

» Tech OnTap Home

April 2006

## HIGHLIGHTS

- » One Customer's DR Tradeoffs
- » The Secret to Smarter Backups
- » April Tool: Nagios
- » Storage Resiliency Best Practices
- » Special Offer:  
FREE Book on Storage Networks

## Fueling the Enterprise Grid: Pursuing the Vision of Enterprise IT

Bruce Moxon, Strategic Technology, NetApp

Server virtualization is increasingly popular but few companies have achieved true application virtualization. Here's why, plus three trends impacting the future of enterprise grids.

[More »](#)



## BLOGGING WITH DAVE

Dave Hitz, NetApp Founder and EVP

"My math says that with a four-disk RAID array using 400GB ATA drives you will lose data in about 10% of RAID reconstructs."

» [Dave's Blog](#)

## DRILL DOWN

- » **A Storage Networking Appliance**  
An intro to the WAFL® file system, Snapshot™ technology, consistency points, NVRAM, and more.
- » **Storage Resiliency Best Practices**  
How to create tiered storage and map applications by tier (26 pages).
- » **Oracle® 10g™ R2 Disaster Recovery**  
Guidelines, commands, and scripts for SnapMirror® Async and Sync (23 pages).



## ADMIN RESOURCES

- » **Tool of the Month: Nagios**  
"We use Nagios to monitor a FAS250 and over 170 other systems and services."

## TIPS FROM THE TRENCHES

### Can YOU Make a Business Case for Storage Networks?

Bill Williams, IT Manager,  
Cisco Enterprise Storage Services

Test your business acumen and get a free Cisco Press guide to showing management the cost advantages of networked storage over DAS (includes best practices for storage migration, TCO calculator, and more).



[More »](#)

### Three Tradeoffs That Impact Disaster Recovery Planning

John Fullbright, Professional Services Consultant, NetApp

Ensuring uninterrupted access for 30,000 Exchange users isn't trivial. See how one company set recovery point and recovery time objectives as part of a multisite DR plan (plus tips for maximizing network bandwidth).



[More »](#)

## ENGINEERING TALK

### When Brute Force Isn't Enough: Using Information to Optimize Backup

Stephen Manley, Technical Director, NetApp

The backup industry suffers from high rates of recovery (and administrator heart) failures. A quick look at how evolving technologies improve backup efficiency by understanding the data and applications they protect.

[More »](#)



**BRUCE MOXON**

Senior Director of Strategic Technology and Grid Guru, NetApp

Bruce Moxon works with enterprise customers deploying grid computing solutions. He brings more than 20 years of experience in scale-out computing architectures for both scientific and commercial applications and writes, speaks, and teaches extensively on the continuing evolution of grid computing. Bruce has architected and developed solutions for a number of high-throughput computing environments, including Perlegen Sciences' SNP discovery system, Bank of America's CRM and analytics systems, and NASA's Earth Observing System.

## Fueling the Enterprise Grid: Pursuing the Vision of Enterprise IT

*By Bruce Moxon*

In an era of flat IT budgets, tight power and space constraints, oft underutilized resources, and increasing competitive pressures that require more agile business practices, virtualization has become the mantra of modern IT architects.

The ability to deploy commodity compute, network, and storage resources; to provision them on demand; and to efficiently manage and effectively share those resources is at the core of the grid—or utility—computing movement.

Over the last few years, we've seen significant growth in the deployment of server virtualization technologies such as VMware, Microsoft® Virtual Server, the open source Xen project, and others. These technologies provide a framework for application consolidation in a multiple-OS environment to more efficiently use shared server resources.

To date, these approaches have been applied successfully in quasi-static production and development environments. For example, production application servers and database servers running on disparate operating environments can be cohosted on a single physical server. And multiple virtual machines (running multiple OSES) can be made available on a shared basis throughout a development organization to reduce physical infrastructure costs.

These approaches have certainly led to increased utilization of server resources. In some cases, application migration capabilities have been leveraged to dynamically balance load across shared computing resources.

However, server virtualization alone cannot fully deliver on the core promise of grid computing—true application virtualization—in which applications can be rapidly provisioned, migrated, and even replicated to support a variety of production and development requirements.

### **"The missing ingredient in the application virtualization 'stack' is storage services."**

For the most part, today's applications maintain their "state" on persistent storage—in file systems or databases. These applications increasingly reside on networked storage (SAN and/or NAS architectures). Moreover, it is becoming increasingly common, especially in scale-out server environments, to also maintain operating system and application software "stacks" on shared storage (network boot configurations).

