



Tech OnTap

Highlights

March 2008

- [Large Capacity Best Practices](#)
- [Tips to Increase Storage Resiliency](#)
- [Defusing the Data Explosion](#)
- [Transform Your Data Center Webcast](#)

netapp.com | Tech OnTap Archive

Defusing the Data Explosion

Brett Battles, Director, Product Management
Mark Woods, Product Marketing Manager



Is your data center running out of power and space? Have you been told that for every new system or application that enters your data center, you must take one out? Discover five NetApp technologies that can help you make more efficient use of your IT infrastructure so you can manage more data with fewer resources in less space and using less power.

➤ [More](#)

Tips from the Trenches

Best Practices for Large-Capacity Storage

Chris Lueth, Technical Marketing Engineer



A single storage system can now exceed 1PB in capacity. Find out how to effectively deploy and manage large systems.

➤ [More](#)

Engineering Talk

Top Five Tips to Increase Storage Resiliency

Steve Lawler and Haripriya
Technical Marketing Engineers



Is your data as safe as it could be? Learn how to increase resiliency through multipath HA, proper sparing, SyncMirror®, and more.

➤ [More](#)

Blogging with Dave



Dave Hitz
NetApp

NetApp Beats EMC

Last month we released new SAN database benchmark results: our system cost less, had fewer disks, and beat the EMC system by 24%.

➤ [Dave's Blog](#)

Tech OnTap Spotlight

➤ Transform Your Data Center

You read the [article](#). Now hear the Webcast and see the technical [Q&A](#) on how Sprint reached 100% automation.

➤ VMware Virtualized Hardware Hotel

Just for fun: You might remember its story from our [previous article](#). Watch NC State's prize-winning video.

Go further, faster™

© 2008 NetApp | [Privacy Policy](#) | [Send Feedback](#)
495 E. Java Drive, Sunnyvale, CA 94089 USA



Quick Links

- netapp.com
- [Tech OnTap Archive](#)

Explore

EPA: Server and Data Center Energy Efficiency

In August of 2007, the U.S. Environmental Protection Agency (EPA) made a report to Congress, "Report on Server and Data Center Energy Efficiency", highlighting the need for concerted efforts to reduce power consumption in data centers. From 2000 to 2006, electricity consumed by U.S. data centers doubled, to 44 billion kilowatt-hours per year. Over the same period, enterprise storage systems tripled their power usage. The EPA estimates that if the current trends continue, by 2011 the United States will need to build an additional 10 power plants just to power data center growth.

[More](#)

Study Puts NetApp at the Top of the Efficiency List

In 2007, NetApp commissioned Oliver Wyman, formerly Mercer Management Consulting, to interview storage and facility managers and IT executives at large enterprises and leading-edge companies to explore how different storage solutions affect power, cooling, and space requirements in the data center.

According to the Wyman study, environments with NetApp storage solutions experience higher system efficiencies than those that have deployed EMC or HP storage solutions.


[More](#)

Diffusing the Data Explosion

By [Brett Battles](#) and [Mark Woods](#)

What would you do if you were told "For every new system or application that enters the data center, you must take one out?"

When a data center runs out of space, power, or the ability to cool equipment, it has to make dramatic changes to continue functioning. Building a new data center is rarely feasible; new enterprise data centers cost tens of millions of dollars and take years to build. Some companies have started requiring their administrators to make sacrifices for every new system or application that they bring into the data center — something has to leave to make room for the newcomer.



Estimated Annual Electricity Use in U.S. Data Centers by End-Use Component

Annual electricity use in data centers is steadily increasing. If this trend continues, 10 new power plants will be required to power this growth in the next three years.

Efficiency is the key to saving power, cooling, and space in the data center. You can make more efficient use of your server resources by pooling them using blades and virtualization software. With the right storage systems, you can realize the same type of efficiencies in your storage infrastructure that you gain from your virtualized server infrastructure. NetApp's alternative approach to meeting rapidly growing power, space, and cooling requirements is simple: We use innovative technology to subtract machines and disks from the power equation by storing data more efficiently.

