

December 2005

HIGHLIGHTS

- **Lessons Learned from Katrina**
(Tech OnTap Exclusive Paper)
- **Five Exchange Issues Solved**
- **Technology behind Simple Storage**
- **Advice from Yahoo!'s Sys Admin**

Transforming SAP's Environment from "Unwieldy" to "Adaptive"

John Wenstrup, Director, Technology Strategy
Mercer Management Consulting

SAP cut storage management over 80%. Learn how SAP's IT team went from screening violent support calls to looking like heroes.

[More ::](#)

BLOGGING WITH DAVE

Dave Hitz, NetApp Founder and EVP

"Three times now we have made quantum leaps to our product line ..."

- **Dave's Blog**

DRILL DOWN

- **Best Practices for Oracle on NetApp**
From aggregate setup to hot backup scripts, this 38-page guide shares lessons from thousands of deployments.
- **Lessons Learned from Katrina**
Step-by-step advice on preparing your NetApp environment for a disaster.
- **Uncompromised Security Initiative**
Podcast (and transcripts) on simplifying data security.

ASK THE SYS ADMIN

Blake Gollhofer, System Administrator, Yahoo!

How to add a disk to the first RAID group in a volume, monitor client workloads, enable SSH with public keys, and use SIO scripts.

- **Got Questions?**

TIPS FROM THE TRENCHES

Plan a Better Backup Strategy

Tim McCarthy, President, Sine-rgy


Veteran Visa CIO, CitiGroup executive, and First Data executive Tim McCarthy tells it like it is: At the end of the day, you need to be confident that both your data and your vendor will survive.

[More ::](#)

Solving the Top Five Exchange Issues Experienced by NetApp Customers

John Fullbright, Professional Services Consultant, NetApp

Tips on dealing with undersized storage, write performance issues, incompatible drivers and patches, and content conversion.



[More ::](#)

Mapping Out a Strategy to Minimize DAS

Nick Triantos, Global SAN Systems Engineer, NetApp

Second in a two-part series: A top NetApp SE shares his checklist for implementing efficient, consolidated storage environments.

[More ::](#)

ENGINEERING TALK

Is NetApp Storage Really Simpler?

Nawaf Bitar, VP of File Systems Engineering, NetApp

Is "storage simplified" just a nice tagline or is there something inherently less complex about NetApp storage? Learn what makes the NetApp architecture unique.

[More ::](#)

Don't Miss: Data ONTAP® GX ::



JOHN WENSTRUP

Director, Technology Strategy, Mercer Management Consulting

During his 12-year tenure as a high-tech strategy consultant to the computing, storage, software, and IT services industries, John Wenstrup has conducted interviews with over 1,400 CIOs and technology decision makers. In this Q&A, John shares insights to help storage managers and systems administrators work smarter based on his findings from interviews with IT executives, system administrators, and other contacts at SAP. These interviews were conducted as part of a recent Mercer report on SAP's move from internal server-based storage to a more "adaptive enterprise" running NetApp networked storage. Per the firm's original agreement with NetApp, [Mercer Management Consulting](#) had complete authority and independence in conducting interviews and compiling the report's findings. Furthermore, all results were verified by the SAP Hosting IT leadership team.

Back from the Edge: Transforming SAP's Mission-Critical Production Environment from "Unwieldy" to "Adaptive"

Q&A with John Wenstrup

Q: Tell us a little about SAP. What was SAP's computing and storage environment like prior to the company's move to NetApp storage?

John: SAP has an extensive software development organization that drives much of SAP's success in the creation of powerful, enterprise-class software. The SAP development organization involves 10,000 developers around the world—from India to most of Europe and Germany, and the United States—all of whom access a complex network of SAP instances to develop new products. The resulting software is in use by many of the world's most complex manufacturing, retail, and financial services environments.

The IT infrastructure supporting these development efforts is one of SAP's most mission-critical areas and is incredibly complex. In 2001 over 1,000 database servers were in use, running on a mix of Windows®, Solaris™, AIX, and HP-UX operating systems and managing over 1,000 unique instances of SAP (referred to as "systems") under development. These had to support several database applications, including MaxDB, SQL Server, and DB2.

Responsibility for managing the infrastructure associated with the SAP global development organization fell to the Product Systems Services (PSS) group within SAP's Hosting organization. SAP had made an early decision to use internal, server-based storage to support its instances of SAP under development. Such an approach was acceptable when SAP only had a small number of instances of SAP to support, but over time that number exploded with new and customized versions of the software. By the time the PSS group was managing more than 1,000 database servers running over 1,000 unique instances of SAP software, the environment had become incredibly unwieldy.

Q: You learned that the SAP Product Systems Services IT Group had reached a "breaking point" by 2001. What would you say were the group's biggest storage management challenges?

John: The PSS server and storage administrators spent about 30% of their time responding to developers who were either running out of space on servers or wanted to migrate a system from one server to another to improve system performance. In effect, the server and storage team was spending a large share of its time reshuffling deck chairs.

RELATED INFORMATION

- [NetApp Storage in SAP's Mission-Critical Production Environment](#)
- [IT Decision-Maker Perspectives on EMC and NetApp Storage in Enterprise Database Environments](#)
- [NetApp Secret Sauce](#)
- [A Roadmap for Networking Storage](#)
- [Mercer Management Consulting](#)

Building an Adaptive Enterprise

Members of SAP's internal IT organization identified three key requirements to transform into an adaptive enterprise:

- Virtualized IT assets
- Flexible and interchangeable modules
- The ability to dynamically alter the IT environment to support changes in the business

Using NetApp storage, SAP was able to fulfill these goals. Other benefits include:

- Increased server and storage utilization. The number of servers required per TB fell 40%.
- Simplified management. The time required to manage storage was cut by over 80%.
- Higher availability. Scheduled and unscheduled downtimes were reduced by up to 90%.

Q. What was the impetus for change at SAP?

John: SAP storage administrators were familiar with NetApp and liked the level of redundancy and data protection they knew NetApp could provide. The administrators teamed with one of SAP's IT executives to do an internal TCO study. They brought this to the attention of the senior management team and essentially told them, "We'll save 26% if we invest in a NetApp NAS system rather than more internal server-based storage."

Management agreed to a NetApp trial. The cost savings achieved in this trial were so impressive and the management improvements so great that they quickly moved to a wider deployment of NetApp. In essence, SAP Hosting decided to use NetApp storage for all new storage needs. The storage administrators and the midlevel executives ended up looking like heroes to senior executives for having driven this large-scale infrastructure change.

Q: What does SAP's environment look like today?

John: SAP now relies on several pairs of clustered NetApp storage systems for the most mission-critical part of its development environment. Each storage system has a hundred or more servers that access data. As a result of moving to NetApp, SAP was able to standardize its server environment using low-cost commodity servers with just enough internal storage to power the database application and perform basic activities.

The impact of deploying NetApp storage has been significant. Due to higher utilization, total storage requirements fell nearly a third, and the time required to manage the storage environment was cut by over 80%. Both scheduled and unscheduled downtime fell significantly, and overall direct internal IT costs were cut in half.

