

» Tech OnTap Home

July/August 2006

HIGHLIGHTS

- » **Optimize Performance w/ FlexShare™**
- » **10 Ways to Cut SAN Complexity**
- » **Three Admin Challenges Solved**
- » **One-Click Management with Firefox**

An Update on SAN Industry Trends

Dave Dale, NetApp Evangelist, SNIA IP Storage Forum Chair and Member of the SNIA Board of Directors

The latest in storage management standards, 10Gb Ethernet's impact on iSCSI adoption, and how SAN boot could revolutionize storage.



[More »](#)

BLOGGING WITH DAVE

Dave Hitz, NetApp Founder and EVP

"When you talk about IT strategies it's important to step back and ask, 'What business problem am I trying to solve?'"

» [Dave's Blog](#)

DRILL DOWN

- » **Thin Provisioning for NetApp SANs**
LUN space usage, Snapshot™ space configurations, and more.
- » **Best Practices for Virtualization**
A road map for using the NetApp V-Series to enable unified block and file access to data on FC SAN storage arrays.
- » **Windows®/UNIX® Authentication Using Microsoft® Active Directory**
A guide to supporting unified Kerberos authentication in a NetApp environment.



ADMIN RESOURCES

- » **'Free' Featured Tool: Firefox**
A user explains how this popular browser simplifies admin across multiple sites.

TIPS FROM THE TRENCHES

10 Tips for Reducing FC SAN Complexity

Nick Triantos, Global SAN Systems Engineer, NetApp

Your approach to switching domain IDs, device naming, and coping with cabling can save (or cost) hours. Are you 10 for 10 with these industry best practices?



[More »](#)

Three Storage Management Challenges, One Simple Solution

Michael Shea, Field Technology Lead, NetApp


Your environment is complicated. Managing it doesn't have to be. See how NetApp can help you track space utilization, identify directories and files outside the backup schedule, manage SnapMirror® copies, and more.

[More »](#)

Getting More from NetApp Storage Part 1: Workload Prioritization

Akshay Bhargava, Product Engineer, NetApp

A unique new NetApp feature enables dynamic traffic prioritization so you can run multiple applications on a single storage system. Best of all, it's bundled in Data ONTAP® and won't cost you a penny!



[More »](#)

[FlexShare Demo »](#)



DAVE DALE

NetApp Industry Evangelist
Chair of the SNIA IP Storage Forum and Member of the SNIA Board of Directors

Dave has more than 25 years experience in the computer industry, including three years as chair of the SNIA IP Storage Forum. At NetApp, he drives product marketing initiatives related to NetApp SAN solutions. He is a regular contributor to industry journals, participates in IT seminars, and is a frequent speaker and panelist at industry events around the globe. Dave's SNIA tutorial, "IP Storage Technologies and Solutions," is consistently one of the top-rated tutorials at the [SNW](#) Conferences.

An Update on SAN Industry Trends

An interview with Dave Dale

The recent Storage Networking Industry Association (SNIA) [Summer Symposium](#) featured a week of Technical Working Group break-out sessions and marketing forum planning meetings, including the biannual IP Storage Forum member meeting. Tech OnTap tapped Dave for his perspective on the various sessions and industry buzz.

Q. Welcome back from SNIA summer symposium. Could you give us a little insight on what the "hot topics" that everyone was talking about were?

There are so many different groups and interests that there wasn't any one topic. Two huge centers of technical activity and education have been around storage management and data management standards. There was a "bigger, better, faster" product focus, which was mostly about infrastructure and focused on 4Gb Fibre Channel and 10Gb Ethernet.

As always, end users at the symposium focused on complexity reduction and on practical issues such as improving storage and server availability and provisioning. I saw a ton of interest in IP SAN boot, which has the potential to completely revolutionize how people think about server provisioning and recovery.

Q. Let's start with standards. What's happening in the area of standards for storage management? Is SMI becoming mainstream?

Heterogeneous manageability is consistently one of the top issues on end-user surveys, and it's also important to vendors who have multiple product lines, each with its own management interface. The [storage management initiative specification](#) (SMI-S) has been a huge push for three or four years. Version 1.0 covered simple discovery and provisioning for FC SAN devices and is now an ANSI standard. Version 1.1 added NAS and iSCSI and richer capabilities such as copy services. The standard is designed to eventually accommodate even the richest of the proprietary storage management capabilities.

Today, all of the leading vendors, including NetApp, participate in conformance testing and offer SMI-S-compliant products, and we are now starting to see end users specify SMI-S compliance in RFPs.

But the big news here is the push to accelerate SMI-S deployment, with initiatives to create reference implementations, APIs, and sample code, all with the goal of accelerating vendor time to market. The Aperi group is creating an open source reference implementation of SMI-S, in cooperation with SNIA, and SNIA is considering expanding its charter from standards specification to include the creation and distribution of code.

Q. Are there key areas in which we can expect to see standards emerging for data management?

A big "ah ha" within SNIA a couple of years ago was the realization that IT is about managing data. Storage management standards are essential, but not sufficient. So

RELATED INFORMATION

- [NetApp SAN Technology Center](#)
- [Consolidation with 4Gb FC SANs \(pdf\)](#)
- [Webcast: A User Perspective on Enterprise-Class SAN](#)
- [Tips for FC SAN Implementations](#)

Consolidate FC SAN Infrastructures with Brocade and NetApp

Consolidating storage on a Fibre Channel SAN offers many advantages, including simplified management, reduced downtime and—most importantly—significant reductions in administrative overhead.

