

HIGHLIGHTS

- » **Tips for Optimizing SnapMirror®**
- » **Simplifying Backup Management**
- » **VMware and NetApp Snapshot Copies**
- » **Resources:**
 - **Podcast: High-Performance NAS**
 - **Demo: Protection Manager™**

Five Ways to Use NetApp Snapshot™ Copies with VMware VI3

Mike Slisinger, Professional Services consultant, NetApp

NetApp Snapshot copies incur no performance penalty and virtually no overhead, even when hundreds are retained. See how IT teams leverage this unique technology to streamline backup and recovery, DR, and more in VMware environments.

[More »](#)

BLOGGING WITH DAVE

Dave Hitz, NetApp Founder and EVP

"There are many ways storage companies can help you reduce power, but—surprisingly—more efficient hardware is low on the list."

» **Dave's Blog**

DRILL DOWN

- » **NetApp Vision and Product Strategy Presentations from Analyst Day '07**
Don't miss Dave Hitz's 3-5 year perspective on virtualization, protecting data, and separating data from storage.
- » **Best Practices for Implementing Scale-Out Storage Architectures**
Outlines a process for characterizing application/performance needs and developing a data distribution strategy.
- » **High-Performance NAS: Oxymoron?**
Podcast discussion on how next-gen NAS solutions scale out, plus best practices.



ADMIN RESOURCES

- » **End-User Review: Hobbit**
NetApp users explain how a tool inspired by Big Brother helps them monitor 100s of servers, switches, and storage systems.

TIPS FROM THE TRENCHES


Three Common Backup and Replication Management Challenges Solved

Jai Desai, product and partner engineer, NetApp

Wish you could manage all of your replication schedules from a single interface and automate routine tasks? See how NetApp Protection Manager enables you to easily:

- Make sure all data is protected
- Eliminate repetitive, manual setup
- Rapidly change global policies

Includes a 13-minute demo plus a mini-FAQ addressing common customer questions.



[More »](#)

ENGINEERING TALK

SnapMirror: The Director's Cut

Stephen Manley, technical director, NetApp Engineering

One of the original NetApp SnapMirror engineers provides an overview of core SnapMirror design and how it is used in a cross-section of real-world deployments. Plus discover the tricks that the pros use to avoid or resolve common challenges.

[More »](#)



Mike Slisinger

Professional Services Consultant, NetApp

Mike joined NetApp in 2003 and serves as a specialist in VMware, Windows® applications, and SAN. During his tenure at NetApp, Mike has been heavily involved in VMware/NetApp integration. He has authored several Technical Reports involving VMware, and has been directly involved in many large deployments of VMware, Exchange, SQL Server, and SAN. At his previous company, Mike led a large Windows/UNIX® consolidation project using NetApp storage that resulted in a 42:1 consolidation ratio.

Five Ways to Use NetApp Snapshot™ Copies with VMware VI3

By Mike Slisinger

Using VMware Infrastructure 3 (VI3) with NetApp storage has proven to be a topic of keen interest to Tech OnTap readers. A recent spot survey on VMware and NetApp identified backup and restore as the area of greatest interest, closely followed by high availability and disaster recovery. The use of NetApp Snapshot copies with VMware figured prominently in your comments, serving as the impetus for this article.

Like Microsoft (via VSS) and Oracle, VMware offers a native snapshot functionality that can capture previous states of data. However, since copy creation is done via the VMware server, the creation and retention of snapshots consumes resources on the VMware server. This additional overhead can impact the performance of the entire VI3 environment.



NetApp Snapshot technology represents an ideal complement to the native snapshot capabilities built into VMware. Because NetApp Snapshot copies occur on the storage system, actions such as backup and restore have no impact on server resources or server performance.

Additionally, NetApp Snapshot technology can protect any datastore used by VMware in a single operation. For instance, if a VMFS datastore contains multiple virtual machines, a NetApp Snapshot copy can capture the state of all of them in a single operation.

Although these advantages are important, the key differentiator of NetApp Snapshot copies is that – unlike competing snapshot technologies that use copy-on-write algorithms or make complete data copies – NetApp Snapshot technology is extremely space and resource efficient.

The NetApp WAFL® file system never overwrites existing blocks, so as blocks are changed NetApp simply keeps the original blocks to maintain a consistent, point-in-time view. This means that additional storage is consumed only as changes are made, and there is virtually no overhead on storage system processors, busses, and so on, even when hundreds of Snapshot copies are retained. For more details, read the NetApp technical report [A Storage Networking Appliance](#).

This article examines five ways in which IT teams leverage NetApp Snapshot technology and related NetApp technologies in their VMware environments to obtain:



RELATED INFORMATION

- [Podcast: Data Protection in VMware Environments](#)
- [Love VMware, Hate Backups?](#)
- [Best Practices for VMware ESX 3.0 on NetApp Storage \(PDF\)](#)
- [SAN Boot with VMware/NetApp](#)

Love VMware but Hate Backups?

While no one disputes the obvious benefits of VMware and server virtualization, the radical change in architecture has significant implications for backup, HA, disaster recovery, and so on.

NetApp VMware expert Vaughn Stewart tackles some of the thorniest questions posed by end users. Examples include:

- The fact that a single service disruption could affect my entire environment is keeping me awake at night. Do I have to live with this level of insecurity?
- My backup processes haven't changed, but suddenly I'm spending almost all my time managing backups. How do I stop the madness?

[Read Love VMware, Hate Backups?](#)

Server Consolidation and Business Continuity in VMware Environments

Do you need to create a complete environment for server consolidation and disaster recovery? In this VMworld presentation, Mike Gower of I.T. Blueprint Solutions Consulting Inc.

1. Near instantaneous backup
2. Fast and flexible virtual machine recovery
3. Accelerated data management
4. Robust disaster recovery
5. Application data protection and management



Figure 1) A comparison of snapshot technologies.

#1 Near Instantaneous Backup of Virtual Machines

Backing up VMware servers can be a challenging proposition because of the sheer number of virtual machines – dozens to hundreds – that must be protected. In addition, a virtualized server environment has less total I/O resources than a physical server environment. This means that you probably need to change the way you do backups when you move to a virtual environment, and this is why it is helpful to offload backups onto storage systems wherever possible.