A Before and After Comparison of NetApp vs. Internal Server Storage at SAP's Product Systems Services Group

	Internal Storage	NetApp Storage
Total Storage:	272TB	192TB
Server Intensity:	7.3 servers/TB	4.4 servers/TB
FTE to Manage:	2.0 FTE	0.25 FTE
Data Migration:	Up to 5 hours downtime	30 minutes downtime (often nondisruptive)
Other Benefits of the NetApp Environment:	• Resources required to manage each terabyte cut ~50% • 4x faster recoveries from downtime incidents • Significantly fewer failures that impact developers • 50% lower data center costs (inc. power, space, environment)	

Source: NetApp Storage in SAP's Mission-Critical Production Environment: Benefits and Total Cost Comparison. [Mercer Management Consulting report](#).

Q: After migrating many of SAP's mission-critical servers to NetApp, you found the NetApp environment required only .25 of an FTE's time to manage nearly 200TB of storage. Where did these time-savings come from?

John: Because SAP doesn't run out of storage space very often, they are not migrating nearly as much data as they did before. Now system administrators just point and click when they need to migrate or allocate more space to a system. Tasks such as data migration are anywhere from eight to 20 times faster than under the previous architecture. The resources saved are now freed up to work on higher level activities that are more strategic and of more benefit to the company.

Q. In your opinion, how much of the benefit that SAP achieved from this deployment was due to the shift to networked storage and how much was specific to NetApp storage?

- Reduced TCO. Cost on a per terabyte basis was reduced by 36% over that of an internal server-based storage environment
- Reduced management cost. Ongoing, direct internal IT costs were lowered by 49%.

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

NetApp vs. EMC: IT Decision-Maker Perspectives

After interviewing storage managers and database administrators at 15 to 20 companies, Mercer Management Consulting identified key differentiators between IT database environments running NetApp vs. EMC storage.

Findings include:

- NetApp implementations generally take 1-1.5 FTE, while EMC implementations take up to 5-6x longer and often require a third-party services company
- NetApp solutions require less training and 3-5 times fewer resources to administer, monitor, and manage
- Database storage administrators using NetApp storage can recover 3-4 times faster from a database crash or corruption

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

John: Frankly, I do not believe SAP could have done a NAS-based consolidation with any vendor other than NetApp. As SAP IT executives pointed out in our discussions, no other vendor had a solution robust enough to make the executive team confident. By “robust,” I mean data protection capabilities like Snapshot™, MetroCluster capabilities for failover, and most of the other functions associated with Data ONTAP®. While difficult to pinpoint precisely, I would anticipate this combination of features in the NetApp implementation was probably responsible for somewhere around half of the actual cost savings SAP experienced.

As an example, SAP experienced significantly lower recovery times and lower incidents of failure, which it directly attributes to the use of NetApp Snapshot functionality. The ability to quickly point and click in data migration was also unbelievably helpful.

More broadly, this study is consistent with other recent research, which suggests NetApp storage is significantly easier to manage than competitor storage. Last year, for example, Mercer conducted a study that found large-scale database environments using NetApp storage to be significantly easier to manage than those based on EMC’s CLARiiON and Symmetrix systems.

Comparison of NetApp and EMC Storage Solutions for a 2TB Database Environment			
	NetApp Storage	EMC CLARiiON	EMC Symmetrix
Full Time Equivalent Staff to Implement	1-1.5 FTE	3-4.5 FTE	5-9 FTE
Ongoing Management and Monitoring	A typical NetApp FC SAN environment requires 1/3 the IT resources of a CLARiiON environment and 1/5 the resources of a Symmetrix environment.		

Source: IT Decision-Maker Perspectives on EMC and NetApp Storage in Enterprise Database Environments. [Mercer Management Consulting report](#).

Q: Based on your research, what advice do you have for other administrators or storage managers facing the same types of issues?

John: I would suggest starting by identifying the areas of your storage environment that require the most investment of administrator time for management and day-to-day tasks such as data migration. In addition, I would suggest identifying the areas of your storage environment growing the fastest in terms of storage capacity, performance requirements, and usage. Those “deep pools” of cost are the typically the easiest places for storage administrators to achieve significant cost savings by considering investments in functionality common to solutions such as the NetApp consolidated storage solution.

Our last few studies have found that the capital cost of implementing a NetApp storage solution paid for itself very quickly by ongoing total cost savings. If you are in any environment where you manage a significant numbers of servers and storage across disparate application areas, consolidating onto an easy-to-manage storage platform will drive cost savings. Often, customers save hard money both on infrastructure—given higher utilization—and on internal costs—by spending less time managing storage. This is particularly helpful for fast-moving companies with changing IT environments.

The final piece of advice I would share is to develop your own basic, back-of-the-envelope total cost of ownership analysis. Such data can be an effective way to convince senior management of the merits of a move to new storage approaches, such as a move to a more scalable, easier to manage storage platform.

Q: SAP IT management realized both tangible and intangible benefits from this project. How can this experience help other CIOs and senior IT managers as they strive to get better results in their own environments?

John: One key message that came through in our SAP study is that “infrastructure matters.” In fact, the performance of the SAP Hosting storage environment contributed to driving some of the metrics that drive SAP’s business—namely, time-to-market and software quality. While SAP’s developers didn’t know or really care about the fact that the storage had moved to a NetApp solution, the dramatic increase in management flexibility and uptime led to considerably higher developer satisfaction. Moreover, the number of frantic calls into the IT organization dropped dramatically. Uptime and availability increased. The systems performed faster. And customer satisfaction skyrocketed.

Even though the financial impact of the business levers I just mentioned were not captured in the original total cost of ownership study, such metrics can have a significant impact on the organization. As the SAP IT executives said, “While the cost savings were well celebrated, it was the improvement in business metrics that really made our decision to deploy NetApp a success for SAP.”

Read the Mercer Management Consulting report, [NetApp Storage in SAP’s Mission-Critical Production Environment: Benefits and Total Cost Comparison](#).

© 2005 Network Appliance, Inc.



Tim McCarthy

President of Sine-rgy

Tim McCarthy, President of Sine-rgy, has 25 years of IT experience at all levels, including CIO for Visa Debit Processing Services plus over 20 years at CitiGroup and First Data Corporation. He currently consults to several Fortune 500 companies.

Build a Better Backup Strategy

by Tim McCarthy

It seems like everybody's evaluating their storage backup strategy these days. Many people I talk to are interested in the benefits of disk-based backup solutions either in some kind of hybrid model with tape or to replace tape entirely. In fact, a whopping 69% of companies surveyed by Goldman Sachs said they plan to deploy disk-based backup solutions this year. That tells me that storage managers are getting the message about the myriad of benefits associated with disk-based backup—including faster data recovery, faster backup, and less risk.

Faster, Cheaper Backup: Priceless

When I was CIO at Visa Debit Processing Services, we implemented a painful data replication process to a remote backup site. The project took 18 months of intense application changes to coordinate and account for all transactions from all sites, and then we still had to run one settlement process for all the client banks. The coordination and the data currency issues were a real challenge, but the 18 months of application changes were a nightmare! If we had had a cost-effective disk-based backup solution available it could have streamlined this process to only a few months.

There's a lot that Visa, and companies like yours, could accomplish with disk-based backup. Multiple usable copies of your data available onsite and offsite present endless opportunities to add value to your organization. Why risk shipping tapes when you can securely replicate the data to create the tapes offsite? We've all faced the challenge of needing to test application patches, QA new applications, or mine databases. Can you risk running these operations on your primary storage? Today some storage vendors enable you to run tests on online database backups on secondary storage, without impact on primary storage or your data protection SLAs.

