



Tech OnTap

Highlights

July / August 2008

- [DR in Virtual Environments](#)
- [Exchange 2007 on VMware](#)
- [The New FAS3100](#)
- [FAS2050 Wins with Network World](#)
- [Tech OnTap Archive](#)

Price/Performance Perfection: The New FAS3100

Mark Insley, Senior Platform Hardware Engineer

Lee Stork, Senior Performance Engineer



The challenge of midrange storage is to deliver the best price/performance solution, and we nailed it on the head with the new FAS3100.

Find out how it performed and the changes we made to deliver up to twice the performance while keeping costs low.

➤ [More](#)

Tips from the Trenches

Case Study: Exchange 2007 in a Real-World Virtual Environment

Joe Jerebic, Information Systems Manager



Considering virtualized email storage? Get tips for success and learn from UBG's recent migration.

➤ [More](#)

Blogging with Dave



Dave Hitz
NetApp
Founder
and EVP

Can a Programmer Be 10 Times More Productive?

I suspect that most highly productive programmers have a knack for spotting a clever approach.

➤ [Dave's Blog](#)

Engineering Talk

Streamlined Disaster Recovery with VMware Site Recovery Manager

Darrin Chapman, Product Engineer



DR planning is challenging enough without a virtual infrastructure to further complicate matters. See how VMware® and NetApp® enable nondisruptive DR plan testing so you can sleep at night

➤ [More](#)

Tech OnTap Spotlight

➤ FAS2050 Takes First Place

NetApp comes out on top in Network World's iSCSI SAN Clear Choice Test product review.

➤ Backup and SLAs: Got Opinions?

NetApp users and consultants discuss service-level agreements and how having them (or not!) impacts backup planning.



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

Engineering Perspective: How the FAS3100 Nailed the Price/Performance Sweet Spot

By [Mark Insley](#) and [Lee Stork](#), NetApp

When it comes to storage systems, low-end storage is focused primarily on achieving the lowest cost, while the high end is optimized to deliver performance and scalability. Between these two extremes is a growing midrange storage sector in which the goal is to optimize both characteristics and deliver the best price/performance.

This is what we were challenged to do with the new FAS3100 and V3100 series: optimize the platform to substantially increase performance for the midrange, while at the same time boost capacity and scalability.

Substantial improvements in performance and price/performance as demonstrated on leading benchmarks were driven by:

- New cost-effective single chassis and a two-power-supply HA-ready design
- A FAS3100 chassis design allowing focus on improvements to the heart of the storage system, the processor and memory subsystems
- Software optimizations

In this article we'll detail the changes to the FAS3100 platform and discuss its performance on the SPC-1, SFS, and Exchange benchmarks.

What Is the FAS3100 Series?

The new series consists of two models: the FAS3140 (which is targeted to replace the FAS3020 and FAS3040) and the FAS3170 (which replaces the FAS3070). Equivalent V-Series systems provide similar benefits while front-ending legacy storage from IBM, Hitachi Data Systems, HP, EMC, Fujitsu, and 3PAR legacy storage platforms.

	FAS3020	FAS3040	FAS/V3140	FAS3070	FAS/V3170
Processors	2 × 32-bit (2 single-core)	4 × 64-bit (4 single-core)	2 × 64-bit (2 dual-core)	4 × 64-bit (4 dual-core)	4 × 64-bit (4 dual-core)
Memory	4GB	8GB	8GB	16GB	32GB
NVRAM	NV5 (1GB)	NV6 (1GB)	NV7 (1GB)	NV6 (1GB)	NV7 (4GB)
Expansion Slots	6 × PCI-X	6 × PCIe	8 × PCIe	6 × PCIe	8 × PCIe
Onboard I/O	8 × GbE 8 × 2Gb FC	8 × GbE 8 × 4Gb FC	4 × GbE 8 × 4Gb FC	8 × GbE 8 × 4Gb FC	4 × GbE 8 × 4Gb FC
Max Disk Drives	168	336	420	504	840
Max Capacity	84TB	336TB	420TB	504TB	840TB
Data ONTAP	7.0 +	7.2.1 +	7.2.5 +	7.2.1 +	7.2.5 +

Table 1) Comparison of FAS3100 series with FAS3000 series. (All numbers shown are for dual-controller, active-active configurations.)

