

HIGHLIGHTS

- » **Pitfalls That Can Affect Recoverability**
- » **Engineering View: Oracle® on FCP**
- » **IEEE Recognizes NetApp Innovation**
- » **Centralized Backup for Remote Sites**
- » **Calculate Potential Capacity Savings**

Real-World Scale-Out Storage: Experiences with Data ONTAP® GX

Reinoud Reynders, IT Manager, Leuven University Hospital

Concerned about scalability? Faced with annual data growth exceeding 100%, one IT team investigated, tested, and deployed Data ONTAP GX. See what they learned.



[More »](#)

BLOGGING WITH DAVE

Dave Hitz, NetApp Founder and EVP



"What exactly does Oracle gain from Direct NFS? You don't have to worry about how to configure NFS, and Oracle bypasses the operating system."

» [Dave's Blog](#)

TIPS FROM THE TRENCHES

Backups Are Worthless, Recoveries Are Priceless

Jot Gill, Strategic Business Consultant,
NetApp Value Engineering

Are you exposed? Check out three commonly over-looked backup problems that can cause recoverability issues in heterogeneous storage environments.

[More »](#)

DRILL DOWN

- » **Webcast: Moving Beyond Tape Backup**
Hear how Swinton Insurance centralized backup for 400+ remote sites.
- » **Overview: SnapVault® for NetBackup™**
A guided tour of the NetApp approach to in-line deduplication and data replication in NetBackup environments.

ADMIN RESOURCES

- » **Exchange Server 2007 Best Practices**
Considering Microsoft® Exchange 2007? A 27-page guide to capacity, sizing, I/O planning, server consolidation, and more.
- » **Estimate Deduplication Savings**
Calculate your potential capacity savings with a new online deduplication tool.

ENGINEERING TALK

In the News: NetApp Cofounders Recognized for Storage Innovation

Data Storage Connection Article

Dave Hitz and James Lau helped alter the landscape of the storage industry through the creation of WAFL® (Write Anywhere File Layout) technology.

[NetApp Cofounders Dave Hitz and James Lau to Receive IEEE Award »](#)

Oracle on NetApp Part 2: Simplified FC SAN Management

Alvin Richards, Technical Director, NetApp Engineering



Learn how NetApp engineers have streamlined Oracle on SAN environments by developing tools that integrate NetApp storage with Oracle RMAN and ASM, any-to-any DR capabilities, and more.

[More »](#)

[FEEDBACK](#)

TECH ONTAP ARCHIVE



Reinoud Reynders

IT manager, Infrastructure and Operations, Leuven University Hospital

After joining [Leuven University Hospital](#) over 10 years ago, Reinoud Reynders worked his way up through the ranks to become an IT manager responsible for strategic planning and design as well as 43 staff members. In his spare time, Reinoud pursues outdoor activities that include hiking 'le Massif Central in France and Mount Kilimanjaro in Tanzania, sailing, and gliding.

Real-World Scale-Out Storage: One IT Team's Data ONTAP® GX Experience

By Reinoud Reynders

Leuven University Hospital (UZ Leuven) first began consolidating existing file servers onto NetApp storage in early 2000. In seven years our environment has expanded from 0.5TB of storage to over 500TB. Today we rely on more than 20 NetApp storage systems to support everything including: mission-critical medical applications, Sybase applications, Exchange, VMware, plus Windows® home directories.

One of our biggest concerns is scalability. If our current environment continues to grow at the same pace, we will need hundreds of storage nodes in the near future, and system management will invariably consume enormous amounts of time.

Naturally, we were intrigued when NetApp Founder Dave Hitz described scale-out architectures as the next leap in storage technology and talked about how the new NetApp operating system, Data ONTAP GX, combines "multiple systems into a common pool of storage that can be managed and accessed as if it were a single system."

Although it was initially targeted at large, compute-intensive environments, our team believed that many Data ONTAP GX system features—modular scalability, a single global namespace, and transparent file migration—could provide significant operational benefits in our data centers.

In late 2006 we began seriously considering where we might deploy a scale-out architecture. If the Data ONTAP GX system performed as promised, not only could we significantly improve operational efficiencies in a few carefully chosen areas now but we would also jump ahead of the curve as scale-out technology matures.

Our team spent nearly a year investigating, evaluating, and testing the Data ONTAP GX system and deployed our first production applications on the platform in September 2007. Scroll down for our full deployment journal; highlights include:

- **Target environment:** two applications that consume over 100TB of storage were selected to test the Data ONTAP GX system in a production environment: (1) Picture Archiving and Communications System (PACS) images and (2) digitized patient record archival.
- **Deployment:** Data ONTAP GX system consisting of four FAS3070 nodes.




Figure 1) Market adoption cycles

RELATED INFORMATION

- [Introduction to Data ONTAP GX \(PDF\)](#)
- [Best Practices for Scale-Out Storage \(PDF\)](#)
- [UZ Leuven: Using Cutting-Edge Technology Today](#)

UZ Leuven: Profile of a Technology Innovator

Leuven University Hospital is one of the largest healthcare providers in Europe and among the most innovative IT teams in the world.

In 2004, UZ Leuven was recognized with a NetApp Innovation Award in the "Visionary" category and today runs its entire operation on NetApp storage.

UZ Leuven continues to push the envelope through the early adoption of cutting-edge technologies, including beta testing the early versions of the following technologies:

- Data ONTAP 7G
- iSCSI
- NetApp FlexShare™
- VMware ESX 3.0
- A-SIS deduplication
- Protection Manager™
- Data ONTAP GX

By aggressively testing new technology during beta, UZ Leuven is able to fast track its deployment of the latest technology to meet specific objectives. [Learn more about this IT innovator.](#)

Testing began June 9, 2007. The patient record application entered production on Data ONTAP GX in early September, and PACS application data was completely migrated to the Data ONTAP GX system a few weeks later. Early results:

- **Stability.** Though minor issues did arise during testing, most were due to the fact that Data ONTAP GX 10.0.2 was still in beta when the system was installed. NetApp was very responsive in addressing this, and we got the released code as soon as it was ready.
- **Successful fault-injection testing.** We tested all of the features, established performance baselines, and simulated every disaster scenario we could think of, including turning off nodes, shutting down power, and disconnecting network cables and system interconnects. Everything behaved as expected.
- **Positive initial impressions.** Our team loves the interface and the ability to manage our four nodes (or more if we had them) as a single large system.