While lots of folks are seriously considering disk-based backup, most are going to a hybrid disk-tape model. Whether you decide to go with the hybrid paradigm or pure disk as you map out your own backup strategy, here are a few crucial considerations to keep in mind:

Survivability Is Number One

Technology is always changing, but one thing remains the same: the number one priority of any backup strategy is making sure you can recover your data as quickly and reliably as possible. It's like Maslow's hierarchy of needs—you can't even think about meeting other needs until you know that your data is protected, secure, and available when you need it.

It pays to do due diligence. Develop a detailed test plan and always ask for proof of concepts using copies of your data (if at all possible). But more importantly, check out your vendor's commitment to survivability. Ask whether backup software is a core, long-term component of their solution portfolio. What is their reputation and how long have they been in this business? No serious enterprise should entrust the backup of

RELATED INFORMATION

- [Achieving 5-min RPO, 15-min RTO](#)
- [Cookbook: iSCSI vs. FC SAN](#)
- [Demand Uncompromised Security](#)
- [Ten Reasons IT Projects Fail \(and Ten Ways to Avoid Them\)](#)
- [Arup Delivers Failproof Access to Decades of Project Data \(Customer Story\)](#)

Tim's Predictions for Backup

1. **Companies will increasingly rely on a series of specialized applications and one overall backup solution.**

The days of having a single backup application are probably gone forever. There are too many platforms, too much complexity, and too few open systems.

With the introduction of solutions like Microsoft® DPM, Oracle 10gTM, Symantec Backup Exec, and NetApp SnapVault®, the industry is clearly moving toward backup solutions designed to understand the application data they protect rather than randomly storing bytes.

In this scenario, a tool like Symantec NetBackup functions as a data management traffic cop and gives you high-level management of the entire process.

2. **Expect a rising emphasis on the value of data.**

Backup applications are increasingly becoming data management traffic cops that understand both the content of the data being managed and business policies.

mission- or business-critical data to an unproven vendor that's not as serious about safeguarding your data as you are.

A Word about Security

I'd argue security is nearly as important as the basic recoverability of your data. With disk-based backup, you're moving supersensitive data from its own private fortress on primary storage and putting it on the shared arrays with all sorts of other data. That can cause real panic in the hearts of security administrators. Here's the good news: today, many vendors offer integration with virtual systems and portioning technology where you can isolate domains—even though it's all sitting on the same secondary storage.

Most security solutions today focus on the network perimeter or specific devices and people and fail to provide security of data throughout its entirety. You need to make sure you get a complete solution to secure data as it is stored and transported across multiple heterogeneous systems. Make sure once you've backed it up or replicated to secondary storage or tape you can lock your data down. Maybe you've gone through a lot of headache and pain to make sure data is locked for compliance purposes on your primary storage. Now you're moving it off to this secondary disk backup or even to tape. Is it still going to be locked? And, equally important, if the encrypted state is maintained on secondary storage will you still be able to access and decrypt it in five, 10, or even 15 years?

One Size Does Not Fit All

Data backups are a necessary evil of every company, no matter the size. Every company has different requirements and needs, which means many organizations need custom data backup solutions.

Two suggestions:

1. **Ignore the hype.** Carefully map both existing and emerging technologies into your own requirements. Tape is a perfect example. I've helped companies eliminate tape because the total cost of ownership is so high relative to disk and tapes take up way too much physical space. Not to mention the painful manual intervention tape requires. But, regardless, I still recommend tape for long-term compliance retention needs.

At some point, it might be cheaper to store compliance data on removable disk drives, but companies have had their existing tape libraries for years and know they work. Instead of completely displacing tape, I've seen many cases where adding disk to the equation has actually extended the life of tape and tape systems that might otherwise have been obsolete. In my experience, once storage managers streamline their tape infrastructure and know they can always go to disk for near-term backups, they are actually much more comfortable with the tape that's left.

2. **Understand the role of snapshots.** Another topic on clients' minds lately is snapshots. With expanding data sets, everybody struggles with keeping backups current, and they worry about ever recovering such massive amounts of data. Snapshots are an efficient mechanism to address both issues. However, the expansion of snapshot technology across so many systems and storage arrays has opened up the question of how, when, and where to use them.

First, you want to understand the reliability, performance, and space efficiency of different snapshot technologies. Not all snapshots are created equal. Remember, the number one priority of a backup is reliable, rapid recovery, so use snapshot technology you can trust. Second, snapshots are space efficient, but they are not free. Determine how much space you can set aside for snapshots. If cost is an issue, explore technologies that allow you to store snapshots on remote, lower-cost systems. Third, understand and test the snapshot recovery mechanism. Will you allow end users to recover their own data from snapshots? How will you recover a database? You should conduct fire drills for snapshot recovery, just as you do for any other backups. After all, snapshots are your first line of backup defense.

Build for the Long Term

In 1990 Citibank backed up cardholder statement images to optical platters and sent

3. Interoperability between vendors will be increasingly important.

You need to know if your vendor's solution is focused on interoperability, and not just among its sister companies. EMC, for example, offers a fairly robust solution and integrates between offerings, but what happens when you want to use VERITAS® instead of Legato?

Sample Vendor Evaluation Questions

1. Is the product a recent addition to the market? How many systems have been installed in production environments?
2. What impact, if any, will the solution have on application performance?
3. How many write operations can the recovery engine sustain? Is this engine on a clustered server?
4. If the hardware fails, can it be quickly and easily replaced?
5. Which operating systems are supported and do the operating system vendors certify the solution?
6. Does the solution support remote replication? If not, can it be integrated with other remote replication products? Will replication affect performance?
7. How is the product priced? By number of agents? Amount of protected storage?
8. What is the upgrade plan?

Not All Snapshots Are Created Equal

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

The cause: different approaches to Snapshot equivalent functionality.

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

them to an offsite storage vault for “safekeeping.” The problem was that disks got lost, dropped, and damaged, and they weren’t encrypted. Some “safekeeping.” And it wasn’t cheap. To further complicate matters, it wasn’t long before we realized that the optical platter technology wasn’t long for this world, and all the manufacturers had stopped supporting and investing R&D money in that product line. Bad move. We had to rewrite all 12,000 platters, some of which we couldn’t read at all, to tape.

So where is the next “optical platter” in your future? I don’t know that, but I do know this: when you develop, or review, your backup strategy, it’s critical to select a vendor that’s likely to be around in 15 years. You also want your data stored in a format that can be easily accessed, searched, and migrated. Someday, for some reason, you will need to read that 15-year-old data and you’d better have the correct technology to do it.

In a Nutshell

Moving to pure-play disk-based backup or to some sort of disk-tape hybrid can deliver great benefits—including faster data recovery, faster backup, and less risk. There are many considerations you need to weigh as you plan your backup strategy: always remember that one size does not fit all and look well beyond the short term.

© 2005 Network Appliance, Inc.

JOHN FULLBRIGHT

Professional Services Consultant, Network Appliance

John Fullbright is a NetApp professional services consultant and “resident expert” on Exchange. Prior to NetApp, John was a rapid response engineer in the Microsoft Global Solutions Support Center. As a member of the Microsoft Rapid On-Site Services team, John was regularly tapped to solve the challenges of premier Microsoft customers. He has also served as an intelligence analyst and IT professional for the U.S. Army.

Top Five Exchange Issues and How to Solve Them

By John Fullbright

The NetApp Global Services organization teams with customers around the globe to help design, implement, and manage consolidated Microsoft® Exchange infrastructures. Based on experience gleaned working with companies across a variety of industries, NetApp professional services consultants (PSCs) have found that issues often arise involve several key areas:

- Storage sizing
- Write performance
- Fibre Channel drivers
- Content conversion
- Windows® service packs

The following paragraphs summarize common Exchange challenges and offer advice on how to approach them.