In this environment, the acts of provisioning new applications, migrating them from one

### RELATED INFORMATION

- [GRIDtoday](#)
- [Enterprise Grid Alliance](#)
- [RedHat "Integrated Virtualization"](#)
- [Oracle IT's Implementation of Grid](#)
- [NetApp iSCSI-Based Diskless Server Farm](#)
- [Distributed Storage for Scalable Grid Infrastructures \(PDF\)](#)

### NetApp Vision for the Grid

This five-minute interview with NetApp senior VP of Products Suresh Vasudevan summarizes our vision and development roadmap for storage grid architectures.

[Listen to the podcast \(MP3 Audio\)](#)

physical server to another, and replicating them for development, test, and QA purposes all require support of shared storage services.

We see a variety of trends impacting the move toward a fully realized application virtualization model including:

- **Nondisruptive operation.** With increasing application consolidation in the server and storage layers, the impact of a system outage is no longer limited to a single application. There will be increasing pressure for true 24x7 operations and nondisruptive upgrades. *Storage Magazine* does a great job of exploring the progress—and work still to be done—in this area in its [March 2006 issue](#) (registration is required, but it's worth it).
- **Global namespace (GNS) solutions and file-level migration in tiered storage environments.** GNS solutions provide a uniform view of files across many physical file servers. Additionally, they support transparent data migration—often at the file level—to facilitate effective storage tiering. These capabilities exist in clustered storage systems such as the recently announced NetApp Data ONTAP® GX, virtual file management solutions for the Microsoft® Distributed File System (DFS), and inband NAS aggregation switches (think of them as "file routers") from the likes of NetApp partners Acopia Networks and Neopath Networks.
- **Diskless boot.** An increasing number of organizations are exploring the use of file server-based system "images" and diskless boot as a way to rapidly provision and switch between different software stacks, especially for dev/test environments. With iSCSI software initiators now readily available for Windows® and Linux® environments and emerging support for software initiator-based diskless boot using standard Ethernet cards instead of costly iSCSI HBAs, I fully expect this trend will accelerate.

(A real-world example of a large-scale diskless boot environment is the 1,000-node [NetApp Kilo-Client project](#), which was profiled in last month's issue of Tech OnTap. Software-initiator-based diskless boot capabilities were demonstrated last August by [Emboot](#) and [String Bean Software](#) using the latter's iSCSI target, WinTarget™. Microsoft [acquired](#) the WinTarget technology in March of 2006, and NetApp has recently announced support for [iSCSI software boot-from-SAN capabilities](#) from Microsoft.)

Over the next few issues, I'll take a deeper dive into the topic of enterprise grid computing and explore ways in which the companies I talk with are leveraging the unique capabilities of Data ONTAP. Upcoming topics will include:

- Use of MultiStore® to develop a dynamic provisioning and load balancing approach for a storage service provisioning model
- Application image cloning and diskless boot for rapid application provisioning in rack-mount server environments
- Database cloning for improved development, test, and QA efficiency and out-of-band reporting and ETL activities
- Whole application migration strategies (server and storage) for application right-sizing

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

In 2006 NetApp Engineering built a 1,120-node diskless server farm. Unique advantages of this architecture include:

- Provisioning time cut 10x to 500x
- Space-efficient copies
- Enables preemption and resumption of system state
- Highly efficient utilization
- Scalable and flexible

Learn more about the Kilo-Client:


[Introduction to the Kilo-Client](#)  
[Kilo-Client Presentation Deck](#) (PDF)

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

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

Oracle was one of the first companies in the world to envision and adopt a grid architecture.

Less than 12 months later the Oracle IT team had cut its overall IT spend by roughly 30%. Today, Oracle can quickly provision and decommission systems from one purpose to another and rapidly repurpose assets.



Learn more about the Oracle implementation. Read the [Q&A with Bill Weils](#), senior director of the Oracle Global IT Command Center.



**BILL WILLIAMS**

Manager, IT, Cisco Enterprise Storage Services

As a program manager for the Network Storage Virtual Team, Bill guided the implementation of SAN storage and infrastructure in Cisco data centers around the globe. In 2003, he began a consolidation project that produced a direct savings of \$1.5M and lowered Cisco's immediate storage costs by more than 20%. In his current role, Bill focuses on total cost of ownership, metrics, and operational efficiencies. He recently authored a Cisco Press book to help project managers and technical architects gain a broader knowledge of financial methods and storage networking technologies.

