

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 ▣](#)

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'es 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 Re-provisioning)

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.

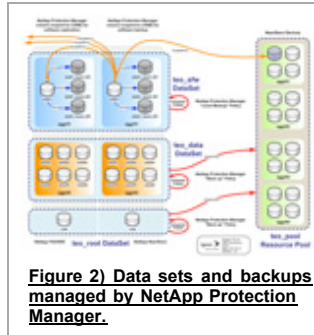


Figure 2) Data sets and backups managed by NetApp Protection Manager.

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.

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.

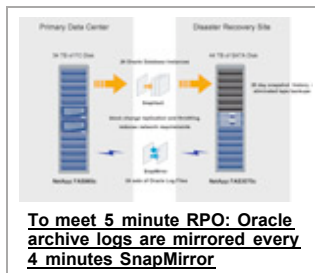
Every night at 1:00 a.m. we put all our Oracle Databases into hot backup mode and create a Snapshot copy of each one using a customized script. This takes about 10 minutes. Then starting at 1:15 a.m., SnapVault runs and vaults all those changes to our DR facility. We throttle the performance to 20MB 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.