Introduction to Data ONTAP™ 7G

NetApp, TR 3356

TABLE OF CONTENTS

INTRODUCING DATA ONTAP 7G

Data ONTAP 7G, the newest software release from NetApp, simplifies customer environments by making it easier than ever for storage administrators to configure and manage systems that meet the exact requirements of their organizations. This document introduces several new concepts at a high level to help readers understand how this new software can enhance their particular environments. Detailed information on these topics is available on the NetApp Web site.

This new release introduces a number of significant features, many made possible through a new building block within WAFL® known as a flexible volume. The FlexVol™ technology decouples the previous direct connection between volumes and their associated physical disks, vastly increasing flexibility and storage efficiency. A new entity termed an "aggregate" now provides the connection between the logical flexible volume and the underlying physical storage and isolates the volume from this connection. A flexible volume is now able to stripe all of its data across the entire aggregate, thereby improving performance for "small" volumes.

FlexClone™, another new technology introduced in this release, makes it possible to replicate volumes instantaneously without duplicating shared blocks. Each FlexClone copy can itself be replicated, allowing many possibilities for branching of storage environments, similar to branching of source code trees within modern source code control systems. Virtualization, common to generations of application programmers when applied to memory allocation and management, is now available to data storage managers.

The purpose of this paper is to help readers understand these new capabilities so they can apply them to their own organizations' storage environments.

Summary of New Features

New features and technologies of Data ONTAP 7G are mentioned briefly here and discussed in detail in the subsequent sections of this paper.

Flexible Volumes

Data ONTAP 7G introduces flexible volume (FlexVol) technology, a breakthrough technology in which volumes are logical data containers that can be sized, resized, managed, and moved independently from the underlying physical storage. This enhances a storage administrator's ability to address a variety of data management requirements while preserving the familiar semantics of volumes and the current set of volume-specific data management and space allocation capabilities.

FlexClone

Data ONTAP 7G introduces a powerful new feature that allows storage administrators to instantly create clones of flexible volumes. A FlexClone volume is a writable point-in-time image of a FlexVol volume or another FlexClone volume. FlexClone technology adds a new level of agility and efficiency to storage operations. FlexClone volumes take only a few seconds to create and are created without interrupting access to the FlexVol volume being cloned. FlexClone uses space very efficiently, leveraging the Data ONTAP architecture to store only data that changes between the parent and the clone. This provides a huge potential saving in time, dollars, space, and energy. In addition, FlexClone volumes have the same high performance as other flexible volumes.

FlexClone volumes also enable administrators to access the destination mirror created through the NetApp SnapMirror® product. Previously, it was necessary to break the mirror in order to make any changes to the destination copy. With FlexClone, an administrator can now clone a Snapshot™ copy held in the mirror and make it available for both reading and writing at the remote site while allowing the mirror facility to continue running unaffected.

Role-Based Administration Controls (RBACs)

With only one set of security credentials available on the controller, it has become common practice for only the system administrator to have access to many of the functions provided by Data ONTAP. This has historically limited the use of many of its truly unique capabilities (such as SnapRestore®) and frequently relegated them to the realm of pure data recovery. If capabilities like Snapshot and SnapRestore were sufficiently constrained so that administrators were comfortable with users availing themselves of these technologies, huge benefits could be realized in the areas of software development, testing, and quality assurance (QA) as well as enterprise-class application deployment testing.

RBACs provide a mechanism for defining separate groups of users and roles and allows them to access defined subsets of controller functions. For example, a QA group can now be granted access to create and destroy cloned volumes without the involvement of the system administrator. The administrator can grant the QA group the permissions to perform only those functions without enabling them to access any controller commands that could put production data at risk.

Regulatory Compliance

Data ONTAP 7G provides new features for regulatory compliance, including LockVault™. LockVault is built on top of existing SnapLock® and SnapVault® functionality. LockVault is an ideal solution for unstructured data, such as home directories, project directories, and document stores that must be retained for compliance with regulations such as SEC 17a4 and Sarbanes-Oxley.