Issue 1: Undersized Storage

Exchange performance problems can often be traced to storage undersizing. For example, I was recently called in to assist a customer that had started an upgrade from Exchange 2000 to Exchange 2003 but ran into outages and I/O latency midway through the implementation.

To uncover the problem, we spent about three weeks helping the IT team assess its existing environment. Data collected over time was analyzed to determine Exchange usage at different times of the day and on different days of the week. The data revealed that the loads differed throughout the day and night, and peak loads were experienced once per quarter. The assessment also showed that the customer had underestimated I/Os per second (IOPS) per user by about 50% and had, as a result, undersized the storage environment.

As part of the assessment, NetApp Global Services helped the customer identify how operational procedures could be improved. It was clear the customer needed a better way to manage load variances—at its Bay Area headquarters and at data centers in Europe and Asia—for thousands of Exchange seats. One procedural fix involved establishing a test lab and testing new releases, patches, and so forth to ensure consistent global deployment.

To address the sizing issue, a NetApp PSC revisited the technical design and then took Exchange online for maintenance, enabling online defragmentation in the middle of the night. The verification of Snapshot™ copies generates quite a bit of I/O, so this process was scheduled over time to avoid a large I/O spike. Using data isolation in the placement of the logs and the Snap Info Directory helped mitigate the impact. And, of course, more storage was added—eight shelves of 72GB drives that not only increased user mailbox sizes, but added IOPS capacity to give additional breathing room on the performance side as well.

So how can you tell if your Exchange performance problems are due to undersized

RELATED INFORMATION

- [Implementing iSCSI to Support Exchange 2003](#)
- [Ask the Sys Admin Column](#)
- [Example Proof of Concept for Microsoft Exchange Running with FC or iSCSI on a NetApp MetroCluster](#)
- [Five Steps to an Economical, Bulletproof Infrastructure for Microsoft Exchange Server](#)
- [NetApp Global Services](#)

storage? Fortunately there's an easy way to spot this: your users will periodically, and simultaneously receive the "requesting data" dialog box in Outlook.

The reason users see that dialog box has to do with the way Exchange works. Exchange is, of course, a transactional database. So it has log buffers in memory. Those log buffers get committed to log files on disk, and eventually those transactions in log files get committed to the database. When the log buffers in memory fill up to a point called the "high-water mark," the Exchange application begins dumping them to the log piles on disk. If the disk I/O is too slow—and the log buffers actually fill up if they can't dump to disk fast enough—then the application will quiesce; it will stop all I/O to clients until the data can be flushed to disk. It does that so you never lose a transaction. That's the nature of a transactional application.

When a client I/O is quiesced (when it halts all client I/Os), all users simultaneously get the "requesting data" dialog box in the Outlook client, and they all call the help desk. So the easiest way to identify disk latency issues is the frequency of the requesting data dialog box.

Issue 2: Write Performance Issues

Write performance problems are another common issue in Exchange environments. A major U.S. military organization recently contacted NetApp Global Services about performance issues it was experiencing. The scope of the deployment was substantial: nearly one million Exchange seats and 40PB of data.

We discovered that the performance issues were due to a reduced spindle count. If you're experiencing performance issues and they do not appear to be capacity-related, it's a good rule of thumb to investigate spindle count.

Issue 3: Fibre Channel Drivers

Another area of potential issues involves Fibre Channel drivers. Not long ago, one top NetApp enterprise account was having major issues with an Exchange deployment using NetApp Fibre Channel SAN storage. Drives would stop responding, drop off, and even disappear. After quickly determining that the issue wasn't related to NetApp hardware, the professional services team found that the problem involved the Fibre Channel HBA driver.

The customer was using an Emulex HBA and had downloaded an incorrect version of the driver from the Emulex site. That particular driver had a resource leak, so it would work initially but would slowly quit working. There were also timeout settings, which caused issues if the host attach kit wasn't used to install the driver. Once the resource leak was identified and Emulex provided the correct driver, the professional services team installed the driver and successfully resolved the issue.

Issue 4: Content Conversions

Content conversions can raise a challenge strictly due to the nature of Exchange itself. A global semiconductor customer encountered an issue with content conversion during a migration from a Sendmail system to Exchange (in other words, a system where the message is stored in MIME format and, when it's imported into Exchange, is converted to MAPI).

Fortunately, there's an easy fix: allocate storage with sufficient IOPS capacity and then relocate the system temp directory and the working directory to that storage. Because content conversion occurs in either the system temp directory or the Exchange working directory, depending on the direction, these directories can easily become bottlenecks.

Issue 5: Windows Service Packs

Another customer recently teamed with the NetApp Global Services team to upgrade from a single Exchange 2000 server supporting a couple thousand users to three four-node clusters running Exchange 2003. As part of the project the IT team wanted to consolidate users from different sources onto a standard platform to support a total of 14,000 Exchange mailboxes.

The IT team had previously worked with NetApp Global Services to replace a legacy Fibre Channel SAN solution with an iSCSI-based NetApp storage solution for Exchange and NAS and Fibre Channel SAN storage for other applications. Faced with

the challenge of completing the upgrade in just one week, the customer tapped the NetApp Global Services team for help.

The implementation was completed on schedule, but during the process the customer experienced major network connectivity problems. With a little sleuthing, we determined that the issue was caused by the service pack the customer was running: Windows 2003 Service Pack 1.

This is a very big potential gotcha. If you're deploying Exchange on Windows 2003 Service Pack 1, you must apply the hot fix or your environment simply won't work. Windows Service Pack 1 includes a "security fix," MS05 019, which breaks network connectivity. This issue is described in the Microsoft support article [KB 898060](#), which provides a link to a hot fix that fixes the issue.

Conclusion

Undersized storage, low spindle counts, content conversions, and compatibility issues are all-too-common problems NetApp PSCs see every day. Every environment is different, and [NetApp Global Services](#) can help you plan the best strategy to secure and extend the benefits of your investments.

© 2005 Network Appliance, Inc.



NAWAF BITAR

VP of File Systems Engineering, Network Appliance

Nawaf Bitar plays a pivotal role in the development of the NetApp corporate strategy for the Data ONTAP® operating system software and, in particular, the NetApp WAFL® file system. Prior to his current role Nawaf was instrumental in developing the integration strategy of Spinnaker technology into NetApp software and systems.

Nawaf Bitar began his career at NetApp in 1998 as a part of the NetApp NetCache® content delivery product group. During this time, he engineered an entire rewrite of the NetCache product. Before joining NetApp, Nawaf Bitar held various management and engineering positions at HP/Apollo, Kubota Pacific, Lockheed Martin, and Silicon Graphics.

Is NetApp Storage Really Simpler?

by Nawaf Bitar

NetApp was built to provide “storage simplified.” From inception, NetApp has worked hard to be a technology leader and along the way has achieved a fairly impressive record of industry firsts. This history of innovation has been possible because of a single basic ingredient that sets NetApp apart from every other major storage competitor: a file system.

Instead of leveraging a general-purpose operating system like UNIX® or Windows®, NetApp wrote the WAFL (Write Anywhere File Layout) storage virtualization technology as part of the Data ONTAP kernel that ships with every NetApp storage system. This highly efficient embedded file system is the foundation of a fully unified architecture that enables natural, building-block evolutions.