Watch the [Webcast](#).
Read [Storage Consolidation with 4Gb FC SANs \(pdf\)](#)

Benefits of Upgrading to a 4Gb FC SAN Infrastructure

4Gb Fibre Channel is a natural evolution for SANs. Key benefits include:

- Up to 2X faster communication between servers and storage
- Reduced need for additional switch and HBA ports
- Bandwidth-intensive transfers happen faster
- Ideal for large block, sequential I/O applications
- Backward compatibility with 2Gb infrastructures

Read NetApp Technical Reports on

another high-profile SNIA initiative was kicked off involving standards, education, and evangelism around data management. [The Data Management Forum](#) is driving this, with considerable technical and end-user involvement.

Key standards around metadata transfer and around data classification are in the works. In particular there is lots of activity around XAM (eXtensible Access Method)—an initiative to establish standard interfaces to coordinate metadata between applications and storage systems—and around the classification of fixed content, which is a prerequisite to making automated tiered storage from primary to archival a reality.

Q. Vendors like NetApp and EMC have announced 4Gb Fibre Channel solutions, and there's been a huge amount of discussion about 10Gb Ethernet and the impact it will have on SAN and NAS. What's the current state of the market?

We're definitely seeing the move beyond just 2Gb Fibre Channel and Gigabit Ethernet infrastructure. Most of the Fibre Channel array vendors are now shipping solutions with 4Gb connectivity, qualified with 4Gb HBAs and switches, and every vendor I talked to at the last SNW intends to ship arrays with 10Gb Ethernet connections by the end of the year.

Q. What do you expect to see in terms of adoption for 4Gb Fibre Channel?

To understand what's happening with 4Gb Fibre Channel, it's useful to look back at the last FC SAN upgrade cycle. When Fibre Channel first entered the mainstream market, devices operated at 1Gb per second. A few years later, 2Gb FC technology became widely available. Because 1Gb FC and 2Gb FC were compatible, and per port pricing was attractive, the majority of companies ended up upgrading to 2 Gb FC over a period of time.

Today you're seeing exactly the same situation with 4Gb Fibre Channel. The majority of companies will eventually upgrade because staying on 2Gb technology in the long term will mean losing out on new features and more power-efficient switches. For a substantial period of time, you can expect most Fibre Channel SAN environments to have a mixture of speeds.

Q. So how does 10Gb Ethernet fit into the picture?

Regardless of what you're doing in the SAN space, a robust, scalable data communications infrastructure is a competitive requirement in today's enterprise. In the past four years the deployment of Gigabit Ethernet has exploded, driven mainly by the need to Internet-enable the corporate intranet and extranet, and in part because many enterprises have implemented Voice-over-IP and storage over IP (both NAS and iSCSI).

10Gb enables IT organizations to scale their LAN infrastructure to accommodate ever increasing amounts of data, and to extend their high-performance LAN to interconnect data centers within the metropolitan area. These are the areas where 10Gb Ethernet is showing up first.

Over the past four years we've seen significant improvements in per-port pricing for 10Gb Ethernet and the proliferation of higher-speed I/O options for servers (such as PCI-X and PCI-express), which makes 10Gb I/O connections more practical for server-to-storage traffic. Over the coming year you can expect to see further significant price reductions and accelerated deployment. Most important, the standards designed to enable server clustering using 10Gb Ethernet have made significant and rapid progress within the IETF and are expected to be finalized in the very near future. This will drive the deployment of 10Gb Ethernet to also address high-performance, low-latency server-to-server traffic.

Q. How will the adoption of 10Gb Ethernet for server-to-storage traffic impact the market for iSCSI?

Since the 10Gb Ethernet and iSCSI standards emerged within about a year of each other, they became somehow linked as being interdependent—each being the "killer app" for the other. This is a misconception that seems to be very widespread, particularly among people with no hands-on experience of iSCSI who think that choosing an iSCSI-based solution means you have to compromise on performance. The reality is different. iSCSI solutions using standard 1Gb Ethernet NICs provide perfectly acceptable performance for the vast majority of enterprise applications.

Other factors have a much more profound effect on the performance of the array. The number, type, and rotational speed of the disks have the biggest impact—large arrays

4Gb:

- [Benefits of Upgrading to a 4Gb FC SAN Infrastructure \(pdf\)](#)
- [Data Center Storage Consolidation with 4Gb FC SANs \(pdf\)](#)

Data Center Proven: Performance and Flexibility for FC SAN Environments

Realize the industry's best total cost of ownership and increase performance with scalable storage systems, industry-leading management capabilities, and high-speed 4Gb SAN infrastructure.

Watch a [presentation by Tom Georgens](#), EVP and general manager of NetApp Enterprise Storage Systems.

SAN Design and Data Migration Services

Designing and implementing an FC SAN environment requires hands-on knowledge of every component as well as their interactions at the system level. This list can include NetApp storage systems switch fabrics, HBAs, software, etc.

Need help? NetApp Global Services can help you design, implement, and migrate data to a robust NetApp FC SAN environment.

Learn more about NetApp Services for:

- [SAN implementation \(pdf\)](#)
- [SAN data migration \(pdf\)](#)

deliver more performance than small arrays, and Fibre Channel drives provide more performance than a similar number of ATA drives. Then the ability of the array to optimize data placement, and its ability to stripe application data over the largest number of disks also have a significant impact on performance. Interconnect bandwidth is a third-order issue compared with these, and is relatively insignificant.

The main thrust of the new 10Gb Ethernet solutions that will start to ship this year will involve supporting large numbers of Gigabit-connected servers (by connecting storage to 10Gb ports on a 1/10Gb Ethernet switch). This will enable the deployment of much larger iSCSI-based SANs and will allow iSCSI to support very high performance applications, which need low latency and more than 1Gb of storage bandwidth.