Mixed Storage

Data ONTAP 7G has the ability to utilize mixed storage including FC and SATA disks for primary storage. Mixed storage provides additional storage flexibility and investment protection. Using the mixed-storage option and leveraging SATA disks, primary storage costs can be reduced by 50% or more.

Upgrading to Data ONTAP 7G

Data ONTAP 7G supports both traditional and flexible volumes. Traditional volumes that existed before an upgrade to 7G will still exist as traditional volumes after the upgrade. By default, newly created volumes on 7G will still be created as traditional volumes. A flexible volume can be created by overriding the default traditional setting during volume creation. After the upgrade to 7G, both traditional and flexible volumes can be used concurrently on the same storage appliance.

Overview of Flexible Volumes

Flexible volumes are a groundbreaking new technology. These volumes are logical data containers that can be sized, resized, managed, and moved independently and nondisruptively from the underlying physical storage.

As shown in Figure 1 below, an aggregate is defined as a pool of many disks, from which space is allocated to volumes (volumes are shown in the illustration as FlexVol and FlexClone entities). From the administrator's point of view, volumes remain the primary unit of data management. But transparently to the administrator, flexible volumes now refer to logical entities, not (directly) to physical storage.

Flexible volumes are therefore no longer bound by the limitations of the disks on which they reside. A FlexVol volume is simply a "pool" of storage that can be sized based on how much data you want to store in it, rather than on what the size of your disks dictates. A FlexVol volume can be shrunk or increased on the fly without any downtime. Flexible volumes have all the spindles in the aggregate available to them at all times. For I/O-bound applications, flexible volumes can run much faster than equivalent-sized traditional volumes.

Flexible volumes provide these new benefits while preserving the familiar semantics of volumes and the current set of volume-specific data management and space allocation capabilities.

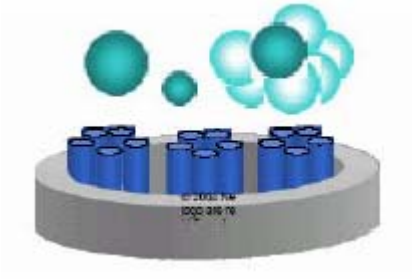

Figure 1) An aggregate consists of a pool of many disks from which space is allocated to volumes.

Improved Performance

With Data ONTAP 7G, disks are still organized in RAID groups, which consist of a parity disk (two in the case of RAID-DP™) and some number of data disks. Previously, Data ONTAP combined one or more RAID groups into a volume. While this still can be done (creating what is known in 7G as a "traditional" volume), RAID groups will now usually be combined into aggregates.

For example, assume that a database called Vol Database is the busiest volume on the system. Before Data ONTAP 7G, only one volume could reside on any given RAID group. Frequently, the RAID group would have only a handful of disks because that was all the capacity the volume required. The consequence was that the performance of the volume was limited due to the low number of disks comprising the volume.

In Data ONTAP 7G, RAID groups are combined to create aggregates. Since volumes are still the usual unit of storage management, it will be common to include all disks on a single NetApp controller in one aggregate and then allocate multiple volumes on that one large aggregate. This makes it possible to tap the unused performance capacity of all the disks, making that capacity available to the busiest part of the system. A FlexVol volume is flexible in changing size because the underlying physical storage does not have to be repartitioned.

Capacity Guarantees in FlexVol

Data ONTAP 7G introduces a new storage management concept of "guarantees." A guarantee differs from the previous management concept of "space reservations," which is familiar to customers using the iSCSI or Fibre Channel facilities of Data ONTAP. Guarantees extend the control by allowing administrators to set policies that determine how much storage is actually preallocated when volumes or files are created. This allows administrators to effectively implement the concept of "thin provisioning." This concept of "thin provisioning" enables the administrator to in effect oversubscribe their storage safely. Provision your storage space once, then grow as needed in the aggregate.

Guarantees, set at the volume level, determine how the aggregate preallocates space to a flexible volume. When you create a FlexVol volume within an aggregate, you specify the capacity and, optionally, the type of guarantee.

