

## HIGHLIGHTS

- Exchange Server 2007: Backups & DR
- Top SnapVault® Questions Answered
- Using NetApp to Manage EMC Arrays
- New Midrange Platform: FAS3040
- Webcast: NetApp/VMware Solutions

## A Mind Shift in Data Protection: Agilent Moves to Disk-to-Disk Backup

Tim Dietrich, Systems Architect, Agilent Technologies

Replacing tape backup with a NetApp SnapVault solution eliminated 24-hour backup windows and freed up server CPU cycles while providing better protection for remote offices.

[More »](#)

## BLOGGING WITH DAVE

Dave Hitz, NetApp Founder and EVP

"This is an unusual blog entry: I'm releasing a 20-page strategy document that I wrote in late 2003."

• [Dave's Blog](#)

## DRILL DOWN

- **New FAS3040 Midrange Storage Platform Provides 126TB Raw Capacity**  
See full specs plus a 14,000-user configuration for Exchange Server 2003.
- **Webcast: NetApp/VMware Solutions**  
A sysadmin discusses his deployment of VMware/NetApp (plus ask questions live!).
- **Vote for Your Favorite NetApp Beers!**  
Would you prefer a WAFL® Ale, RAID-DP™ Red, or perhaps a Hot Swap Hefeweizen?



## ADMIN RESOURCES

- **End-User Review: filestat Shell Script**  
How one IT manager reduced help desk calls by enabling self-service volume info.
- **NetApp NOW™ Community Forums**  
Submit questions and exchange ideas on topics ranging from Data ONTAP® to Decru DataFort® appliances.

## TIPS FROM THE TRENCHES

### Experts OnTap: Picking a Backup Solution How One IT Team is Evaluating Its Options

An inside look at one Tech OnTap member's evaluation of NetApp SnapVault. After grilling a panel of NetApp product managers and technical experts, this customer agreed to share:

- Initial questions and concerns
- Detailed 30-day evaluation plan
- Key learnings and next steps

Plus, submit your topics for *Experts OnTap!*

[More »](#)

### How Exchange Server 2007 Will Impact Your Backup and DR Plans

Robert Quimbey, Microsoft Alliance Engineer, NetApp

Find out how implementing new features such as LCR and CCR will impact your storage design, plus tips for choosing between iSCSI and Fibre Channel SAN for Exchange.

[More »](#)

## ENGINEERING TALK

### Simplifying Data Management for EMC, HP, IBM, Fujitsu, and Hitachi Environments

Steven Lusnia, V-Series Solution Architect, NetApp

The NetApp V-Series product line enables incremental-only Snapshot™ copies, unified SAN and NAS storage, end-user recoveries, and more in EMC and other environments. Includes 5 real-world examples.



[More »](#)

[FEEDBACK](#)

TECH ONTAP ARCHIVE



**Tim Dietrich**

Systems Architect, Agilent Technologies

Tim has been working in IT for over 15 years, with 11 of those years spent at Agilent functioning as an internal IT consultant. Tim was responsible for driving the project described in this article, including evaluating potential solutions, serving as advocate and IT architect, and managing all phases of implementation and rollout.

## A Mind Shift in Data Protection Agilent Moves to Disk-to-Disk Backup

By Tim Dietrich

Agilent Technologies is the world's premier measurement company and a technology leader in communications, electronics, life sciences, and chemical analysis. A spin-off of Hewlett Packard in 1999, Agilent has over 19,000 employees, serves customers in 110 countries, and had net revenue of \$5.0 billion in fiscal 2006.

Agilent has a total of four locations in the United States and three additional locations in other parts of the world. One of our research and development divisions is based in Southern California. This division has about 30TB of storage and about 200 computer-savvy engineers.

### Backup Challenges

In 2005, our backup environment in Southern California consisted of predominantly network-based backup to tape. We had about 300 systems on the network being backed up to centralized tape libraries. To meet our SLAs, we were keeping copies of the backups locally so we could do quick file restores, and we were also duplicating tapes for off-site storage and disaster recovery. The whole environment was getting to be very challenging to support. Our key challenges included:

- Long backup windows—approaching or exceeding 24 hours
- Backups were dragging down system performance
- Complicated management
- Low reliability
- No remote office support
- Labor-intensive tape duplication process
- Increasing costs

Tape just wasn't scaling for us. We were seeing huge data growth year over year, and continuing to address that problem with backup to tape was no longer getting the job done.

### Storage Infrastructure Overhaul

We were already pursuing the idea of centralizing the business data that was spread across our distributed systems down to a few storage devices in order to make it easier for the R&D staff to manage their data and to improve our storage utilization and data availability. However, the real clincher on this strategy came when we spoke to NetApp

**OnDemand  
TechTalk Webcast:**  
Increasing Backup  
Efficiency with a  
NetApp SnapVault®  
Solution

#### RELATED INFORMATION

- [Webcast: Increasing Backup Efficiency](#)
- [Ask the Experts: SnapVault](#)
- [D2D Backup or Faster Tape?](#)
- [Forrester Survey Findings: The Shift to Disk-Based Data Protection Solutions](#)

### A Quick Primer on NetApp Data Protection Software

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

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

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

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

and learned about SnapMirror® and SnapVault technology available with NetApp storage systems. We knew we had found a backup solution that offered the sea change we needed to overcome our existing backup limitations.

We wound up deploying a number of frontline NetApp storage systems—FAS900 and FAS3000 series systems—and centralizing all of our data on them for primary storage. We also deployed a NetApp R200 as a second tier. That allows us to use our FAS storage for what it is intended—high-performance storage—and use less expensive storage for backups.

Client systems primarily access the centralized storage using NFS and CIFS, which provides our engineers the most flexibility in reaching their data from anywhere on the network. We had also been running Fibre Channel SAN in the past for a few servers that required greater I/O performance, but found SAN to be terribly complex to support and expensive to maintain. For the situations where we still needed SAN levels of performance, we started relying on iSCSI, which we've found to be very fast, reliable, and much more cost effective.

We adopted SnapVault for backing up our primary storage systems to our secondary storage; we back up that secondary storage to tape using NDMP for disaster recovery purposes. Because we only do tape backups once a week now, as opposed to backing up to tape every day, we have the luxury of taking as much time as we need to complete those tape backups and have dropped from a couple dozen tape drives down to four.