Consolidated Chassis Design

The FAS3100 series chassis is a new, two-in-one design that houses two controllers in a single enclosure for active-active configurations. This approach was logical given that the vast majority of current FAS3000 series customers purchase their systems in this configuration. The chassis continues to provide redundant power and cooling, but requires fewer power supplies and other components, providing a more cost-effective approach versus the two discrete chassis of the previous generation.

Explore

The New SPECsfs2008 Benchmark

While SPECsfs97_R1.v3 has been the gold standard for NAS performance for many years, time marches on and it is now time for some enhancements. SPECsfs2008 is a follow-on to SFS97, and includes important enhancements such as:

- A CIFS workload
- An updated NFSv3 workload
- Support for additional clients

The establishment of a standard benchmark for CIFS is particularly exciting. Only a few results have been submitted so far, but look for this benchmark to replace SFS97 in coming months.

[More](#)

NetApp Platforms Garner Awards

NetApp was named the top NAS vendor in both the midrange and enterprise categories in this year's Diogenes Labs—Storage Magazine Quality Awards. Winners were chosen based on a survey of actual NAS users.

[More](#)



Figure 1) The new FAS3100 series chassis versus the previous active-active configuration.

The new chassis offers several additional benefits:

- An integrated backplane reduces external cabling. Heartbeat communication (used in active-active configurations) between the two controllers is provided by an Infiniband connection built into the backplane. Two fewer power supplies also lead to a reduction of power cables.
- Single-controller configurations can be easily upgraded to active-active configurations by simply plugging in a second controller.

Processor, Memory, and Chipset Upgrades Boost Performance

The cost-effective system chassis design allowed us to focus on improvements in processors, I/O chipsets, memory, and NVRAM—the components that impact performance. The benefits of faster processors, the latest I/O chipset, and increased memory add up, yielding the significant performance boost we get from the FAS3100 versus the FAS3000 series, as described later.

Processor Selection

For midrange storage, we don't choose the highest-frequency, bleeding-edge processors. Instead, we look for solid, proven processors that make economic sense. Processor vendors offer two road maps: one that includes the latest, fastest processors (where processor generations come and go quickly) and an embedded road map in which they commit to support individual processors for longer periods.

For the FAS3100 series we've chosen dual-core 64-bit processors from the embedded processor road map, enabling us to ship these systems without any changes to the processing subsystem for several years. This reduces our exposure to more frequent processor qualification and additional support complexity along the way. Enterprise customers standardize on a particular storage platform and want to purchase that platform for an extended period of time, so this strategy also helps us better serve these customers.

I/O Chipset

The FAS3100 series uses a newer generation HyperTransport™ chipset than the one originally introduced with the FAS6000 family. (NetApp's use of HyperTransport chipsets was described in some detail in an [earlier Tech OnTap article on the FAS6070](#)). While the newer HyperTransport chipset I/O characteristics are very similar to those of the previous generation, an increased number of buffer credits for PCI transactions in the current chipset provides a performance advantage for intensive I/O.

Memory and NVRAM

The final hardware investment we made for the FAS3100 series was to boost the amount of memory and NVRAM in the FAS3170. A dual-controller system now has 32GB of memory (16GB per controller) versus 16GB for the FAS3070. The memory we chose was also slightly faster, to further reduce memory bottlenecks. While both systems use DDR2 SDRAM, the FAS3070 used DDR2-400 while the FAS3170 uses DDR2-667 DIMMs.

NVRAM was also increased in the FAS3170 to 4GB in active-active controller configurations compared to 1GB in the FAS3070. This was done to maintain balance with the larger memory configuration. NetApp systems use NVRAM to temporarily cache writes to disk. Batching writes allows most data to be written in efficient stripes, including parity, rather than performing a much slower read-modify-write operation.

Optimized Software

With any hardware platform, as you add new features the software has to be optimized to take advantage of them. NetApp® software has now been optimized to take full advantage of the latest multicore chipsets; the FAS3100 series leverages these improvements with Data ONTAP® 7.3. The FAS3170 benchmarks covered in this article all used Data ONTAP 7.3 to maximize the resulting performance.

Enhanced Expandability

In terms of overall expandability versus that of the previous models, the FAS3140 supports an additional 84 disks over the FAS3040 for a maximum of 420 spindles (420TB using 1TB drives), while the FAS3170 supports an additional 336 disks over the FAS3070 for a maximum of 840 spindles (840TB).