Because NetApp Snapshot copies don't actually move any data, they occur almost instantaneously. Users typically retain a specified number of Snapshot copies on their storage systems for convenient restore. NetApp Snapshot copies can also be used as the basis for backing up to tape. SnapVault®, the NetApp disk-to-disk backup solution, can also be used to move Snapshot copies to a secondary disk storage system for greater data protection.

Loyola Marymount University (LMU) recently moved to a virtualized server and storage environment with VMware and NetApp to overcome problems with server sprawl; limited power, cooling, and real estate; and to improve availability for critical applications.

LMU's VMware environment consists of nine ESX 3.0.1 hosts in three clusters connected to a NetApp FAS3050 storage system via iSCSI. Storage for each virtual machine is provided by one or more RDM (raw device mapping) LUNs. LMU uses NetApp Snapshot technology to protect its VMware environment, citing such reasons as:

- Zero downtime
- Zero CPU load on host
- No need for backup agents

LMU developed its own script to automate backups and is working with NetApp to document its implementation. This script uses the VMware snapshot capability to quiesce the virtual machine so that a consistent NetApp Snapshot copy can be created. You can find out more about LMU's VMware environment in [a recent NetApp/VMware Webcast](#). You can learn more about the use of NetApp Snapshot technology for VMware data protection by reading the NetApp technical report [Best Practices for NetApp and VMware ESX Server 3.0](#).

#2 Fast and Flexible Virtual Machine Recovery

Recovery is a critical issue in any data center environment. Studies have shown that 40% of application downtime is caused by operator errors, another 40% is due to application errors, and only 20% is caused by technology failures (Gartner Group, 2000).

If you have dozens or hundreds of virtual machines, problems requiring recovery are almost a certainty. As mentioned earlier, most restores are not due to hardware failure. Instead, they involve an accidental deletion or an event where a virtual machine instance becomes corrupted. Protecting your VMware environment with NetApp Snapshot technology ensures that you have multiple restore options on the existing storage system:

- Use NetApp SnapRestore® to revert to the latest (or most appropriate) saved Snapshot copy. This effectively resets your environment to the state it was in at the time the copy was created.
- Use NetApp FlexClone® to clone the Snapshot copy and create a writeable version, mount the resulting clone, and extract the broken or missing component you need. For instance, if a volume contains multiple virtual



Figure 2) A virtualized server-to-storage environment.

and Chris Fraser of Long View Systems explain how they designed a complete solution for Canadian customer Impark.

By combining VMware ESX with NetApp storage systems and NetApp SnapMirror software for data replication, they were able to craft a complete solution with multiple tiers of service to meet the customer's specific needs.

Technical details of the solution are provided, along with motivations, observed benefits, and future plans.

[Download the full presentation and audio.](#)

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, including all the local snapshot copies, and can be mounted read/write to recover from failure. If a snapshot backup is deleted on the source, it goes away on the mirror at the next replication.

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

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

Get the details. Read the reports:

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

The Core NetApp DNA

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

These basic building blocks of NetApp technology support:

- Multiprotocol environments

machines, this method can be used to restore a single one without affecting other virtual machines.

To return to the example discussed in the previous section, LMU primarily utilizes the second approach. The Snapshot copies that LMU creates for backup purposes include multiple virtual machines. Using FlexClone, LMU can quickly create a writable version of the RDM used by a particular virtual machine and return it to service. This is well illustrated in the [recent VMware Webcast](#). LMU cites the process as being nearly instantaneous (reducing restore time from hours to minutes) and easily automated.

#3 Accelerated Data Management Using Cloning

As mentioned in the previous section, NetApp FlexClone is a useful tool for recovery in VMware environments. The FlexClone capability is built on two unique features of NetApp Data ONTAP®: Snapshot copies and flexible volume technology (FlexVol™ volumes). A FlexVol volume is a virtual and dynamic storage entity that can grow or shrink without being constrained by the characteristics of underlying physical storage. FlexVol volumes are built within aggregates (an underlying set of disks defined by the storage administrator on a NetApp storage system). Any FlexVol volume automatically utilizes all spindles in its underlying aggregate for guaranteed performance, regardless of volume size.

FlexClone uses the same mechanism as Snapshot technology to create writable clones of volumes or LUNs (FlexVol volumes). As with snapshots, clones can be created nearly instantaneously with effectively no storage overhead; they share the same underlying storage blocks as the baseline volume/LUN. As the baseline and the clone diverge (for example, due to updates to the clone's data blocks), those new blocks are written to disk, and the volume accumulates only the changed blocks.



Figure 3) NetApp FlexVol technology.

Naturally, this ability to create space-efficient clones of important data volumes is valuable in any server environment. FlexClone volumes are frequently used in database environments to create copies of production databases for test, development, QA, and so on. In addition to the recovery example discussed earlier, there are several additional uses of FlexClone that are unique to virtual environments:

- **Clone a production virtual machine for testing.** Have you ever wanted to test something with production data? With FlexClone, you can instantaneously clone an entire VMFS and use the cloned volume in your lab environment to duplicate the production server without scrounging up double the storage or waiting for time-consuming data copies.
- **Create a gold image of a virtual machine configuration and clone it to provision new virtual machines for permanent or temporary use.** This works best when you are using an RDM LUN for each virtual machine. Once you have the virtual machine configured the way you want it with appropriate software and so on, you can use FlexClone to duplicate it many times in a matter of minutes. This method is extremely space efficient and avoids the copy-out time required by other provisioning methods.

For example, I'm aware of an independent software vendor that uses the latter cloning method to support its sales staff. The company has multiple software products and serves a variety of markets. It maintains a gold master virtual machine configuration for each application in each market, including appropriate data sets. A salesperson at a customer site can access this facility through the Web. A FlexClone volume is automatically made of the appropriate gold master, providing a virtual demonstration environment tailored to the particular customer and allowing the demonstration to start immediately. When the demonstration is complete, the FlexClone volume is deleted and the original gold master is not affected.

This approach also makes sense in lab and development environments, classrooms, or anywhere you need to rapidly provision a set of virtual machines with defined capabilities for temporary use.

#4 Robust Disaster Recovery

(NFS, CIFS, FC, iSCSI, and so on)

- Clustered failover, mirroring, and disk-to-disk backup
- RAID-DP™ and other software-based resiliency features
- The near-instantaneous creation of writable clones
- Block deduplication using A-SIS

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

NetApp Solutions for VMware Environments

In this Webcast, VMware vice president Brian Byun and NetApp vice president Patrick Rogers discuss the advantages of server to storage virtualization and how NetApp and VMware products compliment each other.

