

✦ [Tech OnTap Home](#)

May 2006

HIGHLIGHTS

- ✦ [A User Perspective on 7G](#)
- ✦ [Improved Database Test/Dev](#)
- ✦ [A Dilemma: Faster Tape or D2D?](#)
- ✦ [Featured Tool: MRTG for Filers](#)
- ✦ [Win a Pass to a Gartner Event](#)

Switching from Reactive to Proactive Management

Dave Kresse, Storage Management GM, NetApp

NetApp's strategy for helping you do less and accomplish more.



[More](#) ✦

BLOGGING WITH DAVE

Dave Hitz, NetApp Founder and EVP

"Sometimes I just feel like bragging: NetApp made the Fortune 1000 list!"

✦ [Dave's Blog](#)

DRILL DOWN

- ✦ [Data Center Proven Launch Page](#)
An intro to the FAS6000 series and NetApp 4Gb FC SAN solutions plus a data center cost comparison kit.
- ✦ [Virtual Tape Library Webcast](#)
How one IT team cut backup times 83%.
- ✦ [Congrats to Recent Winners!](#)
Renault F1 Team Gloves: Adrian Barnes, IT Manager, Alcan Smelting & Power UK
Renault F1 Team Helmet: Michel Khan, IT Director, Specsavers Optical Group

ADMIN RESOURCES

- ✦ [Featured Tool: MRTG for Filers](#)
This open source monitoring tool provides long-term trending data and enables proactive capacity planning.

TIPS FROM THE TRENCHES

A Customer Dilemma: Faster Tape or Disk-to-Disk Backup?

Chris Kim, Professional Services Consultant, NetApp


How one IT team developed a plan for eliminating 3-5 day backups. Includes tests that help identify bottlenecks and tips for calculating change deltas.

[More](#) ✦

Reducing Time-to-Deployment: Improving Agility for Database Test/Dev

Bruce Moxon, Sr. Director, Strategic Technology, NetApp

New technologies streamline database replication, enabling you to create working copies of very large databases in minutes using very little incremental storage.



[More](#) ✦

[Demo](#) ✦

CUSTOMER PERSPECTIVE

The Versatile Storage Platform

Ben Rockwood, Systems Administrator, Homestead.com

Tired of specialized storage systems that eventually become isolated islands? A user perspective on why nothing beats NetApp when you need flexibility plus volume, aggregate, and LUN creation examples.

[More](#) ✦



DAVE KRESSE

General Manager, Storage Management and Integration, NetApp

During his three-year tenure at NetApp, Kresse has talked with hundreds of IT teams and visited customer sites around the world to learn firsthand what life is like for storage administrators. Kresse and his team are devoted to helping organizations derive true value from their use of NetApp manageability products. Akin to Microsoft's Bill Gates, his vision is simple: that such rich NetApp functionality becomes an essential part of every NetApp customer experience.

From Panic to Power: Making the Switch from Reactive to Proactive Data Management

Q&A with Dave Kresse

Q: How does the NetApp approach toward management differ from that of other vendors?

Dave: The most common feedback we get from storage administrators is that they want to spend their time on value-added tasks such as utilization and performance that directly impact their business. However, many find that they spend an increasing portion of every day operating as a "help desk" responding to one-off requests from system administrators, DBAs, and application administrators. This makes their job very complex, and they are often unable to get to the higher value activities because they are stuck in fire-fighting mode. This problem only gets worse in the face of dramatic data growth.

Traditional approaches to management accept this situation and focus on managing complexity at the storage administrator level. Unfortunately, the complexity of today's storage environments extends well past the storage layer into the server and application layers, which are the domain of DBAs and systems administrators.

NetApp has focused on simplifying the problem using an approach we call integrated data management (IDM). This approach enables systems administrators and DBAs to manage their own data and self-serve basic requests. The storage administrator still has ultimate control of the process through the ability to govern what the other data owners are allowed to do. Focusing on building more and more tools solely designed for storage administrators is a losing battle, but you can remove the one-off requests barraging these folks by automating the process across the stack. In that, there's great potential for productivity gain.

Q: If you automate a lot of the work that storage administrators currently perform, do they run the risk of being out of the loop?

Dave: Absolutely not. Our focus on integrated data management doesn't eliminate the need for storage administrators. The difference is that they can now spend time on more strategic activities: bringing real value to the business, improving asset utilization, looking at ways to optimize performance, and thinking about a storage tiering strategy that matches the right level of data to the right-priced storage, instead of needing to be involved in every process.


Q: How does this management focus translate into the solutions and tools that NetApp offers its customers?

Dave: NetApp has invested heavily in developing manageability solutions with this


RELATED INFORMATION

- . [Data Center Proven Launch Page](#)
- . [Productivity Management in the Data Center \(Webcast Series\)](#)
- . [Dave Kresse on Data and Storage Management\(Podcast\)](#)
- . [Integrated Data Management \(pdf\)](#)

Pre-IDM: The Storage Admin Is Outnumbered with Requests



With IDM: System Admins and DBAs Are Empowered to Manage Their Own Data



approach in mind. Many of the products that are part of these solutions are available today and have been proven in many data center customer environments, including:

- Operations Manager (formerly known as DataFabric® Manager), File Storage Resource Manager, and VFM™ (Virtual File Manager) at the storage administrator level,
- The NetApp SnapDrive® product line for UNIX® and Windows® environments at the systems administrator level, and
- NetApp SnapManager® for Exchange, Oracle®, and SQL Server.

On our product roadmap we have committed a team of development resources to ensure both a growing breadth of management functionality and deeper levels of integration between management tools. Throughout the next year, customers can expect to see expanded capabilities in all of the software family suites.

Q: The concept sounds great, but can you be more specific? What is NetApp actually doing with these products to automate the one-off requests that bog down storage administrators?

Dave: I'll give you two examples. The first involves provisioning using SnapDrive, and the second involves streamlining application-specific backup and restore operations with SnapManager.

SnapDrive allows provisioning to take place at the system administrator level, without needing to involve the storage administrator level. Using Storage Suite, a storage administrator implements policies that set the boundaries within which the system administrator can provision the required data containers, be it LUNs, volumes, or file systems, without having to wait for the storage administrator for every new request.


