

» Tech OnTap Home

March 2006

HIGHLIGHTS

- » **iSCSI-Based Diskless Server Farm**
- » **Remote Office Best Practices**
- » **Simulated I/O Tool**
- » **Special Offer:**
Win a Renault F1 Team Helmet

Storage Passes the Qualifying Round

Q&A with Graeme Hackland,
IT Manager, Renault F1 Team

Renault F1 Team's centralized storage strategy plus enter to win a signed racing helmet.



[More »](#)

BLOGGING WITH DAVE

Dave Hitz, NetApp Founder and EVP

"Categorizing products according to their use of virtualization is about as useful as categorizing them by the programming language."

» [Dave's Blog](#)

DRILL DOWN

- » **NetApp Outlines FY2007 Plans**
Execs share our vision for the future.
- » **Benefits of Upgrading to a 4Gb SAN**
Upgrade considerations and benefits.
- » **NetApp Disk-to-Disk Backup Solutions**
Resources include an expert panel's advice on choosing solutions for your environment.



ADMIN RESOURCES

- » **Simulated I/O Tool**
This client-side workload generator provides a much better indication of database performance than just running scripts or doing a DD test.

TIPS FROM THE TRENCHES

Winning the War on Complexity: Oracle IT's Implementation Grid

Bill Weils, Sr. Director, Oracle Global IT Command Center

Does grid really work? Bill can prove it. One of the first companies to envision and adopt a grid architecture, Oracle permanently cut IT spend nearly 50% within 24 months.

[More »](#)

Five Tips for Improving Remote Office Backup and Recovery

Darrin Chapman, Product Engineering Manager, NetApp

With more than 200 remote sites and connections as slow as 2Mb per second, Norsk Hydro is a textbook example of backup and recovery done right. This article includes advice for managing OSSV clients, implementing bare metal recovery, and more.


[More »](#)

ENGINEERING TALK

The NetApp Kilo-Client An iSCSI-Based Diskless Server Farm

Gregg Ferguson, Lab Administrator, and
David Brown, Engineering Support Manager, NetApp

Booting a 1,120-node server farm from an iSCSI-based SAN cut provisioning time up to 500x. Learn how NetApp Engineering clones LUNs in milliseconds and changes OS in the time it takes to reboot a server.



[More »](#)



GRAEME HACKLAND

IT Manager, Renault F1 Team

Born in Scotland and raised in South Africa, Graeme moved to England in the mid-90s and has been a member of the [Renault F1 Team](#) for the past eight and a half years. During this time the team's IT group has grown from five people to 25. Graeme is currently based at the team's England factory and is responsible for virtually all of the company's IT infrastructure, including storage, networking, WANs, security, desktop hardware and software, and IT support.

How Storage Passed the Qualifying Round for the Renault F1 Team

Q&A with Graeme Hackland

Q. First, congratulations on winning the 2005 Drivers and Constructors world championships.

We understand that the Renault F1 Team has implemented a centralized storage strategy with NetApp as your chosen storage supplier. Could you tell us how you reached this point?



Graeme: Before we really embraced a centralized storage strategy, we tended to just throw in any old device. We used to use mostly servers with direct-attached storage. We might buy some storage from somebody and put it in, and then when we wanted to go back and get some more a few months later, we might find the device wasn't available anymore. When the little 1U network-attached storage (NAS) devices came along, they were very cheap and easy to deploy, so if a particular department wanted a bunch of storage, they got one of those 360GB devices.

We worked like that for about three or four years before realizing what a nightmare it was to keep all these little devices with various operating systems backed up, looked after, and highly available. We had a complex environment to back up to tape which consisted, at one point, of around 60 Windows® servers, some Novell servers, and a whole UNIX® environment supporting our CAD/CAM 3D design applications.

Q. What was the first step toward centralizing and standardizing storage?

Graeme: The Renault F1 Team has IT systems in two main areas: at our factories and at the track. We have two factories, one in France (where the engine is made) and one in England (where the chassis is made). At the factories we have our CAD design environment and what we call our office environment. Office data includes e-mail, Microsoft® Office, race videos, pictures, car data—basically everything that isn't involved in actual car design.

We initially made the investment on the CAD side because this is where we have our most critical data. In the 2005 season, for example, over 50 full-time designers created more 13,000 new technical drawings from which we produced approximately 250 engines. That's just in one year. In 2006 we've already created over 7,500 drawings.



In early 2000 we formed technical partnership with Network Appliance. We chose NetApp storage to support vehicle design and development for a variety of reasons. First, it met our performance and availability requirements—we saw a 60% improvement in image loading times after moving to NetApp storage. NetApp is also compatible with the Dassault Systemes CATIA design package. We use CATIA design software on UNIX and Windows, so a multiprotocol

RELATED INFORMATION

- [The Renault F1 Team's Use of NetApp Technology \(presentation\)](#)
- [The Renault F1 Team Speeds Race Car Development](#)
- [The Renault F1 Team Steps Trackside Data Up a Gear](#)
- [NetApp D2D Solutions](#)
- [NetApp iSCSI Storage Center](#)
- [NetApp Product Development Solutions](#)

Which Types of Data Are Stored at the Track?