Senior systems administrator Jason Guibert of Loyola Marymount University describes his real-world experience deploying and using VMware and NetApp. Key topics include HA, backup and recovery, and cloning.

In a follow-up Q&A, Webcast participants asked Jason about specific technical aspects of his deployment, including iSCSI, RDM implementation, cloning, and failover.

[Watch the Webcast](#)

The virtual infrastructure of VMware lends itself to disaster recovery. If you have a copy of your virtual machine data in a remote location, you can rapidly restart one or many virtual machines, as necessary. However, VMware does not have replication mechanisms built in, so this replication has to be done by third-party software either on the VMware host itself or, preferably, externally. This makes sense, given the potential for I/O constraints on server hardware hosting numerous virtual machines.

The most effective way to provide disaster recovery for VMware involves asynchronous replication over a wide area network (WAN) using NetApp SnapMirror® software. SnapMirror uses the built-in Snapshot capability of Data ONTAP to efficiently identify and replicate only changed blocks. Other common replication methods use time-consuming and resource-intensive file system walks to check timestamps and identify files that have changed, and then replicate whole files, even when only a single block is altered. When SnapMirror prepares to replicate, the first thing it does is create a Snapshot copy. By comparing the current Snapshot copy to the one from the previous round of replication, changed blocks (not files) can be quickly and easily identified for replication, saving significant time, storage system resources, and network bandwidth.

At a recent VMware technical conference, VMworld 2006, a Canadian reseller of VMware and NetApp presented a detailed technical explanation of a VMware and SnapMirror business continuance implementation for its customer, Imperial Parking. The company has two main data centers, in Calgary and Vancouver, over 500 miles apart. Different service level categories (gold, silver, and bronze) were defined and each virtual machine was assigned to a category based on business requirements. Appropriate processes were defined for restarting services in the event of failure.

Imperial Parking estimates that in the first year, the new configuration resulted in over \$480K Canadian in avoided costs and/or savings. To learn more, download the presentation and audio: [Server Consolidation and Business Continuity Using VMware ESX Server and NetApp SnapMirror Technology](#).

#5 Application Data Protection and Management

The final piece of the puzzle is the ability to protect and manage data for Enterprise applications running within virtual machines. As with everything else we've discussed, the preference here is to move as much of the heavy lifting as possible off of the VMware server. NetApp accomplishes this through its SnapManager® software, which includes versions tailored for Microsoft® Exchange and SQL Server as well as Oracle® databases.

These software products run within the virtual machine, which currently uses the Microsoft iSCSI software initiator to access NetApp storage directly. A future version of SnapDrive® for Windows will support RDM LUNs, which will allow virtual machines to use either the Fibre Channel or iSCSI initiators of the VMware server itself.

Each SnapManager product is designed to work in concert with the application to ensure that databases are appropriately quiesced before Snapshot copies are created, ensuring consistent and verified backups while avoiding the need for excessive data movement by the server. Restores, tracking of valid Snapshot copies, and other necessary data management tasks are carried out correctly without disruption, and with minimal server overhead.

The Pennsylvania Attorney General's office recently did a complete upgrade of its IT infrastructure, moving from over 130 standalone servers with direct-attached storage to an environment consisting of 40 dual-processor blade servers running VMware ESX Server and a pair of NetApp storage systems—one for primary storage and one in its disaster recovery site. NetApp SnapManager for Exchange with Single Mailbox Recovery and NetApp SnapManager for SQL Server are used to provide backup/restore and data management capabilities within the virtual environment. The customer cited the space-efficient, disk-based, verified backups of Exchange along with the ability to quickly recover entire mailboxes, single messages, and contacts as particular benefits of the overall solution.

To learn more about this deployment [view a recent webcast](#) or read the [customer success story](#).

Conclusion

The unique capabilities of NetApp Snapshot technology form the basis of an ideal set

of data protection and management solutions for a VMware Virtual Infrastructure. Backup, restore, cloning, disaster recovery, and application data management are all facilitated through Snapshot copies or products based on Snapshot technology.

[Comment on this article](#)



Jai Desai

Product and Partner Engineer, NetApp

Jai spent more than four years working in all areas of QA for Cisco Systems prior to joining NetApp as a requirements engineer in April 2005. In this role, he documented deployments at top NetApp accounts and worked with NetApp engineering to simulate these environments (and planned architectural modifications) during new release testing. In his current role, Jai focuses on the Data Suite in the NetApp manageability software family. He writes NetApp best practice guides for products such as Protection Manager™ and Virtual File Manager™ and consults with NetApp SE teams around the globe regarding customer implementations.

Three Common Backup and Replication Management Challenges Solved


By Jai Desai

During a recent [analyst day presentation](#), Dave Hitz explained that the NetApp vision for managing complex IT environments is built around the idea of empowering users to think in terms of data and data management rather than the underlying physical storage.

An example of the practical, real-world implementation of this vision involves NetApp Protection Manager. Protection Manager not only simplifies the configuration, management, and monitoring of Snapshot™ copies, disk-based backup, and replication across an entire NetApp storage infrastructure, it can also incorporate Windows®, UNIX®, and Linux® server backup using Open Systems SnapVault® as part of a broad data protection strategy that can encompass remote offices and virtually any existing storage array.

This month Tech OnTap is pleased to feature a nine-minute [Protection Manager Technical Demo](#) previously shared only in customer presentations. This demo illustrates key components of Protection Manager and provides a step-by-step guide to solving three common administrative challenges:

1. **How do you make sure everything is protected when data is distributed everywhere?**
2. **How do you scale your data protection environment without spending all day on tedious, manual tasks?**
3. **How do you rapidly roll out global changes across all sites and systems?**



A quick summary of how Protection Manager addresses these administrative challenges is highlighted below; to see exactly how these tasks are accomplished [watch the technical demo](#).

Challenge 1: Make Sure All Data Is Protected

Historically, managing backup and disaster recovery across multiple sites, systems, and applications has been a time-consuming and often manual process. Storage admins can easily get bogged down in the details of managing huge numbers of storage systems, disks, and tapes, rather than staying focused on the important data they contain.

RELATED INFORMATION

- [Protection Manager Demo](#)
- [Tech Report: Protecting Data with Policy-Driven Data Management \(pdf\)](#)
- [Experts OnTap: How One IT Team Is Choosing a Backup Solution](#)
- [From Reactive to Proactive Admin](#)

A Quick Primer on NetApp Data Protection Software

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

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

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

At a more technical level, SnapVault takes a point-in-time image based on qtrees, while SnapMirror copies an entire image at the level of a LUN inside a volume.