Q. Will the widespread deployment of 10Gb mean that Fibre Channel goes away?

Nobody wants to rip and replace, so the answer is not really. Fibre Channel's performance advantage will gradually disappear. However, even today, the decision of whether to deploy Fibre Channel or iSCSI usually comes down to the question of whether you already have Fibre Channel SAN infrastructure deployed at that location or not. If the answer is yes, you'll usually choose Fibre Channel and will likely upgrade to 4Gb at some point, since it's fully backward compatible and offers considerable efficiency and speed improvements. If the answer no, iSCSI is likely to be attractive.

Over time, all organizations will be confronted with decisions about their next-generation data center fabric. Most of these organizations will already be using 10Gb Ethernet in their data communications infrastructure. The question at that point will be "Should I standardize on one single interconnect technology for my next-generation data center, or does it make more sense to deploy multiple network types?"

Q. What's on the mind of end users in the SNIA community today?

There is a lot of end-user interest in simplifying management complexity and using storage-based technologies to solve the server provisioning challenges. Thin provisioning is interesting—it definitely has the potential to improve efficiency—but SAN boot has the potential to completely revolutionize the way people think about server availability, provisioning, and recovery.

Q. Why is SAN boot revolutionary?

If you can boot from a highly available shared storage device instead of from a local disk, you can potentially get rid of the disk drives in all your servers. Since disk drives are easily the most failure-prone component in any server, this will significantly improve the MTBF (mean time between failures) of your servers. The additional implication of this is that you keep your operating system image and the different software stack for applications centrally. If you have several hundred servers, this means that instead of having a hundred copies of the operating system and keeping patches up to date on all of them, you just have one "golden" image. This solves all sorts of recovery problems.

At the SNIA Symposium I had an opportunity to talk with an IT director who has implemented this kind of environment. Recently they had a catastrophic power failure in their data center (three redundant power sources all died). Because they had boot from SAN, they were able to bring the entire data center back up in 20 minutes. With conventional means, booting each of hundreds of servers individually, it would have taken days.

With iSCSI software initiators now readily available for Windows® and Linux® environments, and emerging support for software initiator-based diskless boot using standard Ethernet cards instead of more costly iSCSI HBAs, the trend toward using SAN boot is increasing substantially. Tech OnTap previously profiled NetApp Engineering's 1,000-node [Kilo-Client project](#) as a real-world example of a large-scale diskless boot environment.

Q. Great, any final comments?

End-user involvement in SNIA helps keep vendors focused on the fundamental issue of making interoperable heterogeneous storage networking a reality. IT end users can get involved by participating in user groups such as the SNIA End-User Council or one of the storagenetworking.org chapters. I hope some of the Tech OnTap community will be [inspired to join](#).

Comment on this article



NICK TRIANTOS

Global SAN Systems Engineer, NetApp

A member of the elite Global Systems Engineering group, Nick helps top enterprise companies 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). Nick maintains a [blog](#) and previously authored a Tech OnTap [article](#) on choosing between iSCSI and FC SAN.

10 Tips for Reducing Complexity in Your FC SAN Environment

By Nick Triantos

The success of a SAN is highly dependent upon proper planning to accommodate present and future requirements. Since each environment is impacted by the specific configurations it includes, it is difficult to make specific recommendations that will always apply. However, the following represents a list of tips based on industry best practices that can help reduce the complexity of any FC SAN environment.

1. Document Everything

Documentation is critical because it serves as a current and as a future reference guide for anyone who gets involved over the course of the SAN's lifecycle. Two types of documentation are especially important:

- A diagram depicting logical/physical relationships among SAN components
- A document describing different failure scenarios that were tested prior to releasing in production with the results and necessary corrective action

2. Standardize on One FC Switch Vendor

While standards exist for interconnecting connecting switches (FC-SW2) from various switch vendors, these standards are implemented differently by each vendor. In most cases (if not all), critical functionality is lost when switches from multiple vendors are connected. For example, when connecting a Brocade switch to a McDATA switch, trunking is not supported; neither are virtual channels, extended fabrics, or alias servers. Additionally, specific firmware versions are required to support such connectivity.

3. Use Unique Switch Domain IDs

In order for nodes in a SAN to communicate with each other, the fabric assigns each node a 24-bit hexadecimal FCID address (for example, 0x011f00).

This address is carried in the destination portion (D_ID) of the frame destined for a specific node and in the source portion (S_ID) of the frame sent from the node. The switch domain ID is represented by the first 8 bits of the 24-bit hex address.

```
brocade1:admin> switchshow
switchName: brocade1
switchType: 16.2
switchState: Online
switchMode: Native
switchRole: Principal
switchDomain: 1
switchId: fffc81
switchMntr: 10:00:00:60:69:c0:62:f7
switchBeacon: OFF
Zoning: OFF
port 0: id N2 Online F-Port 50:0a:09:84:96:37:c9:3d
port 1: id N2 Online F-Port 50:0a:09:84:86:37:c9:3d
port 2: id N2 No_Light
port 3: id N2 No_Light
port 4: id N2 No_Light
port 5: id N2 Online F-Port 21:00:00:e0:8b:17:c0:23
port 6: id N2 Online F-Port 50:0a:09:81:83:41:44:1d
port 7: id N2 No_Light
```

By default, each switch is assigned a domain ID within a range from 1 through 239. At

RELATED INFORMATION

- [Nick's Blog: Dynamic Queue Management for FC SANs](#)
- [Webcast: FC SAN Consolidation with Brocade and NetApp](#)
- [Storage Consolidation with 4Gb FC SANs \(pdf\)](#)
- [NetApp SAN Technology Center](#)

A User Perspective on Enterprise-Class SAN

In the Folksamerica SAN environment, data, applications, and boot images are mirrored from New York to Colorado. If the connection between sites is broken, the IP addresses of each server can be simulated on identical hardware in the remote facility.