NetApp insiders often describe WAFL as “elegant.” WAFL technology boosts performance by letting us write wherever disk heads happen to be—this minimizes latency and makes it very convenient to do things like RAID striping. And, because we have the ability to manipulate metadata, we can more easily optimize performance.

Without the benefit of a file system technology, our competitors effectively have to start from scratch.

The WAFL file system has served as the foundation for many seminal storage system developments, including the introduction of the industry’s first snapshot technology. A NetApp Snapshot™ copy is a point-in-time, read-only view of a volume that provides easy access to old versions of files, directory hierarchies, and/or LUNs. While other data storage vendors have since copied the Snapshot concept, NetApp was the first company to ship high-volume, open systems and “point-in-time” Snapshot functionality. Independent industry analysts continue to validate the NetApp Snapshot implementation as the most efficient, stable, recoverable, and scalable of all solutions.

Because the entire NetApp product line leverages the Data ONTAP kernel and WAFL file system, NetApp is able to solve a multiplicity of problems within a single architecture. This architecture helps us keep the customer interface simple and provides incomparable extensibility. As a result, NetApp has been able to integrate support for every major compute platform, protocol, class of data, and application: Windows, UNIX, Linux®, CIFS, NFS, FCP, iSCSI, primary, secondary, low end, high end, WORM, and applications such as Oracle®, VERITAS®, and SAP.

As we recognize new customer needs, NetApp is able to build on our core technologies to quickly engineer innovative solutions. After introducing Snapshot capability, NetApp launched multiple derivatives to meet mirroring, rapid data restoration, data permanence, and other requirements. Existing products can even be combined to

RELATED INFORMATION

- [NetApp Gives Storage More Punch](#)
- [Tom Mendoza's Oracle OpenWorld Keynote: Applying the Cisco “Magic” to Storage](#)
- [NetApp Technology Overview ESG Labs Data ONTAP 7G Validation Report \(analyst report\)](#)
- [Data ONTAP 7G Customer Community \(password required\)](#)

A History of Innovation “Firsts”

1993	NAS appliance
1996	Multiprotocol appliance
1998	Content delivery appliance
2001	Nearline storage appliance
2002	Unified SAN/NAS appliance
2003	iSCSI storage system
2004	Commercial RAID-DP™
2004	Flexible virtualized storage

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

create new offerings. Several years ago, for example, customers told us of their need for compliant backups. The resulting NetApp LockVault™ product was a relatively trivial engineering exercise—the NetApp engineering team took an existing product for compliance; an existing product for backup; and, because they both run on the same platform, turned them on together. A project that otherwise might have taken years was delivered in months.

The end result is that NetApp customers do not have to deploy, learn, and support a new platform or develop a new set of best practices every time they change or expand their infrastructures to meet evolving business requirements. They have the added benefits of seamless hardware upgrades, simple data sharing across systems and applications, and a lower total cost of ownership.

In contrast, without the benefit of a file system technology, our competitors effectively have to start from scratch every time they want to solve new customer challenges. This often involves developing or acquiring multiple platforms and having to endure painful engineering contortions to integrate new functionality while at the same time trying to maintain some measure of interoperability.

What's Next for NetApp?

The NetApp storage grid promises groundbreaking performance and scalability, adaptability, and resiliency—all under the mantle of simplicity. NetApp continues to build on our technology foundation as part of a natural progression toward grid storage.

Last November NetApp announced Data ONTAP 7G with FlexVol™ technology to enable storage resource pooling and the automatic creation of multiple flexible volumes on a large pool of disks. Personally, I believe this latest innovation is one of the most exciting leaps forward in our product line because it provides a layer of virtualization that enables customers to manage data, not disks.

Every company using any type of NetApp storage should be thinking about Data ONTAP 7G. This release represents a critical milestone in our move toward grid storage, and understanding what this technology can do for you today is vital to taking full advantage of future innovations.


Our vision is that in 10 years you'll be able to plug a new storage system into a storage grid, and it will simply appear as additional capacity.

In the same way that we've been able to achieve virtualization inside a single box, my team is focused on delivering virtualization outside the box. Our vision is that in 10 years you'll be able to plug a new storage system into a storage grid, and it will simply appear as additional capacity that can be leveraged by any application.

And for those of you who are not yet NetApp customers? Until you've experienced the NetApp secret sauce, you don't know what you're missing. Another vendor may tout its own special sauce—but without a file system such as WAFL, by the time the vendor gets it to the table, it's probably already cold and missing a key ingredient or two.

Simplicity is the ultimate sophistication.

—Leonardo da Vinci—



The NetApp unified storage architecture virtualizes tiered storage.

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

Data ONTAP 7G Overview

Nawaf and the NetApp director of Software Engineering discuss the core technologies of Data ONTAP 7G and their impact on enterprise environments.

Topics include:

- Free-space flexibility
- Precision-tailored data management
- Qtree comparisons
- Thin provisioning
- Migration services
- Potential productivity gains (two real-world examples)

[Listen to the 30-minute presentation.](#)

[See the full presentation including slides on the NOW customer site](#) (password required).

Not All Snapshots Are Created Equal

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

The cause: different approaches to Snapshot equivalent functionality.

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



NICK TRIANTOS

Global SAN Systems Engineer, NetApp

A member of the elite Global Systems Engineering Group, Nick helps top NetApp global enterprise customers solve their toughest technical challenges. Nick has been in systems or support engineering roles for nearly 15 years, including positions at HP as an account support engineer (Server Group) and presales technical consultant (Storage Group).

Mapping Out a Strategy to Minimize DAS Part II: A Roadmap for Networking Storage

by Nick Triantos

Last month I shared feedback on how to evaluate your environment and choose the best technologies—Fibre Channel, iSCSI, or a combination of protocols—to build a consolidated storage infrastructure. In Part II of this series, I'm sharing (for the first time ever!) the checklist I use to help NetApp customers plan and implement efficient networked storage environments.

Most projects typically follow a course through four phases: discovery, planning, design, and implementation. The checklist is arranged accordingly, with an emphasis on the discovery and planning phases—do the right background work, and the last two phases should be straightforward.

1. Discovery Phase

Define the business issues	Get the macro view first—for example, the problem isn't that mail servers fail regularly but that stock traders can't communicate with clients.
Identify and interview stakeholders	In addition to meeting with storage, network, and database administrators, plan to talk with the department, business unit, and group managers who have expectations about deliverables, functionality, availability, and responsiveness. The project owner and security officer should also be on your list.
Document business issues and translate into technical requirements	Summarize the business and technical requirements shared by key stakeholders, plus specific procedures, practices, and policies. Identify corporate retention policies, security issues, and compliance/regulatory requirements. What levels of downtime are acceptable? Who needs access to which systems/applications? Understanding how things like poor performance and downtime impact the business will help you determine what level of spending is justified for systems, tools, applications, and so on.

RELATED INFORMATION

- [Mapping Out a Strategy to Minimize](#)
- [DAS: Understanding Your Options](#)
- [Ten Reasons IT Projects Fail \(and](#)
- [10 Ways to Avoid Them\)](#)
- ["Changing the Economics of Storage" Webcast Series](#)
- [NetApp iSCSI Storage Center](#)

Key Questions for Evaluating Application Requirements

1. Does your application require block or file I/O?
2. What levels of redundancy do your applications require?
3. Do your applications use small or large block I/O, how many IOPS are being pushed, and how much bandwidth do you need?
4. Is your current infrastructure pure Ethernet, or do you have Fibre Channel fabric?
5. How do requirements differ between your main data center, regional data centers, DR facilities, and remote offices?
6. Is getting the "lowest cost" solution a top priority?