NetApp also offers a suite of SnapManager products that are the result of deep integration with enterprise databases and applications, such as Oracle Databases, Microsoft® Exchange, and Microsoft SQL Server. With our solution, DBAs or application administrators can manage their own space-efficient approaches to database backup, recovery, and cloning. This enables them to operate as independent entities and respond quickly to dynamic business requests. For example, when they need to restore a corrupt database, they can do it themselves without having to involve system, storage, and backup administrators. This results in extreme efficiency and high availability for enterprises.

Q. Are there any real-life examples of customers who have successfully adopted this type of approach?

Dave: We have many customers who are already seeing the benefits of integrated data management. A great example is Norsk Hydro, Norway's largest company and a global energy and materials provider. Hydro supports over 50,000 employees in 60 countries and has three core data centers, four hub data centers throughout Europe, and hundreds of remote offices.

The company previously had three independent islands of storage and a wide variety of vendor and in-house management tools. The IT team was having a hard time maintaining uptime, it was highly complex to manage, and management costs were going out of control.

As a result of deploying NetApp storage and data protection solutions, including Operations Manager, the IT team can now centrally monitor and manage its entire distributed infrastructure and provide tiered storage with varied SLAs for its internal IT groups. Over the last three years the company's storage has tripled to over 400TB, and



Source: [Integrated Data Management white paper](#).

ATMs and the Changing Role of the Storage Admin

Before the onset of ATM technology, customers used to have to wait in line for someone to help them do various tasks. With ATM technology, the process is now automated to give customers secure, self-service access to their funds and account information 24 hours a day.

Benefits of automation:

- The bank's controls are defined by preestablished policies and enforced by automated processes.
- Customers are happier because they no longer need to wait in long lines for service and can access their funds and account information whenever they want.
- Bank employees are more efficient and can concentrate on higher value activities.

In much the same way an ATM provides banking customers with self-service access to funds and information, the integrated data management approach can provide application and system administrators with secure, controlled, self-service access to storage capacity and services. This gives them more flexibility in their choices and allows them to respond more quickly to changing conditions.

Get the details. [Read the white paper](#). (pdf)

NetApp Focuses on Data Management, Integration, and Partnerships

Some of the most innovative work NetApp has done in the area of integrated data management has been fueled by strong relationships forged between NetApp and its partners.

NetApp achieved some ground-breaking work in this area, including:

- Being the first storage vendor to integrate with Oracle Enterprise Manager
- Being the first storage vendor to develop a plug-in for use in SAP Adaptive Computing environments

Much of the NetApp commitment to integrated data management is reflected in the integration work found

it hasn't had to add any additional headcount to manage it.

Hydro's deployment is profiled in a [NetApp success story](#) and was recently showcased in an article highlighting [*Tips for Improving Remote Office Backup and Recovery*](#).

Q: What is the long-term NetApp vision for integrated data management?

Dave: NetApp is differentiated in the market because of its software. The value realized through the functionality within Data ONTAP® is well understood, but there is also differentiated value in the approach to manageability NetApp is pursuing. I believe this approach and the tools NetApp offers are valuable for anyone using NetApp storage.

Whether we're talking about a small law firm with a few tens of Exchange users or a global 100 company with hundreds of storage systems, our integrated management tools are proven to deliver real customer value. My vision is that these tools will become an integral part of every NetApp customer environment.

Comment on this article

in its SnapManager product suite:

- [SnapManager for SQL Server 2000 Best Practices \(pdf\)](#)
- [Best Practices for Oracle \(pdf\)](#)
- [SnapManager for Oracle with Oracle Database 10g Grid \(pdf\)](#)
- [SnapManager for Exchange Best Practices \(pdf\)](#)



CHRIS KIM

Professional Services Engineer, NetApp Global Services

Chris Kim joined NetApp in 1999 from Legato Systems. He initially provided application support for Microsoft® Exchange and SQL Server, then joined the NetApp Technical Advisor Team as an "Expert's Expert"—one of four people who traveled the globe teaching customers about NDMP. For the past four years, Chris has been the primary on-site consultant for a major NetApp customer. During that time, the customer has grown from 200TB of storage to nearly 1,000TB and from 25 NetApp storage systems to 110.

A Customer Dilemma: Faster Tape or Disk-to-Disk Backup?

By Chris Kim

Most companies reevaluate their backup environments every few years to determine whether they can extend the life of their existing infrastructure or if it is time to upgrade their backup strategy. NetApp Global Services (NGS) has helped companies big and small assess their environments and select the best products to get the job done today and prepare for the future.

For example, NGS was recently tapped to assist a company with a 1,000-node server farm and a backup environment that included:

- 200TB of electronic design automation (EDA) data
- 20 NetApp storage systems
- Data mirroring between sites using NetApp SnapMirror® software
- Tens of millions of small files
- Three to five days to complete a full tape backup
- Daily change delta of 10% to 30% (20TB to 60TB per day)

The IT team's goals were simple:

- Cut the full backup window to two days or less.
- Find a solution that's easy to implement in the current environment.
- Keep the budget under control.

To meet these goals, the IT team initially upgraded its existing tape backup environment and replaced 15 SuperDLT tape drives with 15 LTO tape drives. Unfortunately, this investment did not shorten the customer's backup windows.

It was at this point that the customer asked NGS for help redesigning its backup environment. In this article, I'll share a little insight into how NGS helps customers like this one improve backups (and recoveries!) by analyzing and evaluating their existing environments.

Step 1: Find the Bottleneck

Hint: Faster tape doesn't always fix the problem.

If you're backing up data from networked storage to tape, you might encounter bottlenecks in any of multiple areas. NGS uses several tests that can help determine whether or not simply upgrading existing tape libraries can solve ongoing backup issues. For example:

- **Test reads by dumping to a null device.** With this test you are essentially running a read of the file system on the NetApp system itself instead of writing to a separate system. This is a simple matter of running the built in **dump** routine to copy the target file system to `/dev/null`. Then look at how quickly the file can be read.