Key benefits include ...

- No data loss and no need to restore from tape since the solution was installed over a year ago
- Full remote recovery in less than four hours
- No additional headcount requirements despite 80% storage growth
- Media costs cut by two-thirds



Watch the [Webcast](#) featuring Folksamerica's VP of Infrastructure. Read the [Folksamerica success story](#).

many companies the same domain ID is used in separate fabrics, but I advise against it. If fabric merge is attempted at some point, domain ID conflicts will prevent that merger unless steps are taken to change the joining switch's domain ID. Keep your domain IDs unique.

4. Disable Unused Switch Ports

Make sure you're protected from anyone who might gain physical access to the switch. Prevent unauthorized port access by disabling unused switch ports.

5. Determine Your Zoning Strategy and Use Aliases

There are various zoning methods; you need to choose one and stick with it for the lifecycle of the SAN. Zoning by initiator/target is the single most widely accepted practice in the industry. Each zone contains an initiator WWN and one or more target WWNs. While in a large fabric this requires the creation of a considerable number of zones, it has the advantage that any disruption in the zone will only affect the smallest number of nodes.

Always use aliases to assign meaningful names to devices connected to the fabric, as it is much easier to identify devices by name rather than by WWN. Aliases, zone and zoneset naming, and content require careful planning in order to accommodate any future fabric merger requirements that may arise.

Here are two things I've had to learn through painful experience:

- The name of a zone in Fabric A should not be used for a different type of zone in Fabric B. For example, if you create a zone named myZone in Fabric A, you should not use the same name as an alias, zone configuration, or zoneset name in Fabric B. In this scenario, merging the fabrics will cause a zone type mismatch.
- If an alias, zone, zoneset, or zone configuration name is the same on both Fabric A and Fabric B, but the content between the two fabrics is different, the fabrics will not merge.

6. Standardize on FC Switch Firmware

In addition to standardizing on an FC switch vendor, you should also standardize, when possible, on firmware running in these switches. While there typically is compatibility between switch FW revisions, managing a single version is always easier.

7. Standardize HBAs and HBA Drivers by OS

It would be difficult to standardize on a single HBA vendor across every operating system in the fabric, but you should choose one HBA vendor, card type, model, and HBA driver per operating system. It is much easier to manage and upgrade a single driver HBA revision than multiple ones. Overall, the less component diversity that exists, the better you will be able to manage and maintain control over your SAN.

8. Use FC Cables

Sometimes standardization can go too far. Some companies standardize on the 62.5µm 850nm cables regularly used by Ethernet for their FC SAN environments. Unfortunately, Ethernet cables have a limited reach of up to 200m.



The standard cable for an FC SAN implementation is 50µm (micron) 850nm multimode cable. Unless you have a unique need for 62.5µm cables, make sure you use 50µm FC cables.

9. Mark the Cables

This seems obvious, but it's so easy to overlook. There is nothing worse than having to trace cables under raised floors to identify connections between the various SAN components. You need a detailed, logical scheme to identify end-to-end connections.

For example, on the switch side, where a host is connected, the connection might be identified as SRV/SRV_Hostname/HBA_Slot_#. This indicates that the connection represents a server (SRV), followed by the server hostname and the host PCI slot number where the HBA physically resides.

On the host side, the cables would be marked as SWT/ SWT_Hostname/Blade/Port_#.

Benefits of Upgrading to a 4Gb FC SAN Infrastructure

4Gb Fibre Channel is a natural evolution for SANs. Key benefits include:

- Up to 2X faster communication between servers and storage
- Reduced need for additional switch and HBA ports
- Bandwidth-intensive transfers happen faster
- Ideal for large block, sequential I/O applications
- Backward compatibility with 2Gb infrastructures

Read NetApp technical reports on 4Gb:

- [Benefits of Upgrading to a 4Gb FC SAN Infrastructure \(pdf\)](#)
- [Storage Consolidation with 4Gb FC SANs \(pdf\)](#)

Data Center Proven: Performance and Flexibility for FC SAN Environments

Realize the industry's best total cost of ownership and increase performance with scalable storage systems, industry-leading management capabilities, and high-speed 4Gb SAN infrastructure.

Watch a [presentation by Tom Georgens, EVP and general manager of NetApp Enterprise Storage Systems](#).

SAN Design and Data Migration Services

FC SAN implementations require expertise with planning, designing, configuring, and application integration. Performance, availability and interoperability are always at the forefront of a SAN implementation.

NetApp Global Services can help you plan and implement a powerful SAN solution, quickly and efficiently that provides maximum protection and peace of mind.

Learn more about NetApp services for:

- [SAN implementation \(pdf\)](#)
- [SAN data migration \(pdf\)](#)

This indicates that the host HBA is connected through a switch (SWT), followed by the switch hostname, the switch blade and the port number where the physical connection takes place. The same scheme can also be used for connecting FC target ports to an FC switch.

10. Beware Patch Panels

A patch panel is a large number of consolidated ports that connect incoming and outgoing lines of a communication system (for example, LAN, FC). While patch panels are important in terms of cable management, they tend to limit distance because of signal loss. As a result, you may see behavior in the fabric where nodes constantly log out and log back into the fabric.



