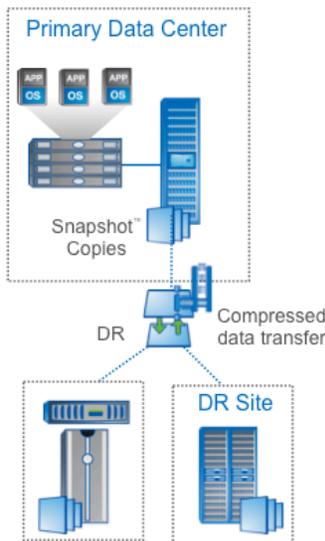


# Back to Basics: SnapMirror

By: Srinath Alapati, Technical Marketing Engineer

*This article is the fourth installment of Back to Basics, a series of articles that discuss the fundamentals of popular NetApp technologies.*

NetApp® SnapMirror® software has been the preferred technology for replication and disaster recovery in a wide variety of NetApp storage environments for years because of its proven efficiency, simplicity, and modest cost when compared with other DR solutions. Over the years, NetApp has continued to enhance SnapMirror with new features and capabilities to make the product fit an even broader range of requirements and to use network bandwidth even more efficiently.



### 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 1) NetApp SnapMirror.

The use of SnapMirror technology offers significant advantages:

- **Efficient.** Block-level updates reduce network bandwidth and time requirements. Starting with Data ONTAP® 7.3.2, volume SnapMirror also offers native network compression to further reduce bandwidth costs.
- **Flexible.** Data can be replicated between dissimilar NetApp storage systems. One-to-one, one-to-many, many-to-one, or

## Explore

### More Back to Basics

Learn the fundamentals of NetApp® core technologies. Installments in this series so far:

- [Chapter 1: Thin Provisioning](#)
- [Chapter 2: Deduplication](#)
- [Chapter 3: FlexClone](#)
- [Chapter 4: Volume SnapMirror](#)

many-to-many replication topologies are supported with async mode.

- **More productive.** When you use SnapMirror in combination with NetApp FlexClone®, you can use the data stored in your DR environment for dev/test, data mining, or other purposes.
- **Consistent.** Through integration with the NetApp SnapManager® suite, application data can be replicated while making sure of full consistency for quick recovery.
- **Safe.** Your DR plan can be tested without affecting production and ongoing replication so you can test more frequently to make sure there aren't any surprises should disaster strike. To protect against application data corruption, your DR site can keep multiple Snapshot® copies on hand and quickly and easily restore to a point in time before the data corruption occurred.

There are two operating modes for SnapMirror: volume and qtree. Volume SnapMirror is generally the preferred mode. Because of its relative popularity, much of our development effort, including integration with the SnapManager suite of products, has focused on volume SnapMirror. As a result, volume SnapMirror offers greater flexibility and efficiency. This chapter of Back to Basics explores how volume SnapMirror technology is implemented, the most common use cases, best practices for implementing SnapMirror, and more.

## How Volume SnapMirror Is Implemented in Data ONTAP

---

**Volume SnapMirror** operates at the physical block level. It replicates the contents of an entire volume, including all Snapshot copies, plus all volume attributes verbatim from a source (primary) volume to a target (secondary) volume. As a result, the target storage system must be running a major 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 and the savings are experienced on the WAN as well.

Volume SnapMirror begin with a baseline copy in which all data in the volume is replicated from source to target. Once the baseline is completed, replication occurs on a regular basis. Should it be necessary, the target 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 NetApp SnapVault®, which is intended primarily for disk-to-disk backup.

Differences Between Asynchronous Volume SnapMirror and Qtree SnapMirror
<ul style="list-style-type: none"> <li>Volume SnapMirror maintains same number of Snapshot copies on source and destination.</li> </ul>
<ul style="list-style-type: none"> <li>Volume SnapMirror integrates with the NetApp SnapManager suite and SnapDrive®.</li> </ul>
<ul style="list-style-type: none"> <li>Target volumes inherit deduplication savings from source volumes. Deduplication on source also reduces network bandwidth requirements.</li> </ul>
<ul style="list-style-type: none"> <li>Volume SnapMirror is not sensitive to dataset characteristics such as number of files and directories.</li> </ul>
<ul style="list-style-type: none"> <li>Volume SnapMirror supports multihop or cascaded configurations.</li> </ul>
<ul style="list-style-type: none"> <li>Network compression is generally supported with volume SnapMirror only.</li> </ul>

**Table 1)** Key differences between asynchronous volume and qtree SnapMirror.

Volume SnapMirror supports asynchronous, semi-synchronous, and synchronous replication; asynchronous replication is by far the most commonly used.