Tech OnTap May 2006 | Page 5

RELATED INFORMATION

- [Information-Optimized Backups](#)
- [Tips for Remote Backup](#)
- [Three Trade-Offs That Impact Disaster Recovery Planning](#)
- [Webcast: How One IT Team Cut Backup Costs 54%](#)
- [NetApp Disk-to-Disk Backup Solutions](#)

Information-Optimized Backups

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

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

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

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

Backup and Recovery Design Services

Designing and implementing a backup and recovery plan require hands-on knowledge of each component affected

- **Test writes using NDMPcopy to copy data to an idle storage system.** This encapsulates NDMP and essentially takes tape out of the equation, showing you how quickly the application could write to another storage system.

If the read and write performance is faster than the existing tape system can handle *and* future data growth is expected to be minimal, it may make sense to simply upgrade your existing tape drives or evaluate a virtual tape library (VTL) solution. However, if your read and write speeds are underutilizing your existing tape environment, neither faster tape drives nor a VTL is likely to improve your backup situation.

In this example, the IT team didn't invest time up front to identify bottlenecks. As a result, the team didn't learn that faster tape drives wouldn't impact backup performance until it had already invested in new LTO tape drives.

The customer then asked NGS to evaluate its environment and propose a solution. Since we already knew that the tape library wasn't the bottleneck, we began an assessment that included checking things like backup server memory and network capacity. We learned that because our EDA customer had literally millions of small files in each storage volume, its file-by-file backup method—which required walking the file system and cataloging each file as it is backed up—turned out to be the source of the problem.

Step 2: Determine the Rate of Data Change

Hint: Backup logs are valuable, but make sure you review all available info sources.

To determine the rate of data change, you first need to understand the file system: how many files are there, where are they located, how often do they change, and so on. We typically begin this process by examining current backup log files to see how much data is being backed up each night. Depending on the customer's version of Data ONTAP®, we might use a tool called SnapDelta to measure data change. If a customer is using NetApp SnapMirror software to mirror incremental changes off-site, we can also compare the value for the nightly block transfer from the SnapMirror log file to the backup value and estimate how much of the data being backed up each night is redundant.

NGS also uses several other types of log files. For example, if you're running tape with NDMP, it records a backup log on the NetApp system. Symantec® NetBackup™ also records log files. The key difference between the backup log and NetBackup log is that NetBackup records when the file header info is transferred over. Although the backup log might indicate the data is transferring quickly, you need to have copied everything (including the header) to the destination for the update to be complete. For a file system that has lots of header info (and thus lots of files), the Symantec logs take this into account in terms of the complete backup time.

In this case, the customer is already using SnapMirror to replicate incremental data changes off-site, so we were able to develop an initial estimate of data redundancy. However, since we knew that this customer's environment includes a lot of little files, we also looked very closely at the Symantec logs and weighted this data heavily in our analysis.

We determined that on average only about 30% of the data being backed up to tape had actually changed; the rest was a duplicate of data from previous backups. This meant that the customer could potentially cut the amount of data it was backing up by 70% by replacing traditional tape backup with a solution that copies only incremental changes.

Step 3: Map the Solution to Your Requirements

Hint: Not all disk-to-disk backup solutions can back up only incremental changes.

Based on our initial assessment, we suggested replacing the error-prone tape systems with a faster and more reliable disk-based backup approach that would eliminate the overhead created by tape loads, unloads, and other activities.

In traditional backup solutions, if a file has changed since the last backup, the entire file is copied to backup media, even if only a single block has changed. In most environments this means that much more data is backed up than necessary. Traditional methods also have to search through the entire file system file by file to identify the backup data set. In large file systems with many files, this process can be extremely time consuming.

as well as their interactions at the system level. This list can include host servers, file systems, application software, network infrastructure, storage devices, and clustering and multipathing applications.

Need help? NetApp Global Services can help you design and implement a systematic process to back up and restore data.

Learn more about [NGS Backup and Recovery Services](#). (pdf)

Real-World Example: How One IT Team Cut Backup Costs 54%

In a recent TechTalk Webcast, systems engineer David Medeiros explained how the Cal State San Marcos IT team automated business continuance, centralized management, and achieved high availability and scalability.



Watch the Webcast.
Read the [Q&A transcripts](#). (pdf)
Read the [Cal State San Marcos success story](#). (pdf)

NetApp [Webcast and Podcast library](#).

A Quick Primer on NetApp Data Protection Software

NetApp customers have two alternatives for data protection: SnapMirror and 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.

Get the details. Read the reports:

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

Traditional solutions also usually require the backups to flow through the media server to the backup media, introducing additional opportunities for bottlenecks.

NetApp disk-to-disk solutions—including mirroring and backup technologies such as SnapMirror and SnapVault®—utilize NetApp Snapshot™ technology to automatically identify incremental changes at the block level without time-consuming file-system searches. Since only changed blocks are transferred, these solutions use backup storage and network bandwidth very efficiently while still providing the effective equivalent of a full backup. These products (and NetApp VTL solutions) also eliminate the data flow issues by backing up directly to backup media.

To learn more about how NetApp makes backups more efficient, read [When Brute Force Isn't Enough: The Value of Information for Backups](#).

The Result: A Manageable Process with Room to Grow

Based on the NGS analysis, the customer decided to back up its existing NetApp primary storage systems to a NetApp FAS3050 system with ATA disks for built-in redundancy. Using NetApp SnapVault software to copy only changed data to the FAS3050 improves efficiency and streamlines the backup process. (See the sidebar for more details on SnapMirror and SnapVault.)

The IT team's next decision involved how best to manage the SnapVault infrastructure. There are a variety of potential methods for managing SnapVault, including:

- Command-line interface
- NetApp Operations Manager (formerly DataFabric® Manager)
- Integration with third-party backup software, including Symantec NetBackup, SyncSort, BakBone, CommVault, and more

Because the customer was already using NetBackup, the IT team decided to manage its entire SnapVault environment using NetBackup. This gives the customer the ability to manage the new infrastructure using familiar graphical interfaces, monitoring, and alerts.

And what about the 15 new LTO tape drives? The IT team currently plans to repurpose them for third-tier storage.