Get the details. Read the reports:

- [SnapMirror Best Practices \(pdf\)](#)

One of the biggest challenges involves simply ensuring that all data is protected. Changes happen rapidly and aren't always communicated clearly from the application team to the storage team or the backup team. How can you be sure that critical data is actually scheduled to be backed up? Unfortunately, all too often it is when you go to restore and the data isn't there.

In a recent well-publicized example, the State of Alaska had to spend about \$200,000 dollars to re-scan over 800,000 document images that had not been backed up. The documents were lost while a technician was trying to fix a problem with the traditional SAN array. The technician subsequently discovered that a critical database file had "inadvertently" been left out of the backup. Without it, the database containing the images could not be recovered. Ouch. ([Read a full account of the incident in Enterprise IT Planet](#))

The traditional approach to this problem involves running frequent reports and spending a lot of time and effort to make sure that all data is protected. When NetApp designed Protection Manager, this was recognized as a critical pain point. Protection Manager includes a simple "dashboard" that is the starting point for most management tasks. The "Data Protection Status" and "Unprotected Data" panes allow admins to quickly identify and address unprotected data.

The "Data Protection Status" section includes detailed failure information. These warnings and errors enable quick identification when SLAs are jeopardized.

Additionally, as new storage or hosts are added to the environment, Protection Manager automatically detects them and summarizes this info in the "Unprotected Data" section. This enables an admin to easily see all unprotected data in one place and drill down to take the necessary actions to back up and/or mirror the data.




Figure 2. NetApp Protection Manager Dashboard

Challenge 2: Eliminate Repetitive, Manual Setup Processes

With storage capacity increasing rapidly, new applications and new data are introduced on a regular basis. Unfortunately, adding them into an existing backup environment is not a simple task. Admins must first figure out if there is enough capacity on the existing destination and if there is enough network bandwidth to accommodate the backup window. Then schedules, retention policies, destinations, and other variables for each individual server or volume must be established. The process is tedious and can be error prone.

Protection Manager significantly streamlines this process by allowing admins to get the job done in just four easy steps:

1. **Create a new "dataset" to contain the data objects to be protected.** A dataset is a collection of data objects with similar protection requirements. If the new data has the same protection requirements as an existing dataset, simply add it to the existing dataset. That's it. All subsequent tasks are automated by the software.
2. **Select a policy to apply to the dataset.** Protection Manager includes a set of pre-defined policy templates. Admins can use an existing policy or create a new policy to accommodate different schedules.
3. **Select a resource pool.** Protection Manager ensures maximum utilization of secondary storage resources by grouping them into "resource pools." Once a resource pool reaches a defined threshold, an alert is triggered to provision more storage for that pool.
4. **Commit changes.**

Protection Manager automatically correlates the datasets and underlying physical storage resources, so admins don't need to worry about all the particulars of the storage infrastructure. The conformance checker built into Protection Manager validates the configuration to ensure that everything is set up properly. Protection Manager then provisions the secondary storage and initializes the backup. Heavy lifting is done *automatically* by the software, and Protection Manager sends automated alerts if any of the policies are violated or any manual changes are made on the storage system. Alerts are also triggered if any of the data protection operations run into trouble.

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

A Better Way to Do Backup and Recovery

In a recent Webcast, Agilent Systems architect Tim Dietrich joined a team of NetApp data protection specialists to discuss how his IT organization reduced backup times by 98%, ensured data protection, and increased the availability of business applications during backups.



[Watch the Webcast!](#)

Read [A Mind Shift in Data Protection](#)

Challenge 3: Make Global Changes

As government regulations, business policies, and corporate requirements change, an organization's data protection policies are adjusted to meet new SLAs. Suppose you have a corporate mandate to change your global backup schedule, perhaps from a minimum of once every 24 hours to once every 12 hours. How long will that take to carry out and how much manpower will it require to implement the new policy in all data centers and all remote offices? Increasing the number of backups may have significant implications for secondary storage, so in this scenario admins would often need to spend significant time checking out secondary disk and tape resources to make sure they won't be exceeded in addition to everything else.

Once again, Protection Manager makes a potentially complicated and error-prone task simple and straightforward. Two simple but powerful features—datasets and global policies—make it possible to quickly implement changes across an entire organization. From the Protection Manager console, admins simply change the policy or policies that apply to key datasets. Backup schedules and retention policies can be changed with just a few clicks.

Because Protection Manager has direct access to primary and secondary storage, there are no surprises. The conformance checker verifies sufficient capacity in the resource pools to accommodate the backup of all the volumes selected in the dataset and generates alerts to provision more capacity as necessary. And, because NetApp Snapshot and SnapVault transfer changed blocks rather than whole files and utilize space very efficiently, this type of change typically has less resource impact than it might with other disk-based or tape-based backup methods.

Watch the Technical Demo

For an in-depth look at the Protection Manager interface and detailed, step-by-step examples showing how key tasks are accomplished using this tool, watch the nine-minute [Protection Manager Demo](#).

Protection Manager Drill-Down: A Mini-FAQ

Q: Is Protection Manager just an upgrade to the Business Continuity Management option of NetApp Operations Manager (DFM)? What's the difference?

Protection Manager is definitely not just an upgrade to the Business Continuity Management option. It adds policy-based management with significant new functionality that makes it possible for administrators to focus less on the details of physical storage by using abstractions such as datasets and resource pools. The tool significantly simplifies the day-to-day tasks associated with data protection.

Q: I have more than 500 SnapVault relationships. Can Protection Manager detect those? How will they translate to the new concepts of datasets, policies, and storage pools?

Protection Manager can import existing SnapVault and SnapMirror relationships. Select the right policy depending on existing requirements and Protection Manager will take over the relationship and maintain the schedule applied to the policy.

Q: What is the role of the Operations Manager server? Who sends alerts?

The Operations Manager server (formerly known as DataFabric® Manager, or DFM) is the core of the NetApp manageability suite and is required for Protection Manager operation. Relevant information is stored in the Operations Manager database and alerts are sent from Operations Manager. The Operations Manager server can be configured on a Microsoft® cluster for high-availability operation.

Q: Can Protection Manager work with more than one Operations Manager server?

Yes. You can connect to multiple Operations Manager servers using Protection Manager. Any configuration changes are stored in the Operations Manager database.