Based on our initial success, we're very excited about the Data ONTAP GX system because of the unique capabilities it gives us, especially the ability to manage huge amounts of data in a single global namespace. We'll be deploying Data ONTAP GX to help us maintain control of our storage in the face of continued exponential growth. We believe Data ONTAP GX will solve a lot of the problems we're facing today; it's all about data management, storage management, and resource management.

For other companies facing similar challenges (and who isn't), I'd advise you to start learning about Data ONTAP GX now. Test it if possible, or check out other companies using Data ONTAP GX. If you have applications where it makes sense—applications that consume significant volumes of storage and are good candidates for new technology—I'd strongly recommend deploying it.

The UZ Leuven Data ONTAP GX Deployment Journal

When we first expressed interest in Data ONTAP GX, Tech OnTap asked me to keep a journal of key learnings, observations, and milestones. For the past year I've saved team notes and internal updates; the following section fleshes out many of the details summarized above. Keep reading to learn more.



Figure 2) Deployment timeline

September/October 2006: NetApp Briefing on Data ONTAP GX

When we started thinking about Data ONTAP GX, all of the existing deployments were in large, compute-intensive environments. It wasn't clear that Data ONTAP GX was ready to meet the quite different needs of a major hospital, and we weren't sure if NetApp was interested in supporting such a deployment. Our first step was a customer briefing at NetApp headquarters in Sunnyvale. Everyone we spoke with (the Data ONTAP GX product manager; technical experts; NetApp VPs; our local storage team; and our storage partner, UpTime) thought the technology would be a great fit for our environment, and NetApp was ready to back the effort, so we decided to move forward and begin scoping the project.

October 2006 to January 2007: Choosing Applications and Defining Project Scope

We have two data centers at opposite ends of the same large building, so a Data ONTAP GX system with two nodes in each data center seemed like the logical choice to give us optimum redundancy. To meet our archiving needs, we decided that a mix of both Fibre Channel and SATA disk would allow us to migrate data to less expensive SATA storage as necessary without application interruptions or changes.

Understanding Data ONTAP GX Technology

For many IT shops, rapid storage growth has led to a proliferation of discrete storage systems. Almost inevitably, this leads to situations where some storage is chronically overutilized or overloaded while other systems are underutilized.

Data ONTAP GX addresses these difficulties by joining storage systems together into a scalable cluster with a global namespace.

A [NetApp technical report](#) describes the unique features and benefits of Data ONTAP GX, including:

- Transparent migration
- Load balancing
- Striped volume sets
- Scalable mirroring
- Tiered storage
- And more

Implementing Scale-Out Storage

For data-intensive applications, it's critical to complement your compute cluster with highly scalable storage. But how do you:

- Determine the appropriate storage technology?
- Identify the best data distribution strategy?
- Evaluate the most critical elements to achieve specific objectives?

This [best practices guide](#) based on the experience of NetApp experts explains the steps you can take to ensure your scale-out architecture produces optimal results.

We quickly settled on two applications for our Data ONTAP GX pilot project:

- PACS
- Digitizing paper records

Both applications consume large amounts of storage and would allow us to test aspects of the Data ONTAP GX technology in a production environment.




Figure 3) UZ Leuven Data ONTAP GX environment

PACS

The PACS application is used for the storage and electronic distribution of all the hospital's radiology and cardiology images. We already have 60TB of stored data and produce more than 15TB of new images a year. We've been storing all our radiology images on NetApp NearStore® since 2002, which has been a great platform for this type of application. The performance has been great (our radiologists can retrieve all exams in less than two seconds), and the price is right. However, there were two potential limitations to our existing configuration:

- **High availability/DR.** We didn't have a second copy of the noncompressed images online. This was considered too expensive given the amount of storage needed, and most images are only really accessed frequently in the first month. It would be possible to mirror only the images for the most recent month using SnapMirror® and Data ONTAP 7G, but that would require a lot of work every month, and we would also have to stop the PACS environment to make the changes.
- **Increasing storage and processing demands.** Image resolution is always increasing, so the images are becoming bigger and bigger, and the number of images in a single exam is also growing. The need to do more postprocessing of images (for example, 3D reconstructions) is coming in the future. The number of client workstations is increasing, and we also do IT for other hospitals. So performance is likely to become a bigger issue in the future. NearStore doesn't offer the capability to easily scale performance.

Data ONTAP GX gives us the ability to address both of these issues. For the first two months, all images will be stored on Fibre Channel disk with a copy in each data center. After two months we'll migrate the data to SATA and keep just a single copy online. All this can be accomplished transparently without disrupting the PACS application because of the global namespace provided by Data ONTAP GX. As our performance needs change, Data ONTAP GX gives us the ability to adapt using features such as horizontal scaling of performance and capacity and volume striping across storage nodes.

Digitizing Paper Records

Today, all hospital record keeping is completely electronic, but we still have 1.3 million paper files. These files take up a lot of space, and managing them creates a lot of headaches. We are planning to digitize all these legacy files and make them accessible online, which we estimate will consume approximately 25TB of storage. Because we will discard the original paper, we will need to mirror this data for a grand total of almost 50TB of storage.

While this might seem like a very simple application to implement on Data ONTAP GX, a major reason we chose it was because 50TB is a significant amount of data. We already feel that Data ONTAP GX is the direction we will be going in the future, so instead of migrating the data from Data ONTAP 7G to Data ONTAP GX in a few years, we thought it would be better to simply start with the data on Data ONTAP GX. Since this is a very straightforward application and not mission critical for the hospital, it also gives us a good opportunity to implement and learn about Data ONTAP GX. We can put this application into production without worrying that pilot problems will affect patient care.