The Renault F1 Team collects hundreds of readings in real time while the team car is running laps or undergoing tests.

These include strategy and safety data pertaining to fuel loads, pitstop timing, and immediate assessments of any damage made to the car.

Over 200 sensors on the race vehicle share information about:

- Ride height (front and rear)
- Pushrod loads
- Steering angle
- Throttle position
- Brake wear
- Brake temperatures
- Fluid levels
- Engine parameters


Get the details. Read about the [Renault F1 Team's Use of NetApp Technology](#).

solution that provided uniform, one-copy access to data was very beneficial.

Q. How did you move from using networked storage to support just your design applications to consolidating storage company-wide?

Graeme: There really wasn't any one tipping point. Everybody knew our environment was becoming very fragmented and difficult to look after, but for a long time cost was the overriding factor when it came to office and track data.

In 2003, we got initial agreement and funding to replace the storage on all these random devices with centralized storage. All the reasons that we'd chosen NetApp storage for our CAD environment still applied, so we decided to upgrade to the next generation of NetApp storage and expand our deployment.



For our office environment, I didn't even bother going to tender. At this point we were very familiar with the product and its overall reliability. You might have the odd disk go out, but then by the time you arrive at work in the morning, NetApp has already shipped you a new disk. We were very comfortable with the relationships we have with people within Network Appliance, and we thought we could have some added benefits of mirroring to the second NetApp system we planned to acquire.

Q. Was iSCSI important in your strategy to centralize office storage?

Graeme: Absolutely. Some Windows applications want to only work with direct-attached storage. That didn't fit with our philosophy of wanting to centralize storage. iSCSI allows us to "fool" the Windows server into thinking that the storage is directly attached. With iSCSI, we were able to effectively map a drive letter on the Windows server and make it think it was a local drive.

This allowed us to implement applications like e-mail much earlier than if we had had to find funding to source new hardware or storage. Not only do we use our existing Ethernet switches, we run both iSCSI and NAS on the same NetApp systems. iSCSI performance is so good that it seems as if it's directly attached.

Q. You mentioned IT systems operating at the racetrack. How do you use storage in that environment?

Graeme: As the car travels around the track, we take data from the car via radio. Our trackside engineers need to see that data as soon as possible to tell if the car is low on fuel; if there's an issue with any parts; if they should be telling the driver to go quicker, go slower, or even to take the car off the track immediately because there's a danger of the engine failing.



A couple years ago we felt that the data wasn't getting to the engineers as quickly as it could. We looked at three solutions, including NetApp FAS270s. The FAS270s beat everything else hands down in terms of performance. They're quite small and portable, and they have the data capacity we need.

The challenge is that in addition to performance, reliability, and manageability requirements, any equipment we deploy trackside has to perform in extreme conditions. Bahrain has a particularly dry heat in the desert. Other circuits are hot but extremely humid. We also go to circuits that are quite cold. We were testing in Barcelona two weeks ago, for example, and the last day of the test was called off because of snow.

We did a trial for about four months where we sent one NetApp FAS250 around Europe in a truck through the winter. It survived remarkably well, and as a result we chose the FAS270c. It's been a huge success. We have not had any downtime for the past two seasons.

Q. Has your decision to centralize storage and standardize on NetApp helped the Renault F1 Team keep up with the fast Formula 1 pace and rapid-fire changes?

Graeme: Absolutely. Last year, for example, we decided we wanted to move toward disk staging. I suppose the more common industry term these days is disk-to-disk backup.

Renault F1 Team Shifts Gears with a New Disk-to-Disk Backup and Archival Strategy

Dissatisfied with the 18+ hours they used to spend restoring test data from tape, it didn't take long for the Renault F1 Team to jump on the next wave of disk-to-disk backup in a big way. The team now uses NetApp FAS960 systems, a NearStore R200, NetApp Snapshot, SnapRestore®, and SnapVault® technologies plus Symantec NetBackup software as the foundation of its multitiered backup strategy.

Highlights of this strategy include:

- Creating local snapshot copies of local data every 15 minutes on two NetApp FAS960 systems
- Mirroring data between FAS960s (each system is divided in half and mirrored to the other via SnapMirror)
- Copying snapshot data to a NearStore R200 in another building four times a day (SnapVault)
- Storing three years of CAD data, along with race and test data, online using the NearStore R200

Said Graeme, "Our 15-minute snapshot intervals work very well for certain departments such as our Aerodynamics department. They run at a wind tunnel for about 20 minutes. At that point, they may stop, make some changes to the wind tunnel model, then start the tunnel up again."

Every now and then, he says, they overwrite files and may need to go back to the earlier file they were using 15 minutes ago. To minimize the storage capacity required for the snapshots, the F1 Team overwrites all local snapshots made on the production system after three days.

Get the details. Read about the Renault F1 Team's Use of NetApp Technology.

NetApp Disk-to-Disk Backup Solution Overview

NetApp's comprehensive family of products can help you address critical backup issues in any data center or remote office environment.