Figure 1: Agilent Network Architecture

For those few systems where for whatever reason we couldn't use centralized storage, we adopted Open Systems SnapVault (OSSV) to allow those legacy systems to utilize the new backup paradigm. We also use SnapVault to back up data from remote office storage systems in Santa Clara and Atlanta to the R200 in our main location. We typically run those backups on the weekend, and just keep enough local Snapshot™ copies on the remote storage systems to allow end users to do any necessary file recoveries during the week.

## SnapVault Impact

The biggest benefit for Agilent is that SnapVault has significantly cut our backup times. In some cases, we went from 24-hour time spans down to minutes. The whole process of replicating data from the primary storage systems is pretty darn fast. Plus, it's nowhere near as system intensive as the old backup to tape, where you always had to walk through the whole file system, creating a big performance drag. SnapVault knows what blocks have to move in advance and just gets it done. By eliminating the backup burden from our servers and primary storage, we're able to truly use those systems to meet the business needs that we bought them for in the first place.

Here's a list of what I see as the key benefits:

- Backup times cut from many hours to minutes
- Eliminated burden on primary storage and servers
- End users can do their own file restores without assistance from IT, which improves their productivity and ours
- Much greater reliability
- Cost savings as a result of ease of management

Get the details. Read the reports:

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

---

## Information-Optimized Backups

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

Over the past five years, however, disk capacities have grown 20x, with a scant fourfold disk speed performance improvement. This has made it more time consuming moving entire copies of data off of disk during a full backup.. Today, many companies are turning to smarter backup solutions that:

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

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

---

## Sneak Peek: A 25-Page Technical Overview of Protection Manager

Protection Manager is an intuitive policy based management application for NetApp disk-based backup and replication technologies including SnapMirror, SnapVault, and Open Systems SnapVault.

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

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

If we had tried to scale our old backup architecture to meet the needs of all three sites, we would have had to invest well over \$400k in tape hardware and software in order to meet our backup SLAs, and we still would have been struggling with the same old labor-intensive support issues that had always caused us so much pain. By contrast, the SnapVault-based backup solution cost us less than \$300k to deploy, and reduced our support burden by about 500 person hours a year over the old solution.

Overall, it's just been a much happier experience dealing with this new backup environment. We have a lot more confidence, and the user population is a lot happier. The IT team is a lot happier. We're saving money on the deal. So I think it's just been a real win all the way around.


#### **Want to Learn More about SnapVault?**

In a recent Webcast, the Tim Dietrich joined a team of NetApp data protection specialists, including NetApp data protection architect Stephen Manley and services engineering director Kelvin Mayo to drill down on SnapVault technologies, elaborate further on the Agilent implementation, and detail deployment best practices. [Watch the Webcast!](#)

[Comment on this article](#)

terms of the storage infrastructure.

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



TECH ONTAP ARCHIVE


## Experts OnTap: Choosing a Backup Solution How One IT Team Is Evaluating Its Options

*This month Tech OnTap is pleased to debut a new feature: Experts OnTap. Experts OnTap goes far beyond traditional "Ask the Expert" forums. Each month, one Tech OnTap member gets direct access to a panel of cross-functional experts and the opportunity to ask his or her toughest technology questions. Learn how to submit your topic for consideration at the end of this article.*

The customer who participated in this article is a senior systems administrator that has been using NetApp storage for over four years. This article provides insight into his team's decision process as they explore moving from tape backup to disk-to-disk backup and includes:

- Background and challenges
- The IT team's technical questions (with NetApp answers)
- Results of an initial 30-day evaluation
- Next steps and recommendations

The technical questions section is very technical—if you're new to NetApp or SnapVault® software, we suggest you read the [NetApp SnapVault Best Practices Guide \(pdf\)](#) for an overview before tackling the Q&A.



### Company Background and Challenges

This sysadmin works for a company that primarily utilizes FC SAN and iSCSI storage, but three of its many storage systems are dedicated to network-attached storage (NAS) and support home accounts, application data, shared directories, and other data. The IT team uses NetApp SnapMirror® software on these systems to provide disaster recovery. Production volumes are synced every 15 minutes, 24 hours a day, and mirrored first to a local storage system and then to a remote system.

While the team is quite pleased with its disaster recovery solution, backups (and, more specifically, restores) are a growing pain point. Currently the team backs up the three NetApp storage systems directly to tape via NDMP using Tivoli Storage Manager (TSM). Fundamental problems with the backup infrastructure include:

1. **Backup windows are simply too long.** The use of differential backups with TSM is forcing the team to back up about 2TB of data each night. With eight tape drives in service, it has been challenging to complete backups in an eight- or even a 12-hour backup window.
2. **Restore processes are painful.** The IT team receives about one restore request a day, and about 70% of these requests can be restored from local Snapshot™ copies. The other 30%, however, required the IT team to load tapes and search for the best restore option.
3. **Longer backup retention is needed.** Backups are now kept for a maximum of three years, but the IT team would like to be able to restore from any point in time for up to seven years.

The team is very familiar with NetApp technology, and initially thought that a disk-to-disk backup solution using NetApp SnapVault software would be a "no brainer." SnapVault utilizes a baseline copy on secondary storage that mirrors the source volume or qtree. (A qtree is a subvolume that has its own quotas and permissions.) When a backup is scheduled, SnapVault creates a Snapshot copy of the primary qtree and transfers only changed blocks to secondary storage, making it fast and efficient. Restores from disk-based storage instead of tape are also significantly faster and more reliable, and because these backups include only incremental changes, a large number of backups can be retained very efficiently.

### RELATED INFORMATION

- [Sneak Preview: Simplify Management of SnapMirror, SnapVault, and OSSV with Protection Manager \(pdf\)](#)
- [SnapVault Best Practices Guide \(pdf\)](#)
- [Webcast: How Agilent Reduced Its Backup Windows by 98%](#)
- [Forrester Survey Findings: The Shift to Disk-Based Data Protection Solutions](#)

### 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:

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

**[Sneak Peek: A 25-Page Technical Overview of](#)**

However, as the team started to develop their SnapVault evaluation plan, questions arose about potential limitations of the solution and whether implementing SnapVault might impact the existing mirroring schedule.