Q: When I change a schedule for a SnapVault relationship I have previously created with Protection Manager, does Protection Manager handle everything?

Absolutely! Any change to the policy schedule is applied across all the relations

within that dataset.

Q: Are NetApp SnapManager® products supported in Protection Manager?

This is a key development area for the next 12 months. NetApp is focused on supporting a centralized model in which Protection Manager is the central console for data and application backup as well as a distributed model in which database admins can continue using SnapManager to perform backup but take advantage of the dataset and policy concepts.

Q: Can Protection Manager automatically detect Open Systems SnapVault (OSSV) clients? Can you automatically apply a policy to a new OSSV client?

Protection Manager can automatically detect OSSV clients, but does not automatically apply a policy because there could be many policies associated with OSSV backup. However, adding new clients to an existing dataset is a drag- and-drop activity from the Protection Manager dashboard.

[Comment on this article](#)



Stephen Manley

Technical Director of Data Protection, NetApp

Stephen joined NetApp Engineering in 1997 after graduating from Harvard University with a BA in computer science. He has helped define and develop NetApp tape and disk-to-disk backup, mirroring, and compliance technologies, including SnapMirror®. Stephen has traveled the world over, working with customers and backup partners to provide open, reliable, and simple data protection solutions one data center at a time.

SnapMirror: The Director's Cut

By Stephen Manley

The famous story about Abe Lincoln writing the Gettysburg Address on the back of an envelope while riding a train is just that — a story. He had actually painstakingly written the speech days before, and just used the envelope to help him focus on the key points.

NetApp SnapMirror demonstrates how years of hard work can change an industry. When you think of SnapMirror, three things should come to mind: low cost, simple, and reliable. SnapMirror enables disaster recovery for lower tiers of data with options for simple, low-cost mirrors. It offers unique configuration choices to minimize new hardware purchases and optimize existing primary storage. SnapMirror copies can become part of daily administrator and business activity, while improving the reliability of disaster recovery (DR). SnapMirror low-cost, simple, reliable mirroring delivers access to your data, anytime, anywhere, for any use.

This article details three different ways that IT teams are leveraging SnapMirror today:

- [SnapMirror for Disaster Recovery](#)
- [SnapMirror for End-to-End Data Protection](#)
- [SnapMirror Beyond DR](#)
- [Q&A: Troubleshooting, Configurations, and More](#)

[Click Here to Read the SnapMirror Q&A](#)

A detailed Q&A at the end of the article addresses common SnapMirror questions.

1. SnapMirror for Disaster Recovery *Lower the Cost of DR and Protect More Data*

SnapMirror is designed to be inexpensive, simple, and reliable, so that customers can deploy mirroring for all of their business critical data.

Let's face it. Disasters don't strike only the most critical, most protected data, and they often reveal the importance of (now unavailable) lower-tier applications and data. Just as IT organizations tier storage, some now offer mirroring tiers to meet varying DR SLAs. The choice should not be either "mirror over an expensive network to an expensive system" or "use tape." With asynchronous mirroring to ATA storage, companies can abandon their "load the tape and pray" DR plans.

Over five years ago, SnapMirror pioneered mirroring from FC-based primary storage systems to ATA NearStore® systems. Today, our largest customers are rapidly expanding their ATA mirror infrastructure. Furthermore, periodic (for example, hourly or nightly) asynchronous mirrors significantly reduce network traffic because most new data is either repeatedly modified or rapidly removed. By designing SnapMirror to transfer only "net-new" data, most customers see a reduction in network load of 25% to 90%. Lower-cost DR means that you can now afford to protect more of your data.

RELATED INFORMATION

- [SnapMirror Best Practices \(pdf\)](#)
- [Simplified Backup/Replication with NetApp Protection Manager](#)
- [Podcast: Data Protection in VMware Environments](#)

A Quick Primer on NetApp Data Protection Software

NetApp customers have two primary 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, including all the local Snapshot copies, and can be mounted read-write to recover from failure. If a Snapshot backup is deleted on the source, it goes away on the mirror at the next replication.

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

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

Get the details. Read the reports:

- [SnapMirror Best Practices \(pdf\)](#)
- [SnapVault Best Practices \(pdf\)](#)

Can you afford not to?

For DR mirrors to expand and protect new tiers of data, mirroring must be simplified. Sadly, managing mirrors today is such a high-stress job that many admins "draw straws" to pick who will suffer. Like snapshots, though, mirroring should be painless; there are no moving parts, no impossible backup windows, and no processes to manage.

To ensure its simplicity, NetApp built SnapMirror on the foundation of Snapshot™ copies. To set up SnapMirror, an admin initializes the mirror (through network or tape), sets a cron-like schedule, and... that's it. With two steps, he or she can protect applications and files, SAN and NAS, remote offices and the data center. DR should always be so easy that one administrator can manage petabytes of mirrors. To further simplify SnapMirror management, NetApp recently introduced Protection Manager™. You can learn more about it in [Three Common Backup and Replication Challenges Solved](#), which includes a demo of the product in action.

Of course, for the new, simple, tiered DR solutions to have value, mirrors must reliably protect against both physical and logical disasters. Unfortunately, historically, "Mirrors mirror everything, including corruption." If a rogue process deletes half your files, that data will also be destroyed on your mirror. Because application and user disasters occur more frequently than physical outages, traditional mirrors leave companies exposed. SnapMirror builds on NetApp Snapshot technology to close this gap.

SnapMirror protects the current data while also preserving the historical Snapshot copies, so administrators can recover from both physical and logical disasters. Your mirror should protect against all disasters; otherwise it's doing only half its job.

One of my favorite customers once said, "I'm glad you like talking about SnapMirror, but I already bought it. Just help me set it up." In his honor, some practical Q&A is included at the end of this article.

Questions for this section include:

- How do I decide whether to mirror to ATA or FC?
- I think my network is my SnapMirror bottleneck. How can I be sure?
- SnapMirror can't meet its schedule because of the network. Can I change my SnapMirror configuration rather than buying a faster pipe or a compression device?
- How can I optimize my SnapMirror performance on ATA?

[View Q&A related to this section](#)

2. SnapMirror for End-to-End Data Protection *If You Can't Trust Your Mirror, What Can You Trust?*

By focusing on cost, simplicity, and reliability, SnapMirror improves the performance of primary storage systems while simplifying end-to-end data protection.