If anyone comes to us and says they want to retrieve data from five years or more back, we knew it would take a fair amount of time to locate the tapes, get them loading, get them indexed, and restore the data. That's also assuming the individual tape is okay. If it's not, you've got to go back to the next previous one and do the same process again. It took us 18 hours, for example, to get an old test back that somebody had wanted. That was too long. So we wanted to move toward disk backup and use tape just for archive purposes.

Of course, we have to keep tape because our auditors prefer to think there's a physical tape somewhere. They like to know there's something they can see and touch.

Q. How are you implementing disk-to-disk backup technology?

Graeme: On the factory side, we just started to do disk-to-disk backup with Network Appliance. Network Appliance has some really good technology there. We brought in a NearStore® R200, and we've been putting all our Snapshot™ copies on it. The original idea was to put 12 months of Snapshot copies on NearStore.

We've actually started to use the R200 for other things as well. We recently added Symantec's Enterprise Vault for archiving our e-mail and files on the R200. And the image files that we take of different laptops and desktops before we deploy them are being saved on the R200 too. There's no need to use primary storage for that kind of information. We've actually seen quite a good cost savings in terms of disk by implementing the R200. As an archiving appliance it's proved very successful for us.

Q. In addition to cost savings, what other types of results have you seen after implementing disk-to-disk backup?

Graeme: First, you have to understand that we only have one administrator to manage 80TB of storage from Network Appliance. He checks that all the devices are available, looks at logs and alerts, and generally makes sure the disks and the heads are okay. Otherwise, there's not much work involved since the systems just stay up.

Still, we have various departments who regularly need data recovery, often the R&D department or our Aerodynamics department. Someone there is constantly overwriting files and then realizing that they want something that was deleted two days ago.



In the past we usually had to go to tape, which took a bit of time. Now we are able to train the people who most regularly need data restored about how to restore it themselves. We have set it up so they have the same rights to these specific bits of data as they have rights to on the original production NetApp system. Then, they just browse to where their file is and click and restore it themselves. You hear people talk about "almost instant" recovery, and it really is like that.

The individual files needing to be restored are not usually that big. But, sometimes someone may want to see data from a whole race or a whole test. There, we are talking about gigabytes of data needing to be restored—more than the size of a DVD. At the moment, our IT team can still go online and restore it as long as it's in the three-year backup window where all data is kept on disk. In that case, we're talking about a matter of minutes, or tens of minutes, to restore the old test rather than the 18 hours it used to take.

Q. In addition to NetApp software and Enterprise Vault, do you use any other tools for data management?

Graeme: Yes. We've been backing up systems to tape with NetBackup since at least 1997. We think we were probably using it even before then, but that's the furthest anybody here goes back. We plan to go to Version 6.0 within the next few weeks.

We're also using the VERITAS® Cluster Server (VCS) from Symantec across our UNIX, Linux®, and Windows environments. This gives us a central console regardless of the operating system. It helps people do basic fault finding on any of the clusters without having to know much about the particular operating system on which the cluster runs.



Finally, we use another Symantec application, i3, for application performance



Learn which type of disk-to-disk solution is right for your environment.

Find webcasts, whitepapers and more at the [D2D Solutions Portal](#).

management. This allows us to see where there are either potential performance issues before we roll something from testing to production on databases or on J2EE. It also allows us to do fault finding on the large environment to see what's going wrong and why there might be particular performance issues.


Q. You have a long history of working with both NetApp and Symantec as technology partners. How do you see the relationship between them impacting you?

Graeme: In the past, we tended to use Network Appliance™ tools for anything to do with storage and Symantec for anything to do with tape-based backup and restore. Now, the lines are becoming blurred. The integration between NetApp Data ONTAP® 7G and NetBackup 6.0 means it won't really matter to us which tools are used. The same functionality will now be available in all the tools—whether they come from Symantec or Network Appliance. We're definitely fans of this shift because with the tools becoming this integrated, we don't have to train people on two different sets of tools.

Q. Finally, what advice would you give to someone who is interested in NetApp but might already have a sizable investment in another technology path or vendor?

Graeme: I'm in a situation where our environment changes a lot. I've replaced the network twice in the last five years. We are constantly evolving, constantly changing, trying to keep up with new applications and the new systems. From my perspective, it is important to work with somebody who's also bringing new products to market and who's there for the long term. We've bought products in previous years from companies who then disappeared. We learned from that that it's best to partner with somebody who's going to be around and will continue to be at the front of the curve.

We're obviously a huge fan of NetApp storage, but I would find it difficult to tell anyone, "You've got to rip everything out and put Network Appliance in." Instead, I would say that when you come to your cycle of replacing equipment, you should definitely include Network Appliance in your consideration process.





WILLIAM (BILL) WEILS

Senior Director, Oracle Global IT Command Center

For the last 11 years, Bill has focused on increasing the efficiency and reliability of Oracle Corporation's IT systems and global IT support teams. As senior director of the Oracle Global IT Command Center, Bill is involved in managing monitoring response and systems restoration efforts. He is also a key contributor to the development and promotion of technology in support of managed Oracle use of information technology. Bill attributes much of his leadership ability and disciplined approach to his time in the US military prior to Oracle, which include nearly 10 years spent analyzing intelligence, running internal IT groups, and implementing technology infrastructures for both the U.S. Air Force and the National Security Agency.