Using NetApp's technology to increase your data center efficiency yields immediate and continuous benefits. You can reduce raw storage expenditures and enjoy corresponding savings of power, cooling, and data center space, while reducing the carbon footprint of your data center. And perhaps the largest benefit is that you can avoid the expense and huge effort that would go into building a new data center.

NetApp's storage efficiency technologies include:

- Deduplication to eliminate redundant data objects
- RAID-DP® to get better data protection than mirroring while using far fewer disks
- FlexVol® thin provisioning to minimize stranded capacity
- **SATA disk drives that consume 40% less power** than FC drives for equivalent capacity
- Snapshot™ technology for point-in-time recovery using very little storage
- FlexClone® writable Snapshot copies for test and development

Storage efficiency starts with the platform

Many data centers have a different storage system for each storage need-business applications, file services, backup, archiving, and so on. This approach increases complexity, wastes storage, and makes inefficient use of valuable data center resources — labor, power, cooling, and space. With NetApp's unified storage platform, featuring the Data ONTAP® operating system, you can eliminate the need for separate storage systems and special gateways.

Remove duplicate data objects with deduplication technology

The average UNIX® or Windows® enterprise disk volume contains thousands or even millions of duplicate data objects. As these objects are modified, distributed, backed up, and archived, the duplicate data objects are stored repeatedly, resulting in inefficient use of storage resources and wasted energy to power them. Deduplication removes duplicate data objects and creates a data pointer to an exact copy that is already stored on disk.

How much space does deduplication actually save?

Deduplication vendors often claim that their products offer 20:1, 50:1, or even greater data reduction ratios. These claims actually refer to the "time-based" space savings effect of deduplication on repetitive data backups. Because data backups contain largely unchanged data, once the first full backup has been stored, all subsequent full backups see a very high occurrence of deduplication.

In non-backup data environments, such as file archiving or infrequently accessed unstructured data, the rules of time-based data reduction ratios do not apply. In these environments, volumes do not receive a steady supply of redundant full backups, but may still contain a large number of resident duplicate data objects. The ability to reduce the space requirements in these volumes through deduplication is measured in "spatial" terms. For example, if a 500GB data archived volume can be reduced to 300GB through deduplication, the spatial reduction is 40%.

Increase storage utilization rates by using FlexVol thin provisioning

According to industry estimates, storage utilization rates average 25% to 40%. Because it is difficult to predict

actual storage requirements, application administrators typically request much more space than they think they will need to protect themselves if they need more storage down the line. This common practice guarantees over-allocation — an estimated 60% to 75% of all storage capacity that is being powered goes unused. Not only is such a low utilization rate a waste of storage, it is a waste of power. With NetApp FlexVol technology, you can create logical pools of storage from physical groups of disks. This pooled capacity can be freely allocated among many data sets, which eliminates the need to assign physical reserves to each data set. With FlexVol thin provisioning, storage utilization typically rises from an average of 40% to 60% — a saving of 33%.

- William Beaumont Hospital, running 25 Oracle® databases, **experienced a 50% increase in utilization** when implementing Data ONTAP 7G with Flex Vol technology and consolidating from nine systems to three primary storage systems and one near-line system.

Reduce power consumption 40% or more with SATA disks

SATA drives offer the highest available storage density per drive and consume, on average, 40% less power than Fibre Channel drives of an equivalent capacity, enabling you to minimize watts/TB in the data center. According to a recent study, Oracle was able to **reduce their power consumption by 40% with SATA drives**.

Increase storage utilization with RAID-DP

To further maximize your efficiency and flexibility, you can pair NetApp SATA disks with our patented RAID-DP. Our dual-parity RAID-DP is a standard feature of the Data ONTAP operating system. When compared to RAID 10 data mirroring, RAID-DP offers up to 46% greater storage utilization. In addition, it enables you to recover from the simultaneous failure of two drives, unlike other RAID levels that can tolerate only a single drive failure. The increased performance of our RAID-DP means that you can use less-expensive SATA storage for your primary storage without worrying about data loss, while also lowering your storage acquisition costs.

Use minimal storage for snapshot copies with NetApp Snapshot technology

NetApp Snapshot copies provide two significant efficiency advantages. First, Snapshot copies consume minimal storage space. Second, these copies let you leverage a single copy of your data for multiple uses, reducing your reliance on special-purpose storage systems. Reducing the number of special-purpose storage systems you use can radically reduce your power requirements.

