



Quick Links

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

February 2011

DataMotion for Volumes Best Practices and Use Cases

By [Richard Jooss, Director-FC/iSCSI Solution Architect, NetApp](#)

The traditional approach to moving data—take a volume offline, spend hours copying it to a new location, and then reconfigure and restart affected servers and applications—has become nearly impossible in an age of shared infrastructure and round-the-clock operations. This is why NetApp is bringing you a range of options to allow you to nondisruptively move data to a new location—with minimal hassles—whenever it becomes necessary.

If you're a regular reader of Tech OnTap, you already know about [NetApp Data Motion](#) (renamed NetApp® DataMotion for vFiler®), which allows you to move MultiStore® vFiler units and all associated data between storage systems in multi-tenant environments. [With the Data ONTAP® 8.0.1 release, NetApp introduced a number of new capabilities](#), including DataMotion for Volumes, which allows you to nondisruptively migrate volumes containing LUNs between aggregates on the same storage controller.

This article describes DataMotion for Volumes, explains how it differs from DataMotion for vFiler, and looks at some of the best practices and possible use cases.

What Is DataMotion for Volumes?

DataMotion for Volumes is a new feature supported in Data ONTAP 8.0.1 running in 7-Mode. It allows you to nondisruptively migrate a volume from one aggregate to another on a single storage controller (you can't move volumes between the two different controllers that make up an HA pair). DataMotion for Volumes is designed to facilitate migration of any volume containing FC, FCoE, or iSCSI LUNs only. It does not support NFS or CIFS volumes.

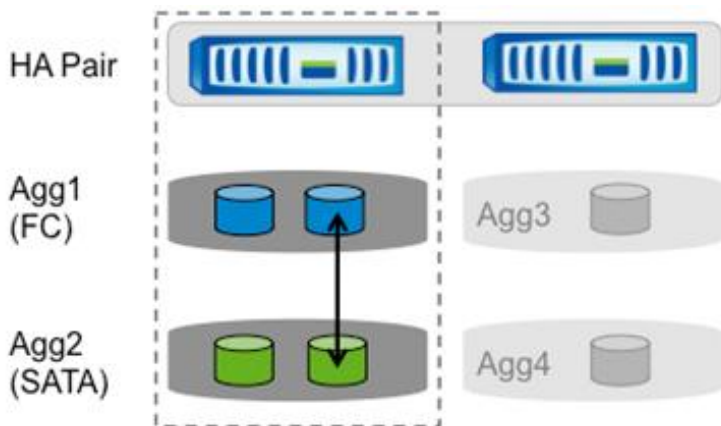


Figure 1) NetApp DataMotion for Volumes lets you nondisruptively migrate volumes containing LUNs between different aggregates on a single NetApp® storage controller.

A key benefit of DataMotion for Volumes is that it maintains all the details associated with a

Explore

MultiStore and DataMotion for vFiler

MultiStore makes secure storage multi-tenancy possible by isolating different tenants or applications in vFiler units. These vFiler units can be migrated from one storage system to another via DataMotion for vFiler.

[More](#)

What's New in Data ONTAP 8?

With the release of Data ONTAP 8.0.1, NetApp introduced a variety of new features. Besides DataMotion for Volumes, Data ONTAP 8.0.1 includes [data compression](#), [Unified Connect](#), support for new platforms, and a variety of other capabilities.

[More](#)

volume after the migration, including:

- Snapshots and snapshot schedules
- SnapMirror®, SnapVault®, and MetroCluster™ relationships
- Thin provisioning settings
- Deduplication state (a rescan is necessary to reestablish the fingerprint database in the new location)

You can move data between all types of media connected to your storage controller. For instance, you can move data from an aggregate containing FC or SAS disks to one consisting of SATA disks or vice versa. For the first release, you can move volumes between 32-bit aggregates or between 64-bit aggregates, but you cannot move volumes between 32-bit and 64-bit aggregates.

Differences between DataMotion for Volumes and DataMotion for vFiler

There are a number of important differences between DataMotion for Volumes and the DataMotion for vFiler technology you may already be familiar with. DataMotion for Volumes works at the level of the FlexVol® volume while DataMotion for vFiler works at the level of the vFiler and will migrate all NFS or iSCSI volumes associated with a given vFiler unit. DataMotion for vFiler allows you to migrate volumes between separate storage systems or HA pairs. DataMotion for vFiler is managed using NetApp Protection Manager. DataMotion for Volumes can only be invoked from the command line by using the `vol move` command.