Winning the War on Complexity

What You Can Learn from Oracle IT's Implementation of Grid Computing

Q&A with Bill Weils

Q: Oracle was one of the first companies in the world to envision and then adopt a grid architecture. What was happening in the Oracle internal IT environment that led to this transition?

Bill: It began back in 1998, when Oracle first saw the transition from client/server applications to Internet/Web interface-based applications. Back then we had a highly distributed environment for ERP- and CRM-based systems. The Oracle IT organization consisted of over 50 data centers supporting everything from education and classroom support to sales and demonstration facilities and even mobile road shows.

Spending in such a highly distributed environment had gotten out of control. Moving to a Web-based architecture enabled us to consolidate our technology to a single source of expertise and helped us minimize the number of facilities we were using.

By early 2000, we were nearing completion of the Oracle® Web-based architecture (applications) deployment. We then began to integrate the applications architecture with a short list of physical (hardware) vendor-based components, systems, and storage. In effect, having first standardized the overall application architecture, it was natural to extend that standardization to include a common set of industry-acceptable hardware components.

Once we had a solid foundation of application and hardware component architecture, we developed a series of processes and procedures that align and realign common resources based upon the current demands of the business. Instead of chewing up a lot of parallel resources while people run around doing just a little bit of their own thing, the grid concept moved us to a position of leveraged resources that allows us to operate in a more focused and precise way.

The Oracle Multilayered Approach to Grid Computing

RELATED INFORMATION


- [Winning the War on Complexity \(Bill Weils Presentation\)](#)
- [Oracle Deploys NetApp Storage for Its Global IT Infrastructure](#)
- [Running on the Grid](#)
- [Distributed Storage for Scalable Grid Infrastructures](#)
- [Oracle Grid Computing](#)
- [Oracle Presentations and Demos](#)

Oracle Principles of the Grid

1. Standardize on low-cost, modular servers and storage
2. Consolidate servers and storage
3. Automate day-to-day management tasks

Oracle Global IT: the Early Dilemma and Current Situation





Source: *Winning the War on Complexity* presentation by Bill Weils.

Q: What was the impact of the shift to a grid architecture?

Bill: Less than 12 months after adopting a grid philosophy we had cut our overall IT spend by roughly 30%. Within a 24-month period we had reached an IT spend reduction of nearly 50%. These savings were not just a one-time return on investment—rather, these savings have been sustainable year over year as a permanent reduction in our IT expense. In short, these costs are gone forever from the budget. In addition to IT cost reduction, we also noticed a more consistent work schedule for IT staff and an increase in application availability, stability, and performance. The flow of daily operations became much more routine, quantifiable, and predictable.

Q: Can you give us any examples of how the Oracle Grid architecture impacts your business today?

Bill: Sure. During standard business activities we routinely provision and decommission applications and data stores that have a short time-to-live value for the business. As these applications and data stores can easily consume large amounts of storage and system resources, once they have reached an effective end of use we need an operations-driven option to recover idle assets for new projects coming on board. We do this in order to reduce computing system waste and increase system reusability.

Because we've brought our storage together with our server pools and take advantage of storage virtualization using NetApp FlexVol™ technology, we were able to set up procedures to quickly provision and decommission systems from one purpose to another and rapidly repurpose our assets. If storage had been split up among 60 different facilities, none of this would have been possible.

Q: Does the Oracle investment in grid computing for its internal infrastructure help Oracle customers in their own efforts to reduce complexity?

Bill: Absolutely. Oracle did a lot of the early legwork. We developed certified architectures and invested in a lot of R&D. We've hardened and certified specific combinations of technologies. More importantly, we've deployed the Oracle Grid architecture in a real-world environment for a multibillion-dollar organization.

We were able to take these practices and principles and put them into a service market through Oracle OnDemand. This same foundational success has been realized by hundreds of Oracle customers.

We can extend this knowledge and share it with the rest of our customers. We're looking to establish a trust-based relationship with our customers that tells them, "If you're willing to adopt the things we've done, you are likely to have a much more successful experience, in a quicker period of time, since we've already proven these combinations in our own environment."

Q: You often talk to people about how Oracle accomplished its internal grid transformation. What are some of the most common questions you hear?

Implementation Snapshot: Oracle on NetApp

Oracle relies on one of the largest deployments of NetApp storage in the world to support its own internal IT operations.

Deployment scope:

- 14,000 servers
- 3000TB of NetApp storage
- 400 servers (60TB) added each month
- 400 OnDemand customers
- \$200B of financial applications



Get the details. [Read the case study.](#)

Grid Essentials: On the Road to Grid Computing

Most of today's evolving grid architectures involve:

- **Distribution of processing nodes.** Breaking the "big server" model, especially the deployment of hundreds or thousands of processing nodes in a blade format.
- **Distribution of storage.** Includes the ability to minimize the amount of data being replicated.
- **Global namespace for state-of-the-art distributed storage access.** Distributed storage appears to each node as if it were "local" storage.
- **Nondisruptive storage management.** Allows processing to continue without interruption while storage management functions are performed in the background.

