



## Quick Links

- [NetApp Community](#)
- [Tech OnTap](#)

February 2011

## NetApp Integrated Data Protection: Making the Right Data Protection Choice

By [Srinath Alapati, Technical Marketing Engineer](#)

One of the great things about NetApp® storage is that all the capabilities to protect your critical data are closely integrated with NetApp hardware and Data ONTAP®. Often, all that's needed is a license key. You never have to buy a specialized appliance or do a complicated software installation to add functionality, and all our data protection solutions take advantage of built-in data management capabilities.




Figure 1) The traditional approach to data protection adds complexity and cost.

The fundamentals of NetApp integrated data protection were explored in [a previous Tech OnTap article](#). In this article, I want to dig into some details of our replication technologies. Most of the important elements of NetApp data protection, such as volume SnapMirror®, qtree SnapMirror®, SnapVault®, and MetroCluster™, use either mirroring or replication. Understanding how these technologies work and how they differ from each other can make it much easier to choose the best data protection strategy. I begin by discussing the various technologies and then provide some guidance on choosing the right options for

## Explore

### Is Tape Necessary?

Is tape still a necessary element of a data protection strategy? In a recent blogpost, Chris Blackwood makes the case that it's not.

[More](#)

### RPO and RTO

**Recovery Time Objective (RTO)** defines how much time it should take to recover if a failure occurs. An RTO of 20 minutes means an application or a dataset will be back online in 20 minutes after a failure.

### Recovery Point Objective (RPO)

In most cases, it's prohibitively expensive to eliminate all chance of data loss when a failure occurs. RPO defines the maximum amount of data you are willing to lose. For instance, an RPO of one hour means you will be able to restore an application or dataset to a point no more than an hour prior to the time the outage occurred. (Note that, for most applications, such an objective would imply backup or replication on an hourly basis.) RPO does not directly affect availability or RTO, except in cases in which an aggressive (short) RPO goal lengthens the time needed to complete the recovery process.

your requirements.

## NetApp Replication Options

Tech OnTap has published quite a bit on SnapMirror, SnapVault, and MetroCluster over the years. However, there are some key capabilities of these products—and some important distinctions between them—that to my knowledge have never been fully explored in a single article. I start with SnapMirror and then explain the other two in relation to it. (Don't be too alarmed if the SnapMirror explanation seems overly lengthy. You won't have to read a similar-length description to understand SnapVault and MetroCluster.) I also include several comparative tables that should help answer any remaining questions you have.

### SnapMirror

Everyone probably knows that SnapMirror is primarily intended to create mirrors in remote locations for disaster recovery. What's less well known is that there are actually two SnapMirror operating modes.

**Volume SnapMirror** operates at the physical-block level. It replicates the contents of an entire volume and all volume attributes verbatim from the source (primary) volume to the target (secondary) volume. As a result, the target storage system must be running a version of Data ONTAP that is the same as or later than that on the source. If deduplication or [NetApp data compression](#) (added in Data ONTAP 8.0.1) is running on the primary system, the destination volume inherits those savings, since the volume is identical.

**Qtree SnapMirror** replicates individual qtrees. Because qtrees are subsets of a volume, qtree SnapMirror operates at a logical level. You can't simply replicate a qtree verbatim, because some of the necessary volume-level bookkeeping information for the qtree would be missing on the target system.

Because replication is happening at a logical level, there are a few important differences versus volume SnapMirror. First, qtree SnapMirror does not inherit deduplication savings. Again, this makes sense if you think about it in the context of the source and the target. On the source, a qtree can contain a deduplicated block that is just a pointer to a block that lies outside the qtree. That block obviously won't exist on the target, and therefore the block must be replicated with the qtree rather than just the pointer.

By default, qtree SnapMirror replicates only the last created Snapshot copy, and so it maintains an asymmetrical number of Snapshot copies at source and target locations. (Volume SnapMirror by definition has the same Snapshot copies on both source and target.) In other words, qtree SnapMirror does not have Snapshot retention capabilities.

Both forms of SnapMirror begin with a baseline copy in which all data in the volume or qtree is replicated from source to target. Once the baseline is completed, replication occurs on a regular basis. Volume SnapMirror supports asynchronous, semi-synchronous, and synchronous replication, while qtree SnapMirror supports only asynchronous replication.

In **async mode**, Snapshot copies of the volume or qtree are created periodically on the source. Only blocks that have changed or have been newly created since the last replication cycle are transferred to the target, making this method very efficient in terms of storage system overhead and network bandwidth.