Not only can mirrors be cheaper to deploy and manage, but they should also reduce the overall cost of tape backups. With explosive data growth and 24x7 operations, admins now run backups during business hours. To compensate for the constant backup load, companies overspend to overprovision their servers and storage. Meanwhile, they maintain an idle mirror as an insurance policy. Why shouldn't the mirror handle the backup, so the primary storage can serve applications and users? Because SnapMirror replicates the active data and Snapshot copies, administrators can back up any application-consistent Snapshot copy while continuing to update the mirror. To reduce cost, tape backups should be run on the mirrors, leaving primary servers and storage for applications and users. Because power and cooling challenges seem to be growing as quickly as data protection struggles, mirrors should not require extra, dedicated hardware. IT departments consolidate to reduce complexity and power consumption. Therefore, mirrors should not double (or worse) the hardware footprint.

SnapMirror
enables admins to
recover from both
physical and
logical disasters.

Mirrors should
simplify and
optimize your
environment not

- [Data Protection for NetApp Storage Systems \(pdf\)](#)

Simplified Backup and Replication Management with NetApp Protection Manager

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

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

Watch a detailed demo that shows how Protection Manager solves three common management challenges:

- How do you make sure that everything is protected when data is distributed everywhere?
- How do you scale your data protection environment without spending all day on tedious manual tasks?
- How do you rapidly roll out global changes across all sites and systems?

[Watch the demo.](#)

Read [Three Backup and Replication Management Challenges Solved](#).

How NetApp 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.

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

SnapMirror uses NetApp Flex[*] technology to minimize its footprint. First, a customer can use SnapMirror to replicate multiple source systems to one destination system.

Second, SnapMirror can replicate System A's data to System B, while also replicating System B's data to System A. User operations can be prioritized over SnapMirror and other system activities using NetApp FlexShare™ software — a standard feature of the NetApp operating system starting with Data ONTAP® 7.2 —to minimize the impact of SnapMirror.

sell more
hardware.

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.

Want to learn more? Read [Workload Prioritization for Optimized Performance](#)

To improve performance, administrators frequently mix primary FlexVol™ volumes and mirrored FlexVol volumes on the same aggregate. Without a dedicated "mirror" aggregate, primary data spreads across all disk drives. More I/O resources are available to primary applications, while SnapMirror runs in the background protecting the data. Mirrors should simplify and optimize your environment, not sell more hardware.

Of course, performance means nothing if, when a disaster strikes, your mirrors are not instantly ready to serve data. Even if a disaster occurs while a mirror is in mid-transfer, the mirror must be available for recovery. Mirrors must also remain consistent while recovering. We've all heard campfire horror stories about an admin incorrectly executing step 23 on page 4 of the DR plan, "And... thwack! The mirror became worthless and it was time to look for a new job." A mirror must always reliably handle a disaster. SnapMirror uses Snapshot copies to always maintain consistency. Regardless of network status, SnapMirror can fall back to the last Snapshot copy for instant DR. SnapMirror DR itself consists of a handful of commands that cannot compromise data integrity. Mirrors must be reliable — if you can't trust your mirror, what can you trust?

Again, I've answered some common questions in the Q&A including:

- Speaking of backup... where does SnapVault® fit into your DR strategy?
- Sometimes my CPU hits 100% when I run SnapMirror. Is that bad?
- You talk about disk I/O a lot. Why are you so obsessed?

[View Q&A related to this section](#)

3. SnapMirror — Beyond DR

Putting Your Mirrors to Work

Mirrors are evolving from expensive DR insurance policies to cost-effective core business tools that deliver DR and more. A single SnapMirror copy can improve DR reliability, optimize test and development, and simplify day-to-day administrator activities.

To reduce the cost of DR, don't let your mirrors idly await a disaster. Put them to work! Offloading tape backup to the mirror is just the beginning; mirrors can meet other challenges. Admins cannot easily test application upgrades, evaluate new database schemas, or run data mining because they fear extra load or data corruption on their primary storage. "Don't touch it, you'll break it." Because mirrors are separate data copies with cycles to spare, they can help solve these problems. SnapMirror uses Snapshot copies and can be used with NetApp FlexClone® technology to safely unify test and development with data protection. To mine data, clients can access the read-only SnapMirror and its Snapshot copies. To run tests, admins can instantly create a FlexClone of their SnapMirror to create a separate, space-efficient, read-write copy. Regardless, the mirror maintains complete DR readiness. With SnapMirror, secondary storage delivers primary business value.

Not only should mirrors deliver business value every day, they should simplify administrators' everyday lives. Data migration has become common in today's dynamic storage environments: applications need more space or performance, data must move to newer arrays, and users or projects change sites. If a mirror can truly handle DR, data migration should be easy. SnapMirror simplifies data migration because it reduces storage downtime to minutes. First, the IT team sets up a SnapMirror relationship to the migration target. Then they disable client access to the old array, update the mirror a final time, split the mirror, and direct clients to the new array.

Simple. Efficient. Fast.

Customers initially worry about using their mirrors for more than DR, but constant use actually improves DR success. Data recovery fails because either the data was protected incorrectly or the recovery process was run incorrectly. Using your mirrors for test and development or data migration helps eliminate those failures. First, there is no better validation of data correctness than having clients or applications use it. It's like running a fire drill every time. Second, there is no better way to make the IT team comfortable with using the mirrors than to simply use them. SnapMirror customers recover from disasters quickly and reliably because they trust the data and themselves.

NetApp didn't build SnapMirror four score and seven years ago, but over the past decade we have protected more tiers of data than any other product. SnapMirror was not created in a moment of genius, but emerged from years of dedication to three design principles: low cost, simplicity, and reliability. SnapMirror was built to ensure that business-critical data shall always be available anytime, anywhere, for any use.

Questions addressed in the Q&A include:

- I'm moving to VMware — how do you work with them?
- After I migrate my data, how do I reconfigure my mirrors?
- Why is NetApp so excited about FlexClone? Doesn't everybody support clones?

[View Q&A related to this section](#)

SnapMirror Q&A

1. SnapMirror for Disaster Recovery

Lower the Cost of DR and Protect More Data

Because SnapMirror is designed to transfer only "net-new" data, most customers see a reduction in network load of 25% to 90%. Lower-cost DR means that you can now protect more of your data. Additionally, SnapMirror reliably protects against both physical and logical disasters. [Read the full details.](#)

Common questions include...