In **async mode**, Snapshot copies of the volume 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 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 forwarding 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 NetApp storage to be synchronized with the target. Consistency point forwarding keeps the on-disk file system images 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).

**SnapMirror network compression** was added starting with Data ONTAP 7.3.2. With SnapMirror network compression, data is compressed only while it traverses the network; data on source and destination systems remains uncompressed. Enabling compression results in two additional steps:

- Compression on the source system
- Decompression on the destination system

On the source system, data blocks that need to be replicated are handed off to a compression engine, which compresses them. The compression engine creates multiple threads corresponding to the number of CPUs on the storage system. The multiple compression threads compress data in parallel. Compressed blocks are then transmitted over the network. On the destination system, compressed blocks are received and decompressed using a similar multithreaded approach. Decompressed data is then written to the appropriate volume.

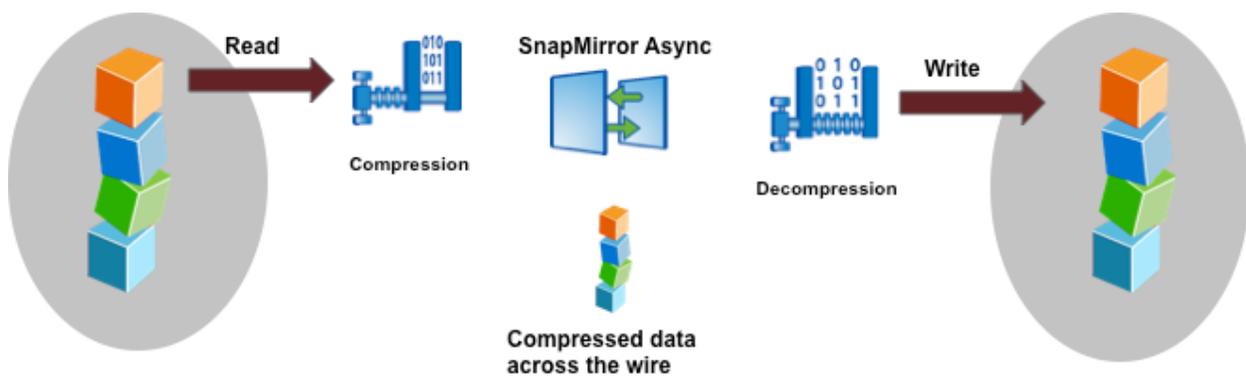


Figure 2) SnapMirror network compression.

The compression and decompression engines can either be configured to conserve network bandwidth or complete a transfer in the shortest time possible, depending on user preference.

SnapMirror network compression is supported on all NetApp storage platforms (including V-Series virtualization systems and the IBM N-series) in the asynchronous mode of operation only. The semi-synchronous and synchronous modes of SnapMirror operation are not currently supported with network compression enabled.

You can learn more about all the functions of volume SnapMirror 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](#). You can also read more about network compression in [a previous Tech OnTap® article](#).

## Use Cases

---

There are two main use cases for SnapMirror:

- Disaster recovery
- Remote data access/data distribution

In addition, the ability to utilize FlexClone volumes and the ability to replicate them is becoming an important emerging use case.

**Disaster recovery.** Using volume SnapMirror, data can be mirrored to another NetApp storage system at a DR facility or secondary data center. If a DR version needs to be made operational, applications can be switched over to servers at the DR site and application traffic redirected to these servers for as long as necessary. When the production site is back online, SnapMirror can transfer the data efficiently back to the production storage systems, and SnapMirror transfers can resume.

Volume SnapMirror supports multihop or cascading configurations. For example, a volume can be replicated from a system in San Francisco to a system in New York City and then from New York City to Singapore.