There are three types of guarantees:

1. **Volume:** A guarantee type of volume ensures that the amount of space required by the flexible volume is always available from its aggregate. This is the default setting for flexible volumes.
2. **File:** With the file guarantee type, the aggregate guarantees that space is always available for writes to space-reserved LUNs or other files.
3. **None:** A flexible volume with a guarantee type of none reserves no space, regardless of the space reservation settings for LUNs in that volume. Write operations to space-reserved LUNs in that volume might fail if the containing aggregate does not have enough available space.

Flexible Capacity Planning

There are virtually no restrictions on the size of a FlexVol volume, and flexible volumes can be resized dynamically. Restrictions in size differ between platform types. Check the system configuration guide on the NOW™ (NetApp on the Web) site for real limits. Administrators can use flexible volumes as a powerful tool for allocation and provisioning of storage resources among various users, groups, and projects. The smallest growth, or shrink increment is 4KB (1 block in WAFL terms). For example, suppose a database grows much faster than originally anticipated. The administrator can reconfigure the relevant flexible volumes at any time during the operation of the system. The reallocation of storage resources does not require any downtime, and it is transparent to users on a file system, or a LUN mapped to a host in a block environment. The effect is nondisruptive to all clients connected to this file system.

When additional physical space is required, the administrator can increase the size of the aggregate by assigning additional disks to it. The new disks become part of the aggregate, and their capacity and I/O bandwidth are available to all of the flexible volumes in the aggregate.

Capacity can also be overallocated where the set capacity of all the flexible volumes on an aggregate exceeds the total available physical space. Increasing the capacity of a FlexVol volume does not require changing the capacity of another volume in the aggregate or the aggregate itself.

Migration to Data ONTAP 7G

To take advantage of all the new features available with Data ONTAP 7G flexible volumes, you must migrate your data from traditional volumes to flexible volumes. Please refer to section 8, below for the details of migrating your data between traditional volumes and flexible volumes.

Overview of FlexClone Volumes

A FlexClone volume is a writable, point-in-time image of a FlexVol volume or another FlexClone volume. A FlexClone volume takes only a few seconds to create and is created without interrupting access to the FlexVol volume being cloned. FlexClone volumes store only data that changes between the parent and clone, providing a huge potential saving in storage capacity, time, and energy. FlexClone volumes have the same high performance as other FlexVol volumes.

FlexClone volumes offer great benefits in situations such as these:

- Using an actual image of the production environment for testing or development without endangering the production volume
- Making improvements to the production environment and "locking them in" incrementally
- Distributing data in changeable form without endangering the integrity of the original
- Initiating independent but parallel operations with identical temporary data

For example, imagine a situation in which the IT staff needs to make substantive changes to a production environment. They need to test in a real environment, but using the production volume for testing would be too risky. By using FlexClone, the IT staff gets an ideal solution: an instant point-in-time copy of actual production data, created transparently and at minimal cost in terms of storage and service interruptions.

IT personnel can make and test their upgrades using the FlexClone volume. Each time they are satisfied with the success of a phase or component of the changes they are making, they can clone their working FlexClone volume to lock in the success. If any change turns out not to work the way it should, they just destroy the working clone and go back to the FlexClone volume that was created at their last success. When everything is working exactly as intended, they can either split off the clone to replace the current production volumes or codify their successful upgrade process to apply to the production system during the next maintenance window. IT personnel can thus test to a higher level of confidence, with far less stress for them and higher levels of service for their customers.

How FlexClone Volumes Work

A FlexClone volume has all the capabilities of a FlexVol volume, including growing, shrinking, and being the source of a Snapshot copy. You can create multiple clones from a single parent volume, and you can create nested clones, using one FlexClone volume as the source for another FlexClone volume. For storing the data that is written to the clone, the writable FlexClone volume uses the same mechanism as a Data ONTAP Snapshot copy to get available blocks from the containing aggregate. But whereas a Snapshot copy simply links to existing data that was overwritten in the parent, a FlexClone volume writes the data to disk (using WAFL) and then links to the new data as well.