Q. How do I decide whether to mirror to ATA or FC?

A. Ask yourself three questions. First, during a disaster, what will be the bottleneck for your performance? ATA is slower than FC, but if you connect to your DR system with a high latency network, disk type won't matter. Second, what RPO do you want? Don't run a synchronous mirror from FC to ATA; I recommend at least 15-minute intervals. Third, what's more flexible — the answers to #1 and #2. or your budget?

Q. I think my network is my SnapMirror bottleneck. How can I be sure?

A. We encounter three SnapMirror bottlenecks: CPU, disk I/O, and network. Although the network throughput or latency can slow you down, don't discount the impact of system load. A system running at 90% CPU and/or disk utilization leaves few resources for SnapMirror, especially because SnapMirror gives priority to user load. If the numbers don't point to the culprit, you can run `snapmirror to tape` (`snapmirror store`). By eliminating the network, you may be able to isolate the bottleneck.

Q. SnapMirror can't meet its schedule because of the network. Can I change my SnapMirror configuration rather than buying a faster pipe or a compression device?

A. Maybe. There are two opposite solutions to this problem. Customers with frequently overwritten data run SnapMirror less often, which eliminates redundant network traffic. Customers with infrequently overwritten data, run SnapMirror more often. By constantly sending the new data to the mirror, they better utilize their network bandwidth (no bursts). There are two easy ways to figure out your workload. First, use the `snap delta` command to monitor the amount of change between Snapshot copies. Second, modify the SnapMirror schedule and examine the `snapmirror logs`. Either way, determine how much more data you would transfer with 4 hourly transfers than with 1 transfer every 4 hours. In the end, though, sometimes you do just need to buy a faster network.

Q. How can I optimize my SnapMirror performance on ATA?

A. Configure your mirrors to utilize all the ATA drives at once using FlexVol volumes. Because the ATA drives are slower than the FC drives on the primary storage, you need to compensate by using them more consistently. Fortunately, FlexVol volumes make it possible to spread your mirrors across the maximum number of spindles, while the SnapMirror simple schedule makes it easy to always run one or more transfers to each destination aggregate.

2. SnapMirror for End-to-End Data Protection***If You Can't Trust Your Mirror, What Can You Trust?***

Because SnapMirror replicates the active data and Snapshot copies, administrators can back up any application-consistent Snapshot copy while continuing to update the mirror. To reduce cost, tape backups should be run on the mirrors, offloading the primary servers and storage. To improve performance, mix primary and mirrored FlexVol volumes on the same aggregate. [Read the full details.](#)

Common questions include...

Q. Speaking of backup... where does SnapVault fit into your DR strategy?

A. SnapVault, the NetApp D2D backup solution for the last 5 years, adds two DR tiers. SnapVault alone offers a tier between periodic mirroring and traditional D2D and tape backup. SnapVault can run CDP-style, with hourly backups and 20:1 to 50:1 deduplication, from which users can recover their own data. However, you cannot make SnapVault writable. Therefore, although recovery will be faster than any other D2D or tape backup, it will be much slower than a DR mirror.

SnapVault plus SnapMirror delivers a unified DR and backup tier. SnapMirror can make any of the SnapVault backups writable for near-instant recovery. One caveat: SnapVault plus SnapMirror can provide only up to 1 hour RPO, while SnapMirror alone can deliver from 0 RPO and up, depending on your schedule. Backup management and deduplication are not computationally free, after all!

Q. Sometimes my CPU hits 100% when I run SnapMirror. Is that bad?

A. Not necessarily. If a CPU runs at 90%, and you have work to do, you're not getting the most out of your system. If it runs at 100%, you worry that the load will slow high-priority operations. Ideally, a system should run as hard as possible, but let high-priority operations take precedence. That's how we designed SnapMirror. It drives the CPU as aggressively as possible, but constantly preempts itself to minimize its impact on user load. If you find that SnapMirror is hurting your application performance, it's more likely due to disk I/O load than to CPU overload. Reduce the frequency of SnapMirror updates, look into FlexShare, and monitor your overall disk I/O loads.

Q. You talk about disk I/O a lot. Why are you so obsessed?

A. Processor performance continues to rapidly increase. Memory steadily becomes cheaper. Bus bandwidth keeps growing. Disk drive capacity explodes—and disk drive performance doesn't. With the vast amount of data and compute power available, disk I/O becomes a precious commodity. NetApp solutions minimize, optimize, and offload I/O for data protection, so your primary storage resources are focused on serving your applications.

3 SnapMirror — Beyond DR***Putting Your Mirrors to Work***

Mirrors are evolving from expensive DR insurance policies to cost-effective core business tools that deliver DR and more. A single SnapMirror copy can improve DR reliability, optimize test and development, and simplify day-to-day administrator activities. [Read the full details.](#)

Common questions include...

Q. I'm moving to VMware — how do you work with them?

A. We love VMware and all virtual server technology. I'll discuss three great opportunities with VMware and NetApp. First, VMware completes our tiered DR story because you can deploy DR tiers for both servers and storage. With virtual machines, the DR site can have fewer servers. We reduce the DR storage cost; they reduce the server cost. Second, VMware works with SnapMirror to further simplify

test and development. FlexClone your SnapMirror, create a virtual server with your application, and start running. That's it. No load on the primary storage and servers; no extra data copies; no complex system/application configuration. Third, VMware highlights the value of SnapMirror for offloading backup. As your server utilization increases, you cannot afford to run heavyweight backups — there's no more excess CPU available! In VMware environments, SnapMirror offloads the backup from both the server and the storage. It's the only way to scale.

(ed note: You can find out more on this topic in [Five Ways to Use NetApp Snapshot Copies with VMware VI3.](#))

Q. After I migrate my data, how do I reconfigure my mirrors?

A. If you migrate a volume, simply change snapmirror.conf to match the source system/volume/qtree name and configure the new source to allow transfers to the destination (use the snapmirror.allow or /etc/snapmirror.allow option). That's it. There's no need to break and/or resynchronize the mirrors. SnapMirror picks up where it left off.

Q. Why is NetApp so excited about FlexClone? Doesn't everybody support clones?

A. NetApp has revolutionized cloning the way it did snapshots, because the two are built on the same principles. NetApp can store hundreds of Snapshot copies per volume without incurring a performance or storage penalty. Our Snapshot copies don't generate extra I/O when data is modified, and they consume space only for changed blocks.