# Can YOU Make a Business Case for Storage Networks?

By Bill Williams

Sure, storage networking technologies can help your organization decrease capital expenditures and increase efficiencies. But how do you justify the investment up front, when it's critical that you convince management and before you commit significant resources?

The first step in building a business case for storage involves determining the current total cost of ownership (TCO). With that kernel of financial truth in hand, you can step through financial metrics such as payback, return on investment (ROI), and net present value (NPV) to better quantify benefits and benchmark the impact of new projects.

As a Tech OnTap member, you're invited to receive a FREE copy of *The Business Case for Storage Networks*. This book addresses the problems of storage growth and increased consumption, the role of the IT department as a cost center, and how SAN technologies can help save money in the long run. It also provides tools to help you and your management team make informed decisions about your storage networking investments.

Already know everything you need to about building a business case? Here's a little challenge to test your mettle.



## RELATED INFORMATION

- ["The Business Case for Storage Networks" \(Cisco Press, FREE Online Access for 14 Days\)](#)
- [Business Case Methodology \(Bill Williams presentation\)](#)
- [How to Talk to a CIO \(If You Must\)](#)

## Cheatsheet

In case you're a little rusty or were distracted by the fancies of youth when first you learned the methods behind these metrics, here are some quick tips:

Payback Period = Initial Investment/Annual Income

ROI = (Returns-Investments/Investments)

$$NPV = CF_0 + \frac{CF_1}{1+Rate} + \frac{CF_2}{(1+Rate)^2} + \frac{CF_3}{(1+Rate)^3}$$

CF is cash inflow and Rate is the cost of capital (Bill's examples typically use 12%)

### Topitoff Draught Supplies

**Just the Facts**

Topitoff Draught Supplies operates two data centers, one each in Stuttgart, Germany, and Golden, Colorado. Topitoff employs 20 full-time equivalent (FTE) staff with a burdened cost per administrator of \$100,000.

Topitoff is considering consolidating storage across its existing data centers. This would involve purchasing \$1.5M worth of new storage and storage networking devices and would reduce headcount costs by 35%.

**Topitoff Needs to Know**

1. The annual dollar value of headcount savings associated with the consolidation project
2. The ROI for the first year
3. The payback period for this investment

# Making a Business Case for Storage

## -- Results --

In a few minutes, Topitoff can determine that, solely based on a TCO perspective, this investment has the potential to offer some long-term benefits to the firm. While negative in the first year, the ROI becomes positive early in the third year. Depending on how Topitoff performs portfolio analysis, however, this project may be sacrificed to other initiatives that offer positive returns in the first year.

### 1. Annual headcount savings = \$700,000

With 20 FTEs at an annual cost of \$100,000 each, Topitoff is currently spending \$2,000,000 on headcount. The new solution is expected to reduce headcount costs 35% ( $\$2,000,000 \times .35 = \$700,000$ ).

### 2. First year ROI = - .53

ROI = (Returns-Investments/Investments)

Topitoff would need to invest \$1.5M initially and would immediately (for the sake of this simple example) see savings of \$700,000 annually. In Year 1, the ROI would be  $(\$700,000 - \$1,500,000)/\$1,500,000$ .

### 3. Payback period = 2.1 years

Payback Period =  
Initial Investment/Annual Income

In just over two years, the investment would pay for itself based on headcount savings:  $\$1,500,000/\$700,000 = 2.1$ .

*Need more in-depth information? Check out Bill's recent presentation on the [Business Case Methodology](#) for strategies, in-depth examples, and multiple real-world case studies. The first 100 people to respond will receive a FREE copy of Bill's book, [The Business Case for Storage Networks](#).*

**Get Your Free Copy of *The Business Case for Storage Networks***



## RELATED INFORMATION

- ["The Business Case for Storage Networks" \(Cisco Press, FREE Online Access for 14 Days\)](#)
- [Business Case Methodology \(Bill Williams presentation\)](#)
- [How to Talk to a CIO \(If You Must\)](#)

## The Business Case for Storage Networks

Bill's book includes:

- A comprehensive business case for storage network adoption and deployment
- How to use EVA, NPV, and ROI metrics to evaluate projects
- Procedures for measuring and tracking TCO before and after implementation
- Best practices for executing a storage migration strategy
- A TCO calculator and decision-making checklist