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

Survey the existing environment	An exhaustive list of every possible question is beyond the scope of this article. Points to consider include: Basic infrastructure: • What are the existing applications, platforms, operating systems (with patch levels), storage software and hardware, and network protocols? Where are each of these in their useful technical/financial lifecycles? • What are your application requirements? Do they use small or large block I/O, how many IOPS are being pushed, and how much bandwidth do you need? • What does your networking infrastructure look like? How many Gigabit Ethernet and FC switch ports are available? Security and data protection: • What types of security features, processes, technologies, and encryption capabilities are in place? Do you have tools for auditing passwords? • What are the current RTO/RPO objectives and backup processes? What resiliency and availability levels are required? Which servers need to be set up for clustering and failover? Where is there a need for replication/mirroring? It can be helpful to determine if volume managers are used. If so, they can help you in no-downtime data migrations—for example, mirror the data after it's synced up, break the mirror, and voila—no extensive restores from tape.
--	--

2. Planning Phase

Determine project scope	Phased approaches generally work best—start with a subset of applications and users, then bring more online as you prove the process and turn on required capacity, and so on.
--------------------------------	---

Create a capacity and scalability plan	Everyone probably already has a spreadsheet with current servers/projections. This is a good starting point, but dig deeper—forecasts are usually not terribly accurate. What are host server and storage performance requirements right now—and in six months, a year, three years? How many users do you support now and will you need to in the future? Do you have enough headroom to handle more transactions? Queuing and load balancing can help you better distribute and manage workload to best use existing resources. What are storage capacity requirements right now—and in six months, a year, three years? How do you scale—for example, horizontally, by adding servers or storage? Or vertically, by maintaining the same footprint and doing add-ons to the existing hardware? If so, plan ahead for management tools (across the entire infrastructure—do you have a single-pane-of-glass product like DataFabric® Manager?), floor space, power and cooling requirements. Incidentally, this has been an area where NetApp FlexVol™ technology provides huge benefits—you don't have to get storage capacity exactly right from the start and have a lot of flexibility to provision and reprovision storage as needs change.
Plan for storage management	Determine how much of your existing infrastructure you will maintain, then develop a plan to manage the end-to-end solution, integrating the existing infrastructure and the new additions. Make sure you have resource management tools that provide information about utilization, trending analysis for forecasting, and so on. While it may not seem intuitive for a storage vendor to be focused on all other areas of your environment, you know the old adage—"if the server slows down, it's always the disk array." You have to be able to prove otherwise! Another word to the wise—today, no single tool provides visibility across storage platforms from all vendors. Be prepared to use each vendor's tools for storage provisioning.
Create a data migration plan	How do you get data out of the old and back into the new infrastructure? Your tolerance of downtime will determine the method you choose. Commonly used methods for data migrations range from backup/restore to and from tape, recursive copy, and dedicated in-band appliances to disk mirroring via the use of host volume managers.
Document failure and recovery scenarios	Develop a plan to inject failures—power off a controller, initiate a failover condition, unplug cables, shut off servers. What recovery scenarios do you expect to see? What are pass and fail conditions?

3. Project Design Phase

- Diagram the storage architecture within requirements gathered during the discovery phase
- Outline backup and disaster recovery procedures
- Design failover and clustering
- Specify recovery methods
- Diagram disk storage layout for each application

4. Implementation Phase

- Install all hardware and/or software initiators, selected software, and disk arrays
- Test functionality with applications running
- Induce failure and recovery scenarios created in the planning phase
- Test for error handling and tolerance (connect/disconnect components and note behavior) under simulated loads

I hope this is helpful to evaluating your existing environment and making plans for moving to a networked storage infrastructure. Teaming with your storage vendor will keep this process as smooth as possible and enable you to maximize the benefits of a consolidated architecture.

Last month: Mapping Out a Strategy to Minimize DAS: Understanding Your Options

© 2005 Network Appliance, Inc.



KEVIN BROWN

VP of Marketing, Decru (A NetApp Company)

Kevin Brown is responsible for worldwide marketing and business development at Decru. Prior to joining Decru, Kevin was a key member of the founding team of Inktomi, a pioneering infrastructure software firm. During his seven years at the company, he served in a number of executive positions, including vice president and general manager of Inktomi's networking business.

Demand Uncompromised Security Podcast Interview for the PodTech.net Infotalk Series

by Kevin Brown

In early November Kevin Brown sat down with John Furrier, the founder of PodTech.net, to discuss the NetApp "Uncompromised Security" initiative and new developments in storage and security.

Download the [full podcast](#) or check out these interview excerpts.

"When they invented the Model T, it didn't have any windows, locks, or alarms. Today there's one button on your key ring that turns your entire security system on and off. Data security has to be that simple."

"Decru has a turnkey appliance. We sit in the middle of the data path, and we can encrypt selectively at wire speed. We are doing up to 10Gb per second of throughput in a single box, and we can do it invisibly without changing any of your existing infrastructure or workflow."

"Decru systems are used on the battlefield today by the military, and now those same boxes are being plugged in at credit unions to protect your credit card number."

"Iron Mountain just announced that they are using [Decru] for all their data and they've now recommended to all 40,000 of their customers that they ought to encrypt data before giving it to the courier."

"Encryption has been around for 4,000 years, but to do it fast without disrupting any of your enterprise applications or systems with full interoperability and with a level of security that has been certified up to the military levels of security...that's unique."

"The idea of simplifying data management now includes not only storage, but the manipulation and the security of that data."

[Download the MP3 podcast.](#)

RELATED INFORMATION

- [Iron Mountain's Five-Step Guide to Protecting Backup Data](#)
- [Compliance Solutions for the Payment Card Industry Security Standard](#)
- [Securing Networked Storage: Decru Technology White Paper](#)
- [PodTech.net Infotalk Series](#)

Security Best Practices

Iron Mountain recently published a 17-page best practices guide that includes a five-step guide to protecting backup data.

1. Assign accountability, responsibility, and authority.
2. Assess storage risk as it pertains to information security.
3. Develop an information protection program that ensures the security of corporate information, regardless of where it is at any point in time.
4. Communicate the processes around information protection and security.
5. Execute and test the information protection security plan.

Want details? [Read the full report.](#)



BLAKE GOLLIHER

Systems Administrator and Occasional Genius, Yahoo!

As a systems administrator at Yahoo!, Blake Golliher works with NetApp storage almost every day (once in a great while they let him take a weekend or evening off). After a frustrated search for a book on management for NetApp storage, Blake decided to leverage his more than five years of experience on the subject and write a book himself.

Got questions? Suspect there's an easier way to accomplish something but not sure how? For the next two months Blake will answer questions and share tips, tricks, and tools for managing NetApp storage environments.

Ask the Sys Admin

By Blake Golliher

To start this monthly column I have identified a few key topics that have come up at Yahoo!, during discussions with friends who run NetApp storage, and at companies I've worked with in the past.

December topics:

- [Adding a disk to the first RAID group in a volume](#)
- [Monitoring clients mounted to a NetApp system and their workloads](#)
- [Enabling SSH with public keys](#)
- [Using SIO to test performance](#)

Q: How do I add a disk to the first RAID group of a volume?

This is quite simple. The "vol add" command has a "-g" option that allows you to add disks to any RAID group you choose.