Lighten the storage load for testing with FlexClone writable Snapshot copies

Testing and development require numerous copies of your data and can stress your storage infrastructure. With NetApp FlexClone technology you can make multiple, instant virtual copies of your data with virtually no storage overhead. With writable copies, your savings are equal to the size of the cloned data set minus any subsequently changed blocks.

- Blackboard ASP **cut their storage needs by 33%** in their test and development environment by using NetApp FlexClone.
- **Oracle experienced a 403% increase in storage utilization** after implementing NetApp FlexClone in conjunction with a move to grid infrastructure.

Adding it all up

The only way to prevent runaway data growth is to stop data proliferation before it can clog up the works. Many of our technologies reduce raw storage requirements and eliminate duplicate data on your systems. Snapshot, FlexClone, deduplication, FlexVol, and RAID-DP are just a few of the technologies that can have a huge impact on your storage footprint.

So, when you add it all up, how does NetApp compare to other storage vendors in efficiency? Because NetApp solutions require less storage per usable terabyte, customers with NetApp systems dramatically reduce the number of hard disk drives and enclosures they need to store their data. According to a study by Oliver Wyman, information from customer deployments showed that as a direct result of NetApp RAID-DP, FlexVol, and Snapshot technologies, customers with NetApp storage systems need 50% less power and generate 50% less heat per usable terabyte of storage than comparable systems from other storage vendors in typical environments.



Brett Battles and Mark Woods

Brett (right) leads product management for storage subsystems and OEM solutions at NetApp. Brett held various positions in engineering and business development prior to joining NetApp. He has a BS degree from the Georgia Institute of Technology and a PhD from Stanford University.

Mark has over 15 years of experience in product management and marketing. Prior to joining NetApp, Mark worked for Hewlett-Packard in server businesses. He earned a BSEE from the University of Colorado and an MBA from the University of Texas.

Go further, faster™



Quick Links

- [netapp.com](#)
- [Tech OnTap Archive](#)

Explore

Sizing Storage for Oracle

Designing a storage system to deliver the necessary performance for Oracle® OLTP applications is not a trivial process. If you pay attention only to the storage capacity you need, without factoring in disk performance, you're likely to be unpleasantly surprised.

A recent Tech OnTap article provided an in-depth look at Oracle sizing, including:

- Disk performance considerations
- Gathering your database configuration
- Measuring I/O load
- Factoring in growth

[More](#)

FAS6070 Features and Performance

When the FAS6070 was released it offered considerable performance improvements and other technical enhancements compared to NetApp's previous storage systems. Tech OnTap presented an engineering round table to discuss the enhancements that yielded:

- A performance increase of 2X
- A capacity increase of 5X
- Improved resiliency
- The ability to mix FC and SATA drives

Read about the features that contribute to performance, scalability and resilience across the FAS6000 product line.

[More](#)

Considerations for Large-Capacity Storage Systems

By [Chris Lueth](#)

Rapid growth in disk size has made it possible to configure individual storage systems with staggering amounts of storage. For instance, you can take a NetApp FAS6080, add 1,176 1TB SATA disk drives, and end up with over a petabyte (1,000 terabytes) of raw storage capacity in a single system.

For many IT shops, the advantages of large-capacity systems are compelling. Large drives give you the lowest price per GB of storage. More capacity from fewer spindles means fewer disks and storage systems to manage, plus less energy consumption and reduced cooling requirements — major concerns in most data centers.

So, is there a down side? Not really, but there are important considerations you should take into account when using large disks. Disk capacity has grown much faster than disk quality or performance. Since the new larger disks are just as likely to fail as smaller ones and performance remains the same, reconstructing a failed 1TB disk is a bit like trying to fill a swimming pool with a garden hose. You have to be prepared to wait a while for the extended reconstruction process.

Longer reconstructs with large SATA disks don't mean that you shouldn't deploy large-capacity systems; you just have to be aware of their unique requirements and the limitations they impose. This article explains the issues you should be thinking about when considering a large-capacity system, including:

- Applications that are (and are not) suited to large-capacity systems
- Data availability
- Data protection
- RAID reconstruction
- RAID scrubs and background media scans
- Provisioning
- Infrastructure complexity

Much of the discussion applies to any large-capacity storage system, but I'll also provide some specifics for large-capacity NetApp systems along the way.