**Sync mode** sends updates from the source to the destination as they occur, rather than

### NetApp Syncsort Integrated Backup (NSB)

NetApp has joined with Syncsort to revolutionize backup by combining Syncsort management software with core NetApp technologies.

Syncsort Integrated Backup uses Snapshot@ copies and block replication to make backups 95% faster, 99% more reliable, and 90% smaller on disk.

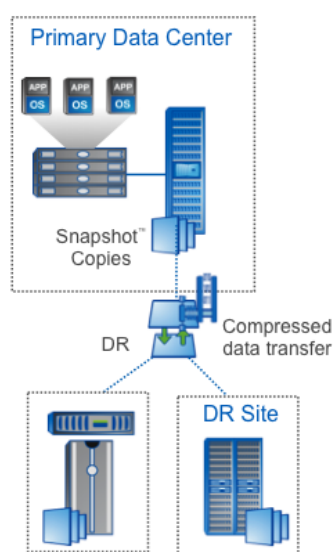
A recent Play by Play video highlights the advantages of this solution, especially in VMware and non-NetApp storage environments.

 [More](#)

according to a predetermined schedule. This helps data written on the source system to be protected on the destination even if the entire source system fails. NVLOG and Consistency Point (CP) forwarding are used to keep the target completely up to date. NVLOG forwarding enables data from the write log that is normally cached in NVRAM on a NetApp storage to be synchronized with the target. Consistency Point forwarding enables the on-disk file system images to be kept synchronized.

**Semi-sync mode** differs from sync mode in two ways. Writes to the source aren't required to wait for acknowledgement from the target before they are committed and acknowledged, and NVLOG forwarding is not used. These two changes speed up application response with only a very small hit in terms of achievable recovery point objective (RPO).

You can learn more about all of these modes by referring to [TR-3446: SnapMirror Async Overview and Best Practices Guide](#) and [TR-3326: SnapMirror Sync and SnapMirror Semi-Sync Overview and Design Considerations](#).



### Satisfy broad set of RPOs

- Sync: Zero RPO, up to 200km
- Semi-Sync: Near-zero RPO; 500km
- Async RPO mins – hours, long distance
- Replicates Snapshot copies for multiple RPOs

### Flexible configuration support

- Different models and disk types
- Many to many
- 3 node DR (sync and async together)

### Extends benefits of storage efficiency

- Up to 90% less capacity required
- DR data stays deduplicated
- Uses network compression transferring up to 70% less data to reduce costs / accelerate RPO

Figure 2) SnapMirror.

Finally, one of the key things to know about SnapMirror is that both volume and qtree SnapMirror result in targets that can be made writable. In other words, if a failure occurs that affects the source or primary systems, you can fail over operations and start writing to the target. Once the failure has been corrected, you can do a failback resync to copy delta changes back to the source and restore normal operation. This capability is a key differentiator versus SnapVault.

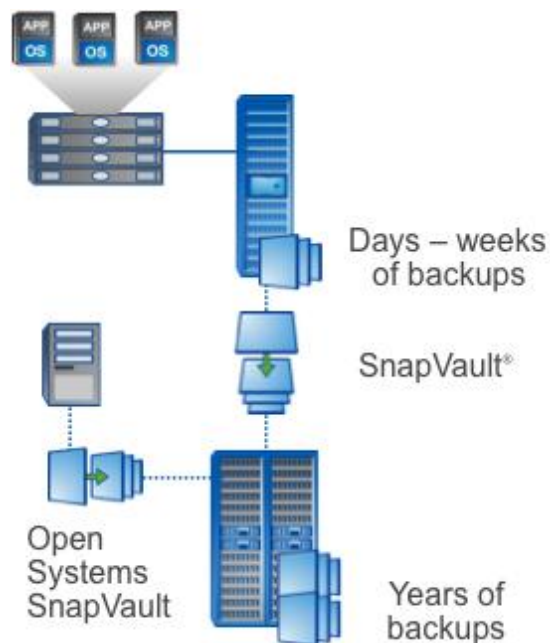
## SnapVault

SnapVault is primarily intended for disk-to-disk backup. Like async SnapMirror, SnapVault leverages NetApp Snapshot technology to back up and restore systems at the block level. Similarly, SnapVault identifies and copies only the changed blocks on a system (not changed files) to secondary storage. This not only increases performance by limiting the amount of data transferred during backup and restore operations, it limits the capacity needed to store backups, allowing you to perform backups more frequently if needed.

In terms of its basic operation, SnapVault is very similar to qtree SnapMirror—it performs replication on a logical basis at the qtree level. Like qtree SnapMirror, therefore, it's not an

exact replica of the source volume and doesn't inherit deduplication or the data compression state from the source. (You can run deduplication and/or data compression on the target as you would with any other NetApp volume.)