February and March 2007: Hammering Out the Details

With the project scope determined, it was time to hammer out some of the details, including finding the budget and convincing the CIO and the hospital board. After those tasks were completed, negotiations with our NetApp account manager could begin in earnest.

It was also time to determine exactly what storage systems would make up our Data ONTAP GX system. We already decided that to do Data ONTAP GX right we needed four nodes: two in each data center. The FAS6000 was ruled out as too expensive for our four nodes. The FAS3070 seemed like a perfect fit. It offered a full 64-bit environment and seemed to have the right price/performance for our needs.

The only problem was the FAS3070 was not supported yet with Data ONTAP GX. That didn't seem like a huge stumbling block, since we knew it was going to take awhile to get our data centers set up with the necessary power and other infrastructure for the new systems. However, the only problems we had during the whole deployment (although they were minor) resulted from this choice. We ultimately settled on four nodes with a total of 12TB of FC disk and 168TB of SATA disk.

April 2007: Infrastructure Hold Ups

Our four FAS3070 nodes arrived, and we had an on-site visit from a NetApp SE who specializes in Data ONTAP GX to go over the setup and project schedule in detail. Unfortunately, our infrastructure was not fully ready to support the new systems. We knew we faced this delay up front. In addition to the Data ONTAP GX cluster, we also added a FAS3070 MetroCluster system, NetApp Protection Manager, and some additional disk shelves for other systems. In April and May, we implemented MetroCluster and installed the new shelves. At the end of May, our resources became available, and the new Data ONTAP GX release was also available. Perfect timing.

June 2007: Installation

Because of the FAS3070 nodes, we needed to install a beta version of the Data ONTAP GX software. Support for FAS3070 controllers was added in Data ONTAP GX 10.0.2, which would not be finalized for another two months. That was not a huge problem because we needed time to learn the system and do our functional testing before going to production anyway. Here are some installation details:

- Every node of a Data ONTAP GX system is connected to every other node through a dedicated Gigabit Ethernet network. All system communication is done through this channel using a custom NetApp protocol (SpinNP). This has to be a private network (it's no problem to use a separate VLAN on your existing switching infrastructure), and it must have jumbo frames enabled.
- We provided all the necessary ports on the switch, and we configured them correctly, but the interconnect network didn't work.
- It was clear that we had a problem with jumbo frames, which we had never used before. Our network guys don't like it, and now I understand why: it's complicated to implement it end to end.
- Because the Data ONTAP GX system is split between two data centers, we used our 10 Gigabit interconnect. It was not enough to turn jumbo frames on for every port; we needed to reconfigure the interconnect too. This is tricky, because all communication between the two data centers uses that link.
- Ultimately, we ended up using separate switches and a separate connection between the two data centers. (This problem was subsequently corrected, and we now use the normal interconnect between the two data centers.)
- The next problem was the additional four-port NIC cards in each FAS3070. Those NICs (with TCP/IP offload engines, or TOEs) were supported by Data ONTAP 7G but not with the beta release of Data ONTAP GX. This took some time to figure out.
- After some quick work by our account manager (to switch those cards ASAP), we were able to complete the installation.
- Overall, the installation went perfectly. However, if you were hoping that the zeroing of disks goes away in Data ONTAP GX, you will be disappointed. Our Data ONTAP GX system has more than 100TB of 500GB ATA disks. Zeroing this storage takes a while. In Belgium, we don't like to wait, so this was a perfect moment to go to a local pub for a real Belgian beer like a Duvel.
- After only four days, the system was up and running, everything was documented, and we had our first training session. We also worked with our SE to make an action plan for the next several weeks.

The installation went very well. The most time was spent migrating from the Data ONTAP 7G configuration to the Data ONTAP GX configuration (which was only necessary because of the beta issues). The team of NetApp Professional Services; our partner, UpTime; and our two storage administrators worked very well together. We only used the manual setup. Data ONTAP GX has an automatic setup feature in

which you only install the first node and it adds the rest automatically. This is a great feature, but we wanted to be able to play with the bits and bytes to start out.

July 2007: Data ONTAP GX First Impressions

When you work with Data ONTAP GX, you recognize the WAFL® file system, and you find a lot of known commands, but it's not exactly the same. There are several things in particular that we really like:

- All the Data ONTAP GX commands have exactly the same structure and syntax. This is really great.
- Data ONTAP GX has autocompletion! This must-have feature works perfectly.
- Because it's a single system, you can manage all your storage from any one of the controllers. With Data ONTAP 7G on the four systems, you would have four telnet windows open at the same time. (One for each system.) Now, with Data ONTAP GX, you can manage everything as a single system from one window.
- With Data ONTAP GX you can use wildcards. So, if you want to change the language setting on all your volumes on all your nodes, you only need one command, not a lot of manual effort or a complicated script.

August 2007: Testing

By August, all the systems were upgraded to the supported production release of the Data ONTAP GX software: 10.0.2. Now it was time for the fun part, doing all the "don't try this at home" tests that you just can't perform under normal circumstances.

We first tested all the features, including:

- Transparent migration
- Mirroring
- User interface
- Global namespace
- CIFS (needed for the patient file scanning project)

We also made some baseline performance measurements and tested DNS. You can access all four nodes using one name, so there's some special stuff there.

Next came our DR testing. We tried everything we could think of, turning off nodes, disconnecting power, networks, the system interconnect, etc. None of these tests resulted in any big problems. The biggest problem we had occurred when we accidentally swapped a system interconnect network cable and a client data network cable. (This happened when the network cards were replaced due to the ordering problem I mentioned previously.) This resulted in some strange behavior, but obviously isn't a problem that's likely to occur under normal circumstances.