With the transition now under way, the company is pleased with the preliminary results:


- **Goal: Cut the full backup window to two days or less.**

Incremental daily backups are currently completed in less than one hour and should be completed in under five hours once the project is done.

- **Goal: Find a solution that's easy to implement in the current environment.**

The new solution will be part of the familiar and well-tested NetBackup implementation.

- **Goal: Keep the budget under control.** The new solution has solved the problems and permanently reduced costs. The IT team expects a two-year return on investment given administrative time savings, lower tape media costs, and so on.



- **Goal: Shorten recovery time (lengthy restores from on-site and off-site tape and occasional media failures were compromising the IT team's ability to meet internal SLAs).** This wasn't an original goal, but ended up being one of the *biggest* selling points of a disk-to-disk solution. The new solution has enabled the team to implement a local disaster recovery solution, thereby substantially improving recovery times and recovery points.

Comment on this article



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.

Reducing Time-to-Deployment: Improving Agility for Database Development, Test, and QA

By Bruce Moxon

Last month I introduced the notion of whole application virtualization, contending that a number of core storage services are necessary to deliver on this vision and the promise of grid computing. This month, I explore two of those services, mirroring and cloning, in the context of a rapid application development model for database-centric applications that increases an organization's agility.

Every database administrator (DBA) has experienced it—a request from the application development organization for a copy of a production database to support application development and testing. For larger databases in excess of a few hundred gigabytes such a request sets off a series of potentially time-consuming planning, provisioning, and implementation activities. (Oh, well, there goes the weekend ...)

The Plan ...

With traditional tools and approaches, a strategy for replicating relevant portions of the production database while minimizing impact on production operations must first be developed.

Depending on an organization's operations, this plan might include one of two fundamental approaches: building a dev/test copy of the database from backup tapes or doing a (partial or full) extraction of data from the production database—using a variety of vendor-supplied and homegrown tools. Backup-based replication is more straightforward and nonintrusive on the production system, but can be a lengthy process—especially for tape-based backups. Extraction-based replication can be faster—especially if partial replication is sufficient for the scenarios under development or test—but may require further planning to ensure that the production system is minimally impacted during the period of data extraction.

Provisioning ...

Once the plan is in place for the replication of the database, adequate resources—including both servers and storage—must be provisioned. The servers must then be loaded with the appropriate software "stack" (operating system, drivers, database, application servers, ...) and the database must be restored or created and populated.

Replication ...

For many organizations, the effort to replicate a large (multi-100GB) database can take hours or even days. Replications requiring restores from tape are particularly lengthy and manually intensive.

In many large-scale environments, multiple copies of production databases may be

RELATED INFORMATION

- [Fueling the Grid: Pursuing the Vision of Enterprise IT](#)
- [A User Perspective on FlexClone](#)
- [ESG Study: Data ONTAP 7G \(pdf\)](#)
- [Data ONTAP 7G: The Ideal DB Application Platform \(pdf\)](#)
- [FlexVol and FlexClone Demo](#)

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 and includes:

- 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](#).

Advantages of NetApp FlexClone Technology in Database Environments

This five-minute swing bench demo shows:

required to support a range of routine activities, including application and database development, test, QA, and perhaps user training. This clearly has a multiplicative factor on the physical and human resources required to support such an endeavor.

Reducing Time-to-Deployment for Database Applications

In the storage industry at-large, a few key technologies have matured over the past few years to help organizations address this growing need for physical database replication and rapid deployment.

Disk-to-disk backup technology has become the norm in many organizations, greatly enhancing the speed at which database replicas can be constructed from backups. Snapshots provide space efficient, point-in-time copies of data containers (file systems, volumes, LUNs) that can be used for rapid recovery and, increasingly, as the basis for creating full writable data copies in a technique known as cloning.

While these technologies greatly simplify the task of replicating the database data, organizations employing these approaches are still faced with two significant operational challenges. First, the act of cloning the production data can adversely impact production application performance. As a result, cloning may need to be restricted to "off" or "nonpeak" hours. Second, this approach still requires a complete copy of the underlying data to be made. This can pose nontrivial provisioning challenges, especially when copying large databases multiple times.

The NetApp Fulcrum ...

Due to its unique "DNA," which includes Data ONTAP® and WAFL®, Network Appliance has developed innovative implementations of these general industry approaches in ways that deliver unmatched speed, efficiency, and simplicity to the task of database replication and deployment.

With Data ONTAP 7G, Network Appliance introduced FlexClone™ volumes. These are writable flexible volumes that are created from NetApp Snapshot™ copies. As such, they are effectively near-instantaneous, space efficient, writable "clones," sharing the same physical blocks with the baseline file system or LUN. As the baseline and clone data containers diverge (for example, due to continuing updates in the production database or development changes to a cloned database), only the divergent blocks need to be written to disk.


The result is that working copies of very large databases can be created within a few minutes and with minimal incremental storage. A typical 300GB database can be cloned and brought up in just a couple of minutes. (NetApp customer Ben Rockwood provides specific examples of FlexClone volume creation in his article "[The Versatile Storage Platform](#).")

FlexClone volumes can only be created on the same storage system on which the Snapshot copy of the original data resides. Hence cloned data shares the same storage resources as the source of the clone. In cases where the production storage system is not fully utilized or where NetApp control of service capabilities can be employed, this may be sufficient to shield the production database from additional load that could be imposed by development activities. More commonly, however, operational requirements will dictate that the development databases be created on a separate storage system. Enter NetApp SnapMirror® software.

SnapMirror can be used to continually replicate production data (synchronously, if desired) to a secondary storage system with minimal impact on the production environment. The cloning and subsequent development activity can then take place on the secondary system, completely "out-of-band" of the production system. This secondary system might also serve as the organization's disaster recovery (DR) target or might be constructed as a low-cost development system utilizing lower cost controllers, iSCSI, and SATA drives.

A blade server-based deployment depicting this scenario is shown below. In this scenario, new database servers would be dynamically provisioned from the blade server pool, using the database containers cloned on the secondary (Dev/Test) system.