We've also moved the NVRAM from a PCIe slot to the motherboard on both models. This makes an additional PCIe slot available—four slots in each controller for a total of eight expansion slots in active-active configurations. These slots can be used to support additional network ports or the new Performance Acceleration Module, an intelligent read cache that reduces latency and improves I/O throughput without additional fast disk drives.

FAS3100: Demonstrated Price/Performance

We've demonstrated the performance of the new FAS3100 platform using three different benchmarks:

- SPC-1, to assess SAN performance using an industry-standard benchmark that generates workloads typical of OLTP database activity
- SFS97, to assess NAS performance (NFS protocol)
- The Microsoft® Exchange Solution Reviewed Program (ESRP), to establish the number of Exchange users the storage system can accommodate

SPC-1 Price/Performance

A complete description of the SPC-1 benchmark was provided in a [recent Tech OnTap article](#) comparing the performance of the previous-generation FAS3040 relative to the EMC CX3 Model 40. In that testing, the FAS3040 achieved 30,985.90 SPC-1 IOPS™ at a cost per IOP of \$13.61 versus \$20.72 for EMC. (With snapshot technology enabled on both platforms the cost/IOP gap favored NetApp even more, with NetApp achieving \$14.89 versus \$54.49 for EMC.)

The new FAS3170 achieved 60,515.34 SPC-1 IOPS, nearly doubling the results posted on the FAS3040. Demonstrated price/performance was \$10.01/SPC-1 IOPS, a substantial improvement over previously published performance numbers.




Figure 2) SPC-1 results for the FAS3040 versus the FAS3170.

Full SPC-1 results can be found at the following links:

- [SPC-1 Benchmark Results \(all\)](#)
- [FAS3040 SPC-1 Results](#)
- [FAS3040 SPC-1 Results with Snapshot Enabled](#)
- [FAS3170 SPC-1 Results](#)
- [EMC CLARiiON CX3 Model 40](#)
- [EMC CLARiiON CX3 Model 40 with Snapshots Enabled](#)

SFS97

FAS3100 performance was also assessed using the SPECsfs97_R1.v3 benchmark, which targets file service workloads by focusing on NFS version 3 throughput (operations/second) versus latency. The FAS3170 reached new heights here as well, demonstrating approximately 60% improvement over the previous-generation FAS3070, with a maximum of 137,306 ops/sec throughput. Latencies for all data points measured were also extremely low, resulting in an overall response time (ORT) of 0.94ms on the FAS3170 that peaked at just 2.1milliseconds for the final data point, compared to an ORT of 1.16ms on the FAS3070 with a peak of 2.9ms. [SFS97 results for the FAS3070 versus the FAS3170](#).




Figure 3) SFS97 results for the FAS3070 versus the FAS3170.

Exchange Performance

Many customers use midrange storage systems to host Microsoft Exchange. The number of Exchange users supported by each storage system was assessed using the guidelines set out by the [Microsoft ESRP](#). Using this assessment, the FAS3170 doubled the number of supported users versus the earlier FAS3070 model, scaling to over 50,000 Exchange users while using 188 disk drives to meet ESRP throughput and latency guidelines. For comparison, the high-end EMC Symmetrix DMX3 supports about 84,000 ESRP Exchange users but uses 448 disk drives to meet ESRP throughput and latency guidelines. From a price/performance perspective, we were very pleased with where the ESRP results for the FAS3170 landed. [See full NetApp results](#).




Figure 4) Microsoft ESRP showing the number of supported Exchange users for each platform.

Conclusion

Optimizing the design of midrange storage systems is more subtle than designing high-end storage. The design implementation is done with customer needs in mind and therefore needs to be a careful balance between performance and price. The new NetApp FAS3100 series—along with its V3100 series companions for front-ending legacy storage—were designed to strike a balance between performance and price. This is born out by the impressive performance and price/performance of the platform as demonstrated on industry-standard benchmarks. The FAS3170 achieves 2x the SPC-1 performance of the FAS3040 plus a 25% improvement in price/performance (\$10.01/SPC-1 IOPS versus \$13.61/SPC-1 IOPS). The same platform also achieves 40% more SFS97 performance and supports 2x the Exchange users versus the FAS3070. Both systems are available and shipping now.