Overall, the testing program went very well. The behavior of our four-node Data ONTAP GX system was exactly as we had hoped, and we are very impressed with its capabilities. All the new features work perfectly. Transparent migration is comparable to VMware VMotion. You can write to a volume while it's migrating. Sadly, the testing came to an end too soon, and we had to get down to serious work.

Late August and September 2007: Production Begins

Paper Records

At the same time that we were testing Data ONTAP GX, we also began testing the new application to present the images of the scanned patient records to physicians. The application shows a thumbnail for every page of patient data that was scanned. The typical case would be about 80 such thumbnails, but some of the larger files might have up to 800 images. For these larger records there were problems on the client side with memory, but never any issues on the Data ONTAP GX side. Of course, we have a total of 3,500 clinical users and up to 1,500 simultaneous users, so the load may be significantly higher when scanning is complete and the application is in full production.

The scanning process started the second week of August. We're scanning about 3,000 patient files per day, each one averaging 40 double-sided pages. An outside company is doing the scanning for us. Once the data is transferred to us, we do some image manipulation and optical character recognition where possible, and then the data is stored on the Data ONTAP GX system. The application is live, and load will slowly increase as more and more patient data is scanned.

PACS

The PACS application is more critical, so we're proceeding more slowly with plans to go live by the end of September. Our current PACS repository consists of 200 to 250 qtrees, each 200GB in size. We began by migrating the oldest qtree from existing storage to new flexible volumes on the Data ONTAP GX system, since the probability of access is low.

Once testing was completed satisfactorily with this data, we migrated one of eight volumes of new images to Data ONTAP GX and carefully monitored activity and performance. Once we were comfortable that everything was working as expected, we migrated everything to Data ONTAP GX, including all new image caches. We plan to maintain two copies of all new images on FC disk (one copy in each data center) for two months. After the first two months, images will be migrated to SATA, and a single online image will be maintained. (The PACS application maintains a tape backup of all images.)

Overall, application implementation was very fast, so there was no problem staying on schedule. Data ONTAP GX is a great system to administer, and both NFS and CIFS work well.

What's Next?

We've just finished our deployment, so obviously we're looking forward to spending the next couple months running this in our production environment. We're also starting internal discussions and planning for next year.

Based on early results, I strongly believe we will find opportunities to begin migrating other applications to the Data ONTAP GX platform. The PACS application consists of just two large Sun™ servers that access Data ONTAP GX, while the patient records application runs on a single server. We'll be looking for similar types of applications to migrate, because that makes it easy to identify any problems that might occur versus applications where 3,000 clients access Data ONTAP GX directly.

In the longer term, I see the combination of Data ONTAP GX and VMware as winning the lottery.



Figure 4) UZ Leuven NetApp storage environment

TECH ONTAP ARCHIVE



Jot Gill

Strategic business consultant, NetApp Value Engineering

Jot has over 10 years of IT experience in the storage, large systems, and open systems arena. Prior to joining NetApp, he ran his own consulting business that performed heterogeneous storage assessments, recoverability analysis, consolidation studies and implementation, and migrations. Jot started his career in the large systems and storage group at IBM where he was part of the IBM Shark validation team and also responsible for the first in-line, large heterogeneous migrations from third-party storage to IBM. He has an engineering degree from the University of Toronto, Canada.

Three Hidden Backup Problems That Can Affect Recoverability

By Jot Gill

Storage vendors like to talk about backup, but storage managers are quick to point out, "Backups are worthless; restores are priceless." Unfortunately no matter how diligent you are in the design and implementation of your backup strategy, relying solely on information from backup hardware and software can result in significant delays or even failures when it comes time to recover. Backups that are reported as successful may still not meet the data consistency or policy requirements of the application using the data. Verifying application recoverability and ensuring that you can meet business availability requirements are essential.

NetApp has helped numerous IT teams across a wide range of industries do a full assessment—from applications to infrastructure—that takes into consideration all elements that can impact backups. These assessments have identified a variety of commonly occurring "hidden" problems that can cause major issues if a disaster or failure occurs.

This article highlights a comprehensive backup and recoverability assessment methodology and identifies three common issues that these evaluations can uncover:

- Hidden scheduling conflicts that increase recovery risk
- Lack of coordination between Snapshot™ copies and backups that can lead to the wrong data being backed up
- Inconsistent application backups with VMware

These examples should help you assess your own backup and recovery environment and identify any hidden vulnerabilities.

A Comprehensive Backup and Recoverability Assessment Methodology

Typical backup assessments look only at the backup environment to determine whether or not key backup hardware and software are operating correctly. These evaluations don't examine anything upstream from the backup process. A more thorough assessment may be necessary to establish true confidence in the recoverability of your environment. This approach should include (in order):

1. Analysis of application alignment down to the infrastructure level to understand how applications, databases, and file systems map to servers, networks, and storage to identify all dependencies on storage and backup

RELATED INFORMATION

- [Information Optimized Backup](#)
- [Moving to Disk-to-Disk Backup](#)
- [Integrating Apps and Storage](#)
- [NetApp Data Protection Strategies \(PDF\)](#)

Information-Optimized Backups

It used to be that a larger library, faster tape drive, upgraded network, or new backup server kept backups running.

Over the past five years, however, disk capacities have grown 20x, with a scant fourfold performance improvement for backup. Today, many companies are turning to smarter backup solutions that:

- Reduce backup storage consumption by eliminating redundancy within backup files
- Shrink backup windows by copying only incremental data changes
- Ensure reliable, recoverable backups by validating backups as they are created

Read [When Brute Force Isn't Enough: The Value of Information for Backups](#).

2. Careful evaluation of the nuts and bolts of backup and replication, including:
 - o Database journaling and transaction logging
 - o Snapshot schedules
 - o Replication schedules
 - o Analysis of the backup catalog, backup schedules, etc.
3. A determination of recovery objectives and needs. The ideal approach involves a formal interview process that includes the line of business owners, application owners, IT staff—anyone who has information about a specific application, including recovery needs, objectives, business impact, available downtime, etc. Every discussion should be broken down by application.