## SnapVault Technical Q&A

To address these questions, Tech OnTap facilitated a phone meeting between the customer IT team and a small team of NetApp experts.

**Sysadmin:** We have a significant number of individual qtrees in each storage volume and would like to keep our backups around for a significant time. We understand that a single storage volume on a SnapVault secondary system has a limit of 251 Snapshot copies. How can we overcome this limit? What are the pros and cons of each possible solution?

**Darrin Chapman:** You're right, there's a limit on the number of Snapshot copies you can store in a single volume regardless of whether you're talking about Snapshot, SnapVault, or SnapMirror.

The first thing that SnapVault does to help with that is called Snapshot coalescing. Say you have 10 qtrees in your primary volume and you vault them all to a volume on secondary storage. Without coalescing, you would be using up 10 Snapshot copies, and—if you were backing up daily—you'd be out of Snapshot copies in that volume in less than a month. In a nutshell, what Snapshot coalescing does is wait for those 10 individual qtrees to be backed up to the secondary volume and then takes a Snapshot of the secondary volume. That Snapshot then contains all the changes and the 10 underlying Snapshot copies can be released, saving you a lot of snaps.

Still, I've looked at your evaluation test plan and you clearly still have the possibility of hitting the limit. There are two possible solutions. The first is to archive the Snapshot copies off to other volumes, or off to tape for longer-term storage. The second option utilizes FlexClone™ clones and is probably more appropriate for your situation from what I can tell.

Let's say you've got volume A on your primary storage system and you vault qtrees within A to volume B on secondary storage. When you hit 249 or 250 Snapshot copies on volume B, you clone the last Snapshot in volume B with FlexClone to create volume C. Now you've got room for another 256 Snapshot copies. Note that because FlexClone clones only save changed blocks, this process doesn't consume any new storage until changes occur, and the rate of growth is no different than it would be in the original volume B. To make this work you simply rename volume B with a name that indicates it is a full archive and then rename volume C as volume B. The final step would be to do:

```
# snapvault start -r
```

This resynchronizes SnapVault without doing a new level zero baseline.

**Sysadmin:** We understand that each model of a NetApp storage system has a maximum limit of simultaneous snap transfers it can perform. We're planning on limiting simultaneous SnapVault threads to three on our secondary storage system so that we always leave threads available for SnapMirror. How do we schedule our SnapVault jobs so they complete quickly without interfering with SnapMirror?

**Darrin:** The first thing you should know about this is that having a limited number of threads doesn't affect the Snapshot coalescing we talked about in the previous discussion. If a volume has 20 qtrees and a limit of three streams, three jobs will start and the rest will be queued up. When one finishes, another starts; coalescing doesn't happen until all jobs complete.

Also, as you've said, it doesn't matter if you're using SnapVault or SnapMirror, the system limit on simultaneous streams applies to both. If you don't limit your SnapVault streams, SnapMirror streams may queue up waiting for SnapVault streams to complete. Your current plan to limit your SnapVault streams to three by limiting the number of simultaneous streams on your secondary system is a good one. This way you can always be sure that you'll have streams available on your primary systems for SnapMirror on virtually any NetApp storage system.


## Protection Manager

Protection Manager is an intuitive policy based management application for NetApp disk-based backup and replication technologies including SnapMirror, SnapVault, and Open Systems SnapVault.

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

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

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



**Darrin Chapman**

Data Protection Subject  
Matter Expert & Product  
Engineering Manager.

*Darrin has been involved  
with every NetApp best  
practices guide involving  
data protection since  
2002.*

There's a pretty good table on page 28 of the [SnapVault Best Practices Guide \(pdf\)](#) that shows these limits by a NetApp storage system.

Even though you're limiting simultaneous SnapVault streams, you want to take some care in scheduling to ensure that your SnapMirror streams that run every 15 minutes aren't held up by SnapVault. The best way to do this is to limit how many streams are kicked off at any one time. If you've got one volume on your secondary storage system with 50 qtrees, that's 50 transfers that will get kicked off at once. On the other hand, if you spread those qtrees evenly across five volumes and stagger the starts for each volume, you're down to 100 transfers at once and so on.

**Sysadmin: What is the most efficient way to establish the relationships between volumes/qtrees on the primary and on the secondary?**

**Cindy Cui:** NetApp recently released a product called **Protection Manager™** that works with NetApp Operations Manager (formerly known as DataFabric® Manager or DFM) versions 3.5 and later.

Protection Manager lets you manage all your SnapMirror, SnapVault, and Open Systems SnapVault® schedules from one convenient interface. It lets you group data with similar protection requirements into datasets. The physical storage system for the SnapVault destination can be grouped into resource pools. Once you create the datasets and define your resource pool, you can apply preset policies and automate backup and replication tasks. The backup relationships will be created automatically and the destination systems will be provisioned accordingly.



**Cindy Cui**  
NetApp Product  
Manager

All of these things minimize the amount of manual steps one has to go through in setting up a backup environment. When you're managing a lot of relationships, this can save you a lot of manual and/or scripting effort. (See the sidebar to find out more about Protection Manager.)

To take advantage of Protection Manager functionality, you need to upgrade to Operations Manager 3.5.

**Sysadmin: During a manual file recovery from secondary storage, what access controls are carried over if we copy/paste or drag/drop from the secondary to the primary?**

**Darrin:** One of the cool things about SnapVault is that you have multiple means of restoring files and datasets. If you've replicated a Windows® file with SnapVault and that file has various attributes and ACLs that have been set on the host, when you copy or drag it back to that host it maintains the same ACLs and attributes and permissions that it originally had on that host. The only caveat is for hard links. SnapVault can't maintain consistent permissions, attributes, and ACLs for hard links that span multiple file systems. If all your hard links are within the same qtree or within the same file system itself, it's fine. But if hard links span file systems and qtrees, you can potentially lose those hard links.

**Sysadmin: If we have to recover an entire volume, do we restore just the latest Snapshot or would we have to go back to the first Snapshot? What is the process necessary to restore an entire volume or qtree?**

**Jai:** If you have to restore an entire volume, you simply locate the Snapshot you want and restore it. It's a one-step process whether you're using the CLI, Operations Manager, or Protection Manager to do the restore. You don't have to restore multiple backups in sequence the way you do with a tape backup scheme that consists of full backups and incrementals.