Get the details. [Read the full report.](#)

Bill: The first question people almost always ask is, "Is it real? Does the grid really work?" They want to know if this combination of technologies actually mechanically, physically works. My answer to the question is, "Yes, but how successful you'll be really depends on your understanding of the requirements of business."

To help us define Oracle requirements, we made the conscious decision to educate our end-user community to apply what we call "truth-based metrics." If the end-user community said they needed a system to respond quickly, we would say, "Great, but how quickly? Is that three seconds? Is it 10 seconds? What is reasonable?" We had to help our end users understand their requirements so we could measure accurately whether or not we'd met them.

We also had to educate people to separate their actual business requirement from the discussion of how we were going to meet that requirement. It took patience and a firm leadership structure that allows you to tell end users, "What you've presented as a requirement is really a discussion of how I'm going to do it. I need to have the freedom to meet your requirements in the way I think is most feasible."

Q: Any tips to help IT contacts "sell" the grid computing vision to their end users?

Bill: At Oracle we started with a series of early adopters who saw the vision and were willing to take early risks in order to get the early benefits. Then we went to the larger community and used those early successes as a reference. Eventually, you reach a magic tipping point where you're no longer dependent on your success references. It simply becomes an adopted standard, and everybody rushes to adopt the standard.

Q: How much of what Oracle accomplished applies to organizations operating at a much smaller scale?

Bill: Every business can benefit from an ability to flexibly repurpose underutilized resources. This is especially true for smaller organizations, which generally do not have the freedom and luxury of spare "standalone" capacity. Unfortunately, many IT professionals could easily focus on the scale of the Oracle operation (several petabytes of storage and thousands of systems) and possibly overlook the more effective principles of the Oracle Grid that could apply to their own business environment.

Since the Oracle Grid architecture is built on a principle of using well-managed iterations of building block components, it is an effective architecture whether a business needs to have one or two iterations of the building blocks or several thousand. The underlying importance of the Oracle Grid is the ability to manage resources effectively, tapping into and realigning underutilized resources, while increasing overall stability and availability of technology spend.

Q: What other types of questions do customers ask you about the Oracle deployment?

Bill: The IT professionals that we speak with are very interested in understanding more about the grid components that Oracle uses in our own data center. In our data center today, we have a strong focus on low-cost Dell® systems running either Red Hat Linux® or SUSE Linux operating systems with Network Appliance™ NAS-based storage.

The next question they ask is how we selected this short list of vendors. We first explain that Oracle certifies a much larger variety of vendors/technologies than those found in our data center, and they need to understand which combination of technologies provides the best return on investment for them. Since we find the most effective technology management when the numbers of vendors and technologies is reduced, Oracle has itself chosen a short list from among those certified that is the most effective for our business.

Our selection of partners was focused on three key areas. First, each partner needed to have existing successful technology with a very strong maturity/improvement vision for the future. Second, each partner needed to understand the vision of the Oracle Grid and be willing to work as diligently on grid development as Oracle. Third, each partner was willing to provide dedicated engineering resources to work cooperatively in ensuring the highest levels of cross-technology compatibility.

The reasons we have partnered with NetApp involve many things important to not only Oracle business—rather, these are items important to any business. NetApp technology provides flexibility of repurposing, ease of overall management, ease of

maintenance, strong performance per megabyte data rate transfer, ease of integration into an existing technology stack, and a commitment to future improvement that will provide increased performance while decreasing the data center footprint.

Stated simply, the key elements of a partnership with NetApp are found in the effectiveness of the technology and the NetApp commitment to customers as a business partner. The fact that this solution costs less, because NAS architecture itself costs less, is simply an added bonus that eases the difficulty of securing a technology spend during a time when IT spend is under so much review.

Comment on this article.



DARRIN CHAPMAN

Data Protection Subject Matter Expert and Product Engineering Manager, Network Appliance

Darrin Chapman is the person you turn to for just about any question involving disaster recovery or backup and recovery. He's been involved with every NetApp best practices guide involving data protection since 2002 and in his spare time designs training courses for customers and NetApp technical staff. Originally schooled as an electrical engineer, Darrin's background includes several years in network engineering for AT&T and Nortel.

Five Tips for Improving Remote Office Backup and Recovery

By Darrin Chapman

Recent articles have suggested that up to 60% of remote backup procedures fail and less than 50% of data stored outside data centers can be recovered in a timely manner. Ouch!

At NetApp, we've helped thousands of customers efficiently manage and protect their remote office data. One of these customers is Norsk Hydro (abbreviated Hydro), which has operations in 60 countries and is one of the world's leading energy companies.

Over the past five years, Hydro has replaced a mesh of storage networks and direct-attached storage islands with a centrally managed storage infrastructure. The team has also implemented standard backup and recovery processes across over 200 remote sites, many of which are mirrored to a central disaster recovery facility. Over 3.5TB of remote data is currently protected at the company's central business continuance center. Worldwide, Hydro has nearly 400TB of NetApp storage.