With this information in hand, you can see if the desired recovery point objective (RPO) and recovery time objective (RTO) can be met for each application and begin to identify potential problems. It may also be valuable to look at the recovery capacity objective (RCO, a relatively new concept) to understand how much storage is needed to recover.

Here are three examples in which applying this methodology uncovered hidden problems. These assessments were carried out by collaborative teams that included NetApp Value Engineering, NetApp Global Services, outside consulting groups, and the customers' internal IT staffs.

Recovery Challenge #1: Meet Strict RPO and RTO Objectives **Hidden Problem: Scheduling Conflicts Increase Recovery Risk**

In most organizations, the company sets universal goals for the amount of acceptable data loss (RPO) and how long a given application can be down (RTO), and the storage team designs a backup plan to achieve these goals. While the plan may look solid on paper, in many large organizations the IT team supports a variety of applications and environments and may find it difficult to be 100% confident in their ability to meet these targets.

For example, I recently visited a major telecom company to help them verify their ability to meet corporate RPO and RTO objectives and identify any potential points of exposure.

Customer Environment

The customer environment consisted of applications running on both Oracle9i™ and 10g™ with Windows®, Solaris™, and Linux® servers. Storage included NetApp storage systems running SnapVault® for disk-to-disk backup and SnapMirror® for replication, plus a variety of other storage products, including EMC storage with SRDF software for replication. The data protection schedule included the following steps:

- The IT team created Snapshot copies every hour, retaining the four most recent Snapshot copies, a daily Snapshot copy, and a weekly Snapshot copy on primary storage.
- Older Snapshot copies were sent to physical tape and recycled on a set schedule.
- Long-term copies for compliance and other retention requirements were sent off site on tape.
- Every four hours NetApp SnapVault was used to copy the latest Snapshot to a secondary storage system on site; this Snapshot was then made available to the backup subsystem.
- NetApp SnapMirror and EMC SRDF replication software was used to mirror data to a remote disaster recovery site every hour.
- Every 24 hours data was backed up to tape at the DR site and was recycled on a scheduled basis.

On paper, it would appear that these steps should have ensured that the IT team could meet or exceed the company-mandated four-hour RPO.

Customer Findings

A thorough analysis (using the method described in the previous section) was extremely revealing. Interviews with the database team revealed that DBAs were not communicating regularly enough with the storage and backup team. Because there was a lot of pressure from the business units and application owners, almost every IT

A Quick Primer on NetApp Data Protection Software

NetApp customers have two potential alternatives for data protection: SnapMirror or SnapVault software.

SnapMirror is replication software intended for disaster recovery solutions. The mirror is an exact replica of data on the primary storage including all the local Snapshot copies and can be mounted read/write to recover from failure. If a Snapshot backup is deleted on the source, it will go away on the mirror at the next replication.

SnapVault, in contrast, is intended for disk-to-disk backup. A separate Snapshot retention policy is specified for the target environment, allowing long-term archival of Snapshot backups on secondary storage. Secondary storage used by SnapVault cannot be mounted read/write. Backups must be recovered from secondary storage to the original or an alternative primary storage system in order to restart.

At a more technical level, SnapVault takes a point-in-time image based on qtrees, while SnapMirror offers both qtree-based replication and replication of a full volume.

Get the details. Read the reports:

- [SnapMirror Best Practices Guide](#) (pdf)
- [SnapVault Best Practices Guide](#) (pdf)
- [Data Protection for NetApp Storage Systems](#) (pdf)

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 allows administrators to apply predefined policies to their data, thereby eliminating ambiguity and the potential for error inherent in manual management.

A simple dashboard depicts comprehensive data protection information at a glance, including unprotected data, longest lag times, etc. The application automatically correlates the datasets and underlying physical storage resources, so administrators don't need to think in terms of the storage infrastructure.

resource (including DBAs) was keeping multiple copies of data to meet recovery and availability objectives. In addition, the majority of test/dev environments were being given the same priority as production environments for backup and replication.

The IT team (logically) assumed that if one replication succeeded every four hours, they were meeting the four-hour recovery point objective. However, because there was no communication from the storage layer upstream to the database layer the IT team had no visibility into production or test/dev and had no way of knowing whether replication or backup operations were missing data at critical points.

In fact, as the result of an extremely busy test/dev schedule, DBAs were setting up and breaking down so many test/dev environments that it was causing SnapMirror and SRDF replication schedules to be missed. (So much data was changing that one replication cycle was not completing before the next cycle was set to begin.) Effectively, this could make it impossible to recover to certain points in time. We were able to identify these issues by mapping applications to their associated databases, then to storage, and finally to specific server backup functions. The discovery and mapping of the entire hierarchy were key.

Solution

The solution in this case was to separate the replication processes for the production and test/dev environments and use information gathered about each application to synchronize replication schedules with important events (such as archive log completion) to create a consistency point during each replication cycle. This ensures that replication of production data will not be affected by test/dev and that the target volumes will always contain the data needed for recovery to a given point in time.

An additional outcome of this analysis was the determination that, given the fast cycle time of test/dev environments, there was no need to back up test/dev data to tape. This decision saved a significant amount of tape resources and time

Recovery Challenge #2: Restore Doesn't Return Expected Data Hidden Problem: Non-coordinated Snapshot and Backup

In most customer environments, it's common to take a Snapshot copy of production data and then mount that Snapshot copy on another server for the purposes of backup, test/dev, reporting, etc. This is certainly true in NetApp environments, but it's also common at sites using IBM FlashCopy, EMC SnapView, etc. Nobody disputes that this is a useful and worthwhile way of doing things, but problems can arise if you don't properly coordinate your operations.

Customer Environment

We recently worked with a large hospital network to debug a problem in which the tape backups didn't always contain the data the IT staff thought they should. The hospital was using applications based on both Oracle® and Sybase and using Veritas™ NetBackup™ and Tivoli Storage Manager to back up data to tape.