For example:

```
filer01*> vol add vol10 -g rg0 2a.103
```

vol10: volume name

rg0: RAID group to which
the disk will be added

2a.103: disk you want to add
to the RAID group

This works for any RAID group in the volume or aggregate until you reach the "raidsz" limit. I'd recommend adding one disk at a time for almost all situations, but sometimes you just want to balance out the RAID groups in a volume. NetApp recommends adding one RAID group of disks at a time to prevent raid I/O imbalance.

Q: How can I tell what NFS clients are mounted on my NetApp system and compare their workloads?

Often companies have multiple NFS clients accessing a single NetApp storage system and want to keep tabs on their workload. Tools that I've found especially valuable include:

```
nfsstat  
- 1 (assess workload by system)
```

-h (review client commands by host)

```
qtree stats
nfs_hist
options nfs.response.trigger
options nfs.response.trace
```

For this article, I'll concentrate on nfsstat's per_client_stats feature which is used by the nfsstat command. If there is interest, I can delve into the other tools mentioned.

First, turn on the nfs.per_client_stats_enable client-side statistics option on the NetApp system to start collecting these statistics.

```
nfs-labfiler1> priv set advanced
nfs-labfiler1*> options nfs.per_client_stats.enable
nfs.per_client_stats.enable on
nfs-labfiler1*>
```

(This feature is off by default because it can cause a spike in CPU activity. I've never personally seen any issues, and I've used this feature on a production system with over 400 Linux® compute nodes hammering the filer cluster. However, I'm usually not one to argue with tribal technical support lore, so please be aware this is a possibility.)

Once the NFS client stats option is enabled, nfsstat will have a new option available: "-l" option. Here's an example:

```
nfs-labfiler1*> nfsstat -l
192.168.1.15  admin.lab.yahoo.com      NFSOPS = 37077 ( 0%)
192.168.1.82  web101.lab.yahoo.com      NFSOPS = 842687044 ( 6%)
192.168.1.83  web102.lab.yahoo.com      NFSOPS = 1677935681 (12%)
192.168.1.84  web103.lab.yahoo.com      NFSOPS = 1681348737 (12%)
192.168.1.85  web104.lab.yahoo.com      NFSOPS = 1670770452 (12%)
192.168.1.86  web105.lab.yahoo.com      NFSOPS = 1414701470 (10%)
192.168.1.87  web106.lab.yahoo.com      NFSOPS = 1670077384 (12%)
192.168.1.88  web107.lab.yahoo.com      NFSOPS = 799282778 ( 6%)
192.168.1.89  web108.lab.yahoo.com      NFSOPS = 795929812 ( 6%)
192.168.1.90  web109.lab.yahoo.com      NFSOPS = 804642376 ( 6%)
192.168.1.91  web110.lab.yahoo.com      NFSOPS = 805806244 ( 6%)
192.168.1.92  web111.lab.yahoo.com      NFSOPS = 789021617 ( 6%)
192.168.1.93  web112.lab.yahoo.com      NFSOPS = 799261522 ( 6%)
nfs-labfiler1*>
```

The command displays statistics in a simple table showing the percentage of workload carried by each of the top NFS clients. From the above example, it becomes clear that web102 through web106 are doing most of the work, while web107 through web112 are doing roughly half of what the other set is doing. Further investigation might show the upstream load balancer may be making some bad decisions, or the Web apps team might be migrating users to new Web nodes with more users moving to those hosts.

Another handy option is the "-h" option.

For example, let's say one host was doing a ton of operations, and you weren't sure why. Part of your investigation might involve figuring out what commands this client was doing versus those being executed by other clients. You can break down the commands that host was doing by issuing:

```
NFS-labfiler1*> nfsstat -h web101.lab.yahoo.com
```

```
Client: 192.168.1.82 (web101.lab.yahoo.com) -----
```

Server rpc:

TCP:

calls	badcalls	nullrecv	badlen	xdr call
0	0	0	0	0

UDP:

calls	badcalls	nullrecv	badlen	xdr call
843147736	0	0	0	0

Server NFS:

```
calls badcalls
843122784 0
```

Server NFS V2:

```
null getattr setattr root lookup readlink read
0% 0% 0% 0% 0% 0% 0%
wrcache write create remove rename link symlink
0% 0% 0% 0% 0% 0% 0%
mkdir rmdir readdir statfs
0% 0% 0% 0%
```

Read request stats (version 2)

```
0-511 512-1023 1K-2047 2K-4095 4K-8191 8K-16383 16K-32767 32K-65535 64K-131071
0 0 0 0 0 0 0 0
```

Write request stats (version 2)

```
0-511 512-1023 1K-2047 2K-4095 4K-8191 8K-16383 16K-32767 32K-65535 64K-131071
0 0 0 0 0 0 0 0
```

Server NFS V3:

```
null getattr setattr root lookup readlink read
1833024 0% 1 0% 0 0% 755764732 90% 64482082 8% 188 0% 4193855 0%
write create mkdir symlink mknod remove rmdir
0% 0% 0% 0% 0% 0% 0%
rename link readdir readdir+ fsstat fsinfo pathconf
0 0% 0 0% 3092232 0% 0 0% 13756658 2% 12 0% 0 0%
commit
0 0%
```

Read request stats (version 3)

```
0-511 512-1023 1K-2047 2K-4095 4K-8191 8K-16383 16K-32767 32K-65535 64K-131071
613279 77320 418644 72377 104536 29076990 0 0
```

Write request stats (version 3)

```
0-511 512-1023 1K-2047 2K-4095 4K-8191 8K-16383 16K-32767 32K-65535 64K-131071
0 0 0 0 0 0 0 0
```

NFS-labfiler1*>

While the output can be very long, it enables you to easily compare the number of ACCESS calls or GETATTR calls from the client's workload to another client. This can be extremely helpful when both systems should be doing the same workload but one is doing considerably more operations.

One caveat: make sure you are comparing similar workloads. The example shown above by using the command "nfsstat -h web101.lab.yahoo.com" shows only one host. Issuing "nfsstat -h" with no host after it will display a breakdown of all clients NFS statistics (similar to nfsstat -h *).

This data is cumulative from the NetApp system's last reboot, so if you are looking for more granular reports, you may use the "-z" option (as you can on most Data ONTAP® monitoring commands) to zero out these counters at will.

For more information on identifying data points and other options, see the manual page for nfsstat on the NOWTM site (access restricted to customers; password required). I also recommend using the "NFS Illustrated" and "Managing NFS & NIS" books published by O'Reilly & Associates to help decode these specific commands.

*** The NetApp Data ONTAP operating system does have a feature request for nfs.per_client_stats to report on more than just 254 nodes. For now, though, Data ONTAP will only report on the top 254 NFS clients. This might be fixed in burt 103532.**

Q. Is it possible to enable SSH to use public keys when accessing a NetApp system?

You should use SSH for almost everything. Not only is it secure, but it can act as a drop-in replacement for rsh if you give the NetApp system your public key in advance. This works very similarly to typical OpenSSH style.

The method I use to turn off the password requirement is based on the O'Reilly Linux Hacks book. I tapped NetApp Global Services for help to get this working, so I have to thank them for this.