- Performance benefits of running OLTP database loads on aggregated storage
- Ability to increase and decrease volumes in seconds
- Ability to create database clones in under a minute for testing



[Launch the demo.](#)

FlexClone Performance

Unlike other implementations of cloning technology, FlexClone volumes are implemented as a simple extension of existing core mechanisms. As a result, the performance of FlexClone volumes is nearly identical to the performance of flexible volumes.

The creation of a FlexClone volume is nearly identical to creating a Snapshot copy and is usually completed within seconds. Clone metadata is held in memory just like a regular volume, so the impact on storage system memory consumption is identical to having another volume available.


Get the details. Read the 32-page [Introduction to Data ONTAP 7G \(pdf\)](#).

Simplifying Data Management: Doing More while Spending Less

NetApp designed its original business plan to include the "Cisco magic":

- An embedded file system
- Commodity hardware
- Native multiprotocol support

Get the details. [Watch Tom's Oracle OpenWorld keynote](#) or read transcript [excerpts](#).



In environments where third-party storage is deployed in the production environment, the same approach can be employed with native database mirroring capabilities, with a variety of emerging heterogeneous replication technologies, including NetApp [SnapDrive®](#) and [V-Series](#) technologies, and potentially with third-party replication tools such as those in [Topio's Data Protection Suite™](#) or [IBM's SAN Volume Controller \(SVC\)](#).

Finally, this same approach can be applied in additional operational scenarios to improve utilization and agility without impacting production system performance. Examples include:

- Offloading end-of-month reporting to cloned databases
- Offloading intensive ad hoc queries and data warehousing extract-transform-load (ETL) operations in support of business analytics
- Distributed development, using SnapMirror to replicate data across geographically distributed sites and employing FlexClone at the remote site to create copies for development, QA, or training

Through the NetApp SnapManager® suite of products, these capabilities are being integrated into standard database management frameworks, such as Oracle® Grid Control and Microsoft® SQL Server Enterprise Manager. This approach provides DBAs with more rapid and space-efficient approaches to database backup, recovery, and cloning, all from the same management suite they're using today.

The implications for many database development organizations are far reaching. NetApp's unique FlexClone implementation, coupled with appropriate replication technology, can significantly reduce the time required to provision database replicas for a variety of development and operational scenarios. Because of the space savings aspect of this approach, many organizations are finding that not only can they do what they were doing before faster, they can also do a lot more of it (many simultaneous dev/test cycles, for example).

Finally, increased agility and reduction in physical storage requirements are accompanied by a simplification of operational procedures—in short, helping to deliver on the promise of grid computing.

Previous article by Bruce Moxon:

- [Fueling the Enterprise Grid: Pursuing the Vision of Enterprise IT](#)

Comment on this article



BEN ROCKWOOD

Systems Administrator, Homestead.com

A self-professed Sun™ zealot, Ben became an advocate of NetApp NAS and iSCSI storage after using heterogeneous NetApp storage to power Homestead.com. In his words ...

"NetApp can serve NFS faster than the people who wrote it. NetApp also nailed iSCSI out of the box. The great thing about Network Appliance™ systems is that they are exactly that, appliances. They just work! And when problems do occur, a NetApp representative contacts me before I know about the problem myself. Of all the storage systems in my data center, the 940s set the standard that I hold other vendors to."

The Versatile Storage Platform: A User Perspective on Data ONTAP® 7G

By Ben Rockwood

NetApp storage systems are best known as network-attached storage (NAS) devices, but I think few people appreciate just what a versatile storage platform they really are. With a single system you can provide NAS, both CIFS and NFS, Fibre Channel SAN (both fabric and DAS), and IP SAN (iSCSI and iFCP). I've found times in my past where I knew I needed storage for a project, but it wasn't clear at the time of purchase just exactly how we needed to access the storage. In these cases, choosing NetApp storage provided me with the widest range of possible solutions to solve the problem and further to continue using that storage beyond its original intent.

To truly appreciate the versatility of the NetApp architecture, it is important to view how the storage is managed and how the storage is accessed as related functionality. Too often storage vendors separate these two concepts, which creates overly specialized storage systems that eventually become isolated islands. Storage consolidation is a constantly reoccurring storage theme, which typically requires migration to yet another storage platform only to repeat the scenario again in the future. Choosing a platform that will not only meet your needs today but also has the flexibility to meet your needs in the future will allow you to continue to leverage your initial investment, scale that investment over time without migrating the underlying platform, and ultimately deliver a return on investment that you can depend on for years to come.

In this article, I'll take a broad look at how the NetApp operating system, Data ONTAP, manages data on disk and how that data is accessed from your client systems. Some of the technologies I'll explore include:

- RAID, RAID-DP™, and traditional volumes
- Aggregates and FlexVol™ volumes
- Snapshot™ functionality and FlexClone™ technology
- LUN creation and masking

In releases prior to Data ONTAP 7G, volumes were defined much like they are with other tools, creating an intimate relationship between volume and disk. A volume (on which a file system sits that contains data) would be created on top of one or more *RAID groups*. Each RAID group is RAID4, containing *n* data disks plus one parity disk. With RAID4 you can lose no more than one disk, which means that when a disk fails, you're completely vulnerable to a second failure during the disk rebuild (and with 300GB drives, a disk rebuild can take a while).

RELATED INFORMATION

- [Cuddletech: The Blog of Ben Rockwood](#)
- [Improving Agility for Database Development, Test, and QA](#)
- [Architecting Storage for Resiliency](#)
- [Intro to Data ONTAP 7G \(pdf\)](#)
- [FlexVol and FlexClone Demo](#)

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 and includes:

- 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](#).

The 6.x and later releases of Data ONTAP include a nifty feature called *double parity*, which allows you to have two parity disks (this isn't just a mirror of parity; [learn more about RAID-DP here \(pdf\)](#)). By using RAID-DP you could lose up to two disks, which covers you during a spare reconstruction and gives you a little extra security and peace of mind.