**Sysadmin: How do we properly size or forecast the storage required on the secondary vault? Do we use current snaps to calculate this data or are there other means available?**



**Jai Desai**  
Product & Partner  
Engineering

## Sneak Preview



A brand-new technical report on Protection Manager will be published in April.

As a Tech OnTap member, you get a [sneak preview!](#)

**Darrin:** Accurate sizing basically requires a good estimate of what your change delta is during the backup interval for each volume or qtree to be backed up. One simple way to do this is using the snap delta command. If you give this command to Snapshot copies, it will tell you how much the data has changed from one to the next. Of course, the longer the time period you cover, the more accurate your estimate of the rate of change of a particular qtree or volume is likely to be, and you'll want to take into account possible future increases in workload. In your case, since you are doing a 30-day evaluation, that should give you pretty good baseline data to estimate how much storage you'll need when you go into production.

## 30-Day Customer Evaluation: Results

To begin the SnapVault evaluation, the IT team configured backups according to a detailed SnapVault Evaluation Test Plan. They chose a range of volume sizes. Some volumes had millions of files and up to 160 qtrees while others were small with few files and qtrees.

Once configuration was completed, the evaluation really just required that the IT team monitor the ongoing backups. They particularly needed to determine:

- If backups occurred at the scheduled times
- How long backups took to complete
- How much data was transferred

The final item was considered key to sizing the secondary storage they would need if they chose to move forward after the evaluation completed. It would also allow them to decide whether the evaluated backup plan and retention policy were feasible from a cost standpoint. In other words, if the plan would require hundreds of terabytes of storage over the next few years, it might be a deal breaker.

Overall, the team was very pleased with the results they saw during the 30-day period:

- No jobs failed.
- Everything scaled as expected.
- All SnapVault transfers completed within a few hours (even with the limited number of simultaneous streams). Some nights the entire backup took an hour, other nights it took up to 2.5 hours, depending on how much data was created and changed that day. That compared very favorably with the 12+ hours required for the current tape backup scheme.

The only real surprise of the evaluation occurred as a result of changes made on the server chosen to host the NetApp Operations Manager software that was used to manage SnapVault. (For the evaluation, the customer used an older version of Operations Manager that did not support the more fully featured Protection Manager software.)

Because this was an evaluation, the server that was chosen to run Operations Manager wasn't technically a production server. As a result, several unapproved patches were applied to the server during the evaluation by other members of the IT group. One of these included changes to support the new Daylight Savings Time rules. At one point, someone accidentally set the date a month ahead, throwing the SnapVault log information out of whack. However, things went back to normal when the date was reset, and the team learned a valuable lesson. If you are using Operations Manager to control SnapVault, the server on which it runs must be considered a production server and provisions must also be made to allow another server to take over if a problem occurs on the primary server.

## Customer Conclusions and Next Steps

The IT team's goal is to implement a plan that supports fast, disk-based restores from any specific day within 90 days and any specific week up to 7 years beyond those 90 days. This plan would include the following retention and frequency of backups:

- 90 nightly backups at 7pm Monday through Saturday
- 161 weekly backups at 7pm Sunday

The growth trends seen during the evaluation indicate this plan would require about 4TB of secondary storage over the next 3.5 years using a NetApp SnapVault solution. Now that the evaluation is complete, the IT team's next step is to perform a cost

analysis and compare it to the current TSM backup environment, then present the conclusions to senior management.

Based on the results of the 30-day evaluation, the team feels confident recommending a disk-to-disk backup solution based on NetApp SnapVault software and anticipate using Protection Manager once they upgrade to Operations Manager 3.5. In fact, the team is already talking about undertaking a similar evaluation of Open Systems SnapVault to accommodate direct-attached storage.

The team had two key recommendations for anyone considering a move from tape to disk-to-disk:

1. **Identify your exact requirements.** Know what to expect from the environment; otherwise your scope can get completely out of control. If you have established SLAs to work to, that's even better. We struggled with things such as how many different iterations of the data we were going to keep, how much of the data we were going to keep, how long we were going to keep the data, and how frequently we were going to do backups. Figuring this out was the longest part of creating our test plan.
2. **Do a long and careful evaluation.** In any environment there's going to be something that you didn't expect or didn't plan for. In addition to our adventure with Operations Manager, we also discovered that creating the initial baseline for our large volumes with 20 to 30 million files in them takes a long time. That's the kind of information you'll be able to use in the planning process as you move forward.

### Experts OnTap April Topic: SnapMirror Environments

---

#### Submit your Toughest Questions for Experts OnTap!

Have a challenge your team is brainstorming or thorny questions that remain unanswered? Submit your challenge to [techontap@netapp.com](mailto:techontap@netapp.com). Please include your key issues, the types of questions you need answered, and a description of your environment.

Each month, Tech OnTap will select one member to conference with a panel of multifunctional NetApp panel product and deployment experts. Highlights of the resulting discussion will be featured in future Tech OnTap articles.

---

[Comment on this article](#)

TECH ONTAP ARCHIVE



**Robert Quimbey**

Microsoft Alliance Engineer, NetApp

Robert recently joined NetApp after 8 years as a member of the Microsoft® Exchange product team, where he was responsible for storage and high availability. While at Microsoft, Robert authored [Optimizing Storage for Exchange Server 2003](#) and [Configuring, Validating, and Monitoring your Exchange 2007 Storage](#). His role at NetApp involves creating best practices for Exchange 2007 and working on various projects, including an in-depth analysis of the impact of Microsoft's Clustered Continuous Replication (CCR) with I/O.

## How Exchange Server 2007 Impacts Backup, Recovery and DR Strategies

By Robert Quimbey

The introduction of Microsoft Exchange Server 2007 is raising a lot of questions in the minds of Tech OnTap readers. This was made clear in a recent TechTalk Exchange live chat session, where questions about 2007 were prominent. From a storage perspective, there are several new features in Exchange Server 2007 that will affect some storage environments. The most significant ones are:

- Local Continuous Replication (LCR)
- Cluster Continuous Replication (CCR)
- Changes in backend I/O traffic to storage

The introduction of these features raises a number of questions:

- How will Exchange Server 2007 affect my backup strategy?
- How will Exchange Server 2007 affect my disaster recovery (DR) strategy?
- Which storage networking protocol should I use? FCP or iSCSI?