[Receive a FREE copy of The Business Case for Storage Networks.](#)



## JOHN FULLBRIGHT

Professional Services Consultant and Microsoft MVP, Network Appliance Global Services

John Fullbright is a NetApp Professional Services consultant and "resident expert" on Exchange. In April of 2006, John won a [Microsoft Most Valuable Professional \(MVP\)](#) Award recognizing his competency in Exchange Server technical communities. Prior to NetApp, John was a rapid response engineer in the Microsoft Global Solutions Support Center. As a member of the Microsoft Rapid On-Site Services team, John was regularly tapped to solve the challenges of premier Microsoft customers. He has also served as an intelligence analyst and IT professional for the U.S. Army.

# Three Key Trade-Offs That Impact Disaster Recovery Planning

By John Fullbright

Implementing a disaster recovery (DR) solution requires making tough choices. NetApp Global Services (NGS) brings a quantitative approach to solution design that has helped customers around the world understand the trade-offs associated with different approaches, make hard decisions, and implement solutions that efficiently address their specific priorities.

Recently, for example, a major American insurance firm began finding it increasingly difficult to complete tape backups in a 24-hour window. The firm engaged NetApp Global Services to help it design and implement disaster recovery for a Microsoft® Exchange environment that included:

- One major satellite operation approximately 1,000 miles from headquarters
- 30,000 Exchange seats (15,000 at each site)
- 14TB of Exchange data (7TB at each site)
- Fibre Channel SAN infrastructure
- OC3 connection between sites

This article highlights three significant trade-offs impacting the project and how this particular customer chose to proceed.

## Trade-Off 1: Is Your Goal High Availability, Ongoing Backup and Recovery, or Remote Recovery in the Event of a Site Disaster?

There are basically two approaches to online data protection: replication (mirroring) and disk-to-disk backup.

Replication creates an exact duplicate of your data set on another storage system, which may be local or at another site. This is the best solution if your goal is to ensure immediate failover for high availability or to remotely restart operations in the event of a site disaster. This doesn't take the place of backup, because if something disappears from one side of the mirror, it will be gone from the other side at the next replication cycle. Distance is a key factor: while asynchronous replication can be done over great distances, there are technical limitations associated with synchronously mirroring data between sites more than 80 kilometers apart.

Disk-to-disk backup and archival methods supplement or replace failure-prone tape systems with more robust secondary disk. Data is backed up across a network to a remote location and restored from there if a disaster occurs. Unlike replication, hosts can't attach to the secondary storage directly. Instead, you have a historical record of iterative changes to your data set, so you can recover data from any point in time.

In this case, the IT team decided to implement a combination of methods.

## RELATED INFORMATION

- [NetApp Disk-to-Disk Solutions](#)
- [NetApp Technology Overview](#)
- [NetApp Data Protection Strategies](#)
- [Tips for Remote Backup](#)
- [Data Protection Best Practices Portal](#)  
(password required)

## 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 that can be mounted read/write to recover from failure. If a 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. It retains all backup copies as they appeared at the time they were created on primary storage for a user-specified period of time. 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 copies an entire image at the level of a LUN inside a volume.

Get the details. Read the reports:

Implementing remote disaster recovery was a top priority. The company has two major sites located about 1,000 miles apart with relatively good network bandwidth connections, so the decision to use NetApp mirroring software to asynchronously replicate Exchange logfiles and databases between sites at specified intervals was relatively straightforward. The IT team also chose to address its existing tape backup issues by moving to a disk-to-disk-to-tape environment using NetApp SnapVault® software.

## Trade-Off 2: If the Math Doesn't Work, Can You Increase Network Bandwidth or Compromise on How Much Data You Can Lose?

To size your DR infrastructure, you have to determine the amount of data that is changed—and therefore that must be replicated or backed up—every day. Once you determine your change delta, the next step is to figure out the maximum amount of data you can afford to lose if a disaster occurs. Dividing your change delta by the replication interval enables you to estimate how much data has to be transferred in each interval. This is where trade-offs start to come in. If the math doesn't come out in your favor, you'll either have to increase network bandwidth or consider a longer recovery point objective (RPO)—and more potential data loss.