The disk space associated with the Snapshot copy and the FlexClone volume is accounted for separately from the data in the parent FlexVol volume. When a FlexClone volume is first created it needs to know the parent FlexVol volume and also requires a Snapshot copy of the parent to use as its base. An existing Snapshot copy can be used, or a Snapshot copy can be created automatically as part of the cloning operation. The FlexClone volume gets a copy of the Snapshot metadata and then updates its metadata as the clone volume is created.

Creating the FlexClone volume takes just a few moments because the copied metadata is very small compared to the actual data. The parent FlexVol volume can change independently of the FlexClone volume because Data ONTAP uses the Snapshot copy to keep track of the changes and prevents the original parent's blocks from being reused while the Snapshot copy exists. The Snapshot copy is read-only, and can be efficiently reused as the base for multiple FlexClone volumes. Space is used very efficiently, since new disk space is used only for the small amounts of metadata, or for changes to either the parent FlexVol volume or the FlexClone volume.

From the point of view of the storage administrator, a FlexClone volume is treated just like a FlexVol volume for most operations. The storage administrator can use the Data ONTAP interface to manage volumes,

Snapshot copies, and FlexClone volumes—including getting their status and seeing the relationships between the parent, Snapshot copy, and clone. The main difference is that Data ONTAP forbids operations that would destroy the parent FlexVol or base Snapshot copy, as long as there is one (or more) FlexClone volume depending on them. A FlexClone volume can be split from the parent volume to create a fully independent volume as long as there is adequate free space in the aggregate to copy shared blocks.

Splitting a clone from its parent removes the connection between clone and parent. When the split has been completed, there are no longer any shared blocks, and the cloned volume no longer has any Snapshot copies. If a split operation is interrupted—for example, by a restart—the split automatically resumes from where it left off when the restart is completed. A split can fail if aggregate lacks sufficient space. If this occurs, the administrator can simply free up unused space or add new storage and restart the split.

These technical details should be noted:

- Management information in external files (e.g., */etc* files) associated with the parent FlexVol volume is not copied and modified.
- Quotas for the clone volume get reset rather than being added to the parent FlexVol volume.
- LUNs in the cloned volume are automatically marked offline until they are uniquely mapped to a host system.

To learn more about how FlexClone volumes work, please see TR3347 titled "A Thorough Introduction to FlexClone Volumes" on http://www.netapp.com.

Practical Applications of FlexClone Technology

FlexClone technology provides dramatic improvements for application test and development environments and is tightly integrated with the NetApp patented file system technology and microkernel design in a way that makes competing methods archaic. The technology enables administrators to generate multiple, instant dataset clones without any storage overhead. FlexClone volumes are ideal for managing production data sets, allowing effortless error containment for bug-fixing and incremental development. FlexClone volumes simplify platform upgrades for enterprise resource planning (ERP) and customer relationship management (CRM) applications and provide data for multiple simulations against large datasets for Electronic Computer Aided Design (ECAD), Mechanical Computer Aided Design (MCAD), and seismic applications—all without unnecessarily duplicating data or wasting physical space. The ability to split FlexClone volumes from their parents lets administrators easily create new permanent, independent volumes for forking project data.

Table 1 lists some of the applications and benefits of FlexClone technology.

Applications	Benefits
Application Testing	▪ Administrators can make infrastructure changes without worrying about crashing production systems. ▪ IT personnel can lock in infrastructure changes on the FlexClone volume and designate them to be applied to production volumes during the next maintenance window. ▪ Infrastructure change can be managed with less risk, less stress, and more reliable results.
Data Mining	▪ Data mining operations and software can be implemented more flexibly because both reads and writes are allowed.
Parallel Processing	▪ Parallel processing applications across multiple servers can use multiple FlexClone volumes of a single milestone or production dataset to produce results more quickly.