These and other differences are summarized in Table 1.

Table 1) Comparison of DataMotion for Volumes with DataMotion for vFiler.

	DataMotion for vFiler	DataMotion for Volumes
Granularity	vFiler	FlexVol
Scope of Operation	Migration between HA pairs	Migration between aggregates on the same storage controller
Protocols	NFS, iSCSI	FC, FCoE, iSCSI
Management	Protection Manager, CLI	CLI only
Data ONTAP Release	7.3.3 or later Future 8.x release	8.0.1 or later

How Does DataMotion for Volumes Work?

DataMotion for Volumes leverages proven NetApp SnapMirror technology for volume migration. Knowing a little bit about how the process works can help enable success. A migration occurs in three phases:

- Setup
- Data copy
- Cutover

In the setup and data copy phases, the source volume continues to serve I/O requests without interruption.

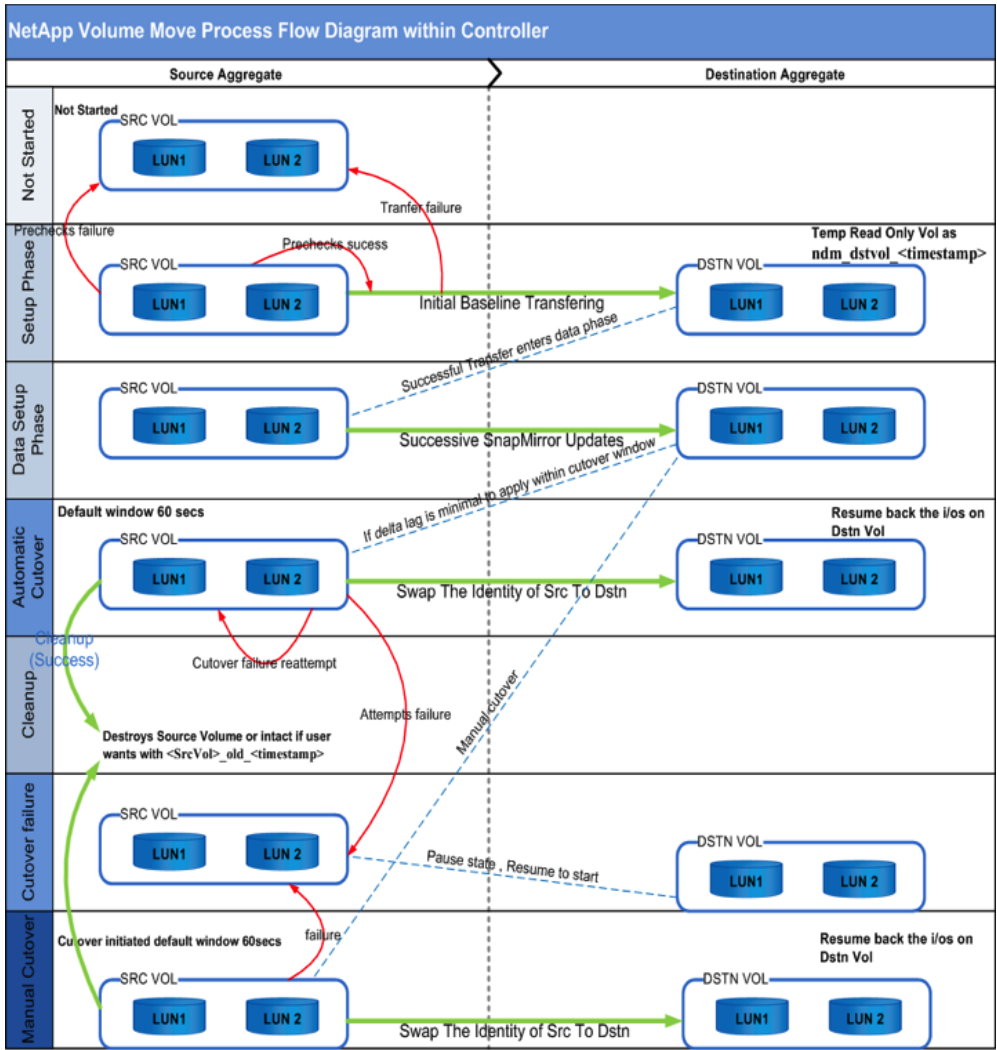


Figure 2) Illustration of the full DataMotion for Volumes process.