This article answers the questions listed above. To learn more about the other new features of Exchange Server 2007, check out the [Microsoft Exchange Server 2007 Web site](#).

### Question #1: Does LCR Replace Backup? Answer: Definitely Not!

Local Continuous Replication replicates Exchange databases to another set of disks on the same physical server. The question that this immediately raises is, "Does that mean I don't need to do backups any more?" The resounding answer to that question is "No!" Regular backups must still be performed.

The purpose of LCR is not backup; it is high availability (HA). LCR creates a copy of the Exchange database, which provides two sets of the same data. The LCR copy is just slightly behind the primary. Data is written to the primary Exchange log file first, and then, slightly later, that log file is replicated to the LCR copy or target. The trigger for replication is the closing of the log file. The log file is 1MB in size, so after 1MB is written to the primary, it is replicated to the LCR target, and then played into the target database.

If something goes wrong with the primary data store, there is another copy available for use (although slightly behind), but *that copy isn't a replacement for backups!* Suppose



#### RELATED INFORMATION

- [NetApp Solutions for Exchange Environments](#)
- [NetApp/Microsoft Exchange 2003 Pilot Involving the FAS980 \(pdf\)](#)
- [Webcast: Simplifying and Scaling Storage for Exchange Server 2007](#)

### NetApp and Exchange Server 2007

NetApp is a long-term Microsoft partner and has been involved with Exchange Server 2007 for some time. NetApp is ...

- Involved in the Exchange Server 2007 Technology Adoption Program.
- A member of the Rapid Deployment Program and a global launch partner for Exchange Server 2007.
- The only storage provider with equipment in all 12 Microsoft Technology Centers.

### Sample Configuration: Supporting 14,000 Exchange Users with the FAS3040

This report describes the testing of a 14,000-user configuration for Exchange Server 2003 with the new NetApp FAS3040 storage system, based on the Microsoft Exchange Solution Reviewed Program (ESRP) Storage program.

The report includes:

there is a logical corruption in the primary database. As soon as the log fills, the secondary database copy becomes corrupted as well. Similarly, if something is deleted from the primary database, a short time later it is deleted from the secondary. The deleted item is stored in the database (dumpster) by default for 14 days before being expunged during online maintenance. To recover deleted data more than 2 weeks old, a backup is required.

Be aware that LCR does not address database verification. Only incoming log files are inspected. Performing regular backups with verification—such as those performed with NetApp [SnapManager® for Exchange](#)—can verify that the production database is healthy. This ensures the success of the backup and the ability to use a consistent database (and log file) image for restoration when required.

## Question #2: Does CCR Replace My DR Strategy? Answer: It Might, But There are Tradeoffs.

As previously discussed, since LCR is focused on storage resiliency, there is no benefit if the server goes down. Instead, if the primary storage fails, the failover is not automated. This means that the administrator has to manually fail over to the target copy. That brings us to Cluster Continuous Replication, which provides Exchange server resiliency. Instead of keeping a copy of the Exchange database on the same machine, the copy is kept on another machine.

The second machine that stores the Exchange database copy is deployed as part of a two-node Microsoft Cluster (MSCS) that includes the original machine. (Right now, only a two-node cluster is supported.) To be in an MSCS cluster, network latency must be below 500ms (heartbeat) to ensure that the cluster nodes can communicate with each other. Next, with CCR, latency and throughput must keep up with the log generation. This is bandwidth-dependent, so the infrastructure may support a distance of anywhere from 1 to 200 miles. Typically, this translates to 60 to 100 miles.

If the primary cluster node fails, the system automatically uses the secondary node running against the replica of the Exchange database. In a controlled failover, where the primary node is still available, all log files are copied to the target, and no data is at risk. In the event of a catastrophic failure of the primary node, CCR (not LCR) attempts to recover all mail sent through transport from the hub transport server, which may not have been replicated at the time of failure.

Messages that had been applied to the primary Exchange database, whose log files had not been replicated, are still potentially available from the [transport dumpster](#) on the hub transport servers in the infrastructure. Assuming that the hub transport server was configured properly and didn't also fail when the primary CCR node went down (think site disaster), this data can be recovered. (Note: More I/O per message will be required if transport dumpster is enabled on the Hub Transport server; see [Capacity and Transaction IO Requirements for Exchange 2007 Edge Transport and Hub Transport servers](#)).

The storage infrastructure question raised by CCR is, "Does CCR eliminate the need for disaster recovery mechanisms such as NetApp SnapMirror®?" The answer is ? maybe. It depends on a number of things:

- **Recovery objectives.** How much loss of e-mail data is acceptable? Remember that SnapMirror replications are triggered as a result of a SnapManager for Exchange backup, while CCR replication occurs as a result of a filled 1MB log file.
- **Ability to support the extra I/O from CCR.** CCR imposes additional I/O overhead to that of normal Exchange 2007 I/O. Microsoft recommends that CCR clusters use isolated storage from other servers; see [Continuous Replication LUNs](#). In an isolated environment, the additional I/O overhead of the target LUN would keep up with the source and would not require additional disk performance, although at higher latencies. CCR clusters that use shared storage (shared between clusters) require that the storage (both source and target) be over-provisioned to handle 100% more I/O than the source LUN for CCR deployments. For example, if the source LUN requires 1,000 I/Os, in a shared storage environment the source and target LUNs would need to be provisioned with storage to handle 2000 I/Os each.
- **Distance to the disaster recovery site.** Remember that Microsoft Geographically Dispersed Clusters require network latency below 500ms, and Exchange 2007 requires a disk latency below 20ms, which typically restricts

- Simulated Exchange configuration
- Hardware/software used for the test
- Disk configuration, including mailbox store disks and transactional log disks
- Streaming backups
- Replications
- Overall best practices

Full results from the [Microsoft ESRP Test](#). ([pdf](#))

distances less than 100 miles between systems in the cluster. If the DR site is more than 100 miles away, then a solution other than or in addition to CCR is required.