Online Backup	▪ If administrators discover corruption in the production dataset, they can continue production immediately by mounting the FlexClone volume instead. ▪ They can use database features like DB2 write-suspend or Oracle® hot-backup mode to transparently prepare database volumes for cloning (necessary because databases need to maintain a point of consistency).
System Deployment	▪ Administrators can maintain a template environment and use FlexClone volumes to build and deploy identical environments or environments with defined variations. ▪ They can create a test template that can be cloned as needed for more predictable testing. ▪ They can use the Data ONTAP SnapMirror feature in combination with FlexClone volumes to perform migration faster and more efficiently.
IT Operations	▪ IT staff can maintain multiple copies of production systems for live production, development, testing, reporting, etc. They can refresh working FlexClone volumes regularly to work on data that is as close to live production data as possible ▪ They can use FlexClone volumes for special runs that can now happen in parallel with regular operations.

Table 1) Sample applications for Data ONTAP 7G FlexClone volumes.

FlexClone Performance

Because Data ONTAP cloning operations are tightly integrated with WAFL and the controller architecture, the performance of FlexClone volumes is nearly identical to the performance of flexible volumes. Unlike other implementations of cloning technology, FlexClone volumes are implemented as a simple extension of existing core mechanisms.

With FlexClone technology, the impact of cloning operations on other system activity is relatively light and transitory. The creation of a FlexClone volume is nearly identical to taking a Snapshot copy, and is usually completed within seconds. The clone metadata is held in memory just like a regular volume, so the impact on controller memory consumption is identical to having another volume available. Ongoing access to the clone is nearly identical to accessing a regular volume.

When a clone is split to create a fully independent volume, free blocks in the aggregate are used to copy blocks shared between the parent and the clone. This incurs disk I/O operations and could potentially compete with other disk operations in the aggregate. The copy operation also uses some CPU and memory resources, which could impact the performance of a fully loaded controller. Data ONTAP addresses these potential issues by completing the split operation in the background, setting priorities in a way that prevents significant impact on foreground operations. If a critical job requires the full resources of the controller, it is possible to manually stop and restart the split operation.

The final area to consider is the impact on disk usage from frequent operations where FlexClone volumes are split off and used to replace the parent FlexVol volume. Data ONTAP obtains free blocks from the aggregate and assigns them to the split volume, using contiguous chunks if they are available. If there is abundant free space in the aggregate, the blocks allocated to the split volume will be mostly contiguous. If the split volume is used to replace the original volume, the blocks associated with the abandoned original volume become available and can create a potentially large free area within the aggregate. That free area will most likely be largely contiguous. In cases where many simultaneous volume operations have reduced contiguous regions for the volumes, Data ONTAP 7G implements brand-new block reallocation functionality. The new reallocate command makes defragmentation and sequential reallocation more flexible and effective than ever before. It reduces any impact of frequent clone split-and-replace operations, and it also optimizes performance following other disk operations that might unbalance block allocations (e.g., adding disks to an aggregate). For additional information on these capabilities, please see the “Data ONTAP 7G Command

Reference” and “Storage Management Guide.”

SIZING VOLUMES AND AGGREGATES FOR PERFORMANCE

One of the major benefits of Data ONTAP 7G, with its introduction of FlexVol and FlexClone technology, is the ability to distribute I/O operations across currently deployed disks in a much more efficient manner than was previously possible with traditional volumes.

Storage Utilization in Midtier Enterprise Systems

An examination of the usage patterns of several thousand midrange enterprise controllers revealed that while half of the systems containing traditional volumes have an average capacity utilization of over 60%, the majority of user I/O is serviced by one or two volumes that consist of only a very small subset of total disks available to the system. Figures 2-A through 2-C show the distribution of file- and block-based I/O operations across midtier enterprise systems that have three, five, or eight volumes.

A significant number of the systems examined have a traditional volume that services more than 50% of the entire I/O served by the controller. One out of 10 systems has a traditional volume that services more than 85% of all the I/O served by the controller.