Here's an example of the creation and configuration of a traditional volume:

```
nova> vol create trad_vol0 -t raid_dp 5
Creation of a volume with 5 disks has completed.
nova> vol status
      trad_vol0 online   raid_dp, trad
nova> sysconfig -r
Volume trad_vol0 (online, raid_dp) {zoned checksums}
Plex /trad_vol0/plex0 {online, normal, active, pool0}
RAID group /trad_vol0/plex0/rg0 {normal}

RAID Disk Device  HA  SHELFBAY CHAN Pool Type  RPM  Used (MB/blks)  Phys (MB/blks)
-----
dparity  7.48  7      6  0  FC:A  0  FCAL 10000 34500/70656000 34732/71132960
parity   7.41  7      5  1  FC:A  0  FCAL 10000 34500/70656000 34732/71132960
data     7.33  7      4  1  FC:A  0  FCAL 10000 34500/70656000 34732/71132960
data     7.49  7      6  1  FC:A  0  FCAL 10000 34500/70656000 34732/71132960
data     7.42  7      5  2  FC:A  0  FCAL 10000 34500/70656000 34732/71132960

...
nova> df -h trad_vol0
Filesystem                total    used    avail capacity  Mounted on
/vol/trad_vol0/           72GB    132KB   72GB      0%   /vol/trad_vol0/
/vol/trad_vol0/.snapshot  18GB      0KB   18GB      0%   /vol/trad_vol0/.snapshot
```

When you create a traditional volume (simply referred to in releases before Data ONTAP 7G as a "volume"), the volume sits on top of the RAID groups you created. If you want to increase the size of the volume, you must add more disks. For performance reasons you typically want to add additional full RAID groups at a time.

This method of allocation is fairly inflexible: you can't shrink volumes, and your volume sizing is based on the physical disk sizes, not simply needed capacity. To make up for this, you can create within a traditional volume a *quota tree* (qtree), which by and large can be treated like a volume itself. For example, you can replicate individual qtrees using NetApp mirroring software (SnapMirror®) and define different access types (UNIX®, NTFS, or mixed). But qtrees are only so flexible.

Enter Data ONTAP 7G, in which NetApp introduces a new concept called *aggregates*. Aggregates add a layer of abstraction between the physical disks and the volumes. The idea is simple: allocate a certain number of disks to a big aggregate, which is built on RAID groups just like traditional volumes. However, there is no file system within an aggregate; it is just allocatable space.


Now within this big aggregate we can create something NetApp calls *flexible volumes* ("FlexVol"). Because a FlexVol volume is abstracted from the underlying disk, you can create a volume to meet your capacity needs without regard for physical layout. Furthermore, you can grow or shrink a FlexVol volume easily at any time. This truly is flexible. One advantage of flexible volumes is that as your needs and requirements change over time, you can manipulate your storage by sloshing things around within the aggregate.

This is an example of aggregate creation with RAID-DP enabled:

Advantages of NetApp FlexClone Technology in Database Environments

This five-minute swing bench demo shows:

- Performance benefits of running OLTP database loads on aggregated storage
- Ability to increase and decrease volumes in seconds
- Ability to create database clones in under a minute for testing



[Launch the demo.](#)

FlexClone Creation

Unlike other implementations of cloning technology, FlexClone volumes are implemented as a simple extension of existing core mechanisms.

The creation of a FlexClone volume is nearly identical to creating a Snapshot copy and is usually completed within seconds. Clone metadata is held in memory just like a regular volume, so the impact on storage system memory consumption is identical to having another volume available.

Get the details. Read the 32-page [Introduction to Data ONTAP 7G \(pdf\)](#) or the 35-page [Thorough Intro to FlexClone Volumes. \(pdf\)](#)

RAID-DP vs RAID4

RAID-DP significantly increases data protection with zero to minimal impact to capacity utilization or performance versus RAID4. And, since RAID-DP is an integral part of Data ONTAP, there are no hidden costs associated with it.

RAID-DP offers:

- Protection against up to two disk

```
nova> aggr create aggr0 -r 7 -t raid_dp 14#36
Creation of an aggregate with 14 disks has completed.
nova> sysconfig -r
Aggregate aggr0 {online, raid_dp} {zoned checksums}
Plex /aggr0/plex0 {online, normal, active, pool0}
RAID group /aggr0/plex0/rg0 {normal}
```

RAID Disk Device	HA	SHELF	BAY	CHAN	Pool	Type	RPM	Used (MB/blks)	Phys (MB/blks)
dparity	7.48	7	6	0	FC:A	0	FCAL 10000	34500/70656000	34732/71132960
parity	7.41	7	5	1	FC:A	0	FCAL 10000	34500/70656000	34732/71132960
data	7.33	7	4	1	FC:A	0	FCAL 10000	34500/70656000	34732/71132960
data	7.49	7	6	1	FC:A	0	FCAL 10000	34500/70656000	34732/71132960
data	7.43	7	5	3	FC:A	0	FCAL 10000	34500/70656000	34732/71132960
data	7.34	7	4	2	FC:A	0	FCAL 10000	34500/70656000	34732/71132960
data	7.50	7	6	2	FC:A	0	FCAL 10000	34500/70656000	34732/71132960

```
RAID group /aggr0/plex0/rg1 {normal}
```

RAID Disk Device	HA	SHELF	BAY	CHAN	Pool	Type	RPM	Used (MB/blks)	Phys (MB/blks)
dparity	7.35	7	4	3	FC:A	0	FCAL 10000	34500/70656000	34732/71132960
parity	7.44	7	5	4	FC:A	0	FCAL 10000	34500/70656000	34732/71132960
data	7.51	7	6	3	FC:A	0	FCAL 10000	34500/70656000	34732/71132960
data	7.36	7	4	4	FC:A	0	FCAL 10000	34500/70656000	34732/71132960
data	7.45	7	5	5	FC:A	0	FCAL 10000	34500/70656000	34732/71132960
data	7.52	7	6	4	FC:A	0	FCAL 10000	34500/70656000	34732/71132960
data	7.37	7	4	5	FC:A	0	FCAL 10000	34500/70656000	34732/71132960

```
nova> vol create flexvol0 aggr0 40g
Creation of volume 'flexvol0' with size 40g on containing aggregate
'aggr0' has completed.
nova> vol create flexvoll aggr0 1g
Creation of volume 'flexvoll' with size 1g on containing aggregate
'aggr0' has completed.
nova> vol size flexvol0 -20g
vol size: Flexible volume 'flexvol0' size set to 20g.
nova> vol size flexvoll +9g
vol size: Flexible volume 'flexvoll' size set to 10g.
nova> df -h
Filesystem                total      used    avail capacity  Mounted on
...
/vol/flexvol0/             16GB       92KB     15GB      0% /vol/flexvol0/
/vol/flexvol0/.snapshot    4096MB      0KB    4096MB     0% /vol/flexvol0/.snapshot
/vol/flexvoll/             8192MB     96KB     8191MB     0% /vol/flexvoll/
/vol/flexvoll/.snapshot    2048MB      0KB    2048MB     0% /vol/flexvoll/.snapshot
```