In addition, you can't make a SnapVault volume writable (for immediate recovery) as you can with SnapMirror; as a result, recovery times with SnapVault may be much longer than with SnapMirror if you transfer a lot of data across a network. If you also own SnapMirror, it is possible to make a SnapVault volume writable, but keep in mind that SnapVault is one-directional; it doesn't have failback resync to bring the source back to currency.



**Figure 3)** SnapVault. Open Systems SnapVault (not discussed in the text) allows third-party storage to be integrated into the backup framework.

The key weapons in the SnapVault arsenal—because it operates at a logical level—are Snapshot retention and Snapshot coalescing. You can retain as many Snapshot copies as you want (up to the limit of 255 per volume) on a SnapVault volume and expire Snapshot copies automatically according to a schedule that you set. Coalescing allows you to run multiple SnapVault processes from multiple sources to a single target and then create a single Snapshot copy on the target that includes all the different sources. This reduces the number of saved Snapshot copies; if you run deduplication on the target system, you can then deduplicate identical blocks across all qtrees in the backup.

You can learn more about all aspects of SnapVault from the [SnapVault Best Practices Guide](#).

**Table 1)** Comparison of SnapMirror and SnapVault.

Feature	Volume SnapMirror	Qtree SnapMirror	SnapVault
Replication type	Physical	Logical	Logical
Replication network	FC or IP	FC or IP	IP only
Multiple paths for replication	Yes	Yes	No
Data ONTAP version sensitive	Yes	No	No
Network Compression	Yes	Yes with approval	No
RPO (How much data can I afford to lose?)	1 minute <sup>1</sup>	1 minute <sup>2</sup>	1 hour
Failover capability	Yes	Yes	Yes, when combined with SnapMirror
Snapshot retention for backup use	No	Possible but tedious	Yes
Snapshot coalescing	N/A	No	Yes
Failback resync	Yes	Yes	No
Deduplication	Destination inherits deduplication savings; network savings as well	Destination does not inherit deduplication savings	SnapVault and deduplication are integrated; destination does not inherit deduplication savings

<sup>1</sup> Although 1-minute updates are possible, NetApp does not recommend them. Use SnapMirror Semi-Sync for low RPO (<3 minutes).

<sup>2</sup> Although 1-minute updates are possible, NetApp does not recommend them. SnapMirror Semi-Sync cannot be used on standalone qtrees.

## MetroCluster

The NetApp solution for continuous data availability is MetroCluster. This solution is an outlier relative to SnapMirror and SnapVault because it works in a very different way, but conceptually it's very easy to understand. As the name implies, MetroCluster provides "stretch" clustering. It lets you take a standard NetApp HA pair and separate the nodes by up to 100 km. MetroCluster uses a fully mirrored active-active configuration that maintains two complete copies of all mirrored data—one on each side of the cluster. These copies are called plexes and are continually and synchronously updated each time Data ONTAP writes data to disk.

Each controller owns storage volumes (plexes) on both nodes. This not only allows deduplication to occur on both nodes, it allows read operations to be split across both disk sets, which increases read performance by up to 80%. You can read more about MetroCluster in a [recent Tech OnTap case study](#) or you can watch a [complete video explanation](#).

Table 2) Comparison of MetroCluster and SnapMirror Sync.

	SnapMirror Sync	MetroCluster
Replication network	IP or FC	FC
Limit on concurrent transfers	Yes	No limit
Maximum distance	200 km (>200 km semi-sync)	100 km
Replication between HA pairs	Yes	No
Failover	CLI	CLI (Single command), System Center
Use of replica	Yes	
Primary deduplication support		Yes

## Which Option Should I Choose?

Tables 1 and 2 in the previous section are designed to help you choose the best replication option for your particular needs. There are a few considerations that can help you choose from the various technologies discussed above. The first—and most obvious—question to ask yourself is whether what you need is backup or DR.

### Backup

If what you're going for is backup, most people find that a regular snapshot schedule on primary storage (hourly is common), possibly combined with a nightly SnapVault copy to secondary storage (either local or remote), meets their backup needs. Most file restores can be satisfied from Snapshot copies on primary storage, while SnapVault provides the ability to reach further back in time, plus the ability to do big restores in the event of more serious failures.

See the sidebar to view a video on NetApp Syncsort Integrated Backup, which combines the benefits of Syncsort data management and NetApp replication for a variety of important application environments.