Snapshot copies were being created on a four-hour basis. These Snapshot copies were retained for recovery mounted on a secondary server to be backed up. Because the database and associated file system were supposed to be quiescent at the time the Snapshot was taken, it was assumed that the Snapshot was consistent, and hence the backup was also consistent.

Customer Findings

A detailed backup and recovery analysis revealed a lack of coordination between Snapshot copies and backups. Sometimes backups would get delayed because of a tape problem and a new Snapshot copy would be taken and mounted before the backup actually started. This information was gathered by building a dashboard view of all application and database-specific backup steps and correlating it to Snapshot copies and backups with complete logging and audit tracking.

While this may not sound like a big problem, in this case the scheduling was such that databases were being purged in the intervening period, resulting in potential data loss and a backup catalog that didn't reflect reality:

[Read the Protection Manager Technical Report \(pdf\)](#)



Time	Event
6:00 p.m.	Snapshot created and mounted
6:15 p.m.	Tape backup scheduled to begin
6:45 p.m.	Purge database record
7:00 p.m.	New Snapshot created and mounted on same mount point
7:15 p.m.	Tape backup actually begins (purged record is not backed up)

Solution

The solution for this customer was to improve coordination and monitoring between the Snapshot copy creation and backup processes. This included ensuring that the second Snapshot copy doesn't get mounted until the backup occurs, which is made possible through a notification feature built into the NetBackup schedule.

Because this customer is primarily using NetApp storage, we also recommended NetApp SnapManager® products for appropriate applications and databases to provide more visibility and consistency across Snapshot copies and NetApp Operations Manager to provide a visual dashboard for monitoring events.

Recovery Issue #3: Inconsistent VMware Application Backups

Hidden Problem: VMware Backups May Not Be "Application Aware"

The final problem is one that I've seen a lot recently: backups of VMware environments performed using VMware Consolidated Backup (VCB) or other backup methods are not application aware and don't result in consistent backup images for active applications. In addition, they can be extremely resource intensive, especially if they leverage the standard VMware snapshot capability.

Customer Environment

The customer site where I first saw this issue was running Oracle applications and Exchange on VMware and running VCB at regular intervals for backup. (NetApp Global Services has seen the same thing since then at a number of other customer sites.)

Customer Findings

A detailed, application-centric backup assessment discovered that when VCB is kicked off it simply quiesces the running virtual machine and makes a copy without regard to application state. The resulting backup is therefore "crash consistent" (the virtual machine can be restarted from the backup), but not application consistent.

NetApp Global Services uses tools that can look into running virtual machines and map out information relative to the virtual and physical hardware and associated software. Applications are generally not affected by running VCB because it only takes a few seconds for the VMware snapshot or storage snapshot to occur, but the state of the application data on disk is generally inconsistent because applications typically cache critical data and/or metadata in memory for performance reasons rather than writing it to disk immediately. It may be impossible to restore an application to a working state from such a backup.

Solution

The solution is to put the application in "hot backup" mode prior to running VCB, shut the application down entirely (usually undesirable), or run some form of application-aware backup inside the virtual machine for application backup.

With NetApp storage you can run SnapManager for Exchange, SnapManager for SQL Server™, or any of the other SnapManager products on individual virtual machines to properly quiesce the application, flush application data to disk, and take a consistent Snapshot for backup without significantly interrupting application work.

Products such as NetApp SnapDrive® also provide the basis for OS-aware Snapshot copies. SnapDrive scripting capabilities can be very easily combined with VMware ESX/Server API-based scripting. Many customers have used these capabilities to write scripts to create consistent Snapshot copies and SnapMirror copies for disaster

recovery. These "copied" virtual machines can be mounted in standby VMware environments to confirm consistency and availability.

A recent Tech OnTap article describes [multiple uses of NetApp Snapshot technology with VMware](#). NetApp is working closely with VMware to increase the coordination between VMware and NetApp storage to further enhance these capabilities.

Reduce Your Recovery Risk

These are just three of the "gotchas" that can lead to nasty surprises when it comes time to recover. A thorough backup assessment—such as the approach described here—is the best way I've found to uncover hidden problems with your data protection environment before they have a chance to bite you.

TECH ONTAP ARCHIVE



Alvin Richards

Technical Director for Database Integration, NetApp Engineering

Alvin joined NetApp in 2004 after more than 16 years at Oracle and a further 2 years at a database archiving startup. His wealth of database knowledge and experience has allowed him to serve as the leader of a virtual startup organization within NetApp focused on developing applications that allow databases to take full advantage of NetApp storage technologies. Other contributors to this article include Steve Daniel, NetApp director of Database Platforms and Performance Technology, and Vasu Subbiah, NetApp director of Product Management.

Oracle on NetApp Part II: Simplified Fibre Channel SAN Management

By Alvin Richards, Steve Daniel, and Vasu Subbiah

This three-part series on Oracle® storage architectures began with a discussion of how NetApp engineering optimized the use of [Oracle with NFS](#). Next month, look for a technical case study involving Oracle on NetApp using NAS, iSCSI, and FC SAN.

Many IT organizations rely on Fibre Channel SAN architectures as a proven solution for Oracle deployments. Complex deployment and management processes have traditionally been viewed as an unavoidable trade-off associated with high performance SAN environments.

It doesn't have to be that way.

NetApp Fibre Channel SAN offerings deliver not just high performance but also a broad range of manageability features that lower the complexity associated with deploying and managing Oracle environments. This article explores NetApp engineering's work in the FC SAN space, including

- Developing unique management tools tailored to Oracle environments
- Enhancing FC SAN management and efficiency
- Delivering flexible and cost-effective DR solutions
- Ensuring full compatibility with SAN switches, HBAs and file systems
- Integrating with a broad suite of SRM tools

Developing Unique Management Tools Tailored to Oracle Environments

NetApp and Oracle have collaboratively developed and deployed solutions since 1995. In 2004 NetApp began a program of sustained investment in developing unique software tools designed specifically for Oracle environments, including SnapManager® for Oracle and SnapValidator™.