Conclusion

During its lifecycle, a SAN will undergo multiple changes as new devices are added, existing devices become obsolete and are removed, and new personnel is assigned. Maintaining a consistent process to manage the SAN lifecycle and following widely accepted best practices are some of the easiest and most effective ways to reduce complexity and minimize unplanned downtime.

Comment on this article



MICHAEL SHEA

Field Technology Lead, NetApp

Michael Shea joined NetApp in October 2004 and brought a wealth of experience gained from over 10 years spent working in the storage industry. For the past five years he has focused on teaming with customers to craft storage management solutions and processes to realize the business benefits of integrated storage solutions. When not in the data center, Michael explores other types of transport protocols. He reports that he's ridden camels, alpacas, horses, elephants, and even an ostrich and was once bitten by a spider monkey (all true).

Three Storage Management Challenges One Simple Solution

By Michael Shea

My mother raised me with a "get-it-done" mentality and a strong disdain for anything that makes for more work or adds complexity. Conceptually, I think most IT shops would agree with this philosophy. Everyone wants to reduce complexity, and yet we see article after article about the lost promise of value in SAN. Many IT teams are coping with an overwhelming degree of complexity that is keeping them from realizing the very value that motivated them to move to centralized storage in the first place. The cure is worse than the disease, and complexity rules supreme.

A big part of the problem is that SAN technologies such as iSCSI and Fibre Channel have created a new level of sophistication that single-purpose tools simply cannot scale to address. Managing complex SAN environments requires an integrated tool that shows you the big picture and allows you to zoom in to varying levels of depth and detail to manage complex storage environments.

NetApp was founded with the goal of simplifying storage. (Check out a previous Tech OnTap article, [Is NetApp Really Simpler?](#) for details on how the folks in Product Engineering have accomplished this.) In this article, I'll share three specific examples of common management challenges and demonstrate how you can easily solve them in a NetApp storage environment using a tool called NetApp Operations Manager (formerly known as DataFabric® Manager or DFM).

If you're not currently using NetApp storage, this article provides a quick guided tour of how NetApp helps simplify basic management tasks.

Challenge 1: How do I know exactly where all my storage resources are and how they are being utilized?

I don't know much about herding sheep, but one thing is pretty clear: if a shepherd doesn't know where his sheep are, it's a lot harder for him to take care of them. The same holds true with storage management. How can a storage administrator control his environment if he isn't sure what resources he has? A clear and up-to-date view is critical, especially when the environment begins to scale up and out.

Although there are any number of different tools you can use to manage your environment, my favorite is NetApp Operations Manager (OPS MGR). OPS MGR provides a single point of view that starts at a high-level overview of your entire NetApp environment and enables you to quickly drill down to detailed management tasks at any level.

RELATED INFORMATION

- [Guided Tour: NetApp OPS MGR](#)
- [From Reactive to Proactive Admin](#)
- [Is NetApp Really Simpler?](#)
- [Productivity Management in the Data Center \(Webcast Series\)](#)
- [Integrated Data Management \(pdf\)](#)

Is NetApp Storage Simpler?

NetApp solves a multiplicity of problems within a single architecture.

Our "secret sauce" consists of the Data ONTAP® kernel and WAFL file system, both of which are leveraged across the entire NetApp FAS family of tiered storage. This architecture helps us keep the customer interface simple and provides exceptional extensibility.

As a result, NetApp has been able to innovate freely, creating unique storage solutions designed to satisfy emerging needs and address the thorniest storage problems such as backup and disaster recovery.

See for yourself what makes [NetApp storage simpler](#).

Switching to Proactive Data Management

Ever feel like you spend much of your day fighting fires?

NetApp is committed to simplifying management using an approach we call integrated data management (IDM).

Click the image to the right for a [brief introduction](#) to how you can drill down from a high-level overview of your environment to a view of how disk space is utilized inside an aggregate with just a few clicks.

The best part is that the same tool works the same way using the same screens for NAS or SAN—or both at the same time.



A Guided Tour of NetApp OPS MGR

Challenge 2: How can I be sure my data is scheduled for backup?

Ask that simple question in a room full of administrators, and nervous smiles break out. Despite the most rigorous processes to ensure that new directories are added to the backup schedule, people forget or make mistakes, and things are missed. Often the way an administrator finds out (in fact, often the only way) that a backup is not occurring is when a restore is required.

Unfortunately, it happens.

Traditional storage vendors have no way to quickly and accurately flag directories and files outside the backup schedule, leaving storage administrators to coordinate with end users and do much of the work by hand.

If you're using OPS MGR, however, a few clicks can display a list of directories that are not scheduled to be backed up. It makes no difference if the data is sitting on a NAS device, an iSCSI LUN, a Fibre Channel LUN, or even in a directory on a locally attached hard drive. OPS MGR treats each of these scenarios equally. Hence an administrator needs to learn just one simple thing and is therefore able to validate the backup environment without concerns about how the data is accessed. This is part of the beauty of what we call "unified storage."

From the directory backup listing, you can easily add items to the backup list while ignoring items, such as scratch space, that don't need to be backed up. A similar report is available at the qtree level.

Simple, yet very powerful. Mom would approve.

Challenge 3: Is there such a thing as simplified DR for my applications? (Or, how do I manage multiple SnapMirror® copies?)

Disaster recovery (DR) is a lot like life insurance. Everyone needs it. No one ever wants to use it. You want enough of it so that you're protected, but not so much that your wife's next spouse gets a free ride.