**Remote data access/data distribution.** SnapMirror also facilitates the distribution of large amounts of data to geographically remote locations, allowing local read-only access to data. FlexClone technology can be used when locally writable replicas are required. One-to-many and many-to-one configurations are supported with async SnapMirror.

Remote data access not only provides faster access to data for local clients, but also results in a more efficient and predictable use of expensive network and server resources. This allows you to replicate source data at a chosen time to minimize overall network load. The ability to control when data is replicated is also valuable in cases where you need to make sure that a dataset is in a consistent state.

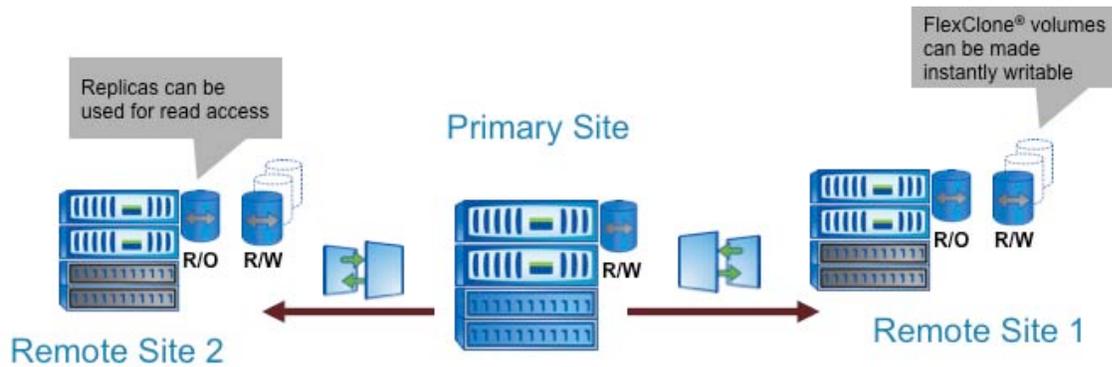
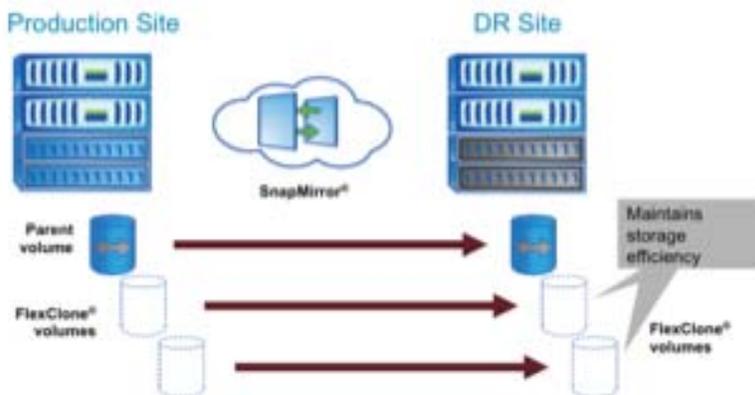


Figure 3) Using volume SnapMirror for remote data access.

**Use cases in conjunction with FlexClone.** SnapMirror provides particular benefits when used in conjunction with [FlexClone technology](#) to support application dev/test environments and for DR testing. Performing application dev/test on your DR storage allows you to get more use out of resources that might otherwise sit idle much of the time. This was described in some detail in [the FlexClone chapter](#).

Testing your DR processes without interfering with ongoing replication mechanisms can be problematic. With FlexClone you can easily clone your DR volumes and fully test your DR processes without interfering with ongoing SnapMirror replication processes.

Some environments make use of FlexClone volumes to provide space-efficient copies for virtual desktop infrastructure (VDI), data warehousing, and local development and testing. In many cases, it might be desirable to replicate such clones to protect them. Before Data ONTAP 8.0.1 (7-Mode), when a FlexClone volume is replicated using volume SnapMirror, space savings are lost. The FlexClone volume on the target requires capacity equal to the size of the parent volume. Starting with Data ONTAP 8.0.1, when operating in 7-Mode, FlexClone volumes can be replicated using volume SnapMirror without the need for additional capacity on the destination system as long as the parent of the FlexClone volume is also replicated.