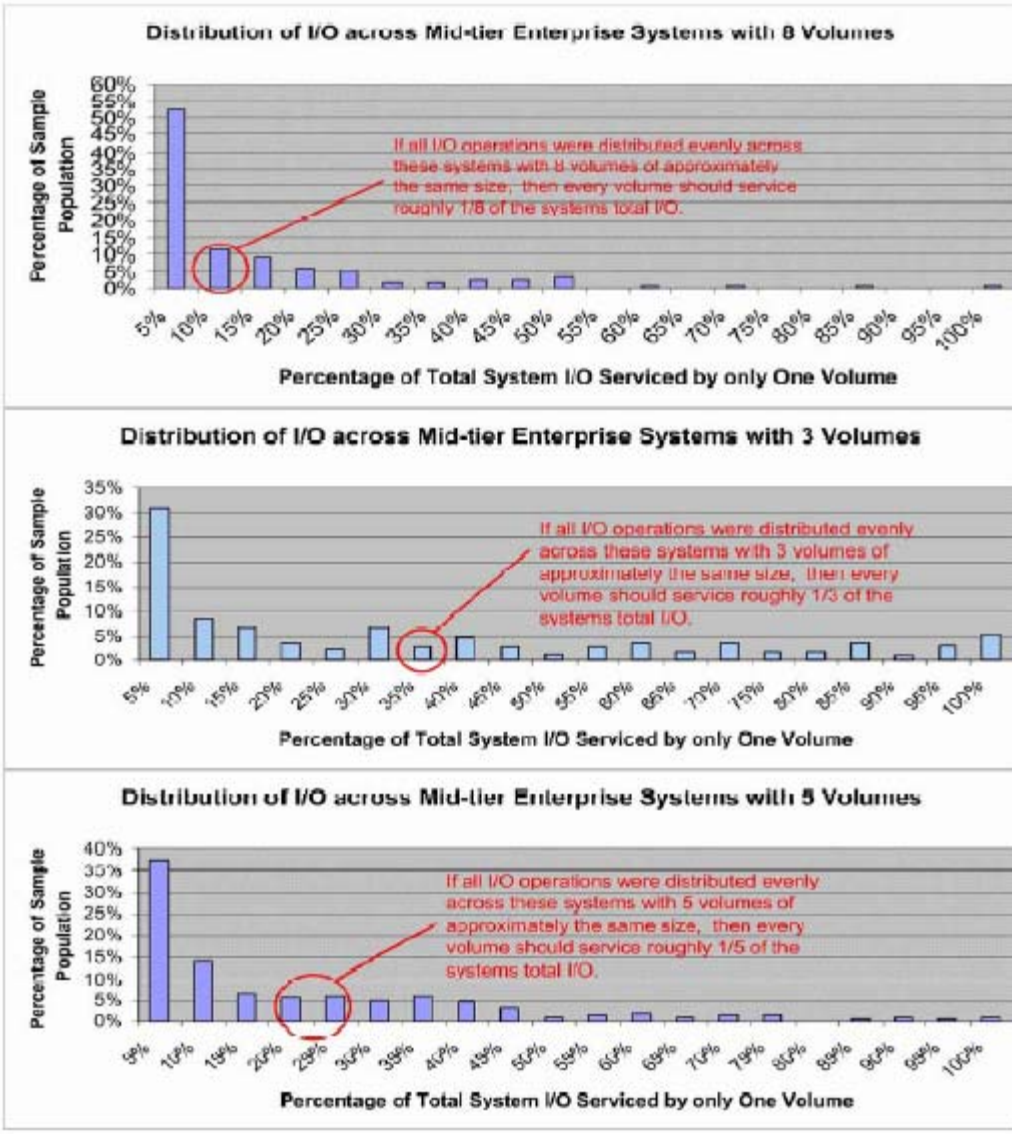

Figures 2-A, 2-B, and 2-C) Distribution of file- and block-based I/O operations across midtier enterprise systems.

If applications hosted on these very busy traditional volumes are spindle-limited, considerable performance gains can be realized by migrating the traditional volume in question to a flexible volume hosted on an aggregate containing all of the storage appliance's disks.