SnapManager for Oracle

Backup, restore (getting the files back from a backup) and recovery (getting a database back into a consistent state) are often complex tasks for Oracle Databases, especially in FC SAN environments where there's a whole storage fabric to deal with.

The goal of the SnapManager for Oracle (SMO) development team was to closely integrate the strengths of a NetApp storage system—including backup, restore, and cloning—with Oracle's RMAN, RAC, and ASM (automatic storage management) products. This integration greatly simplifies the day-to-day management of test/dev clones, Snapshot® copies, and backups. SMO provides a variety of advantages

RELATED INFORMATION

- [Oracle on NetApp Part I: NFS](#)
- [Best Practices for Oracle](#) (PDF)
- [Oracle Deploys Oracle Enterprise Manager Grid Control and NetApp FC SAN](#) (PDF)
- [Integrated Data Management: Improving IT Productivity and Flexibility](#) (PDF)

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 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 technical director 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

discussed at length in a [previous article](#), including three unique benefits no other SAN storage vendor can provide:

- **Ability to restore a single file within an ASM disk group.** SMO understands the ASM file layout, making it possible to restore a single file from a Snapshot of an ASM disk group rather than the entire disk group (which may contain files from different databases).
- **Ability to clone ASM disk groups and make them available to the same host.** SMO understands the ASM label and integrates with ASM to create a unique new label for disk groups cloned with FlexClone®.
- **Ability to create a consistent point-in-time Snapshot copy across multiple storage systems.** SMO can create consistent Snapshot copies of an ASM or Oracle instance, even when the instance spans multiple store systems, and even when the storage shares the underlying disks with other applications.




Figure 1) File restoring with partial single file restore

SnapValidator for Oracle

On rare occasions, faulty hardware can corrupt data in flight between server and storage. Because Oracle keeps a checksum on each block, such corruption will be detected, but data corrupted during a write may not be read back and checked for many months. By then, the operational difficulties of recovering the data may be substantial.

NetApp 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, ensuring that corrupted blocks are never written to disk.

Enhancing FC SAN Management and Efficiency

NetApp offers all the standard features of mission-critical SAN storage, including high availability storage systems, non-disruptive upgrades, and a complete compatibility matrix. (More on compatibility later.) If you're out of I/O capacity in one controller, for example, you can scale across controllers and still create consistent backups of your data. NetApp also offers a variety of unique features, including [dual-parity RAID](#) (RAID-DP®), which provides the data reliability of mirroring (RAID1) at the price/performance of RAID5, and [dynamic queue management](#), which shifts queue depth management from the host to the array target port.

NetApp also has in built protection against so called "lost writes". Most vendors need to verify written data because they are unable to detect situations where the drive does not write the data but says it did. In those situations, most storage systems will see an old block with correct ECC and other checksums yet the RAID parity will be inconsistent. There is no way of recovering lost writes in these environments because it is impossible to tell which block is old. Because the NetApp WAFL file system writes new data to new places and uses a logical ID embedded in the checksum, NetApp offers the unique ability to identify lost writes.

However, the greatest operational improvements arise from the integration of NetApp data management features with SAN technology to provide a superior operating environment for Oracle databases.

NetApp Snapshot Technology

NetApp Snapshot copies are space-efficient, point-in-time copies that can be created in a matter of seconds with minimal impact on performance. 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. With SnapManager for Oracle, the entire process becomes a single mouse click.

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

Applying Innovation to FC SANs

In a typical SAN environment, manually setting host queue depths can be a complicated process. Providing a standard queue depth setting is simply a "shot in the dark" approach that hopes the assigned queue depth will provide adequate application performance.

To address this problem, the NetApp engineering team implemented dynamic queue management, which shifts queue depth management from the host to the array target port. Simply put, dynamic queue management controls queue depths from the array side.

By monitoring QFULL conditions, application response times on a per-LUN basis, and queue lengths, dynamic queue management automatically adjusts queue depth based on actual application I/O requirements.

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

Driving the Fibre Channel Standards Roadmap

NetApp technical experts have taken [leadership roles](#) in the Fibre Channel Industry Association (FCIA) and other related organizations:

- Jeff Eckard (Senior Manager of SAN Product Management) is the FCIA board.
- Fred Knight (Senior SAN Architect) sits on the T10 (SCSI standards) and T11 (interconnect standards) technical committees and is currently contributing to the FCoE standard and evolving security standards.
- Mike McNamara (Senior FC Product Marketing Manager) is the Chairman of the FCIA marketing committee driving all FC messaging and positioning statements as well as the FC roadmap.
- Denise Ridolfo (Industry Marketing Operations Manager) sits on the SNIA and FCIA Alliance Working Group to ensure that technology initiatives are consistent across organizations.

Jeff, Mike, and others have written a number of recent articles on [Fibre Channel standards and trends](#).

Because of the very brief and minimal performance impact, most DBAs schedule multiple Snapshot copies throughout the day as a safety net. NetApp systems can retain up to 255 Snapshot copies of each LUN and revert to any previous copy should a problem occur. Taking more frequent backups can significantly reduce mean time to recovery. With daily backups, the worst case scenario is that a storage manager might have to apply 23 hours and 59 minutes of logs to recover the database. Using Snapshot copies allows backups to be created hourly, which results in a worst case scenario involving 59 minutes of logs.

A data protection solution for a production database begins with Snapshot copies on the primary storage; however, a typical implementation will also copy the Snapshot copies to long-term secondary storage using NetApp SnapVault®, mirror data to a remote disaster recovery site using SnapMirror®, back up to tape, or some combination of these options. Because all of these data protection operations occur on the storage system, the servers are free to keep running Oracle applications at full speed.