**Figure 4)** Starting in Data ONTAP 8.0.1, FlexClone volumes can be replicated with SnapMirror without losing storage efficiency as long as the parent volume has been replicated.

## Using SnapMirror Technology

---

Volume SnapMirror can achieve recovery time objectives (RTOs) ranging from seconds to minutes and recovery point objectives (RPOs) as low as a few minutes. If you need a more aggressive RPO than async SnapMirror can achieve, you must then choose from either MetroCluster™ or synchronous or semi-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 100km since it offers continuous data availability and automatic failover and recovery. SnapMirror Sync doubles the supported range to 200km, and SnapMirror Semi-Sync can reach further than that to achieve the lowest RPO over a longer distance. Sync and semi-sync SnapMirror do not support the same feature set as async SnapMirror; for instance, network compression and SnapManager integration are not supported when using these modes. You can find more information on the use of MetroCluster in conjunction with SnapMirror in a [recent Tech OnTap article](#).

A few general considerations are important when you are getting started with volume SnapMirror:

- Pay attention to the Data ONTAP version requirements for the operating mode you are running.
- Async volume SnapMirror: Destination must be of same or higher major or minor version.
- Sync or semi-sync volume SnapMirror: source and destination systems must be running the same version.

**Table 2)** Data ONTAP source and destination requirements for async SnapMirror.

Volume SnapMirror Source	Volume SnapMirror Destination	Support?
Data ONTAP 7.2*	Data ONTAP 7.2*, 7.3*, 8.0*	Yes
Data ONTAP 7.3*	Data ONTAP 7.3*, 8.0*	Yes
Data ONTAP 7.3*	Data ONTAP 7.2*	No

- Volume SnapMirror only supports replication between aggregates of the same type: that is, 32-bit to 32-bit or 64-bit to 64-bit aggregates.
- SnapMirror operates over both Ethernet and Fibre Channel. Refer to the [switch support matrix](#) (requires NOW™ access) for Fibre Channel requirements.
- Sync and semi-sync modes are sensitive to distance and round trip time (RTT). RTT should be less than 2 milliseconds for sync and less than 5 milliseconds for semi-sync.
- There are limits on the number of concurrent SnapMirror transfers that can be performed. These limits are dependent on the type of NetApp system you have and the Data ONTAP release you are running. Refer to this [link](#) for more information. (Requires NOW access.) For more detailed information, refer to the appropriate technical report:
- [TR-3446: SnapMirror Async Overview and Best Practices Guide](#)
- [TR-3326: SnapMirror Sync and SnapMirror Semi-Sync Overview and Design Considerations](#)

## SnapMirror and Other NetApp Technologies

---

Because of the central importance of SnapMirror in many NetApp deployments, we've taken significant care to make sure that it interoperates with the vast majority of NetApp software solutions. Here are a few specifics you might want to be aware of:

- **SnapManager suite.** The SnapManager suite is designed to provide data protection and DR services for important applications, including Microsoft® Exchange, SQL Server®, and SharePoint®; Oracle®; and SAP®. The VMware® and Microsoft Hyper-V™ hypervisors are also covered. When using the appropriate SnapManager product (or Virtual Storage Console for VMware), you can make sure that application and/or hypervisor data is replicated in a consistent state so that operations can be restarted at the remote site.
- **FlexClone.** See the earlier section on use cases for information on using SnapMirror and FlexClone. Also refer to the [FlexClone chapter of Back to Basics](#).

In some instances, the space-efficient volume clones will contain critical data that warrants replication.

- **Deduplication.** When you replicate a deduplicated volume with volume SnapMirror, the destination volume inherits the space savings.

## Conclusion

---

NetApp SnapMirror technology is an important disaster recovery and general-purpose replication tool that can be used alone or in conjunction with other solutions such as the NetApp SnapManager suite. To learn more about NetApp SnapMirror, be sure to refer to [TR-3446: SnapMirror Async Overview and Best Practices Guide](#) and [TR-3326: SnapMirror Sync and SnapMirror Semi-Sync Overview and Design Considerations](#).

---

### Got opinions about SnapMirror?

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 has 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.

---



*Quick Links*

[Tech OnTap Community](#)

[Archive](#)

[User Groups](#)

---