Target Applications

The first thing you should understand about large-capacity storage systems is that they can be limited by the performance characteristics of the underlying drives and are not suited for all applications. The largest-capacity drives on the market are all SATA disks rather than high-performance Fibre Channel. Regardless of capacity, modern SATA disks spin at the same rotational speed and provide the same throughput. In addition, to achieve the desired size for a given storage container (file system, LUN, etc.) you will be deploying fewer disks, and fewer disks generally mean lower maximum performance from the storage container.

Storage systems and/or host operating systems may also impose size limits that in turn limit the number of spindles you can use for a given storage container. For instance, Ext3, the default Linux® file system, has a maximum size of 16TB, so when using 1TB disks a single file system would be limited to 17 or so spindles when taking into account some lost capacity for such things as formatting.

When you think about large-capacity systems you should think about secondary storage; these systems are not well suited for Exchange, databases, or other applications that require low response times and high throughput. Ideal applications include:

- Disk-to-disk backup
- Target for data replication (using NetApp SnapMirror®, for instance)
- E-mail archiving
- File or document archiving
- Compliance storage

Secondary storage also lends itself well to applications that have large, sequential data streams, including:

- Image capture
- Live video capture
- Seismic data

Data Availability

Because they may have hundreds of SATA spindles, there are some important considerations for data availability on large-capacity systems, including:

- RAID
- High-availability configurations
- Multipath HA

Failure rates for SATA disks are typically higher than those for Fibre Channel, making RAID protection critical. NetApp typically recommends using the NetApp high-performance, dual-parity RAID 6 implementation, RAID-DP™, to protect against the data loss that can arise from double disk failure within a RAID group. Other vendors may offer dual-parity RAID 6 solutions, depending on the storage product. Regardless of the vendor, any large-capacity storage system will benefit over the life of the solution from the higher data resiliency offered with RAID 6.

Although large-capacity storage systems are typically used as secondary storage, NetApp customers who have deployed them usually opt for full high-availability configurations with active-active controllers and no single points of failure to ensure that large data stores remain accessible at all times. An important consideration for a large-capacity HA solution is the time required for one controller to take over disks from the other controller or give them back. This time can be slightly extended with high numbers of SATA disks in the solution versus what would typically be the case with only Fibre Channel disks. This occurs because the SATA disks themselves are slower and the health-checking processes take longer than they do with Fibre Channel drives.

Data ONTAP® 7.2.4 introduces some specific optimizations for takeover and giveback with SATA disks and can improve the performance of large-capacity SATA systems during failover and giveback, making the solution on par with solutions using only Fibre Channel disks. To benefit from these optimizations, we recommend using Data ONTAP 7.2.4 or later for any NetApp SATA-based large-capacity HA storage solution.

One NetApp storage configuration option that has been underutilized is multipath HA. Multipath HA ensures that there are two separate I/O paths from each controller to every disk so that cabling issues or other hardware problems don't disrupt access to disk drives. In HA configurations, such problems can cause failover to occur. Multipath HA reduces the chances of a failover by providing redundant data paths from each controller to its storage. Multipath HA can also help improve performance consistency by spreading the storage workload across the two data pathways.

Data Protection

Backing up the data from a large-capacity storage system also presents unique challenges. Disk-to-disk methods are preferred where possible to minimize backup time. However, with tools such as NetApp SnapVault® and SnapMirror, the time needed to create a baseline copy of a large-capacity storage system may be prohibitive. NetApp offers two tools: LREP (logical replication) and SnapMirror to Tape to assist in creating baselines that can then be seeded on remote systems. Afterwards, only changed blocks are replicated, reducing the impact to the source and destination controllers as well as the network in between the two.

RAID Reconstruction

As with most other system maintenance activities, RAID reconstruction times are extended with large SATA drives. For instance, if a 1TB disk fails, RAID reconstruction on a NetApp system will take approximately 10 to 12 hours when no other load is present. This time will be extended as system load increases.