NetApp FlexVol Technology

In addition to the efficiencies provided by NetApp Snapshot technology, NetApp offers a variety of other features to improve storage utilization including Netapp FlexVol® technology.

FlexVol allows volumes of varying sizes and usage patterns to share a common pool of physical space. By mixing together volumes with differing performance and capacity needs, you often require far fewer disk drives than they would if each volume had dedicated disk resources. Furthermore, as requirements change, flexible volumes can grow and shrink, or the underlying space pools can grow as required.

NetApp FlexClone Technology

NetApp FlexClone® enables the creation of logical copies of production data for test and development, data mining, or other purposes. FlexClone only consumes additional disk space as changes are made to the original LUN or the clone, which makes it possible to clone production data without massively increasing your storage requirement. If operational requirements dictate that development and test databases reside on separate storage from the production system, a physical copy can be maintained using SnapMirror, and space-efficient FlexClones of the SnapMirror destination may be used for development or test.

Blackboard, a NetApp customer, combined the space efficiencies of FlexVol with NetApp FlexClone technology to cut its storage needs by one third and streamline database cloning for test and development. It used to take 36 hours to make a working copy of a 2.5TB database. Obviously, any process that takes 36 hours and consumes 2.5TB of storage is not one that many people are anxious to perform. After implementing a NetApp FC SAN with NetApp FlexClone technology, the IT team can now accomplish the same task in less than an hour without using twice the disk space up front. ([Read the full story.](#))

Delivering Flexible Disaster Recovery Solutions

Implementing disaster recovery in traditional SAN environments is difficult, expensive, and comes with many limitations. Most solutions are complex enough to require the vendor to set them up. Oracle disaster recovery solutions using NetApp SnapMirror replication software overcome these hurdles.

Setting up a mirroring relationship in SnapMirror is an easy process; it's as simple to create a mirror as it is to create a volume. SnapMirror also eliminates many of the constraints that limit other solutions:

- SnapMirror can mirror between different storage systems; no need to have the same platform—and the same disk drives—on source and target.
- SnapMirror can mirror between different hardware tiers (e.g. primary storage for the source and less expensive secondary storage for the target.)
- You can use different storage protocols on source and target. For example, Oracle might use FC SAN on the source and iSCSI to access clones on the target.

Most customers find that the simplicity of creating an asymmetric disaster recovery

Improved Data Protection with RAID-DP Technology

RAID-DP significantly increases data protection, with zero to minimal impact on capacity utilization and performance. And, because RAID-DP is an integral part of Data ONTAP®, there are no hidden costs associated with it.

RAID-DP offers:

- Protection against up to two disk failures in the same RAID group
- Protection against single disk failure plus uncorrectable bit error during the reconstruction time frame
- No significant read, write, or CPU consumption differences
- Larger allowable RAID groups, which mean that capacity utilization stays about the same (one in eight disks dedicated to parity)

Learn more about the [NetApp Implementation of RAID Double Parity](#)

solution significantly reduces the cost of disaster recovery.

Ensuring Full Compatibility with SAN Switches, HBAs and File Systems

A robust FC SAN solution must ensure compatibility with the wide variety of SAN hardware and software products. A single operating environment can include host operating system releases and updates; native I/O stack components (volume manager, multipath I/O, file system, host cluster); Veritas Storage Foundation I/O stack components (VxVM, DMP or vDSM, VxFS, VCS, CFS); server hardware; HBAs; HBA firmware and drivers; host virtualization; other identified features of an I/O stack release; fabric and director-class switches; and firmware. Obviously the number of possible permutations is staggering.

NetApp's compatibility matrix today runs to over 200 pages and 1,600+ entries covering tens of thousands of configuration possibilities. NetApp customers can access the full [compatibility matrix](#) on the NOW website.

Our efforts to achieve the broadest possible compatibility are bolstered by close partnerships with major switch vendors and HBA vendors, as well as with Symantec for support of the VERITAS Volume Manager, VERITAS File System (VxFS), and related products.

Integrating with a Broad Suite of SRM Tools

Compatibility with the tools used to manage all of the elements of your SAN environment is as important as compatibility at the switch and HBA level. For many customers, NetApp Operations Manager software provides an excellent solution for managing NetApp storage.

Customers with existing storage resource management (SRM) infrastructure will find that NetApp SAN solutions easily integrate with existing SRM solutions, including:

- Symantec CommandCentral Storage (resold by NetApp)
- HP Storage Essentials
- EMC ControlCenter (ECC)
- IBM TotalStorage Productivity Center (TPC)

NetApp storage solutions are also supported by a wide variety of up and coming storage management solutions from vendors such as Onaro, Brocade, Akorri, MonoSphere and more.

Summary: Choose a Solution that Fits Your Environment

If your objective is to implement a high performance storage environment that has the flexibility and manageability features required for success, a NetApp FC SAN solution may be the right choice. NetApp has a proven track record of Oracle deployments on FC SAN, and the combination of NetApp storage features such as Snapshot and FlexClone with Oracle-specific software like SnapManager for Oracle and SnapValidator can make the lives of DBAs and storage administrators much simpler.

Ultimately, of course, your choice of storage architecture depends on a wide variety of factors that include existing infrastructure, processes, and skill sets. NetApp offers excellent implementation options and tools for deploying and managing Oracle environments for Fibre Channel, NFS, 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, attend one of four NetApp special sessions including:

Hear a Real-World Data Management Success Story

Exploring the Future of Data Management
Tuesday, November 13 – 3:15 pm

Join Network Appliance co-founder and corporate philosopher, Dave Hitz and Niel Armstrong, VP of IT from lead game manufacturer, Activision, as they examine real-

world issues associated with managing enterprise data, discuss how data storage solutions have evolved to handle these challenges, and explore the future of storage and data management. Get the information you need and enter to win a video game package from NetApp and Activision!

[See the complete list of NetApp speaking sessions at OpenWorld](#)
[Sign up by November 9th for a free VisitorPlus pass to OpenWorld](#)