Got opinions about the FAS3100 platform?

Ask questions, exchange ideas, and share your own opinions in the NetApp online community.



Mark Insley

Senior Engineer, Platform Hardware
NetApp

Mark (left) is a key engineer in the midrange platform group responsible for the FAS3000 series. He was responsible for the electrical design of the FAS3020/FAS3050, the FAS3040/FAS3070, and the FAS3140/FAS3170. Prior to joining NetApp, Mark held a variety of hardware and software design positions, including work on the original Sun™ SPARCStation-1.

Lee Stork

Senior Engineer,
NetApp

Lee's career in performance engineering began over 15 years ago at Tandem Computers, where Lee worked on processor, operating system, and database problems. Over the last three years, at NetApp, he has worked extensively on the FAS6080, the FAS3000, and the FAS3100 series of platforms. He is a graduate of the University of Nebraska, with a bachelor's degree in electrical engineering.



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

Case Study: Exchange 2007 on VMware and NetApp

By Joe Jerebic, Unity Builders Group

There are a lot of tangible benefits to running Exchange in a virtualized environment rather than a physical one, but while a lot of people are evaluating such solutions, it seems that many still remain skeptical about performance and reliability.

My company, Unity Builders Group (UBG), has run Exchange in a virtualized environment for a while now. We recently went through a migration to Exchange 2007 and VMware® VI3, and we are extremely pleased with the results. Tech OnTap asked me to describe the migration process and the advantages of running Exchange 2007 on VMware and NetApp® storage.

Migration to VI3 and Exchange 2007

A year and a half ago, we had roughly 300 mailboxes on Exchange 2003. Since then we've grown to about 450 mailboxes (seats). While that's not big in absolute numbers, it represents 50% growth in our Exchange environment (and other associated infrastructure) over a relatively short time frame.

Our Exchange 2003 server was originally hosted on a VMware server running ESX 2.5, while our other VMware servers were running ESX 3.0.1. Our NetApp environment was and still is on the FAS270 model appliance. Because Exchange 2003 was hosted on ESX 2.5, we were unable to benefit from the advanced features of VMware Virtual Infrastructure 3 (VI3) such as VMotion®, DRS (Distributed Resource Scheduler), or HA.

At the same time, our Exchange 2003 server was in need of a rebuild, and I needed to allocate more storage for Exchange data stores. Since Exchange 2007 was shipping and seemed to offer some significant benefits over Exchange 2003, we opted to migrate to Exchange 2007.

Since Exchange was already virtualized and delivering good performance, we started by upgrading the VMware server hosting Exchange to ESX 3.0.1 like the other two servers. This also gave us the opportunity to rebuild the server to accommodate increased space requirements based on the growth we were experiencing.

From there it was a natural progression to move to Exchange 2007. Because Exchange 2007 is less I/O intensive than Exchange 2003, it is actually better suited to virtualized environments than Exchange 2003.

Exchange 2003	Exchange 2007
32-bit Windows®	64-bit Windows®
900MB database cache	Multi-GB database cache
4Kb block size	8Kb block size
High read/write ratio	1:1 read/write ratio
Storage is common pain point	Eliminates storage pain point

Table 1) Comparison of Exchange 2003 versus 2007.

At this point, we started by doing our homework; we watched Webcasts, read books, and so on until we were comfortable with Exchange 2007. Next we created a few test mailboxes in Exchange 2003 and installed an Exchange 2007 server in our virtual environment. We then migrated the test mailboxes from Exchange 2003 to Exchange 2007 and tested the send and receive functions both internally and externally.

Once we felt comfortable with the success of these test migrations, we proceeded to migrate the mailboxes of our IT staff. When we felt our Exchange 2007 environment was working 100%, we began the migration of the remaining active users, migrating 400 to 450 users over a two-week period.

For those who might not be aware, Exchange 2003 and 2007 can coexist together in your environment while you're performing the migration. We operated in this fashion for several weeks. You just have to do your homework and make sure you don't leave any bases uncovered.

The only real problem we ran into as a result of the migration was that we forgot about our scanning to e-mail functionality. When we finished migrating all mailboxes and public folders to Exchange 2007, we shut down Exchange 2003. As a result, all scan to e-mail jobs went into limbo; people were no longer receiving their scans. This was easy to resolve; the lesson is to be sure you map absolutely everything that depends on your e-mail systems.