Current mean time between failure (MTBF) data suggests that on a storage system with 1,176 1TB disk drives, a system could be engaged in reconstruction as much as 5% of the time under normal operating conditions. Once again, the percentage of time spent in reconstruction gets bigger as the overall workload on the storage system increases.

Media Scans and RAID Scrubs

NetApp uses regular media scans and RAID scrubs as a way to ensure the integrity of stored data, and I assume that other storage vendors provide similar capabilities to detect and correct problems. Much like painting a large bridge, where you start at one end, paint every day for months until you reach the other end, and then start all over, these two NetApp utilities simply keep track of their progress and continue working through the storage subsystem until all the storage has been checked. Background media scans run continuously at a low rate, using built-in diagnostics to detect media errors. RAID scrubs run by default on a weekly basis for six hours, using parity data to check data integrity.

On large-capacity storage systems, NetApp recommends increasing the data rate for media scans and increasing the frequency and duration of RAID scrub execution to ensure that infrequently accessed data (which is typical on secondary storage) gets checked on a timely basis.

Storage System Provisioning

When it comes time to provision a large-capacity system, you need to first find out what limits are imposed by your storage system (and host operating systems for SAN environments) and plan accordingly. For instance, on NetApp systems you can define a maximum of 100 aggregates or traditional volumes on a single storage controller, while the total number of aggregates, traditional volumes, and flexible volumes (FlexVol® volumes) cannot exceed 500. While these might seem like high limits, there are situations in which they can be exceeded. For instance, if your host operating system limits you to 2TB file systems, or if you have standardized on a high number of FlexVol volumes per aggregate, you could potentially run into the 500-container limit before you fully provisioned a maximum-capacity system.

The point is that when you're dealing with large-capacity systems, you can't just jump in and start provisioning. You have to understand what the various storage limits are and do the necessary up-front planning to ensure that you are able to use all your capacity while leaving room for unforeseen future requirements.

Infrastructure Complexity

A factor you shouldn't overlook when planning the deployment of a large-capacity system is the sheer complexity of the entire disk infrastructure. I recently worked with a customer who had 1,008 disks in 72 disk shelves. These were further subdivided into 12 storage loops, each with 6 shelves.

With an active-active environment using multipath HA storage connectivity, each storage loop requires 4 connections, resulting in 48 connections between storage and storage controllers across numerous storage cabinets. If it sounds like the cabling is complex, it is. You can't simply start cabling a maximum-capacity storage system without a plan and expect everything to come out okay. There's a lot of work to do up front to ensure that

everything works properly when you are done. Up-front planning, wire diagramming, and labeling are critical in large-capacity storage deployments.

Conclusion

By becoming aware of potential limitations and gotchas up front and choosing your applications wisely, you can safely deploy storage systems with capacities that just a few years ago would have seemed impossible. If you carefully consider your availability and data protection needs relative to the capacity-versus-throughput capability of the latest SATA drives and plan for provisioning and physical requirements up front, you can avoid the unpleasant surprises that come from pushing the envelope with any technology and reap the benefits of simplified management, reduced direct storage costs, and reduced power and cooling requirements.



Chris Lueth

Technical Marketing Engineer
NetApp

Chris has over 17 years of industry experience. Since joining NetApp five years ago, he has gained incredible technical breadth, including work on NearStore® deployments, RAID-DP, SnapLock®, mid-range and high-end platforms, and storage resiliency. He was previously a chip design engineer and worked on the first multiprocessor motherboard chipset before switching to UNIX® system administration and then, ultimately, to storage.

Go further, faster™



Quick Links

- [netapp.com](#)
- [Tech OnTap Archive](#)

Explore

Disk Drive Resiliency: How NetApp Protects You

You may be surprised by some of the "secret" problems that still lurk inside disk drives despite their remarkable dependability. This article describes five of the most troublesome disk problems:

- Drives fail suddenly.
- Drives slowly degrade.
- A bad drive can lock an FC loop.
- Firmware bugs can corrupt data.
- Committed writes can get dropped.

Learn about the resiliency technologies that NetApp engineering has developed to protect against these problems.