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

Learn more about the NetApp Implementation of RAID-DP. (pdf)

Flexible volumes let users do some really cool things. One of my favorites involves cloning. We can take an existing FlexVol volume and create a clone from it, either using the volume itself (right now) or using an existing Snapshot copy from some point in the past (like before your server blew up). A FlexClone volume looks exactly like the volume it was created from (its parent FlexVol volume or another FlexClone volume), but it uses no additional physical storage!

This code shows you how easy it is to create a FlexClone volume:

```
nova> vol clone create flexclone0 -b flexvol0
Creation of clone volume 'flexclone0' has completed.
nova> df -h
Filesystem                total      used    avail capacity  Mounted on
...
/vol/flexvol0/             16GB      104KB     15GB      0% /vol/flexvol0/
/vol/flexclone0/           16GB    1552KB     15GB      0% /vol/flexclone0/
nova> vol size flexvol0 -10g
vol size: Flexible volume 'flexvol0' size set to 10g.
nova> df -h
Filesystem                total      used    avail capacity  Mounted on
...
/vol/flexvol0/             8192MB     104KB     8191MB     0% /vol/flexvol0/
/vol/flexclone0/           16GB    1552KB     15GB      0% /vol/flexclone0/
```

Snapshot Copies + FlexClone Volumes = A Powerful Tool

Snapshot copy management itself is really easy. The operation takes milliseconds. The combination of robust NetApp Snapshot capabilities and the WAFL® file system means that operations such as volume creation, Snapshot copy creation, and cloning require almost no work at all. No more sitting around waiting an hour for volumes to become usable. This is the real magic.

Here's an interesting example. Say your PostgreSQL system got attacked and the database was corrupt. Your first priority is going to be getting the databases back online, but it would also be nice to have a copy of that corrupt data for further analysis when some of the DBAs get into the office in the morning.

If we were using NetApp storage with Data ONTAP 7G, we could create a clone of the

affected volume based on a prior Snapshot copy that is known to be good, make that available to the system (export the NFS share, CIFS share it, or remap the LUN for iSCSI or SAN), and be back online in the time that it took you to type the commands and to roll forward the logs. Then we could mount that corrupt volume off to a test system for analysis. When the analysis is complete, we could split off the clone that we created so that it would be its own unique volume.

At this point we have two options: we can keep using the clone as is (not losing the changes between the time we recovered and now), or we could delete the clone and revert the whole volume back to the last known good Snapshot copy. If you choose to revert back to a previous Snapshot copy, all of the Snapshot copies that were created after that point in time will be destroyed. Sometimes you need to revert to an older Snapshot copy but don't want to lose more recent changes: in this scenario, cloning based on a point-in-time Snapshot copy and then being able to split off that clone provides some flexibility.

Reviewer note: Depending on the scenario, the process for rolling changes back into a live database can vary. For more information and specific examples, NetApp recommends reading [A Thorough Introduction to FlexClone Volumes \(pdf\)](#).

If you're interested in cloning, I'd recommend Bruce Moxon's article about developing a [rapid application development model](#) for database-centric applications.

Managing LUNs for Fibre Channel SAN and iSCSI

Just managing data is only part of any storage solution and doesn't provide any real value if you can't access that data. NetApp storage systems are legendary NAS devices and well known and trusted in that space. The FlexVol volumes that we created earlier needed only to be exported via NFS ("exports") or shared via CIFS ("cifs shares -add") to be made available on your network as NAS resources. However, many users may be unfamiliar with the ability of Data ONTAP to serve block storage LUNs for iSCSI and Fibre Channel SAN access.

To demonstrate just how easy and painless NetApp is to manage for block environments, let's expand on the concepts we explored earlier to build LUNs and control how they are accessed by clients.

When dealing with iSCSI or FCP (Fibre Channel Protocol: that is, SAN), we are making block allocations, which we call LUNs, available. A LUN is just a big allocated file sitting on a volumes file system (or so it appears). In order to create a LUN, you must first create a volume, as discussed earlier. Then, when creating the LUN, you'll specify the path within that volume. Here is an example:

```
nova> vol create qa_luns aggr0 100g
Creation of volume 'qa_luns' with size 100g on containing aggregate
'aggr0' has completed.
nova> df -h /vol/qa_luns
Filesystem                total      used    avail capacity  Mounted on
/vol/qa_luns/             80GB       100KB    79GB      0%   /vol/qa_luns/
/vol/qa_luns/.snapshot    20GB        0KB    20GB      0%   /vol/qa_luns/.snapshot
nova> lun create -s 45g /vol/qa_luns/qa_solaris1
nova> lun create -s 20g /vol/qa_luns/qa_linux1
nova> lun show
          /vol/qa_luns/qa_linux1      20g [21474836480]  (r/w, online)
          /vol/qa_luns/qa_solaris1    45g [48318382080]  (r/w, online)
nova> df -h /vol/qa_luns
Filesystem                total      used    avail capacity  Mounted on
/vol/qa_luns/             80GB       65GB    14GB      81%   /vol/qa_luns/
/vol/qa_luns/.snapshot    20GB       68KB    19GB      0%   /vol/qa_luns/.snapshot
```

You can see how easy it is to create LUNs. You'll notice that while I created a 100GB volume, 20GB of it is allocated to Snapshot copies. We can limit that easily:

```
nova> snap reserve qa_luns 5
nova> df -h /vol/qa_luns
Filesystem                total      used    avail capacity  Mounted on
/vol/qa_luns/             95GB       65GB    29GB      69%   /vol/qa_luns/
/vol/qa_luns/.snapshot    5120MB     68KB    5119MB    0%   /vol/qa_luns/.snapshot
```

Once you've created a LUN, you make it accessible by *mapping* it to an *initiator group*. Initiator groups contain one or more iSCSI IQNs or FCP WWNs, which map to a LUN on the NetApp storage system. This is traditionally known as LUN masking, because you're making sure that only the initiators (clients) that should see certain LUNs can.