In the past few years, NetApp has done a lot to simplify DR. Our platforms virtualize

Like customers going to an ATM to make withdrawals or deposits, this approach enables systems administrators and DBAs to manage their own data and self-serve basic requests. The storage administrator retains control over the process, but doesn't have to participate in every exchange, freeing up valuable time for other work.

Learn more about the NetApp strategy for helping you shift from [panic](#) to [power](#).

A Quick Primer on NetApp Data Protection Software

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

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

SnapVault, in contrast, is intended for disk-to-disk backup. It retains all backup copies as they appeared at the time they were created on primary storage for a user-specified period of time. Secondary storage used by SnapVault cannot be mounted read/write. Backups must be recovered from secondary storage to the original or an alternative primary storage system in order to restart.

Get the details. Read the reports:

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



Figure 1: Directories Not Scheduled for Backup

An administrator can dole out storage in amounts required and add more in the amount desired, quickly and nondisruptively, without having to think a lot about disk I/O. Since all disks are combined into one unit (an aggregate), the applications accessing it utilize all the spindles in the array at all times. This is very different than how traditional storage works. You don't have dozens or hundreds of disk components to coordinate in a DR replication scenario. Traditional storage forces you to ride herd on all the discrete components, making DR harder.

A common management question involves how you manage and monitor dozens or hundreds of SnapMirror copies across multiple systems, sites, and time zones.

Enabling this option adds a Disaster Recovery tab to the top level view of the main OPS MGR screen. Simply click this tab for a complete overview of your SnapMirror configuration. Your environment. No scripting. In units. You can create new schedules and policies, and You can even set up secure



Delivering Manageability, Not Just Management

Ronald Reagan said it best: "They say the world has become too complex for simple answers; they are wrong." Mom would agree.

Interested in seeing more? Contact your account team to schedule a demo of OPS MGR capabilities or [click here](#) to be contacted by NetApp.



AKSHAY BHARGAVA

Product Engineer, Network Appliance

Akshay Bhargava is a product engineer on the NetApp core system team with responsibility for Data ONTAP® and the WAFL® file system. Prior to joining NetApp, Akshay served in a variety of engineering management and software development roles working on distributed computing, clustering, storage virtualization, and network security. He holds a degree in computer science from Carnegie Mellon University.

Getting More from NetApp Storage

Part 1: Workload Prioritization for Optimized Performance

By Akshay Bhargava

For years the most common approach to storage deployment has been to dedicate one or more storage systems for each major application or workload. That way, if one application overloads a storage system, no other application or group of users suffers the consequences.

Today—with storage system capacity and performance continuing to grow—this traditional approach has become less and less cost-effective. The one-storage-system-per-application approach can lead to a huge number of storage systems, each of which may be underutilized, tying up budget dollars, consuming valuable floor space, wasting energy, and complicating management.

A Simple Solution: Workload Prioritization

To minimize the need for dedicated storage, NetApp has introduced a new priority of service feature—FlexShare™—that enables you to dynamically prioritize storage traffic at the volume level and control service levels. This means you can give higher priority to the volumes used by more critical workloads and safely run multiple applications on a single, consolidated storage system.

The best news is that FlexShare is a standard feature of the NetApp operating system starting with Data ONTAP 7.2. **There is no additional cost to take advantage of this advanced feature.** No other mainstream storage vendor offers non-disruptive dynamic tuning of storage service levels.

Three Independent, Tunable Parameters

Each workload in your storage environment has a different level of importance to your business, so a "one-size-fits-all" approach to storage resource allocation doesn't make sense. FlexShare offers three parameters that can be configured independently on each volume to tune your storage to the needs of every application:

- **Relative priority.** Assign a priority from "VeryLow" to "VeryHigh" to each volume. A higher priority gives a volume a greater percentage of available resources when a system is fully loaded. If higher priority applications aren't busy, lower priority applications can use available resources without limitation.
- **User versus system priority.** Prioritize user workloads (application and end-user traffic) versus system work (backup, replication, and so on) or vice versa. For instance, give your OLTP application priority over SnapMirror®.
- **Cache utilization.** Configure the cache to retain data in cache or reuse the cache depending on workload characteristics. Optimizing cache usage can significantly increase performance for data that is frequently read and/or written.

RELATED INFORMATION

- [FlexShare Design and Implementation Guide \(pdf\)](#)
- [FlexShare Demo](#)
- [Is NetApp Storage Really Simpler?](#)
- [Architecting Storage for Resiliency](#)

FlexShare in Action

The following graph shows how a high-priority application can benefit from FlexShare:



Click to launch the FlexShare demo.

To perform this test:

- A single NetApp system was configured with two identical volumes.
- SIO was used to simulate an OLTP workload against each volume (R/W 40:60, random: 70%, block size: 8KB).

With FlexShare off, each volume delivered a workload of about 6,000 disk ops per second for a total of 12,000. Priority on one volume was then set to "VeryHigh," while the other was set to "VeryLow."

- The high-priority volume climbed to an average of around 9,000

These parameters can be configured dynamically, so you can change all or some of the settings on-the-fly as your needs change, even if they change daily or hourly. Once FlexShare is configured, it directs the way storage system resources are used to provide an appropriate level of service to each application. See the sidebar for more details.

Four Ways You Can Get Started with FlexShare

By assigning an appropriate priority to the volumes used by each application, you can mix workloads on the same storage system without having lower priority work impact critical workloads. Some of the ways FlexShare might be used to increase storage efficiency and utilization in your environment include ...

Consolidating Different Workloads

Most organizations have a variety of workloads—such as databases, business applications, e-mail, and home directories—that could be consolidated to simplify storage administration, increase storage ROI, and reduce the consumption of valuable resources and physical space.