Explore

About Unity Builders Group

Unity Builders Group is a Calgary-based home builder with construction volume of approximately C\$330 million annually—just over 1,200 homes. There are 13 companies under the UBG Unity umbrella, with Canadian offices in Calgary, Edmonton, and Canmore and U.S. locations in San Antonio, Texas, and Colorado Springs, Colorado.

[More](#)

Virtualizing Exchange with NetApp and VMware

Want to learn more about virtualizing your Exchange environment? A recent NetApp Webcast focused on the technical advantages of both VMware and NetApp for Exchange.

[More](#)

Virtualizing Exchange: Chat Transcript

In conjunction with the [Webcast](#) there was also a very active chat session in which participants asked questions about specifics of virtualizing Exchange for their environments.

Review the 16-page transcript for the details.

[More](#)

Reference Architecture for Exchange 2007 on VMware and NetApp

Are you ready to start planning your Exchange 2007 on VMware deployment? This joint tech report from VMware and NetApp gives you the information you need to get started, including:

- Full reference architecture using iSCSI and blade servers
- Sizing guidelines for both VMware virtual machines and NetApp storage
- Complete performance results

[More](#)

Advantages of the New Environment

VMware Advantages

We currently have three ESX servers, all now running ESX 3.5. To support our Exchange 2007 environment as well as a busy Citrix environment, we utilize a combination of VMware VMotion, DRS, and HA. These tools dramatically increase our flexibility and decrease the time it takes to deploy new servers.

VMotion allows us to migrate virtual machines (VMs) from one server to another without interrupting service. For example, if I need to do a BIOS update or other firmware update on a server, I can migrate all the VMs off of that server and take the server down for the upgrade without the users even noticing that anything has happened.




Figure 1) VMotion.

VMware DRS gives us automatic load balancing across our ESX servers so that we get optimal utilization such that a server is never overloaded or underutilized. VMware HA protects us by automatically restarting VMs on other servers should an entire ESX server fail.




Figure 2) VMware DRS uses VMotion to automatically migrate VMs to balance load. VMware HA automatically restarts VMs on another server if a physical server fails.

One unique thing we've done in our environment is to create a rule that keeps the VMs that run our Exchange 2007 mailbox server, our hub and client access server, our Symantec™ Enterprise Vault™ Archive server, and our BlackBerry server on the same ESX server. That way all the communication that occurs between those servers actually happens on the virtual switch within the ESX server and doesn't traverse our physical network switches.




Figure 3) Traffic through network switches has been reduced by making sure that the mailbox, HUB/CAS, Enterprise Vault, and BlackBerry servers always reside on the same ESX server.

NetApp Advantages

On the back end, we use a NetApp FAS270 storage system with 3TB of storage capacity. The NetApp system was installed with the assistance of our partner, Long View Systems, an IT services provider with over 400 technical consultants and offices in Calgary, Edmonton, Vancouver, Denver, Dallas, and Houston. We work with Long View not only because of its technical expertise, but also because of its ability to analyze our business requirements and propose solutions that we can afford and manage today that won't paint us into a corner when it comes time to expand in the future.

The best practices for Microsoft® Exchange suggest using RAID 10 (striping and mirroring) for back-end storage. However, mirroring—which requires double the storage capacity—was an expensive option for us. NetApp gave us the ability to use RAID-DP® to provide the high I/O performance required by Exchange at much less cost without sacrificing resiliency. In fact, NetApp's Double Parity RAID implementation actually

provides better resiliency than RAID 10. RAID-DP is 163 times less likely to lose data than RAID 10 and 4,000 times less likely to lose data than RAID 5.

	RAID 5	RAID 10	RAID-DP
Cost	Low	High	Low
Performance	Low	High	High
Resiliency	Low	Med	High

Table 2) Comparison of RAID 5, RAID 10, and NetApp RAID-DP.

Our ESX servers have Fibre Channel HBAs, which connect them to our NetApp storage system. While we’re using Fibre Channel today, one advantage of the NetApp approach is that we can add NAS protocols (NFS and/or CIFS) and iSCSI in the future if we need to, all on the same storage system.