NetApp can now create hundreds of FlexClone copies without incurring a performance or storage penalty. A FlexClone copy can be instantly created from any Snapshot copy. At that moment, the clone shares all of its blocks with the Snapshot copy; it uses no extra space and generates no extra load. As the application modifies the clone, our clones consume space and I/O only for the changed blocks.

FlexClone copies transform our industry-leading read-only Snapshot copies into industry-leading read-write Snapshot copies. It's easy to get excited about a technology that—combined with Snapshot and SnapMirror—makes data available anytime, anywhere for any use.

[Comment on this article](#)

TECH ONTAP ARCHIVE



Antti Velinen and Francesco Duranti
Systems Administrators

This month Hobbit was nominated by sysadmins in two countries. Antti Velinen (pictured left) has worked for his current employer for 30 years and has managed NetApp storage for two of them. He's worked his way through major technology changes from tape machines to personal computers to Linux and the Internet. Francesco Duranti has been working as a sysadmin since 1995. He worked for a Sun Microsystems partner until he moved to his current position at Kuwait Petroleum Italia in 2000. Francesco started working with NetApp storage in 2004, after his company discovered that NetApp was the solution to its Windows® and Linux® data-sharing problems.

Featured Tool: Hobbit

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: The primary Hobbit developer is Henrik Störner.

What it is: The Hobbit monitor is a system for monitoring hosts and networks, inspired by the Big Brother system. It provides real-time monitoring, an easy Web interface, historical data, availability reports, and performance graphs. Unlike Big Brother, it is designed to work well whether you need to monitor small networks with just a handful of hosts or large networks with thousands of hosts.



How it works: Hobbit monitors hosts and network services and can monitor almost anything else using extensions. All monitoring results are collected by a central server and used to build a set of regularly updated Web pages that show status information on monitored devices, with drill-down functionality to check up on problems. Through the use of agents installed on servers, you can also monitor local disk utilization, log files, and processes.

Why Do You Like Hobbit?

Francesco: Hobbit is simple. All configuration work, including alerts, can be done on the server. Hobbit can send alerts by e-mail, SMS, or whatever method you want via scripts. People receiving alerts can acknowledge they're working on them by responding to the e-mail or via a Web page. Everything is simple and well documented. The client can be compiled on one machine and deployed on the same architecture by copying it and can auto-update itself. Because Hobbit is compatible with Big Brother, there are a lot of extensions available for use with the tool.

Antti: Hobbit has a great user community, and there are many useful extensions available. A few weeks ago I found a nice extension to check databases, BEA WebLogic servers, and NetApp storage systems: the [hobbit_perl_client-1.15](#). (Note: This extension was created by Francesco Duranti, this month's other contributor. Editor.)

How Do You Use Hobbit?

Francesco: We're in the process of finishing the implementation of Hobbit server. At this moment we monitor about 700 objects (switches, routers, storage systems, servers, and databases) divided into 70 Web pages with SMS or e-mail alarms on important services. We also have created some views based on the type of object so that our network admins can view just the networking devices or our database admins can check the status of database instances and servers.

RELATED INFORMATION

- ✦ [Hobbit download](#)
- ✦ [Hobbit information](#)
- ✦ [NOW™ Tools Discussion Forum \(NOW user login required\)](#)
- ✦ Previously Featured Tools:
 - ✦ [NTOP](#)
 - ✦ [Toasters](#)
 - ✦ [Swatch and Kiwi](#)
 - ✦ [Firefox](#)
 - ✦ [Cacti](#)
 - ✦ [MRTG for Filers](#)
 - ✦ [Nagios](#)
 - ✦ [SIO](#)
 - ✦ [Toasterview](#)

Sample Hobbit Output

From Francesco:



Overall status of storage systems and Brocade switches.



We also use Hobbit to monitor our peripheral sites (servers, switches, and routers), and we will deploy Hobbit to monitor servers and storage at our disaster recovery site. At the moment I don't know if we will put a local Hobbit server there or if we will just use the one at our primary site. Most likely we will put a Hobbit server at the disaster recovery site and link it to our primary server.

We save a large amount of data so that we can create various graphs (such as users logged onto systems or CPU/memory utilization). I've written some custom Perl scripts—for example, [hobbit_perl_client-1.15](#)—to monitor Network Appliance™ storage systems, databases, and BEA WebLogic servers. They work well with Hobbit, and there's also a patch to save most of the data collected. Some of those scripts also work with Big Brother server.

By default, Hobbit collects data and does checks every five minutes. One of the nice features is that you can get a "snapshot view" of a monitored object at a specified time in the past. We had a problem in which Oracle® RAC servers were rebooting during the day without any evidence of a problem on the servers or database. While we were trying to analyze the problem we got a snapshot at the moment when the RAC nodes were rebooting and discovered that the two switches where the machines were connected were not reachable. Analyzing uptime (also obtained with Hobbit), we confirmed that the switch had not been rebooted and that the network interfaces were up. This helped us zero in on the network as the cause of the problem. In the end our network admins discovered a problem related to the use of spanning tree on the switches.

Antti: We started using Big Brother about five or six years ago, including extensions to draw graphs and see history more clearly. When I found out that there was a successor to Big Brother—Hobbit—that incorporated these and other extensions, it was a logical transition to make. Hobbit was still safe and familiar, but more efficient in a bigger environment.

We use Hobbit to monitor about 125 critical servers and two NetApp storage systems. The majority of servers are running Windows, but we have some 30 Linux machines as well. Hobbit has agents for both operating systems, so it fits our environment nicely. I currently have Hobbit installed on two separate machines on different network segments. It doesn't require a powerful machine to run (I use a Dell that's about five years old), and it's very easy to install another system once you've got one configured.

Rather than use alarms, I've created a home page for our support team where I collect information from different monitoring and reporting systems. The Hobbit nongreen display is one part of the page, providing a quick overview of issues that may need attention. That way everyone can quickly and easily identify and investigate any system that is having problems.

Caveats

Antti: Hobbit is stronger for UNIX® and Linux than it is for Windows, although it works okay in the Windows world. Some people say it's ugly and too simple, but I really like the way it works.

❖ **Hobbit download**

❖ **Hobbit information**

What do YOU think?

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

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

❖ **Nominate a tool**



Graphical view of storage system operations per second by protocol.



Disk utilization.


From Antti:



Page identifying potential problem (nongreen) systems.



Availability over a defined time period.



Disk usage on a few NetApp volumes.