Some NetApp customers are considering running both CCR and SnapMirror. These IT teams plan to use CCR in the local area for high availability as well as offloading the backup and verification activities, and SnapMirror to replicate Exchange data to a remote disaster recovery site. (An unfortunate lesson of Hurricane Katrina is that it's probably a good idea to have the DR site more than 100 miles away if at all possible.)

The key takeaway is to deploy those infrastructure components that make the most sense to meet the availability needs. NetApp software products, including SnapManager for Exchange 4.0 and Single Mailbox Recovery 4.2, include support for Exchange Server 2007 today and integrate with both LCR and CCR to make this support easier. For example, running CCR to replicate Exchange data to a secondary node and running SnapManager for Exchange on the secondary node to do the backup and verification tasks ensure that these tasks don't affect the primary node.

### **Question #3: Should I Use FC SAN or iSCSI for Exchange Server 2007?**

**Answer: It Depends on the Infrastructure, but iSCSI is Usually Appropriate.**

The final question that a lot of IT teams are asking involves which protocol is best suited to supporting an Exchange infrastructure. This isn't strictly limited to Exchange Server 2007, but some of the changes in Exchange Server 2007 should increase the comfort level with iSCSI.

The first thing that concerns some people about iSCSI is that it typically runs at only 1GB/sec, while Fibre Channel runs at 2 or 4GB/sec. Further, Fibre Channel networks are predictable with regard to bandwidth consumption and network access. The typical Ethernet network deployed for iSCSI is not predictable in this way due to collision avoidance. There's less chance to send data on a busy network, but the way people normally implement these networks—with a dedicated connection—it's not really an issue.

Both Microsoft and NetApp best practices for deploying iSCSI involve using dedicated connections.

The second thing about iSCSI versus Fibre Channel is that on FCP a lot of the protocol handling is done on specialized hardware, while most people running iSCSI tend to use the Microsoft iSCSI software initiator. This processing consumes CPU cycles on the server to process the protocol, and those cycles may be important for other activities on the server.

There are, however, iSCSI HBAs that offload iSCSI protocol processing from server CPUs in a manner similar to Fibre Channel HBAs. People usually choose to use these for one of two reasons:

- **Booting over iSCSI typically requires an iSCSI HBA.** (The network protocols need to be resident so the server can talk to storage to get the server boot image.) The iSCSI HBA looks to the server just like any other disk controller.
- **The server CPU cycles that iSCSI protocol processing consumes are needed for something else.** An iSCSI HBA doesn't necessarily provide better network performance, but it does give back CPU cycles for running more Exchange users. With newer, faster servers with 64-bit CPUs, multiple cores, and back-end buses, there is typically more than enough CPU power to do both jobs, so this reason is becoming rarer and rarer. Keep in mind, an iSCSI HBA is needed, not a TOE (TCP/IP Offload Engine) card, which offloads only TCP/IP processing and still leaves iSCSI processing to the host.

Still, some people voice concerns about using iSCSI because it does not provide the raw performance of Fibre Channel. For small-block I/O applications, however, the critical factor is IOPS with low latency, not bandwidth. Nick Triantos explored this topic in detail in [Choosing between iSCSI and Fibre Channel SANs](#).

Testing has shown that bandwidth doesn't really matter with transactional I/O in Exchange because a single machine can't generate enough load to flood a 1GbE pipe, and that pipe is dedicated between the server and storage. It's possible that if there

are a lot of servers sharing a single 1GbE/sec connection to storage (possibly a bad idea), the connection could become saturated, so that could become a bottleneck. However, even in this case it is possible to add more 1GbE connections to increase the bandwidth to the storage device.

Additionally, with Exchange Server 2007, the actual amount of disk I/O is significantly decreased over prior versions of Exchange. This change in I/O is due to a couple of factors: on 64-bit hardware, additional memory is available to use for database caching, thus reducing I/O; and changes to the Exchange database and the internal I/O activities of Exchange further reduce I/O.

A 1GbE pipe can run into trouble with non-transactional I/O. Streaming online backups, VSS integrity checks, and offline repair or defrags all depend on disk throughput. For these activities, if 1GbE does not meet required SLAs (for backup, restore, or repair), more HBAs or network cards should be used with MPIO to increase throughput. Microsoft v2 iSCSI initiator supports MPIO and can be found [here](#).

NetApp, of course, supports both protocols so feel free to choose whichever option makes the most sense in your storage infrastructure.

## **Exchange Server 2007 and Your Storage Infrastructure.**

The main thing to remember about Exchange Server 2007 is that although LCR and CCR are valuable additions to the overall toolset and can greatly increase Exchange availability, there is still an absolute requirement for Exchange backups. In addition to basic data protection, many companies have internal requirements to maintain off-site backups, and some government agencies and legal regulations require it as well. Exchange database verification is another critical component to a healthy messaging environment that occurs during both VSS and streaming online backup.

Additionally, although the new features associated with Exchange Server 2007 help provide high availability, HA alone can't ensure that Exchange is always up. A complete disaster recovery plan is still required and, depending on requirements, a mirroring solution may be required. Ultimately, these decisions impacting the length of downtime, will be influenced by the level of acceptable risk, and the amount of money available for the solution.

---

[Comment on this article](#)

TECH ONTAP ARCHIVE



**STEVE LUSNIA**

V-Series Solution Architect, NetApp

Steven Lusnia joined the NetApp Professional Services team over six year ago and in mid-2003 started working with large NetApp customers who wanted to leverage NetApp features in their traditional SAN environments. This was the beginning of the V-Series product line. In early 2006, Steve accepted a new role with engineering and product management to leverage his field implementation experience to help drive the evolution of the V-Series product line.

## Simplifying Data Management for EMC, HP, IBM, Fujitsu, & Hitachi Environments

By Steve Lusnia

The NetApp V-Series product line is a storage virtualization solution that enables IT teams to take advantage of the unique features and functionality NetApp offers while leveraging existing SAN investments. The V-Series product line is unique in its ability to unify block and file access to storage array products from Hitachi, HP, IBM, Fujitsu, and—most recently—EMC.

This article provides a quick introduction to the V-Series product line plus five real-world examples of how IT teams have deployed V-Series systems to enhance their traditional SAN environments.