SETUP PHASE

In the setup phase, important prechecks are performed to make sure that the whole operation can be completed successfully. These prechecks are run each time DataMotion for Volumes is started and performed a second time when the cutover phase begins.

Prechecks verify the state of all the objects (such as source volume, source aggregate, and destination aggregate) involved in the move. If any of the prechecks are not successful, you'll be notified and the process will not start. All errors from prechecks are reported on the console and logged to the log file in the root volume of the controller (/etc/log/ndvm).

After all prechecks have successfully completed, DataMotion for Volumes:

- Creates a temporary placeholder volume in the destination aggregate using the naming convention:


```
ndm_dstvol_<timestamp>
```
- Triggers a baseline transfer to establish a SnapMirror relationship between the source volume and the placeholder volume. This is the longest phase; its duration

is directly proportional to the size of the source volume since the entire volume contents are transferred.

DATA COPY PHASE

After the baseline transfer is complete, the data copy phase begins, during which successive SnapMirror updates are initiated to synchronize the destination volume with the source, which remains active. At the end of each successful SnapMirror update, DataMotion for Volumes estimates the delta lag between the two volumes. You receive a notification of this delta lag, which becomes the estimated time for the next update operation. DataMotion for Volumes remains in the data copy phase as long as the delta lag is high. When the delta lag is small, it enters the cutover phase.

CUTOVER PHASE

The cutover phase can be either manual or automatic (the default).

Precutover. DataMotion for Volumes enters the cutover phase if the destination volume can be synchronized completely within the cutover window. The precutover phase is a transient period in which the initial prechecks mentioned above are reverified to make sure that nothing has changed and other important checks are run, including monitoring the state of the NVLOG and system CPU load.

You can set the CPU and disk threshold limit for DataMotion for Volumes. (The default is 100 for both of the following options):

- `vol.move.cutover.cpu.busy.limit`
- `vol.move.cutover.disk.busy.limit`

Automatic Cutover. Once the precutover checks are completed, the cutover phase begins automatically. The cutover must complete within the cutover window, which is by default a maximum of 60 seconds. Once the destination volume is synchronized completely, the identity of the source volume is transferred to the destination and I/Os are then serviced from the destination.

The following processes take place during the cutover phase:

- All LUNs in the source volume are put in a suspended I/O (quiesced) state. The quiesced state prevents new I/Os from being scheduled on the LUNs, and it drains pending I/Os on the LUNs.
- The source volume is quiesced. As part of the quiesce operation, the final delta lag is captured in a Snapshot copy, which is named according to the convention `ndvm_final_<timestamp>`.
- The destination volume is then synchronized completely with the source volume with the delta from `ndvm_final_<timestamp>`. This is the last SnapMirror update between the two volumes before servicing the I/O from the destination volume.
- The placeholder volume is stamped with the name and file system ID (FSID) of the source volume, and vice versa; that is, the identities of the source and placeholder volumes are swapped.
- The migrated volume at the destination is brought online with the identity of the original source volume and the LUNs are made active.
- The source volume is deleted (unless you specify that it be retained).

Cutover Failure. If cutover fails, DataMotion for Volumes reenters the data copy phase and tries again at a later time. By default, it makes three cutover attempts. After that, the cutover

is aborted and I/O continues from the original location. User intervention is required in order for DataMotion for Volumes to continue.

Manual Cutover. In manual cutover, DataMotion for Volumes remains in the data copy phase, performing successive SnapMirror updates. It continues to estimate the time of the next SnapMirror update operation and logs this information. You should review this information before triggering a manual cutover.

Possible Use Cases

There are a variety of possible use cases in which DataMotion for Volumes makes sense.

- **Load balancing.** Move a volume from a busy aggregate to one with less activity.
- **Capacity balancing.** Move a volume from a full aggregate to one with more available free space.
- **Migrate data from third-party disk to NetApp disk.** If you have a NetApp V-Series system as a front end to a disk array from a third-party vendor, you can use DataMotion for Volumes to migrate data from that array to NetApp disk shelves.
- **Change media type.** Move a volume from one type of disk media to another. For instance, move from FC disks to SAS or SATA disks.