NetApp provides close integration with both Exchange and VMware, which can make the life of an administrator a lot easier. I recently participated in a [NetApp TechTalk in which the specifics of many of these technologies was discussed](#). That gave me the opportunity not only to talk about what UBG was doing, but also to hear a lot more about NetApp and VMware solutions.

NetApp’s integrated Snapshot™ technology makes it easy to capture consistent, point-in-time images of Exchange data for backup, and it puts the backup workload on the storage system rather than on Exchange servers.

Snapshot is the foundation technology on which NetApp builds a variety of related services. For instance, NetApp [SnapManager® for Exchange](#) integrates with Exchange to coordinate actions between the Exchange server and storage to create backups that are consistent and fast. NetApp SnapManager for Virtual Infrastructure provides a [similar capability to back up running VMware virtual machines](#). Because backups occur very quickly and put no load on ESX servers, you can create many more backups per day for greater data protection.

For disaster recovery, NetApp SnapMirror® technology builds on NetApp Snapshot to create a simple and reliable mechanism for asynchronous replication to a remote site. Unlike many other replication solutions, SnapMirror allows you to have different storage infrastructures at your primary and remote sites. You can use Fibre Channel HBAs at one and iSCSI at the other, and you can use faster but more expensive Fibre Channel disks at your primary site and less expensive SATA disks at your DR site.




Figure 4) NetApp SnapMirror provides flexible replication for virtualized environments.

NetApp also offers deduplication technology that works with your primary storage (rather than just on backups) to eliminate duplicate copies of blocks and reduce the amount of storage you need. This deduplication technology works particularly well in virtual environments because virtual machines are typically almost identical and thus have a large amount of duplication. The storage savings you get from deduplication of your primary site gets passed along to your DR site, reducing the storage needed there.




Figure 5) When NetApp deduplication is used at a primary site, the space savings are automatically passed on to the DR site.

NetApp has worked closely with VMware to integrate the features of SnapMirror with VMware Site Recovery Manager (SRM), a new tool from VMware that simplifies DR planning, implementation, and testing for virtual environments. ([A companion article in this issue describes the integration of NetApp and VMware technology in SRM in more detail.](#))

Conclusion

The combination of VMware and NetApp technologies with Exchange 2007 has been a real win for us. Virtualizing our Exchange environment gives us greater reliability and flexibility with good performance. VMotion lets us move virtual machines easily for maintenance or other needs. DRS makes sure that Exchange servers are always properly resourced. VMware HA makes sure that Exchange and other important services get restarted right away should an entire ESX server fail.

NetApp storage is a good match for the capabilities of VMware, offering greater reliability at a lower cost than mirroring. We haven't yet implemented all of the NetApp capabilities I've described above in our environment, but having a range of options will give us flexibility to meet our needs. We intend to look seriously at implementing VMware SRM using NetApp SnapMirror in late 2008 to provide greater protection for both our Exchange and Citrix environments.

Got Opinions About Exchange on VMware and NetApp?

While many IT teams are interested in running Exchange on VMware using NetApp, others are skeptical. What are your thoughts about combining e-mail with virtualized servers and storage?



Joe Jerebic
Information Systems Manager
Unity Builders Group

As IS manager of a dynamic company, Joe is responsible for maintenance and support of much of the IT infrastructure, including Citrix and Exchange running under VMware, Cisco network infrastructure, business continuity, and DR. Prior to joining UBG, Joe served as a technical consultant with a large IT firm specializing in VMware implementations. He is a Microsoft Certified Systems Engineer (MSCE) and VMware Certified Professional (VCP).



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

Using VMware Site Recovery Manager to Simplify DR

By [Darrin Chapman, NetApp](#)

Nothing is scarier than the prospect of having to recover an entire site after a disaster, and the addition of virtual infrastructure to your environment may further complicate the situation. VMware® Site Recovery Manager (SRM) is designed to simplify and accelerate disaster recovery for VMware infrastructures, and it also includes nondisruptive testing so you can make sure your site recovery plan will work before you need to use it.

SRM provides automated virtual infrastructure failover for virtual machines (VMs) and servers; it relies on the existing replication capabilities provided by storage vendors—such as NetApp® SnapMirror® technology—rather than provide its own data replication mechanisms for moving VM data to the DR site. NetApp has worked closely with VMware to enable the advanced capabilities of SnapMirror and other NetApp technologies to be fully leveraged by SRM.