### The History of the V-Series Product Line

The V-Series project started in 2003 when an aspiring group of NetApp engineers extended a FAS controller to storage LUNs from a storage array connected through a SAN. Though the "back end" was different—the V-Series system aggregates LUNs that are RAID protected by the storage arrays—the front end brought all of the rich file system access features of Data ONTAP® to the combined solution.

This virtualization solution was originally introduced in 2003 as a Hitachi-branded gFiler™ appliance and was eventually productized directly by NetApp as the V-Series product line.

While the initial gFiler appliances offered only network-attached storage (NAS) functionality for HDS storage arrays, the V-Series product line was expanded to support Fibre Channel and iSCSI protocols for arrays from all major storage providers.

### V-Series Benefits

Deploying V-Series systems in SAN host environments brings the power of Data ONTAP to traditional SANs. The V-Series system functions as both a NAS gateway and a SAN virtualization engine that provides a rich set of storage features and unifies block and file access into a single platform.

V-Series systems are identical to NetApp FAS storage systems in their support for:

- A complete suite of front-end protocols, including CIFS, NFS, iSCSI, and FCP
- Space-efficient NetApp Snapshot™ functionality that writes only changed blocks to disk
- NetApp replication software, including SnapMirror®, SnapVault®, SyncMirror®, and NDMP
- All host-side attach software, including SnapDrive®, SnapManager® for Oracle®,



V-Series Network Diagram  
[Click to Enlarge](#)

### RELATED INFORMATION

- ❖ [V-Series Product Overview](#)
- ❖ [Best Practices for V-Series \(pdf\)](#)
- ❖ [Simplified DR and Cloning for Non-NetApp Storage Environments](#)
- ❖ [Case Study: Cadence Design Systems \(pdf\)](#)

### The Core NetApp DNA

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

These basic building blocks of NetApp technology support:

- Multiprotocol environments (NFS, CIFS, FC, iSCSI, and so on)
- Clustered failover, mirroring, and disk-to-disk backup
- RAID-DP and other software-based resiliency features
- The near-instantaneous creation of writable clones
- Block de-duplication using A-SIS

Learn about [The Core NetApp DNA](#).

### 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 disk speed performance

The benefits for traditional SAN environments are significant. Due to the unique NetApp approach to snapshots, for example, Snapshot copies can be created very frequently with virtually no performance impact. NetApp also eliminates the redundancy inherent in traditional backups, shortens backup windows by eliminating the need to identify changed data, and provides 100% backup validation. To learn more, check out [The Core NetApp DNA](#) and [The Value of Information for Backups](#).

In addition, the V-Series product line enables the use of NetApp host-based manageability software, which not only reduces backup and recovery time when compared to traditional storage array solutions, but also allows application owners to recover their own data with minimal storage administrator involvement.

### How is V-Series Different from Using NetApp FAS Storage?

V-Series and FAS products use the same hardware controller and run the same Data ONTAP operating system. The key difference between a V-Series system front ending a storage array and a FAS system with NetApp disks is that the V-Series controller no longer runs RAID 4 or RAID-DP™. Instead, the V-Series system offloads the raid protection to the storage array. V-Series storage pools are large RAID 0 stripes of back-end LUNs.

### Supported Storage Arrays

The V-Series product line supports storage subsystems from Hitachi, HP, IBM, Fujitsu, and now EMC. Customer demand for unified storage and simplified management in existing EMC environments drove NetApp to invest in supporting the EMC CX array. A V-Series/EMC system was first demonstrated at SNW 18 months ago, and today EMC CX arrays have been fully qualified by NetApp and are now listed on the public V-Series support matrix.

A detailed and current list of [supported arrays](#) can be found on the NOW™ site. The support matrix currently includes IBM DS8000 and DS4000 families, Fujitsu Eternus, HDS Tagmstore, AMS, Thunder, Lightning, and the rebranded equivalent arrays from HP. NetApp has committed to continuing to expand the list of arrays supported behind the V-Series product line and to provide additional virtualization functionality.

### Real-World Examples

The following are a series of examples of how real-world customers leverage the V-Series product line with a variety of traditional storage arrays. Although each example references a specific vendor's storage array, these approaches could easily be implemented with any supported array.

#### V-Series and EMC CX Arrays: Centralized Backup for Over 100 Remote Sites

A leading retailer had standardized its central data center on EMC CX storage arrays, but wanted to implement a centralized backup strategy to support over 100 remote offices running Windows® file sharing and home directories. When NetApp first met this customer, each remote site was backed up to local tape, with support provided by either part-time local administrators or corporate-based administrators.

By standardizing on a unified data management system and leveraging NetApp replication technology, the IT team realized it would be able to leverage its existing SAN investment to centralize backups across its remote locations.

NetApp solution:

- Central data center: V-Series system fronting EMC CX arrays
- Remote locations: NetApp FAS270s for file sharing and home directories
- Replicate FAS270s to the central V-Series system using NetApp SnapVault software
- Centralize tape backups at the central data center and eliminate remote tape



Centralized Backup Environment  
[Click to Enlarge](#)

Benefits:

- Reduced hardware and admin costs due to the elimination of remote tape
- Rapid online file recovery both locally and from the central data center
- Tape restores kept to an absolute minimum
- Standardized backup schedules and processes company-wide

#### V-Series and HP XP12000 Arrays: Simplifying Backup and Application Testing

improvement. This has made it more time consuming moving entire copies of data off of disk during a full backup. Today, many companies are turning to smarter backup solutions that:


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

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

### 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.](#)

Another NetApp customer using a V-Series system supports its SQL Server environment with storage from two HP XP12000 arrays. The IT team wanted to simplify storage provisioning and implement a more flexible storage model for its SQL Server environment.

Introducing the V-Series product line into this traditional SAN environment and creating a virtualized storage environment enabled the IT team to do the following:

- Create a storage pool that is synchronously mirrored between the two HP XP12000 arrays. An entire array can be offlined without any impact on SQL Server database availability. A separate pool of storage that is not synchronously mirrored to another array is also available for other applications.
- Eliminate the need to take SQL Server databases offline for backups. The IT team uses NetApp SnapManager for SQL Server to streamline backups. Transferring these backups to tape is performed from the V-Series controller and no longer a host function.
- Significantly improve restore times. Restoring data from online disk-based storage is much faster and more reliable than the previous restores from tape.
- Reduce the cycle time for testing and development. Databases can be cloned in minutes rather than hours using NetApp FlexClone™ technology.
- Simplify storage provisioning by expanding/contracting volume sizes on the fly.