In mid-2005, I flew to Oslo, Norway, and spent two days meeting with the Hydro IS Partner team. In addition to sharing corporate best practices, my goal was to document how Hydro had successfully designed and implemented its large-scale infrastructure transformation.

Hydro is an example of remote office backup and recovery done right. Below are five things any company interested in protecting remote office data can learn from Hydro's implementation:

#1: Consolidate on Networked Storage Where It Makes Sense

When Hydro initiated this project, all of its remote offices relied on various forms of local storage. Some of Hydro's remote offices are tiny, but others have 200 to 300 users and dozens of small file servers that need to be backed up. With data volumes increasing as much as 80% a year, the situation was becoming unmanageable. One of the first things Hydro did was identify the locations where it made economic sense to consolidate multiple file servers with a single networked storage system.

#2: Replace Remote Tape with Centrally Managed Backup

When Hydro started this project, nearly all of its remote offices were dependent on local tape backups. Roaming IT staff serviced the needs of many locations, meaning it might take days or weeks for something as simple as a file restore. One of Hydro's top priorities was to centralize backup information at major regional data centers.

Unfortunately, low-bandwidth connections between these sites and remote offices (which include distant oil rigs) made significant data transfer over wide-area networks (WANs) impractical. Working with NetApp, Hydro rolled out a solution that enabled it to back up remote office data over connections as slow as 2Mb per second to nearline

RELATED INFORMATION

- [Norsk Hydro Implements Company-Wide NetApp Solution](#)
- [SnapVault Deployment and Configuration](#)
- [Enabling Rapid Recovery with SnapVault](#)
- [Leveraging NetApp SnapVault for Heterogeneous Environments](#)
- [SnapMirror Best Practices Guide](#)
- [Data Protection Best Practices](#) (password required)

Implementing Efficient Backup over Low-Bandwidth Connections

NetApp SnapVault software streamlines backup and restore operations, making network backup and restore possible, even over low-bandwidth, high-latency connections.

Most backup applications work at the file level. For example, if even one byte changes in a 10MB file, the whole file is backed up. By comparison, SnapVault only sends the block containing the change (4KB versus 10MB in this example).

SnapVault provides this level of efficiency even when backing up third-party systems, including Linux®, Windows, and UNIX. In addition to great bandwidth efficiency, SnapVault also achieves great storage efficiency. Since only changed blocks are stored, incremental SnapVault backups use very little disk space and can be maintained on disk longer than other disk-based backup solutions and at a fraction of the cost.

secondary storage systems located in regional data centers.

To centralize backups for remote offices with very low-bandwidth connections, Hydro chose to:

- **Use block-level backups that copy only incremental changes.** Hydro uses NetApp SnapVault® software to back up only changed data from NetApp storage and open systems SnapVault (OSSV) to back up Windows® and UNIX® servers in smaller offices. This is an area where NetApp has found a unique niche because we have a product that intelligently moves only data blocks that have changed (rather than whole files). This makes extremely efficient use of network bandwidth, allowing backups over low-bandwidth, high-latency WAN connections. (See sidebar for more details.)
- **Assign every OSSV client (primary system) its own storage volume.** An OSSV file system or backup is written inside a quote management tree called a qtree (similar to a directory in a file system). Qtrees live within volumes on NetApp systems. Although NetApp advises against it, it is possible to create a single large volume supporting, for example, 50 or more OSSV primaries. An OSSV schedule applies to an entire volume, so in this scenario when backups are initiated the NearStore® secondary storage system will start pulling data from the 50 primary systems and sending it to the 50 qtrees as 50 concurrent backup streams. This can impact secondary system resource utilization and may approach concurrent session limits. In order to avoid this, Hydro assigned each OSSV primary its own volume where only the qtree belonging to that primary lives. A second issue is that one link may be slower than the others. For example, let's say one client has a small change rate and a fast link while another has a very large change rate and a slower link. Even if the speedy client finishes in five minutes, the backup cannot be completed until the slowest system is done. Again, by implementing a process where each OSSV primary is backed up to its own individual volume, this issue was avoided at Hydro.
- **Stagger schedules.** By placing each primary system in its own volume, Hydro ensures it only has one backup stream at a time. The IT team staggers its backup streams by distributing the OSSV backup schedule throughout the day to make the most efficient use of available network bandwidth and resources.
- **Move all data that is backed up off to its own drive.** Hydro implemented a policy in which all data that is backed up regularly on each Windows server is stored in one logical drive (E:), which the IT team refers to as the data drive. That way the team doesn't have to look for the data that needs to be backed up; it just points SnapVault to the E: drive on each server. This was something I hadn't seen before, and it is a very clever approach to simplifying backups

#3: Use Disk Volumes Instead of CDs For Bare Metal Recovery

A less frequent recovery scenario involves restoring an entire server to operational status after a hardware failure or disaster takes it down—commonly known as bare metal recovery (BMR). Given the amount of critical data being stored in some remote offices, Hydro had identified this as a priority for key offices.