[More](#)

Five Little-Known Tips to Increase NetApp Storage Resiliency

By [Steve Lawler](#) and [HariPriya](#)

Over the years, NetApp storage has built a reputation for being simple, easy to manage, and resilient to the problems that can affect data availability. To achieve the highest levels of resiliency, a variety of best practices should be followed.

NetApp recently released a technical report that provides the complete details of [storage best practices for resiliency](#). In this article we provide a few tips you can use to enhance the resiliency of your NetApp storage:

- Use multipath high availability (multipath HA)
- Provide the right number of spare disk drives
- Use SyncMirror® for even greater resiliency
- Bulletproof your HA configurations for nondisruptive upgrades
- Verify your storage configuration using NetApp's automated tools

Tip #1: Use Multipath High Availability

Multipath high availability provides redundant paths between storage controllers and disks for both single-controller and active-active configurations. Having a second path to reach storage can protect against a variety of possible failures, such as:

- HBA or port failure
- Controller-to-shelf cable failure
- Shelf module failure
- Dual inter-shelf cable failure
- Secondary path failure in HA configurations




Figure 1) Multipath HA in an active-active controller configuration.

Even with clustered NetApp storage systems (active-active or HA configurations), multipath HA reduces the chance of a failover occurring and improves availability.

Multipath HA also offers potential performance benefits in situations in which Fibre Channel paths to disk shelves are overloaded by providing twice the bandwidth to your storage. This can be especially valuable when reconstruction is taking place and on older systems that use 1Gbit/sec Fibre Channel connections.

In many cases, open FC ports are already available on storage systems, so multipath HA can be added at the cost of a few cables. That's a small price to pay for a big potential payoff in resiliency.

Tip #2: Provide the Right Number of Spare Disk Drives

On NetApp storage, disk failures automatically trigger parity reconstructions of affected data onto a hot standby (spare) disk, assuming that a spare disk is available. If no spare disks are available, self-healing operations are not possible. The system will run in degraded mode (requests for data on the failed disk are satisfied by reconstructing the data using parity information) until a spare is provided or the failed disk is replaced. During this time, your data is at greater risk should an additional failure occur. (With NetApp RAID-DP™, a RAID group operating in degraded mode can undergo one additional disk failure without suffering data loss.)

The number of spares you need varies based on the number of disk drives attached to your storage system. For a lower-end FAS200 or FAS2000 with a single shelf, one spare disk may suffice (configure two if you want to use Maintenance Center). On the FAS6080, with a maximum spindle count of 1,176 disks, more spare disks are needed to ensure maximum storage resiliency, especially with larger SATA disks that have longer reconstruction times.

NetApp recommends using two spares per disk type for up to 100 disk drives, where disk type is determined by a unique interface type (FC, SATA, or SAS), capacity, and rotational speed. For instance, if you have a system with 28 300GB 15K FC disks and 28 144GB 15K FC disks, you should provide four spares: two of the 300GB capacity and two of the 144GB capacity.

For each additional 84 disks, another hot standby disk should be allocated to the spare pool. The following table provides some additional examples to illustrate this approach. (The table assumes all the disks are of a single type.)

Number of Shelves	Number of Disks	Recommended Spares
6	84	2
8	112	3
12	168	3
24	336	4
36	504	6
72	1,008	12
2	28	2

Table 1) Choosing the right number of spares for a given number of disks of the same type.

Note that if you are using NetApp Maintenance Center, you will need a minimum of two spare drives of each type in your system. Maintenance Center performs proactive health monitoring of disk drives and, when certain event thresholds are reached, it attempts preventive maintenance on the suspect disk drive. Two spare disks are required before a suspect disk drive can enter Maintenance Center for diagnostics.

Tip #3: Use SyncMirror for the Greatest Possible Resiliency

If you need even higher levels of resiliency than HA and RAID-DP offer, consider using SyncMirror in either a local or MetroCluster configuration.

Local SyncMirror provides synchronous mirroring between two different traditional volumes or aggregates on the same storage controller to ensure that a duplicate copy of data exists. This feature is available starting with Data ONTAP® 6.2. The mirroring provided by SyncMirror is layered on top of RAID protection (RAID 4, RAID-DP, or RAID 0 in V-Series).