With FlexShare, you can consolidate two or more such workloads together on the same system. For instance, you might consolidate your Oracle® Database with user home directories, setting the priority of the database higher than the home directories to ensure that home directory traffic doesn't interfere with important database operations.

A database is a prime example of an application that may also benefit from the ability of FlexShare to tune cache utilization. Volumes that store log files, which are typically written but not reread, can be set to *reuse* cache space immediately, while volumes containing data files storing the most active indices and table spaces that are frequently read and updated can be set to *keep* data in cache. (For optimal results, the *keep* setting must be used with care. The active data set configured with this policy should be smaller than the cache.)

Mixing Fibre Channel and SATA Storage

Many NetApp storage systems support mixed storage—combining Fibre Channel disks and SATA disks on the same system. The new NetApp FAS6000 series, for example, can provide hundreds of terabytes of either or both types of storage. (See last month's [discussion on FAS6070 design](#).) Customers typically put higher priority data on Fibre Channel disks and lower priority data on SATA disks. Using FlexShare you can ensure that access to the data on Fibre Channel volumes has priority over access to SATA volumes.

Incidentally, if you're considering mixed Fibre Channel and SATA storage, I highly recommend checking out the best practices outlined on page six of the [Storage Best Practices and Resiliency Guide](#) (pdf).

Setting the Priority of Backup and DR Operations

One of the unique features of FlexShare is that it gives storage administrators the ability to set the priority of system operations (backup, replication, maintenance tasks, and so on) relative to end-user and application traffic.

ops, 50% more than without FlexShare and 3X the low-priority volume.

- The low-priority volume dropped to an average of about 3,000 ops.

Total operations remained at 12,000, indicating that FlexShare has no significant system overhead.

Watch the [FlexShare demo](#).

How FlexShare Controls System Resource Allocation

FlexShare works by controlling the allocation of critical system resources based on the volume priority level and buffer cache policy set by the storage administrator.

- **CPU.** Disk operations on higher priority volumes are preferentially scheduled for processing.
- **Disk I/O.** Higher priority volumes are allowed more concurrent disk reads than lower priority volumes. FlexShare maintains a maximum number of concurrent disk reads allowed per volume.
- **NVRAM.** A volume's priority dictates how much NVRAM it can consume relative to other volumes to ensure high-priority applications do not get write starved.
- **Memory.** Memory consumption is dictated by the configured buffer cache policy, which provides hints to the buffer cache manager:

Keep. The cache manager preferentially keeps items in the cache.

Reuse. The cache manager preferentially flushes these items when clearing space in cache.


Default. The cache manager gives data items normal priority. Keep items have higher priority, and reuse items have lower priority.

Want to learn more? Read the [NetApp technical report](#).

The Versatile Storage Platform

To truly appreciate the versatility of the NetApp architecture, it is important to view how the storage is managed and

how the storage is accessed as related



Many NetApp sites have active user volumes that are backed up or replicated at frequent intervals using NDMP, NetApp SnapVault®, or SnapMirror. It may be advantageous to give higher priority to user access during the day (peak hours). This way, backup and replication occurring during that time won't impact users. During off-peak hours, you can grant higher priority to these critical tasks to ensure that they complete in as short a time as possible.

For organizations that have a primary site and a backup or disaster recovery site, you may choose to grant higher priority to user access at the primary site, while setting higher priority for system access at the secondary site.

Multiple Instances of the Same Application

A single storage system might also be used to host multiple instances of the same application, such as several database instances. For example, Database 1 might be executing an OLTP workload while Database 2 is doing data mining. Using FlexShare to give higher priority to volumes used by the OLTP database ensures that a storage system can support both activities without the data mining work affecting critical OLTP transactions.

The [FlexShare demo](#) demonstrates this scenario using a simulated OLTP workload and two identical volumes on a single NetApp system (R/W 40:60, Random: 70%, Block Size: 8KB).

Get the Most from NetApp Storage with FlexShare


FlexShare gives storage administrators a new level of control over storage system resource allocation and performance. By bundling advanced features such as FlexShare into its base operating environment at no additional cost, NetApp is helping you get more from your NetApp storage than ever before.

You can learn more about FlexShare from my recently published [FlexShare Design and Implementation Guide](#). This guide provides detailed information on how FlexShare works, how it is administered, best practices for deploying FlexShare, and more.

Coming in September:

Part II: Managing Complementary Workloads for Exchange

Comment on this article



functionality. Too often storage vendors separate these two concepts, which creates overly specialized storage systems that eventually become isolated islands.

In his May Tech OnTap article, NetApp user Ben Rockwood provided an overview of how the NetApp Data ONTAP operating system manages data on disk and how that data is accessed from client systems.

Highlighted technologies include:

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

Read a [User Perspective on Data ONTAP 7G](#).



[Click to launch the FlexShare demo.](#)

TECH ONTAP HOME PAGE



Bryan R. Bell

Enterprise Data Architect, VECO

With over 5,000 employees, **VECO** provides engineering, procurement, and management services to the energy, resource, and process industries in both the private and government sectors. Bryan has over 15 years of experience in IT and been employed at VECO for 10 years. Throughout his career, Bryan has focused on enterprise storage, including architecture, implementation, and administration. In addition to his work at VECO, Bryan consults on Web development and graphic design, operates a mountain bike tour business, and is an accomplished musician. He plays trumpet, piano, and guitar and runs a personal recording studio.

Tool of the Month: Firefox

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: Mozilla Firefox was developed by the Mozilla Foundation.

