesting result occurs when the snapshot feature is enabled on each platform. Performance on the NetApp system only drops to about 30,000 SPC-1 IOPS (97% of maximum). On the EMC system, performance drops to approximately 9,000 SPC-1 IOPS, or 36% of the performance level without snapshot enabled. The NetApp system used for the Snapshot tests has a list price of \$446,210.49, while the EMC system has

a list price of \$535,251.02. The NetApp system therefore costs \$14.89 per IOP versus \$59.49 per IOP for the EMC system.

	NetApp FAS3040	EMC CLARiiON CX3-40
Baseline IOPS	30,985.90	24,997.48
Baseline List Price	\$421,730.49	\$517,851.02
Cost per IOP (without snapshots)	\$13.61	\$20.72
Snapshot IOPS	29,958.60	8,997.17
Snapshot (% of Baseline)	97%	36%
Snapshot List Price	\$446,210.49	\$535,251.02
Cost per IOP (with snapshots)	\$14.89	\$54.49

Table 1) Comparison of NetApp versus EMC with and without snapshots.

Conclusion

While we were careful to ensure that we tested both systems under the same conditions-configured to achieve optimal performance in each case-these results demonstrate some significant advantages for NetApp technology. The NetApp Snapshot implementation is clearly more efficient than the EMC implementation and has a far smaller impact on performance. We believe this translates into direct benefits for busy production environments.

The result for the NetApp system is the only SPC-1 result published in the last five years not using RAID 1/0 and the first ever result with double-parity RAID 6. Despite that, compared to the EMC CLARiiON configuration we tested, the NetApp FAS3040 demonstrated higher performance with fewer disks (140 versus 155).

Because of the use of RAID-DP instead of mirroring, the NetApp configuration also demonstrated the highest storage utilization of any recent SPC-1 result. With the FAS 3040, 61% of raw storage was utilized versus 38% for the mirrored CLARiiON CX3 model 40 configuration. (Note that both systems had volumes filled to 82%.) This high rate of utilization in turn leads to the lowest cost per benchmarked gigabyte (the gigabytes that are actually available for use) ever reported for an SPC-1 benchmark: \$33.50 per benchmarked gigabyte.