Typical bare metal applications create a data volume that contains all the information necessary to restore a server. This data is usually written to a tape or bootable CD. If a server fails, you roll in replacement hardware and boot from CD. The server can be quickly rebuilt and returned to operation. One of the challenges with this approach involves updating and managing the resulting CDs and images.

Hydro uses Symantec® Live State Recovery (LSR) to create images, but instead of writing the resulting image to CD Hydro writes the image to a separate disk volume (on the same server or a separate server). The IT team then uses NetApp SnapVault software to move this image file up to a NearStore system at the data center so the images can be centrally managed.

If a server fails and the system disk must be restored, Hydro has four options:

- Restore the image locally from the local data drive if available.
- Restore the image from a NearStore at a regional data center directly to the failed server over a LAN connection.
- Create a new system disk with the image files from the NearStore and send it to the site where the failed server is located.
- Create a CD from image files on the NearStore and send it to the site where the failed server is located. Then simply restore from the CDs using LSR.

Get the details. Read the reports:

- [SnapVault Deployment and Configuration](#)
- [Enabling Rapid Recovery with SnapVault](#)
- [Leveraging SnapVault for Heterogeneous Environments](#)

Not All Snapshots Are Created Equal

After interviewing storage managers and database storage administrators at 15 to 20 companies, Mercer Management Consulting found that NetApp environments require less disk space than EMC for the same size database.

The cause: different approaches to Snapshot™ equivalent functionality.

Get the details. [Read the full report.](#)

#4: Empower End Users

After finding that about 80% of restores were single files, Hydro identified end-user restores as one of its remote office data management goals. Using the .snapshot directory from Windows Explorer, Hydro end users are now able to browse through backups using fully GUI-based standard tools to locate and restore single files without IT intervention.

#5: Don't Lose Sight of the Big Picture

Implementing centralized backup was only one element of a larger, more ambitious master plan. After data backups and BMR images are copied from remote offices to NetApp storage at regional data centers, these systems are then mirrored to a recently established disaster recovery facility where Hydro has implemented a lights-out type operation. Data for critical regional and corporate applications is mirrored to the same facility.

The result is that all critical data is available at the disaster recovery facility. If any remote office goes down, its critical operations can be restarted at the DR facility using current data. The same goes for regional and corporate operations. Large tape libraries are available at the disaster recovery site for any data that needs to be archived to tape.

Hydro has successfully designed a global and highly scalable environment that enables it to protect data at many different levels.

Comment on this article



DAVID BROWN AND GREGG FERGUSON

Engineering Support Manager and Laboratory Administrator, NetApp Engineering

David and Gregg helped design, build, and manage the NetApp Kilo Client. This project is one of several firsts for them. As an engineer at Lawrence Livermore labs, David was a founding member of the CIAC, Computer Incident Advisory Council. David also holds a patent for 10 Gbps mobile data center design and helped start the Engineering Support program at the Research Triangle Park (RTP), North Carolina NetApp office. Gregg, a 22-year industry veteran, held a variety of engineering, systems administrator and IT manager positions before joining NetApp in 2000. After opening the original NetApp RTP facility, Gregg spent over four years in the field as a systems engineer before returning to the NetApp Engineering IT team.

The NetApp Kilo-Client

The World's Largest iSCSI-Based Diskless Enterprise Server Farm

Over the past few years, a variety of new technologies have emerged that promise to change the way data centers are designed and managed. Among these are large computing grids with hundreds or thousands of hosts;

IP SANs running the iSCSI protocol for inexpensive, consolidated block storage; the ability to boot diskless servers over either FC or IP SANs; and sophisticated cloning technologies for advanced storage provisioning.


In early 2006, NetApp Engineering combined these technologies to build what we believe is the world's largest iSCSI-based diskless server farm. All server nodes boot from NetApp storage using iSCSI and can be rapidly rebooted to run Linux®, Windows®, or other operating systems. The system leverages NetApp FlexClone™ technology to rapidly create system images without making full physical copies of those images. The baseline image is shared among all of the nodes of the farm; only host-specific "personalization" needs to be added to the core image for each of the provisioned servers. This unique approach affords near-instantaneous image provisioning with near-zero footprint (only the blocks of the images that differ need to be added to the storage system, which keeps track of the individualized images).

The Kilo-Client project began in 2005, when the NetApp Engineering IT team was challenged to create a testing environment capable of meeting or exceeding the most extreme conditions NetApp products encounter in the field.

The NetApp team designed a 1,120-blade server farm and originally planned to use diskfull server blades in which each blade booted from a local disk. However, as the project progressed, it became clear that the time and administrative overhead required to copy a boot image to over a thousand local disks would result in more time spent configuring and managing the cluster than running actual tests.

Instead, the team chose a diskless approach. Although they weren't aware of anyone using iSCSI-based SAN boot for 1,000+ nodes, earlier that year Gregg

The Kilo-Client Architecture Supports Multiple OS



RELATED INFORMATION

- [Future of IP Storage Q&A](#)
- [NetApp Technical Report Library](#)
- [NOW™ Customer Site](#)
(password required)

BayLISA Presentation: The NetApp Kilo-Client

At the [BayLISA](#) March general meeting, David and Gregg described the server farm in detail and explained how particular technologies were chosen.