For example, roughly 30% of examined systems service read-intensive workloads comprising approximately 80% reads and 20% writes. A typical deployment of a midrange server has four to five volumes, with each volume composed of approximately 10 disks. If a 70% random, 80/20 read/write workload mix was hosted on a FAS940 traditional volume of 10 disks, then migrated with all the other traditional volumes to flexible volumes hosted on an aggregate containing 40 disks, a potential increase in throughput of up to 110% could be realized by the previously spindle-bound application. If the application performed only reads and no writes, an increase of up to 200% could be realized.

Workloads that are biased toward sequential access of data, or workloads composed primarily of writes, would usually not realize such a dramatic increase in throughput. For example, a 40/60 read/write workload mix accessing data randomly in the same scenario stated above would realize only a 10% gain in throughput on a midrange, enterprise-class storage appliance such as a FAS940.

5-VOLUME DF OUTPUT				
File System	**Kilobytes**	**Used**	**Available**	**Capacity**
/vol/vol0/	56384920	1220588	55164332	2%
/vol/vol1/	789388836	653007668	136381168	83%
/vol/vol2/	789388836	586607788	202781048	74%
/vol/vol3/	394694420	299552752	95141668	76%
/vol/vol4/	394694420	186284920	208409500	47%

Table 2) Example of 5-volume DF output.

Table 2 is an example of a typical five-volume midenterprise storage appliance. If the majority of volumes are almost full of data, but only one or two are serving most of the data in the system, the system is probably a very good candidate for migrating all of the traditional volumes to flexible volumes hosted on a single large aggregate. Striping data across all the disks in the system would benefit applications that frequently access this data and would have little or no impact on applications that infrequently access this data.

Best Practices for Migrating to FlexVol Volumes

Following these general rules will help avoid some common pitfalls in deploying aggregates and FlexVol volumes in Data ONTAP 7G. In general, it is safe to migrate all existing traditional volumes on a single controller to FlexVol volumes on a single, large aggregate on the same controller.

NetApp Professional Services is available to provide a comprehensive analysis and consultation regarding your storage architecture, including the best way to deploy Data ONTAP 7G.

Create the Largest Aggregate Possible: When possible, try to maximize the number of disks in any aggregate, especially during the creation of the aggregate.

Use RAID-DP instead of RAID-4: Always use RAID-DP on large aggregates, especially on systems that include ATA disk drives. RAID-DP offers more than 10,000 times the data protection of RAID-4 and comes at almost no capacity or performance premium. RAID-DP does not require a license.

Use the Default RAID Group Size: Default RAID group sizes have been carefully selected for each platform to balance RAID reconstruction times and useable capacity. Consult with a NetApp systems engineer before changing these values.

Increase Aggregates in Increments of RAID Group Sizes: For example, if the default RAID group size for the platform is 14, plan to add 14 disks to the aggregate at once instead of one or two disks at a time.

Stagger Recurring Activities Like Snapshot copies and SnapMirror Transfers: When possible, stagger recurring system activities like FlexVol Snapshot schedules and SnapMirror transfers so that they start at different times. This prevents performance degradation from undue stress on disk operations.

Use Homogenous Disk Drive Capacities within RAID Groups: Write allocation in Data ONTAP is more efficient if all disks within a single RAID group are of the same capacity. Try to avoid single RAID groups consisting of mixed-capacity drives.

Avoid Creating Small Aggregates: Try to create aggregates that contain at least one to two shelves of disk drives. Smaller aggregates can become disk-bound for even sequential workloads.

Separate Aggregates for Separate RPM Drives: In order to fully realize the lower response times of 15,000 RPM drives, create separate aggregates for 15,000-RPM drives and 10,000-RPM drives.

Allow Data ONTAP to Choose Disks and Adapters during Aggregate Creation: Because Data ONTAP automatically spreads aggregates across disk adapters, let Data ONTAP choose the member disks of an

aggregate.