#### **V-Series and HDS Lightning Arrays: Rapid Recovery with NetApp Snapshot Copies**

Major disasters rarely happen. But when they do, NetApp features enable dramatically reduced recovery times. Data center hardware often needs to be refreshed at end of lease or upgraded when larger application servers are required. For example, during a host replacement in an HDS environment where V-Series systems were deployed, a small server was replaced with a larger server to handle an unexpected increase of processing needs. The host data was provisioned from a V-Series system serving FCP LUNs. Due to a host administrator error, the LUNs were accidentally destroyed as presented from the V-Series.

In this situation, the most recent backup was five days old, and the application was a critical path of a larger corporate project. Prior to the introduction of the V-Series product line, the only option for recovery would have involved going back to tape and five days of log replay. Using NetApp Snapshot copies, however, the recovery took only a few minutes, and the company's Oracle Databases were back online within the hour.

#### **V-Series and IBM DS8000/DS4000 Arrays: Outsourced Disaster Recovery**

Another example of how V-Series systems are used can be seen at a company that provides outsourced disaster recovery (DR) services. This provider has deployed IBM DS8000 and DS4000 storage arrays as part of its IT infrastructure.

Introducing a V-Series system at the disaster recovery site has expanded the company's range of customers. Now an IT department running either FAS storage with NetApp disks or a V-Series system with any storage array can replicate to V-Series systems at the outsourced disaster recovery site.

#### **V-Series and HDS Lightning Arrays: Business Continuity for Oracle**

A final example involves an existing SAN supporting Oracle applications running on Solaris™ hosts with storage served from an HDS Lightning array and local storage. The company needed to replicate all of its data to a different state in the event of a complete site outage, but did not want to duplicate its existing HDS environment at a disaster recovery site.

Instead, the IT team inserted a V-Series system into its environment and migrated its SAN-based Oracle application and Windows shares to the V-Series environment. This enabled replication anywhere on the customer's Ethernet network. Data from the V-Series front end on the HDS Lightning solution at the primary location was mirrored using NetApp SnapMirror software to an IBM-branded NetApp FAS system at the remote site. All planning, implementation, migration, and testing were completed using NetApp Professional Services.

Using a non-HDS replication strategy enabled the IT team to keep costs down and still achieve their disaster recovery objectives.

---

▪ **Comment on this Article**

TECH ONTAP ARCHIVE



**RAJEEV K. KARAMCHEDU**

Computer Systems Manager at a Leading Bioinformatics Research Institute

For the last seven years, Rajeev has managed the UNIX® and storage environments of a leading research institute. Rajeev and his team are constantly evaluating new storage and computational technologies in order to facilitate the Institute's mission of deciphering and analyzing genomes. His team's latest projects include NetApp GX evaluation, FAS6070 platform testing, and working with other parallel file system solutions. The team also coordinates a local NetApp User Group. During his "spare" time, Rajeev contributes to a few open source projects and lives to [blog](#) about it.

## Featured Tool: filestat shell script

*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:** Filestat was created by the author.

**How it works:** Filestat uses the SNMP MIB to query a storage system and returns a simple text output to the user.

**Why it's cool:** For optimal storage performance, it is imperative that one understands the file system usage and its layout. Variables such as number of files in a volume and the average size of each file are very useful in determining potential bottlenecks. The Data ONTAP® native commands that are available in Data ONTAP require console access. In contrast, any user with SNMP read capabilities can run this shell script to query the storage for this critical information. This means that end users can get the information they need on the volumes and adjust their usage patterns accordingly. The information that filestat supplies on the number of files and average file size in each volume is also very useful for fine-tuning backups, replication, trend analyses, etc. I'm also planning to feed the information that filestat gathers into MRTG and Cacti tools so we can do historical trending.

**How Rajeev uses filestat:** In our data center, we currently have two NetApp FAS6070 clusters with 50TB+ of usable storage, representing about half of our total storage environment. These systems host our scientific data repositories for individual research needs and for GRID computational needs. The NetApp appliances transfer data to the GRID clients at about 3Gbps. Our current GRID capacity is over 600 compute slots, with total memory amounting to 1TB. As we did during the last few years, we expect to double our capacities and demands next year as well.

The Institute has a UNIX-savvy user base. Power users who tend to write complicated and large-scale computational programs that utilize the GRID heavily also want to know, at regular intervals, the appliances' various performance statistics so they can adjust their workflows as necessary. Filestat, along with other similar tools, provides such information. The result has been fewer help desk requests and substantial time savings for both users and administrators.

Having ready access to this type of information has also had other benefits in our IT group. In the past, not knowing the number and size of files that were accumulating in a volume over time significantly affected our backups. A backup of a volume with a few large files goes faster than a backup of a volume with many small files. Now we can easily identify the volumes in which we expect problems and deal with them. Instead of backups over NFS, we use NDMP backups for those volumes and get better performance. We've also changed the way we've been laying out files and volumes



### RELATED INFORMATION

- ❖ [NOW™ Tools Discussion Forum \(NOW user login required\)](#)
- ❖ [Filestat script](#)
- ❖ Previously Featured Tools:
  - ❖ [Toasters Discussion Group](#)
  - ❖ [Swatch and Kiwi](#)
  - ❖ [Firefox](#)
  - ❖ [Cacti](#)
  - ❖ [MRTG for Filers](#)
  - ❖ [Nagios](#)
  - ❖ [SIO](#)
  - ❖ [Toasterview](#)

### Sample Filestat Output

```
# filestat -f toaster -v pvol1

capacity (GB):      681.12
usage (GB):         614.26
files (m):          17.96
avg file size      35.84
(KB):
```

based on the output. We could see that what we were doing clearly wasn't working and made changes to better accommodate our work. As a storage admin, you need to know this information. Once you have it, you can't live without it!

**Caveats:** Filestat requires SNMP read access on a NetApp storage system in order to execute.

▪ [Filestat](#) script

---

### What do YOU think?

Ask questions, share feedback, and nominate your own favorite tools on the [NOW tools forum](#) (*NOW user login required*).

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

▪ **Nominate a Tool**