LUN masking sounds a lot more difficult than it is. Here is an example:

```
nova> lun show
      /vol/qa_luns/qa_linux1      20g [21474836480]  (r/w, online)
      /vol/qa_luns/qa_solaris1    45g [48318382080]  (r/w, online)
nova> igroup create -f QA-Solaris 10:00:00:00:c9:2b:51:2c
nova> lun map /vol/qa_luns/qa_solaris1 QA-Solaris 0
nova> igroup show QA-Solaris
      QA-Solaris (FCP) (ostype: default):
      10:00:00:00:c9:2b:51:2c (not logged in)
nova> lun show /vol/qa_luns/qa_solaris1
      /vol/qa_luns/qa_solaris1    45g [48318382080]  (r/w, online, mapped)
```

iSCSI and FCP are handled in exactly the same way, except for a single flag. Notice the *-f* in *igroup create* in the preceding code. This means FCP; change it to *-i*, and you're doing iSCSI instead.

This is only a taste of what you can do with a NetApp storage system. I hope you can see that there is a huge amount of flexibility that is either unavailable from other vendors or so expensive as to be impractical (replication using EMC software vs. NetApp SnapMirror software is a good example). Personally, I view NetApp storage systems as the Swiss Army knives of the storage world: they do everything you need to do and a ton of stuff that might come in handy someday (like that little toothpick).

When you need flexibility to roll with whatever comes your way, nothing beats NetApp: the versatile storage platform.

Comment on this article

Talk to the Tech OnTap editor about submitting an article



LELLA BENNETT

Professional Services Consultant, Network Appliance Global Services (NGS)

Lella was formerly the senior systems architect at NAREX (now part of Fair Isaac Corporation), where she was responsible for one of the largest Microsoft® SQL Server installations in the world. Lella joined NGS to provide full-time, on-site management for a new high-availability customer deployment after a competitor's storage system caused a 10-minute outage (estimated cost: \$10 million). The company has achieved 100% uptime since inception (two-plus years) and today relies on over 40 NetApp storage systems. Lella's wealth of experience in ensuring storage system SLA adherence has made her a valuable member of the NGS extended team, which develops managed service offerings worldwide.

May Tool of the Month: MRTG for Filers

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: NetApp created a custom version of the open source MRTG tool.

What it is: The Multi-Router Traffic Grapher—MRTG—was originally developed with routers in mind and can be used to monitor any networked Simple Network Management Protocol (SNMP)—capable device. MRTG for Filers offers Web-based monitoring for NetApp storage systems and provides a current and historical graphical view (last day, week, month, and year) of statistics for network traffic (per interface), CPU load, disk usage, and NFS and CIFS operations.

How it works: MRTG consists of a Perl script that uses SNMP to read the value of any SNMP variable and a fast C program that logs the data and creates graphs. These graphs are embedded into Web pages, which can be viewed from a browser. To create the version specific to NetApp, NetApp modified the MRTG config file to reference the correct SNMP object identifiers (OIDs) as defined by the NetApp Management Information Base (MIB). The makefile was modified to automatically query a NetApp storage system to discover its volumes and interfaces and construct the configuration.

Why Lella thinks it's cool: Using MRTG for Filers you can quickly establish performance baselines and do capacity trending for proactive management. It's free, simple to use, and—since it's an open source tool based on SNMP—very flexible. It's an added benefit that lots of companies are already familiar with MRTG and this tool can integrate into existing MRTG frameworks.

How Lella uses MRTG: When a customer engages NGS in my region, one of the first things we do is install a suite of tools that includes MRTG for Filers. This way, the customer doesn't have to buy any specialized tools right off the bat. Within an afternoon we can be set up to gather a range of useful information that we use to diagnose existing issues and proactively manage the customer environment.


Here's an example. My team provides managed storage services to a customer where a single minute of downtime costs approximately \$1 million. We use MRTG for Filers to monitor bandwidth and proactively track performance patterns to make sure there isn't any overloading across the customer's environment that might impair our team's ability to deliver on our contracted service-level agreements.

The long-term trending data provided by MRTG for Filers enables us to project when operating capacity will be exceeded months in advance and detect unusual spikes in disk space consumption. For instance, I noticed a 25% spike in disk utilization during a single month last spring that turned out to be the result of an annual data load from a supplier. Since purchase cycles for new storage often take 60 days or more, long-term, proactive monitoring ensures that the customer has plenty of time to react.

RELATED INFORMATION

- [Official MRTG Website](#)
- Previously Featured Tools:
 - [Nagios](#)
 - [SIO](#)
 - [ToasterView](#)
 - [Simulate ONTAP™](#)
- [NOW™ Toolchest](#) (password required)

Sample MRTG for Filers Monitoring Screens



Other tools that NGS installs as part of its standard suite include:

- SWATCH (syslog watcher)—proactive log file monitoring for Linux® and UNIX®
- Kiwi Syslog Daemon—proactive log file monitoring for Windows®
- NetApp Operations Manager—comprehensive monitoring and management for all aspects of NetApp storage (purchase required)

Caveats: MRTG is not a replacement for sophisticated monitoring and management tools such as NetApp Operations Manager (formerly DataFabric® Manager), which provides a full suite of alerts, reports, and configuration tools. The MRTG for Filers package is written in perl, but uses one binary executable called "rateup." Versions of rateup are included in the package for Solaris™, HP-UX, and Linux.

Download MRTG for Filers 1.4
(password required).

Want to recommend a tool?
If we feature it, you get a FREE NetApp jacket.

Nominate a Tool