SyncMirror stripes data across two mirrored storage pools known as plexes, which can result in read performance improvements on disk-bound workloads. It provides greater protection against multiple simultaneous failures across mirrors. SyncMirror with RAID-DP is so fault tolerant that it can ensure data availability with up to five simultaneous disk failures across mirrored RAID groups. Because SyncMirror uses native NetApp Snapshot™ technology to maintain synchronized checkpoints, resynchronization after loss of connectivity to one plex takes much less time. Only data that has changed since the most recent Snapshot checkpoint has to be synchronized.

SyncMirror also provides geographical disaster tolerance when used in conjunction with MetroCluster. SyncMirror is required as part of MetroCluster to ensure that an identical copy of the data exists in the remote data center in case the original data center becomes unavailable. When used in active-active configurations, SyncMirror provides the highest resiliency levels, ensuring continuous data availability.

Tip #4: Bulletproof Your HA Configurations for Nondisruptive Upgrades

Configuring your storage systems in an HA configuration with active-active storage controllers is a great way to eliminate single points of failure and increase resiliency. In addition to eliminating potential unplanned downtime, these configurations can also reduce planned downtime through nondisruptive upgrades.

Nondisruptive upgrades (NDUs) give you the ability to upgrade transparently any component in an active-active storage system (software, disk and shelf firmware, hardware components, etc.) with minimal disruption to client data access by doing a rolling upgrade. In order to perform a nondisruptive upgrade, the two storage controllers must be identical at the outset in terms of a variety of factors, including licenses, network access, and configured protocols. You can learn more about NDUs in [a recent Tech Report](#).

The best way to ensure that an upgrade goes smoothly is to check your systems well in advance to ensure that they meet NDU requirements. By meeting these requirements, you also ensure that your HA systems are optimally configured to provide the greatest possible resiliency and data availability. NetApp provides a set of automated tools to make this possible, as described in the following section.

Tip #5: Verify Your Storage Configuration with Automated Tools

Whether you have clustered HA storage systems or single-controller configurations, it's important to ensure that you have the right hardware, firmware, and software installed, especially before undertaking an upgrade. You may have dozens of disk shelves and hundreds or even thousands of disks, so this is no small task. Fortunately, NetApp Global Services (NGS) has developed a set of tools designed to automate processes that would otherwise be tedious and error prone. Running these tools periodically can increase the resiliency of your storage systems and simplify your operations.

➤ Cluster Configuration Checker

This tool detects and identifies the most common configuration causes of failover problems:

- Inconsistent licenses
- Inconsistent option settings
- Incorrectly configured network interfaces
- Different versions of Data ONTAP on the local and partner nodes
- Differences in the cfmode configuration settings between the two nodes

Cluster Configuration Checker is also available as part of NetApp Operations Manager.

Upgrade Advisor has been designed as a one-stop solution to qualify a storage system for a Data ONTAP upgrade. The tool uses live AutoSupport data to first automate the normally painful manual process of documenting every caveat and requirement associated with determining a system's eligibility and then generate a step-by-step upgrade plan for use in upgrading as well as backing out an upgrade.

1 **Current OS**

Target Version

Please select warnings you would like included in the upgrade and downgrade plans, specify method, and click continue.

 - Upgrade warning,  - Back-out warning

* Description	cvnetapp13
 The firmware on the motherboard is 4.3_i2 and is not up to date. The latest firmware for this platform is 4.3.1. Upgrading this firmware is recommended.	☒
 This system has ESH modules that are below the latest release. It is recommended to upgrade these modules to 95.	☒
 This system has ESH2 modules that are below the latest release. It is recommended to upgrade these modules to 18.	☒
 This system has ESH modules and at least one Qlogic 2312 HBA. It is recommended that the jumpers be checked. See Qlogic 2312 HBA .	☒

Method

[PREVIOUS](#) [CONTINUE](#)

Conclusion



Steve focuses exclusively on high-availability storage configurations. With over 15 years of industry experience, he previously worked in telecom, where he gained broad experience supporting enterprise-level customers.

HariPriya specializes in storage resiliency, including disk drives and shelves. She previously worked at Hewlett-Packard, where she focused on RAID and storage issues. HariPriya holds a master's degree in computer science and is currently working on an MBA.