First, I run the following:

```
[golliher@admin.lab golliher] eval 'ssh-agent'
Agent pid 94030
[golliher@admin.lab golliher] ssh-add
Enter passphrase for /homes/golliher/.ssh/id_rsa:
Identity added: /homes/golliher/.ssh/id_rsa (/homes/golliher/.ssh/id_rsa)
Identity added: /homes/golliher/.ssh/id_dsa (/homes/golliher/.ssh/id_dsa)
Identity added: /homes/golliher/.ssh/identity (golliher@admin.lab.yahoo.com)
[golliher@admin.lab golliher]
```

This is the magic voodoo that lets SSH utilize public key encryption and lets the storage system's SSH daemon look for a matching key on the host to which you attempt to connect.

Data ONTAP 7.0 releases come with SecureAdmin™ already installed. All you have to do is set it up and enable it. If you're running a release based on Data ONTAP 6.5, you can install SecureAdmin 3.0 without having to reboot your storage systems. For large deployments, I'd recommend one of my very favorite commands, "software get".

The command "software get" is usually used to grab Data ONTAP software releases, but it can also be used as a sort of wget-like utility that grabs things and brings them to the storage system. The drawback is that "software get" only puts files in /etc/software. The combination of "software get" with AMD (an automounter package for Linux) results in a very cool way of making this easy to install.

Another hint: the package for SecureAdmin 3.0 just extracts one file to /etc/, sa30.mod. So I usually unwrap the package and just take the sa30.mod file and put it on the Web server on my admin host. Then I can easily automate the installation of the package like this:

```
sudo rsh nfs-labfiler1 'software get http://admin.lab.yahoo.com/OnTapSoftware/Other/sa30.mod'
sudo rsh nfs-labfiler1 'priv set -q diag; mv /etc/software/sa30.mod /etc/sa30.mod'
sudo rsh nfs-labfiler1 'secureadmin setup ssh'
(hit enter 3 times, y on the last prompt)
sudo rsh nfs-labfiler1 'secureadmin enable ssh2; secureadmin enable ssh1'
```

So line by line. Line 1: get the module. Line 2: changes the mode to diag (the -q means be quiet about it-that is, don't print the warning about this being a special command that you can only run with NetApp support on the phone), then uses the mv command to move the module from the default dumping location, the location to which it would normally be extracted. Line 3: runs secureadmin setup ssh, which generates the keys. You'll need an expect script for this. Finally, line 4 enables SSH, with the options I've specified above.

Now we have to get your SSH public key over there as well. So you'll have to copy the /.ssh/authorized_keys file to the storage system. If you are using AMD, this again is easily scripted. The main point is you have to get the file to the following path:

```
/etc/sshd/root/.ssh/authorized_keys
```

Once it's there, you should be able to use passwordless authentication to your storage systems. Use the following syntax:

```
ssh -a -2 root@nfs-labfiler1.lab 'options ssh'
```

To force the use of ssh2, use the command "secureadmin enable ssh2" and do not enable ssh1. This enables ssh2 (the only active ssh version) to connect to the NetApp system.

You can use this trick with other users in Data ONTAP. You simply replace /etc/sshd/root/.ssh/ with /etc/sshd/golliher/.ssh/ for the user golliher.

Please be aware of the following known issues and make sure you are at least on the 6.5.6 release of Data ONTAP, or 7.0.2 for best results:

138781 (access restricted to customers; password required)
102525 (access restricted to customers; password required)
90855 (access restricted to customers; password required)

Q. How do I use SIO to test performance?

SIO is great for testing several options for your application, but be wary of using it to test Oracle® Database instances (see discussion at the end of this entry).

Although SIO looks scary, it is actually easy to use. Once you understand its syntax, SIO is a kitten to work with.

To install SIO, grab the package from the [NOW site](#) (access restricted to customers; password required) and extract it; the package will create its own directory called SIO. Change directory into the SIO directory and run “make” for whatever platform you have. In the below example, I’m using a Debian Linux box:

```
golliher@cow:~/SIO$ make linux
gcc -o SIO_linux.o -c SIO_ntap.c -g -D_GNU_SOURCE -D__LINUX__
gcc -g -D_GNU_SOURCE -o SIO_ntap_linux SIO_linux.o -lpthread
rm SIO_linux.o
```

And that’s the output. It created a binary in the SIO directory called “SIO_ntap_linux.” While the packaged “README” doesn’t list support for FreeBSD, the makefile command does list it as a buildable option. I’ve used it on Yahoo! gear before, and it works fine. After that, I copy the binary to /usr/bin and start using it.

One of the things SIO needs before it can do work is a target file. The target file should be roughly the same size as the data-base file you are looking to test with. I use “dd” to do this.

An example test might be to see how many threads help or hurt performance. I ran three tests, on a 100% read workload, 0% write workload using an 8kB block size on a 90MB file for 30 seconds. The only change I made was to the number of threads. I went from 16 to 8 to 4. Here are the results:

```
[golliher@admin.lab golliher] sio/sio_ntap_freebsd 100 0 8k 90m 30 16 /nfs-labfiler1/sio-test-file
Version: 3.00
```

```
SIO_NTAP:
Inputs
Read %: 100
Random %: 0
Block Size: 8192
File Size: 94371840
Secs: 30
Threads: 16
File(s): /nfs-labfiler1/sio-test-file
Outputs
IOPS: 130380
KB/s: 1043036
IOs: 3971490
Terminating threads ...
```

```
[golliher@admin.lab golliher] sio/sio_ntap_freebsd 100 0 8k 90m 30 8 /nfs-labfiler1/sio-test-file
Version: 3.00
```

```
SIO_NTAP:
Inputs
Read %: 100
Random %: 0
Block Size: 8192
File Size: 94371840
Secs: 30
Threads: 8
File(s): /nfs-labfiler1/sio-test-file
Outputs
IOPS: 131407
KB/s: 1051254
IOs: 3969931
Terminating threads ...
```

```
[golliher@admin.mud golliher] sio/sio_ntap_freebsd 100 0 8k 90m 30 4 /nfs-labfiler1/sio-test-file
Version: 3.00
```

```
SIO_NTAP:
Inputs
Read %: 100
Random %: 0
Block Size: 8192
File Size: 94371840
Secs: 30
```

```
Threads: 4
File(s): /nfs-labfiler1/SIO-test-file
Outputs
IOPS: 130609
KB/s: 1044871
IOs: 3932635
Terminating threads ...
[golliher@admin.mud golliher]
```

As you can see from the above output, the number of IOPS seemed to peak with eight threads. Four seemed too few and resulted in fewer I/Os and overall throughput. Sixteen didn't do as many IOPS either, but that's probably caused by local host resource contention more than anything (a fancy way of saying too small a box).

There is one major gotcha with SIO that applies to testing NFS performance for Oracle workloads: SIO generates workload, but it's not an Oracle Database instance. People often attempt to use SIO to gauge performance of an NFS server for Oracle. They dutifully follow NetApp best practice guides for Oracle over NFS, using no attribute caching and using a `DIRECT_O` kernel module.

In those circumstances, normal application workload other than Oracle tends to suffer. Therefore SIO performance will suffer. Or any workload generator, because it's unaware of the `DIRECT_O` (or force direct I/O on SolarisTM) and doesn't use uncached I/O to its advantage. During development of this article I learned about an option for supporting direct I/O, the `-direct` flag. It's available for Linux, Solaris and AIX. Be sure to use this flag for these tests. I've seen many an admin confused by SIO results they weren't expecting using `DIRECT I/O`.

When it comes to Oracle over NFS, testing with an actual Oracle instance is usually the best way to obtain realistic performance numbers.

Reviewer Comment: An alternative to installing and running an Oracle database (which requires a license) is to use [ORION](#) to simulate Oracle I/O workloads.