After reviewing requirements with various organizations, the team established a five-minute RPO. Due to the short cycle, we had to consider peak spikes when calculating the change delta for the logs. This was done by creating a data set using perfmon sampling of average writes per second for the process store.exe at a five minute interval. The team measured a change delta of about 200GB every 24 hours. With a replication interval of five minutes, this translates to about 700MB of data to be replicated every five minutes (200GB/day ÷ 288 replication cycles/day). Given other network traffic, peak traffic would have exceeded the available OC3 network (155 Mbps, or about 19MB/sec).

The IT team couldn't compromise its five-minute RPO objective or upgrade its network infrastructure. Instead, NGS pointed out that for any transactional application where writes are first committed to a log and then from the log to the database, there are two discrete components to the change delta. New data is written to a log file before it is committed to the database, so half of the change delta (100GB) is going to log files.

**"By replicating only the logs every five minutes, the firm was able to cut bandwidth requirements in half while still meeting its five-minute RPO objective."**

The Exchange database is replicated every four hours, and peak traffic does not exceed 13MB. (See sidebar for details.) This not only provides a high level of data protection for the log files, but also spreads the load out more evenly throughout the day, which helps reduce the impact on networks and primary storage.

The potential downside of this approach involves the recovery time associated with replaying these logs. After testing the process, NGS determined that replaying logs added only about five minutes to the recovery time. The benefits associated with spreading the use of bandwidth over time and cutting overall bandwidth requirements easily outweighed the impact of this approach.

**"Throttling spreads the I/O out over time and ensures that the capabilities of the network are never exceeded."**

NGS helped the customer further minimize the impact of replication and backup traffic by throttling the two NetApp products. Both SnapVault and SnapMirror® can be tuned so that they don't exceed a specified I/O rate. Not all DR applications support this capability, but it is a good idea to set a threshold where possible, so that unusual peaks in activity don't lead to unexpected results.

## Trade-Off 3: How Long Can You Afford to Be Down?

The next step is to determine how long it can take to return to operation if a disaster strikes. This is the recovery time objective (RTO). RPO is fairly straightforward, but RTO can be tricky because you have to consider all the steps that it's going to take to

- [SnapMirror Best Practices Guide](#)
- [Enabling Rapid Recovery with SnapVault](#)
- [Data Protection Strategies for NetApp Storage Systems](#)

## Drill Down: Calculating the Change Delta

In the scenario outlined in this article, NGS had to consider peak spikes when calculating the log change delta due to the short cycle.

After creating a data set with 302 samples using perfmon, NGS derived the average data transfer and added five standard deviations. The resulting spread includes 99.945% of all cases. The peak value for logs in this case was 5.36MB/sec each direction, or 10.72 MB/sec peak over the OC3. Add to that another 2.3MB/sec for database traffic, and the result was about 13MB/sec peak.

If the customer had not chosen a longer replication cycle for databases, NGS would have also had to account for peaks on database replication. The total would have exceeded the capacity of the OC3, resulting in failed replication cycles. Each time a replication cycle fails, the solution potentially could not meet the five-minute RPO.

The customer's requirement was to meet the RPO with 99.9% accuracy. By using an asymmetric asynchronous replication plan, NGS was able to statistically prove that the customer would be able to exceed this objective.

## Disaster Recovery Planning and Implementation

DR planning should start with a complete assessment of your existing infrastructure, including vulnerabilities and risks.

Use your application requirements to create a comprehensive plan for disaster recovery. Once you understand your requirements and have a plan, you can begin to integrate your existing systems with new DR software and infrastructure.

If you need help, NetApp Global Services has service offerings to provide professional, on-site assistance to manage any or all of the phases of DR deployment.

Learn more about [NGS disaster recovery services](#).

return to operation.

To establish a realistic RTO, NGS worked with the firm's IT team early in the process to document everything involved with getting Exchange back online. This includes:


- Breaking the replication links and making the replicated LUNs read/write
- Connecting all the LUNs to the hosts at the sister site
- Starting Exchange services
- Replaying the logs

After evaluating all of the steps involved in making the necessary infrastructure changes, starting Exchange, and replaying hours of log files, the customer established a four-hour RTO that dictated that the replication interval for the Exchange databases be no more than four hours.