In this article I'm going to discuss the challenges of DR planning and explain how VMware SRM—in conjunction with NetApp storage functionality—can greatly simplify DR for virtual infrastructure.

Disaster Recovery Planning

Planning and execution are the most crucial aspects of a disaster recovery scenario. There are many natural, human-caused, and computer-driven disasters that can affect data availability. Here are a few of the most common problems with DR planning today.

No plan at all. For some, the cost and complexity of DR is simply too much to address given current resource and budget constraints. There is no time available for planned downtime, and the process of making—and, more importantly, testing—a real DR plan is continuously delayed.

Inability to meet your recovery point objective (RPO) and recovery time objective (RTO). For many of you the RPO and RTO your business operations require are not being met because your plan depends on expensive infrastructure, time-consuming restores, and/or system installations from scratch. The resources to implement and run a test plan are just too demanding, especially if a true disaster has never occurred at your company.

Administration vs. RPO/RTO. When failing over business operations to recover from a disaster, there are many steps that are manual and time consuming. Often, though custom scripts are written and utilized to simplify some of these processes, it is the processes that must be followed that affect the real RTO that any DR solution can deliver. Consider the flow of a typical disaster recovery process:

1. A situation occurs that requires failover to a DR site. (This could result from a power outage that is too long for the business to withstand without failing over or a disaster that causes the loss of data and/or equipment at the production site.)
2. The DR team takes the necessary steps to confirm the disaster and makes the decision to failover.
3. Assuming that the necessary testing was done to confirm that data replication was successful and that the DR site is in a usable state:
 - a. Replicated storage must be presented to ESX systems at the DR site.
 - b. The systems must be attached to the storage.
 - c. The correct virtual machines must be added to the inventory of the ESX systems.
 - d. If the DR site is on a different network segment than the production site, each VM may need to be reconfigured for the new network.
 - e. IT staff must ensure the environment is brought up properly, with systems and services being made available in the proper order to prevent naming/IP conflicts.
4. Once the DR environment is ready, business is supported by the equipment at the DR site. Eventually, the production site will be available again or lost equipment will be replaced.
5. Changes applied to data while the DR site was supporting business operations will sometimes need to be replicated back to the primary site. (Replication must be reversed to accomplish this.)
6. The processes that took place in step 3 above must be performed again.
7. Once the production environment has been recovered, the original replication schedule and relationships must be reestablished (from the production site to the DR site).

In general, DR processes can be lengthy, time consuming, and prone to human error. Essentially the same processes—with the same problems—have to be repeated to return operations to the primary site. Nevertheless, a DR solution is an important insurance policy for any business. Periodic testing of the DR plan is a must if your solution is to be relied on.

Explore

Five Tips to Increase NetApp Resiliency

A recent Tech OnTap article examined five little-known things you can do to improve the resiliency of your existing NetApp storage systems, including using:

- Multipath HA
- Proper sparing
- SyncMirror®
- Nondisruptive upgrade
- Automated tools

➤ [More](#)

NetApp Snapshot Copies and VMware

NetApp Snapshot™ technology is ideally suited for use with VMware. A recent Tech OnTap article explored five uses of Snapshot and other NetApp technologies derived from it in VMware environments:

- Near-instantaneous VM backup
- Fast and flexible VM recovery
- Accelerated data management through cloning
- Disaster recovery
- Application backup and management

➤ [More](#)

Such testing typically involves repeating steps 1–7 above on a scheduled basis. This can be costly (in terms of man hours and downtime), and can be a disruption to normal IT operations. Because of physical environment limitations and the difficulty of performing DR testing, most sites are able to do so only a few times a year at most, and some can't do testing in a realistic manner at all. SRM can automate the DR process and reduce the chance that human error will affect disaster recovery.

SRM Basics

The most difficult and time-consuming part of DR failover in a VMware environment is the execution of the steps necessary to connect, inventory, reconfigure, and power up virtual machines at a DR site. VMware has solved these problems with the introduction of Site Recovery Manager, which simplifies the management of the entire DR process.