Watch BayLISA videos of the Kilo-Client presentation:

- [Project Overview](#) (10 mins.)
- [Architecture](#) (35 mins.)
- [Q&A Session](#) (20 mins.)


Read the [presentation slide deck](#).

Kilo-Client Overview

- **1,120 Blades**
 - QLogic iSCSI SAN initiator hardware in each blade
- **Network Infrastructure**
 - Five Cisco 4948 switches (boot infrastructure)
 - Five Cisco 7608 switches (test infrastructure)

had helped deploy a 250-node diskless environment based on FC SAN technology. The team also had access to internal expertise managing a large grid environment (a 400+-node diskfull testbed already in use at the NetApp Pittsburgh Technology Center).

The result is the NetApp Kilo-Client, a 1,120-node diskless server farm with all nodes booting from a back-end network running iSCSI. This architecture packs massive performance into a footprint of just over 235 square feet and offers many unique advantages:



- **Provisioning time reduced 10x to 500x.** A set of "golden" boot images is created (as iSCSI LUNs) for each operating system and application stack required in the server farm. NetApp FlexClone technology is used to nearly instantaneously create space-efficient copies of those golden images, which are then "personalized" (individual server information such as IP address, hostname, etc.) and used to network boot the blades—potentially across the entire 1,120 host system.
- **Space-efficient copies.** The cloned images (LUNs) only consume additional disk space as boot images change, providing over 1000-fold efficiency of capacity.
- **Enables preemption and resumption of system state.** Tests can be preempted by halting a server and creating a space-efficient, derived clone of that system (using FlexClone). Test configurations can be preserved or shared with other users and re-run months or years later, perhaps even on an alternative (albeit same architecture) system.
- **Highly efficient utilization.** Instead of having equipment sitting idle during an extended setup, saving LUNs allows the team to quickly reinstate previous tests. Most of the configuration work has been automated using scripts, so the team can switch from one operating environment to another in little more than the time it takes to reboot the blades.
- **Scalable and flexible.** By partitioning the grid into smaller subsets, a variety of smaller tests can be run simultaneously (with a different OS for each test as needed), or the team can harness the full power of all 1,120 nodes for the most massive workload tests.

For details, see the BayLISA presentation [The Kilo-Client Host Swarm](#).

- 10 Gigabit Ethernet
- **iSCSI Boot Storage**
 - Five NetApp FAS980s (active boot images; one per 224 blades)
 - NetApp NearStore® (master boot images)
- **Boot Image Management**
 - NetApp FlexClone technology
 - NetApp SnapMirror® software (replicate LUNs from secondary to primary storage)
 - Homegrown scripts

[See the slide deck.](#)



TONY POWELL

Storage Administrator

Tony manages CIFS, NFS, and iSCSI storage for a privately held American company that employs over 2,000 people worldwide and has annual revenues of nearly half a billion dollars. One of Tony's primary responsibilities involves ensuring 100% availability for a 2.5TB SAP environment. Originally trained as a software engineer, Tony's previous experience includes working with storage at CSC and Nortel.

March Tool of the Month: Simulated I/O Tool (SIO)

Every month Tech OnTap showcases a free tool that just might make your life a little easier. Recommend a tool and get a free NetApp cycling jersey.

Author: Network Appliance Engineering

What it is: SIO is a client-side workload generator that works with any target. It generates I/O load and does basic statistics to see how any type of storage performs under certain conditions.

How it works: SIO enables I/O performance testing without having to create large application structures (such as databases). For example, SIO can "approximate" a TPC-C-like workload by specifying a 2 to 1 read-write ratio of 4kb transfer sizes with the appropriate number of threads. A recent "Ask the Sys Admin" column details how to install and run SIO.

What Tony likes about SIO: This tool gives me a much better indication of database performance than just running scripts or doing a DD test in UNIX®. Unlike DD, for example, SIO uses multiple threads and allows for more realistic stress testing. And, since it's a command-line tool, it can be easily scripted.

How Tony uses SIO: I use it quite a bit to tune disks and applications. I first started using SIO about a year ago when we were doing a big SAP migration from HP to AIX. I had written a script that would create and read files, but it wasn't getting good randomization. Using SIO helped me make decisions about how to lay out the file system to get the best performance. Not having to do additional tuning halfway through the migration was a major benefit—although it was originally supposed to take 45 hours, we completed the whole project in about 20 hours.

I also used SIO to test performance running our Oracle® Databases on iSCSI in a UNIX environment. Performance was pretty good, so we ended up putting our database files on NFS using NAS and putting the redo logs and archival on iSCSI.

Caveats

This tool can help eliminate various factors, but no tool can simulate all possible types of conditions.

SIO Download (Includes Full Requirements)

(Password required, customers only.)

RELATED INFORMATION

- [Using SIO for Performance Testing](#)
- [Dec Tool: ToasterView](#)
- [Feb Tool: Data ONTAP Simulator](#)
- [NetApp Technical Report Library](#)
- [NOW Customer Site](#)
- [NOW ToolChest](#)
(password required)

Sample SIO Screen

[illegible]