### DR

For protection from site-wide disasters and to enable business continuity, you'll probably want to choose from either MetroCluster or SnapMirror. By far the most popular alternative in terms of the number of deployments is volume SnapMirror with asynchronous replication. People tend to choose this because it offers simplicity and great economy with efficient use of storage and network resources. NetApp invested a lot of development effort in SnapMirror, creating valuable features such as bandwidth throttling, [network compression](#), and integration with the SnapManager suite of products for application integration.

Both qtree and volume SnapMirror can achieve recovery time objectives (RTOs) ranging from seconds to minutes and recovery point objectives (RPOs) as low as one minute (this requires replicating data every minute), although NetApp does not typically recommend

asynchronously replicating every minute. For recovery times between one and three minutes, SnapMirror in semi-sync mode is a better choice. (If you aren't familiar with RPO and RTO, see the sidebar.)

If you need a more aggressive RPO than async SnapMirror can achieve, you can choose from either MetroCluster or synchronous SnapMirror. Keep in mind that synchronous solutions typically require much greater network bandwidth and specialized network equipment to implement, so this makes them significantly more expensive.

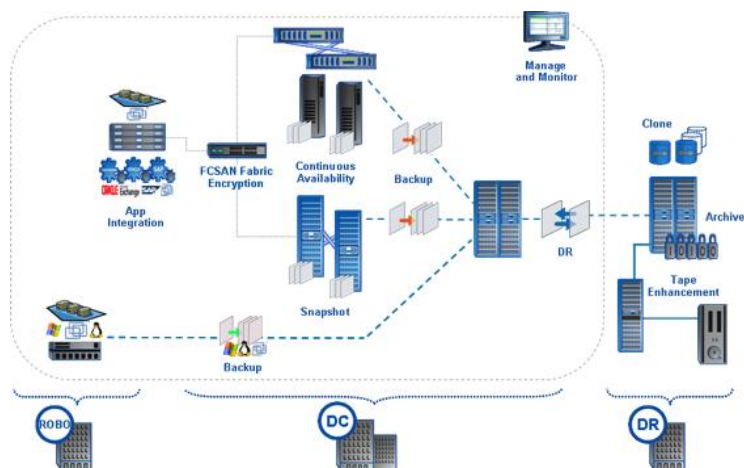
MetroCluster is the preferred solution for distances up to 100 km, since it offers continuous data availability and automatic failover and recovery. SnapMirror Sync doubles the supported range to 200 km, and SnapMirror Semi-Sync can reach further than that should you need the lowest possible RPO over a longer distance.

### Corner Cases

The approaches I outline above should cover most of the situations out there, but, naturally, there are always corner cases. Some people use SnapMirror for backup, usually because they want the ability to quickly and easily make a backup volume writable should that become necessary. Conversely, others use SnapVault for DR because it lets them recover to any point in time. SnapVault volumes cannot be made writable by SnapVault alone, but, as I mentioned (although I haven't explained how), this is possible using SnapVault and SnapMirror.

## Getting Started

Naturally, many NetApp users implement a combination of the solutions I discuss in this article to cover both backup and DR needs. A fairly common scenario is SnapMirror for critical volumes to a remote site combined with a regular SnapVault schedule at the remote site for backup purposes. Some sites even deploy a combination of MetroCluster, SnapMirror, and SnapVault to address data protection needs.



**Figure 4)** The NetApp integrated data protection portfolio (includes features not discussed in this article).

You can read more about advanced configurations, all the topics I cover in this article, plus topics I didn't have space for, such as data protection planning, in the [NetApp Data Protection Handbook](#). You can also check out the other resources I mention in this article for more details. NetApp has developed a lot of expertise with all sorts of data protection solutions. You shouldn't hesitate to go online to the NetApp community or ask questions of

your NetApp team if you need help making the right decisions.

---

➤ [Got opinions about integrated data protection?](#)

Ask questions, exchange ideas, and share your thoughts online in NetApp Communities.



**Srinath Alapati**  
Technical Marketing Engineer  
NetApp

Srinath joined NetApp in 2004 and has been a member of the Data Protection group for over four years. He has 10+ years of experience in IT, managing servers and storage infrastructure. Srinath authored or coauthored multiple technical reports on SnapMirror, MetroCluster, VMware®, and Exchange and speaks at various technical conferences. He is also a core team member involved in NetApp IT's disaster recovery implementation.

---

## Tech OnTap

### [Subscribe Now](#)

Tech OnTap provides monthly IT insights plus exclusive access to real-world best practices, tips and tools, behind the scene engineering interviews, demos, peer reviews, and much, much more.

Visit [www.netapp.com/us/communities/tech-ontap/](http://www.netapp.com/us/communities/tech-ontap/) to subscribe today.