SRM has three main components:

- Discovery/configuration
- Failover
- Test

The most critical part of DR is planning the response you will provide when a disaster strikes. SRM allows you to preprogram your disaster response; this is one of the most powerful benefits of the solution. Powering hundreds of servers up or down can be an enormous challenge especially in the middle of a crisis. Bringing systems online in a specific order is even more of a challenge.

The recovery plan you create during the setup phase of SRM configuration allows you to preconfigure the entire plan. SRM allows you to create the plan in a short period of time because of its built-in discovery capabilities and close integration with Virtual Center. Once a plan exists, it can be executed flawlessly with minimal user intervention. (You simply start the recovery process if and when a disaster strikes.) Creating such a plan using traditional DR approaches can take months—even years—to construct, and include a lot of error-prone manual steps.


The SRM automated test solution provides nondisruptive DR testing without disrupting ongoing replication (for DR) and without affecting SLAs even while various virtual machines are booted and data sets are tested. Replicated data sets can also be cloned and used for development testing and so on.

Using NetApp's cloning technology in combination with the SRM tool to execute virtual machine operations, you can bring a site online and run tests there without impacting production—all from a central location with little intervention. This is a powerful way to run a set of tests to validate a specific application or data set without impacting the production site or agreed-upon SLAs.

SRM and NetApp Storage

SRM enables two separate VMware environments, the Primary and the DR sites, to communicate with each other. VMs at the primary site can be quickly and easily collected into Protection Groups that share common resources and can be recovered together. These Protection Groups are configured into Recovery Plans at the DR site.

In a NetApp storage environment, SnapMirror can be configured to replicate VMs from local NetApp storage to a remote system, resulting in read-only mirrors at remote locations. An advantage of SnapMirror is that it gives you great flexibility in configuring the storage for your DR site, significantly reducing the cost of DR storage. Many replication solutions require you to have an identical storage configuration in both locations. With SnapMirror—while you are required to have NetApp storage at both sites—you can mirror high-end to low-end platforms, FC disk to SATA disk, and Fibre Channel SAN to iSCSI.



VMware SRM discovers these relationships and allows you to preprogram a response, including promoting mirrors to make them writable, mounting file systems, and booting systems.

Upon execution of a DR plan in a NetApp storage environment, SRM will:

- Quiesce and break NetApp SnapMirror relationships.
- Map the LUNs to existing igroups (igroups are tables of worldwide port names—WWPNs—of hosts that have access to LUNs).
- Trigger DR ESX hosts to rescan and detect storage.
- Reconfigure VMs as defined for the network at the DR site.
- Power on VMs in the order defined in the Recovery Plan.

Figure 1) SnapMirror alone or in conjunction with SRM, you can flexibly mirror high-end storage configurations (high-performance platforms, FCP disks, FC SAN configurations) to less-expensive solutions (lower-cost storage platforms, SATA disks, iSCSI).

Upon execution of a DR test in a NetApp environment, SRM will:

- Create FlexClone® volumes of the FlexVol® volumes on the DR storage system.
- Map the LUNs contained within these FlexVol volumes to existing igroups.
- Trigger DR ESX hosts to rescan and detect the storage.
- Connect VM network adapters to a private Test Bubble network.
- Reconfigure VMs as defined for the network at the DR site.
- Power on the VMs in the order defined in the Recovery Plan.

Conclusion

In summary, by automating the resource-intensive and/or manual parts of DR planning, failover, and test, such as mapping VMs to storage, booting in the right sequence, and taking care of IP addressing and naming schemes, SRM greatly simplifies DR in virtual environments. The incremental cost of protecting a VM is almost zero from an operational perspective. The only real costs are the disk space at the destination site and bandwidth to handle the data change rate of that VM. Virtual servers and storage are joined with minimal administrative headache. Your DR plan and recovery procedures can be created in the least amount of time using the least amount of resources.

Got opinions about DR for VMware environments or Site Recovery Manager?

Ask questions, exchange ideas, and share your own opinions in the NetApp online community.



Darrin Chapman

Data Protection Subject Matter Expert and Technical Marketing Manager
NetApp

Darrin Chapman is the person you turn to for just about any question involving disaster recovery or backup and recovery at NetApp. He's been involved with almost every NetApp best practices guide on data protection since 2002 and in his spare time designs training courses for customers and NetApp technical staff. Originally schooled as an electrical engineer, Darrin spent several years in systems architecture at AT&T, Nortel, and EMC.