### The End Result: A Multitiered Storage Architecture

Working with NGS, the IT team was able to stay within its original budget while building an infrastructure that provides multiple levels of protection and allows recovery from both routine failures and a sitewide disaster:

- Redundant servers protect Exchange against server failure.
- Up to 30 Snapshot™ copies\* of each volume (five days' worth) and 48 copies of the logs (four hours' worth) kept on primary storage for quick recoveries.
- Disk-to-disk backup using SnapVault offloads disruptive tape backup from primary servers and storage. Data can then be backed up to tape without affecting primary storage or the operation of Exchange.
- Remote replication using SnapMirror protects against sitewide disasters.



[Click to Enlarge](#)

\* There is no performance impact associated with keeping more than 250 Snapshot copies on NetApp storage. This is not necessarily true of solutions based on copy-on-write.

Other Tech OnTap articles by John Fullbright:

- [Solving the Top Five Exchange Issues Experienced by NetApp Customers](#)



## STEPHEN MANLEY

Technical Director of Data Protection, NetApp

Stephen joined NetApp Engineering in 1997 after graduating from Harvard University with a BA in computer science. He has helped define and develop NetApp's tape and disk-to-disk backup, mirroring, and compliance technologies. Stephen has traveled the world over, one data center at a time, working with customers and backup partners to provide open, reliable, and simple data protection solutions.

# When Brute Force Isn't Enough: Using Information to Optimize Backup

By Stephen Manley

We've all heard Ben Franklin's aphorism, "The definition of insanity is doing the same thing over and over and expecting different results." In the face of expanding data sets, more stringent SLAs, and escalating costs, many companies continue to adhere to traditional backup methods. What might Ben Franklin have to say about that approach?

Simply put, the brute force approach to backup no longer works. In the past, a larger library, faster tape drive, upgraded network, or new backup server kept backups running. Over the last five years, however, disk capacities have grown by 20x, but with a scant fourfold performance improvement for backup. Disks are why backup times have increased by 500%. Throwing more tape infrastructure at the backup problem will no longer solve it.

Fortunately, by deploying backup solutions that leverage available information that they have long ignored, you can meet your business needs while reducing costs. This article highlights three areas where knowing something about the data being backed up can yield major benefits.

## Benefit 1: Reduced Backup Storage Consumption

Every week around the world, companies replay the final scene of *Raiders of the Lost Ark*: stashing enterprise secrets in a warehouse, never to be seen again. With the low rates of data change on primary storage, new backups usually have over 95% redundancy with prior backups. While retaining mountains of redundant backups was an expensive but necessary "best practice" for tape, the industry cannot and should not accept this amount of waste on disk-to-disk backups.

Single-instance storage (aka data deduplication) technologies eliminate unwanted redundant copies across multiple full and incremental backups. Some companies have been able to store one year of full backups in the space that previously would have been consumed in two to three weeks! By exploiting the significant commonality across backups, disk-to-disk products such as Data Domain and Diligent VTL can dramatically increase the storage utilization of disk-based backups.

By understanding the contents of the backups, however, companies can create a more scalable, space-efficient solution. Have you ever played the game Concentration? Most single-instance storage technologies work like a 10TB game of Concentration, trying to find random matches across all the data. Not surprisingly, such brute-force algorithms are expensive and slow.

On the other hand, if the backup system recognizes another full backup that contains the file *connor.doc*, it can begin its search for duplicate elimination with the previous copies of *connor.doc*. After eliminating such obvious matches, the brute-force calculation takes place on a much smaller set of data. Products such as Avamar Axion and Symantec® NetBackup™ with NetApp SnapVault® can dramatically improve the



## RELATED INFORMATION

- [Is NetApp Storage Really Simpler?](#)
- [Enabling Rapid Recovery with SnapVault](#)
- [The Private Lives of Disk Drives: How NetApp Protects Against 5 Dirty Secrets](#)
- [Tips for Improving Remote Backup](#)
- [Data Protection Best Practices Portal](#)  
(password required)

## A Storage Networking Appliance

In the early 1990s, Network Appliance revolutionized storage networking with a simple architecture that relied on NVRAM, integrated RAID, consistency points, and a unique file system to do things that the file servers of the time could not.

This technology is still the basis of every product that NetApp offers, including solutions for optimized backup.

The original technical report about the NetApp storage architecture was recently updated and expanded by NetApp co-founder Dave Hitz and Akshay Bhargava to describe today's NetApp storage appliance, including:

- High-level architecture
- The WAFL® file system
- Snapshot™ copies
- Consistency points and NVRAM
- FlexVol™ and FlexClone™ technology
- RAID and RAID-DP™

If you only read one paper about NetApp technology, read [A Storage Networking Appliance](#).

efficiency of space-optimized backups.