What it is: You're probably familiar with the Firefox Web browser, which offers unique features, including tabbed browsing, Password Manager, and ability to integrate Google Desktop Search. Firefox also provides a fast, easy, and free interface for managing a diverse NetApp environment and can be a big help when it comes to all sorts of HTTP-based SAN and system management.

How it works: With Firefox's tabbed browsing feature, it's simple to create a folder full of links to related items—such as FilerView® for each NetApp system—and place that folder on your toolbar. Then when you want to access those items, all you have to do is click your right mouse button on the folder and select "Open in Tabs." (See Figure 1 in the sidebar.) All the pages in the folder are automatically opened in separate tabs. And, with the Firefox Password Manager, you don't have to remember and manually enter numerous passwords. Various Firefox extensions add more useful functionality.

Why Bryan thinks it's cool: Firefox lets me manage all my NetApp systems from one session and one launch of Firefox. I manage 10 primary storage systems and two NearStore® R200s, all with one-click tabbed browsing, using this interface.

How Bryan uses Firefox: I'm running the latest version of Firefox (currently version 1.5.0.4), but the features I use have been available since I started using Firefox about a year ago.

The feature I use most is tabbed browsing. The first item on my Bookmarks toolbar is a folder called NetApp. You can create this type of folder very easily using Firefox's bookmark management features. My NetApp folder has links to everything I need to manage my NetApp systems, including:

- NetApp Operations Manager (formerly known as DataFabric® Manager, or DFM)
- FilerView for each system
- NOW™ Web site

By right-clicking this icon and clicking "Open in Tabs," I get automatic access to these tools without having to manually log in to each one. With these tools, I can do almost every management task I need to accomplish, including monitoring and managing SnapMirror® copies.

`http://Filer_Name/api?api=/util/true&onsuccess=fv`


I use the same method with other tools that I use all the time. For example, if I need to troubleshoot a problem, one of the first things I do is look at traffic graphs from [MRTG](#)

RELATED INFORMATION

■ **Previously Featured Tools:**

- [Cacti](#)
- [MRTG for Filers](#)
- [Nagios](#)
- [Simulate ONTAP](#)
- [SIO](#)
- [ToasterView](#)
- [NOW Tool Chest](#)
(password required)

Screen Captures of Firefox in Use for NetApp Storage Management



This screen capture shows how Firefox looks after I've opened my NetApp folder in tabs. I have a separate tab with FilerView to each system, NetApp Ops Mgr, and a tab connected to the NetApp NOW site.

(previously featured in Tech OnTap) and PRTG. I have a separate folder on my Bookmarks toolbar for those. When I open that folder in tabs, it gives me one tab for each of the graphs that I need. I also launch other management tools, including HP OpenView, this way. I also frequently open "bit-to-byte" conversion Web sites for proper LUN and volume resizing across my globally based SAN environment.

My company has 15 major locations in Alaska alone. Each machine room is equipped with a Weathergoose climate monitor that provides video, temperature and humidity information, sound, and more for each remote location. Again, I use Firefox to launch connections to my remote locations with a single click if I need to check what's happening.

Along with Firefox, I use the following additional software and extensions:

- Google Desktop Search
- Google Desktop Search Plug-in for Firefox
- Morning Coffee
- Internote

Google Desktop Search will index your desktop system in the same way that Google indexes the Web, including your e-mail and all of your files. I have NetApp Operations Manager configured so that all trapped events are e-mailed to me. I use Outlook to filter all the events into separate folders by category: critical events, errors, information, warnings, and so on.

This gives me a very fast way to search for events using the Google Desktop. I can search for a keyword such as a volume name and get back every event that ever affected that volume in a matter of seconds rather than the hours it might otherwise take to search logs. (The second figure in the sidebar illustrates a search on the string "filer down.") It's not essential that you integrate Google Desktop with Firefox, but doing this gives me one-stop shopping for everything I need to get done. I can manage my entire SAN environment in one place and one window. It sure beats opening consoles, sessions, and other thick clients just to keep things running.

The Morning Coffee extension allows you to schedule a particular set of pages to look at on a particular day. This helps remind me to do routine tasks. For instance:

- Monday: check volume status
- Wednesday: check status of remote site OSSV jobs
- Friday: check status of all SnapMirror and DR operations

Finally, the Internote extension lets you add persistent notes to Web pages. You can tie notes to particular Web pages for future reference, and it beats having sticky notes all over the sides of your monitor!

Caveats: Having all the passwords for my NetApp systems cached in the browser does create the possibility of a security breach. To minimize the risk, I set a master password within Firefox. This feature ensures that stored passwords are encrypted. You also have to enter the master password once in any Firefox session in which you access a password-protected site, so no one can get on your machine and access your storage unless they know the master password. Once you do this, additional "launches" opened during the same logon session are cached, and no passwords are required. In addition, my profile is set to autolock my desktop if I forget to lock it when I leave my desk.

Also, NetApp Operations Manager is specific to NetApp and is not a free tool. If you don't have the budget for Operations Manager or aren't a NetApp user, an alternative option would be to create a similar event log-in e-mail using free tools such as Swatch (Linux® or UNIX®) or Kiwi Syslog (Windows®).

Editor note: Swatch and Kiwi will be profiled in an upcoming issue of Tech OnTap.

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

✱ **Nominate a tool**



This screenshot shows Google Desktop Search (GDS) running in Firefox. I have all DFM traps sent to me via e-mail, where they are filtered into separate folders. I've configured GDS to index my e-mail so I can quickly search my archive of 175,000 messages for a particular keyword (in this case "filer down"). Note that the highlights are created by Google based on the search.

Veco Storage Infrastructure, Disaster Recovery Plan Flowchart, and More