I have heard considerable interest from users in all of these use cases, especially cases needing the ability to change media type. With [the industry move from FC to SAS that's taking place](#), DataMotion for Volumes provides a relatively painless way to retire older media when the time comes.

I've also talked with users who believe the ability to change media type nondisruptively will allow them to be more aggressive in the deployment of SATA disk, especially in conjunction with [Flash Cache performance acceleration](#). These users plan to deploy more applications on less-expensive SATA disk, knowing that if performance ever became an issue they could quickly migrate volumes to FC or SAS disk without disruption.

Best Practices

Attention to a few best practices will help enable your success with DataMotion for Volumes:

- You can perform only one DataMotion for Volumes migration at a time.
- Don't run `vol rename` or change volume attribute settings and LUN attribute settings in the source volume during the DataMotion for Volumes operation.
- Do not start any conflicting operations until the DataMotion for Volumes operation is complete. If any operation starts when the process is not in the cutover phase, DataMotion for Volumes initiates an abort. In the cutover state, DataMotion for Volumes fences all storage administrator operations.
- To maintain data integrity, initiate a pause before a manual reboot or halt is triggered.
- In an HA pair, be sure that both nodes are running Data ONTAP 8.0.1 7-Mode.
- Be sure to check the system health and load before you initiate DataMotion for Volumes.
- To prevent interference with cutover, don't schedule long-running data management or data protection operations during the DataMotion for Volumes process.
- Do not change any configuration for the source volume.

SNAPMIRROR AND SNAPVAULT

- During DataMotion operations, don't make changes to your SnapMirror

configuration, such as the `snapmirror.conf` file.

- If DataMotion for Volumes has not reached the cutover state, SnapMirror or SnapVault updates will be successful.
- If an update occurs during the cutover phase, the transfers may not succeed, because cutover fences the operation. After cutover is complete, the transfer will succeed on the next update.
- In case of a manual cutover, do the cutover when no updates or transfers are in progress.

SNAPDRIVE AND SNAPMANAGER

During cutover:

- LUN provisioning and Snapshot copy management operations in SnapDrive® for Windows® (SDW) will be unsuccessful.
- Any backup and restore operations will be unsuccessful for SnapManager® products (SME, SMSQL, SMVI, SMHV, SMO, SMSAP, and SMOSS). The next scheduled backup job will succeed.

HOST OPERATING SCSI TIME-OUT SETTINGS

NetApp recommends installing the NetApp Host Utilities Kit to avoid time-out errors. Refer to the "NetApp Host Utilities Installation and Setup Guide" for your host operating system at <http://now.netapp.com/NOW/cgi-bin/software>.

FLEXVOL AND FLEXCLONE

- DataMotion for Volumes moves only the source FlexVol volume; FlexClone® volumes are not moved. After DataMotion for Volumes is complete, the original FlexVol volume is retained in an offline state and the child-parent relationship is untouched.
- The FlexClone volume must be split from the parent FlexVol volume before moving the FlexClone volume.

DEDUPLICATION and COMPRESSION

- If deduplication is active on the source FlexVol volume, then it must be stopped for a successful cutover.
- DataMotion for Volumes does not move the fingerprint database and change logs of a deduplicated FlexVol volume. After the DataMotion for Volumes process is complete, users must execute `sis start -s` to rebuild the fingerprint DB at the destination.
- Compressed volumes cannot be moved.

Additional best practices for using DataMotion for Volumes in conjunction with Oracle® Database and Microsoft® Exchange can be found in [TR-3881](#).

Conclusion

DataMotion for Volumes is a powerful data management tool that can increase the flexibility of your operations. Instead of having to carefully plan and initiate volume movements that can only be performed during long-planned downtimes, you can move volumes as necessary to achieve your objectives of load balancing, capacity balancing, or changing media type. To learn more about DataMotion for Volumes, check out [TR-3873](#) and [TR-3881](#).

› [Got opinions about DataMotion for Volumes?](#)

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



Richard Jooss

Director and SAN/iSCSI Architect, NetApp



Richard Jooss is the senior manager of SAN Product and Partner Engineering at NetApp. Rick is responsible for defining technical and business requirements for the SAN ecosystem and storage, and business solution integration with NetApp SAN solutions. Rick has 15 years of experience in the storage industry. He holds a bachelor of science degree in Electrical Engineering and Computer Engineering from the University of Wisconsin.

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/ to subscribe today.