Understanding the redundancy inherent in traditional backups helps reduce disk-to-disk backup costs. Utilizing the information within the backup stream enables efficient, scalable, and space-optimized disk-to-disk backups.

### Benefit 2: Smaller Backup Windows

With disk capacity growing faster than backup performance, companies are creating full backups less frequently. As SLAs become more aggressive, companies need to create backups more often. As a result, incremental and differential backups have become increasingly business critical. Unfortunately, their performance often fails to meet expectations.

If you don't know anything about a data set, figuring out what has changed is like looking for a bunch of needles in a haystack. You're simply looking through every file or block to ensure that you've protected all the new data. Searching for changed data on a system with tens or hundreds of millions of files can be just as agonizing as running a full backup.



How can you avoid searching the entire primary system for changed data? If you already know what's happening on the host system—what parts of what files are being modified and where—you can avoid a brute-force search. Eliminating the search for changed data is at the heart of the "continuous data protection" (CDP) movement. Offerings such as Microsoft® Data Protection Manager, Oracle® FlashBack, Syncsort BackupExpress, Symantec BackupExec, and NetApp SnapVault rapidly find and protect the changed data within a system. Whether they integrate deeply with the application, the system, or snapshots, these solutions use knowledge of the data they protect to provide lightweight, fine-grained backups.

### Benefit 3: Reliable, Recoverable Backups

There's nothing more agonizing than the desperation of attempting a mission-critical restore, only to be stymied by a misplaced tape, corrupt backup image, or critical software bug. Unsurprisingly, the backup industry suffers from high rates of both recovery and administrator heart failures. Solutions that truly understand the backed-up data, however, can calm your recovery anxiety.

Unfortunately, most new technologies actually exacerbate the reliability challenges. Consolidating hundreds of backups by eliminating redundant blocks seems like a fantastic TCO option ... until the first failure. When your consolidated, space-optimized backup system suffers a spectacular hardware or software failure, you lose last year's backups. It's even more disturbing that the cutting-edge system might incorrectly eliminate data that was not actually a duplicate. In either case, a warehouse of hundreds of redundant tape backups suddenly feels very comforting.



Block-level incremental backups help meet recovery point objectives, but increase the likelihood of recovery failure. One lost, corrupted, or incorrect incremental backup can compromise a year's backups. Furthermore, CDP solutions that merely save off streams of changed block data have no safety net. Not only can a corruption on the primary application or system propagate to the backups, but the backup application is more likely to introduce undetected errors if it does not understand the data that it protects.

Never forget "backups are worthless; restores are priceless." Some products have used their understanding of the primary data to increase the reliability of the backups. Products such as Oracle FlashBack are deeply integrated into the application, so that the backup data can be easily validated. The incremental backup data is not just a stream of blocks to be stored; to Oracle, they are a well-understood set of database operations and data. Still, even with Oracle, you cannot be absolutely sure of data recovery until you try.

If only recovery can validate a backup, why not recover every backup as it runs? That way, each completed backup is 100% validated. In fact, why not make those backups securely available to users, so there is no question about their correctness?

Microsoft DPM and NetApp SnapVault, including the Symantec NetBackup integration, offer this high standard of reliability. Backups transition from unreliable, proprietary

## Not All Snapshots Are Created Equal

Mercer Management Consulting found that NetApp environments require significantly less disk space than HP EVA, EMC CLARiiON, or EMC Symmetrix for the same size database.

The cause: different approaches to snapshot equivalent functionality.

Get the details. [Read the full report.](#) (See Page 6 for a comparison of overhead capacity and disk capacity requirements)

opaque images into online, user- and application-accessible data. Online backups then enable tools such as the NetApp SnapManager® suite and Syncsort BackupExpress to leverage their application integration and further validate the contents of application backups.

## Conclusion

The backup challenges facing your company may seem insurmountable. You struggle to meet your SLAs, and every day just brings more data and more requirements. You cannot afford to completely upgrade your backup infrastructure. You feel like you should adopt disk-to-disk backup, but you're not sure what solution to choose. The challenges can feel overwhelming. Thankfully, the situation is not as bleak as it seems.

The backup industry has enjoyed unprecedented innovation over the past five years. Disk-to-disk backup, single-instance storage, and CDP are changing data protection. These advancements have all been built on utilizing information. When brute-force techniques ceased to scale, the industry began to build more intelligent solutions.

As you consider the technologies at hand, consider three important questions. First, will it efficiently eliminate costly data redundancy? Second, does it reduce my backup window? Third, and most importantly, can I *trust* it? It's time to start investigating how you can make your backup solutions work for you. After all, Ben Franklin also pointed out that "an investment in knowledge always pays the best interest."

## Hear Stephen Manley speak live at Symantec Vision 2006 in May:

### [From Vision to Reality: NetBackup and Disk-to-Disk Backup in the Real World](#)

One year ago, Network Appliance and Symantec launched a joint initiative to simplify backup management, optimize data recovery, and minimize the cost of data protection. One year later, listen to our experiences with customers deploying all these different solutions. You'll hear about the laughter, the frequent flier miles, the successes, and the challenges of deploying disk-based backups. Come share in the lessons we've learned about protecting the world's data—one disk-based backup at a time.

## Ready to Learn about NetApp Solutions for Disk-to-Disk Backup?

The NetApp D2D Backup Center includes a demo showing how you can optimize backup and recovery for any storage environment.



Visit the [NetApp D2D Backup Center](#) for resources including a SnapVault demo.



## DANNY DU PREE

Assistant IT Manager, Intersafe Groeneveld

For over 55 years, Intersafe Groeneveld has been supplying a complete range of high-grade personal protective equipment to shield workers from on-the-job hazards. As assistant IT manager, Danny considers Nagios to be part of his "personal protective gear." He routinely interfaces with both end users and suppliers to ensure that service levels meet requirements and is closely involved with all major application and hardware deployments.

# April Tool of the Month: Nagios

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 jacket.

**Author:** Nagios was primarily developed by Ethan Galstad. Many other people in the open source community have contributed to the project over the years.

**What it is:** Nagios is an open source host, service, and network monitoring program.

**How it works:** A monitoring daemon runs periodic checks on the hosts and services using external agents, or "plugins," which return status to the main application. Notifications can be sent via e-mail, instant message, SMS, and so on. Current status information, logs, and reports are easily accessible via a Web browser. Nagios runs on Linux® and most versions of UNIX®. Since the tool is open source, it is possible to get it running on almost any platform.

**Why Danny thinks it's cool:** Nagios is generally simple to use, and it has all the components we were looking for, including monitoring of servers and services. Because the tool is open source, the only cost is the time to install it. Lots of information and user forums are available at [www.nagios.org](http://www.nagios.org).

**How Intersafe uses it:** We use Nagios to monitor our Windows NT®, Windows® 2000, and Windows 2003 systems, plus our switches, routers, network printers, and the NetApp FAS250 system that our SQL Server databases access via iSCSI. It helps us ensure key systems are alive and monitor critical services such as DNS, SQL Server, Oracle®, print spoolers, and so on. In total, we monitor 52 systems and 121 services using standard Nagios agents plus a few special ones we created in PERL.

The most important thing we monitor with Nagios is our Citrix environment and its subsystems. Our offices use Citrix as the main application platform, so we need to know ASAP when things go wrong. Nagios makes it easy to drill down to find the source of problems. Without it, we would have to check a lot of possible causes manually.

We use the following tools to complement the capabilities of Nagios:

- [Ping Plotter](#) to pinpoint network problems
- [Vncviewer](#) to display and interact with the user interface of a remote computer
- [Bmail](#) for a command-line SMTP mail sender (used from within Windows NT scripts)

**Caveats:** The Web-based configuration tools are difficult to use. We do it with the help of a text editor and after changing have to reload all the config files.

*Ed. Note: By the way, the NetApp IT team uses Nagios to monitor about 300 UNIX, Linux and NT servers. A basic set of services is monitored on each system including disk availability and capacity, CPU and memory utilization, etc. Different thresholds for notification can be set for different servers. Specific servers have additional monitoring according to their functions. For instance, on mail servers such things as mail queues, SMTP ports, DNS, and NIS are continuously monitored to ensure smooth operation.*

## RELATED INFORMATION

- [Official Nagios Web Site](#)
- [Mar Tool: SIO](#)
- [Feb Tool: Data ONTAP Simulator](#)
- [Jan Tool: ToasterView](#)
- [NOW Customer Site](#)
- [NOW ToolChest](#) (password required)

## Sample Nagios Screens: Status Overview and Service Problems Screens



[More Nagios Screenshots.](#)