Perform Trend Analysis Frequently: Software management tools like DataFabric® Manager 3.0.1 and the Performance Advisor can greatly assist in tracking resource utilization of hundreds of NetApp storage appliances.

Caveats

FlexVol volumes are made possible by adding an additional layer of metadata between the data stored on disk and the logical volumes presented to the client or host. This added layer of storage virtualization requires more CPU resources than traditional volumes. While this additional overhead is relatively small compared to traditional volumes, storage appliances that are currently running with over 90% CPU utilization or back-to-back consistency points during peak times are not recommended for migrations to FlexVol volumes. Please contact NetApp Professional Service for comprehensive analysis. This will ensure that the adoption of FlexVol volumes and large aggregates will not negatively impact system performance.

Migrating traditional volumes to flexible volumes hosted on an aggregate that contains all of the volumes of a storage appliance greatly reduces the possibility of application performance becoming spindle-limited in the future. Additionally, the added granularity of data management features such as Snapshot copies for FlexVol volumes, storage provisioning, FlexClone writable Snapshot copies, and logical volume migration across multiple storage appliances can be realized only by the adoption of FlexVol volumes and aggregates.

Migrating traditional volumes to flexible volumes hosted on an aggregate with more disks is no guarantee of better performance. Applications that are not limited by the number of disks servicing their I/O will not witness any noticeable change in performance. You can quickly evaluate if there is contention for disk resources during peak application traffic times by looking at the output of `sysstat –u 1` from the command line.

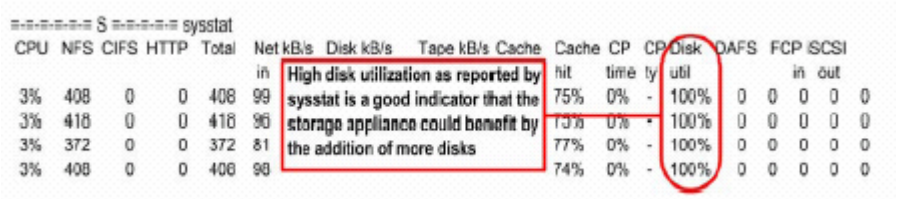

Figure 3) Sample output of sysstat –u 1 command showing high disk utilization.

If the output of the `systat –u 1` command does not show disk utilization of higher than 90%, applications hosted on the NetApp storage appliance will most likely not realize an immediate improvement in either response times, throughput, or job completion time. A comprehensive service is available with NetApp Professional Services to help you determine whether traditional volumes on a NetApp storage appliance are good candidates for migrating to FlexVol volumes on a large aggregate. Please talk to a member of your NetApp technical team to find more information about this service.

BLOCK MANAGEMENT WITH DATA ONTAP 7G

Data ONTAP 7G includes many new features and utilities that make it easier to create and manage LUNs, such as simplifying volume and LUN planning and allowing LUNs to be optimized for efficiency and performance. Data ONTAP 7G also includes enhancements for iSCSI support.

FlexVol and FlexClone

With the release of Data ONTAP 7G, administrators can choose to use traditional or flexible volumes. (FlexVol and FlexClone are described in detail earlier in this document.) With proper planning, these capabilities greatly increase flexibility and efficiency. Some of the benefits are:

- Multiple LUNs with separate Snapshot schedules can share in the performance provided by a pool of disks. This allows optimal use of disk spindles
- Groups of related LUNs on a volume can be cloned by simply cloning the volume

- Destinations that are mirrored by SnapMirror can have writable LUNs created on a different volume simply by cloning the volume.

LUN Reallocation

Data ONTAP 7G introduces a new “reallocate” command to manage reallocation activities. Reallocation scans ensure that blocks in a LUN or a volume are laid out sequentially. This improves the performance of read and write commands, thereby improving the read/write performance of LUN-based applications that access data on the controller. A reallocation scan:

- Runs in the background so applications can continue to run.
- Evaluates the layout of blocks in a LUN to determine the current level of optimization.
- Rewrites blocks as necessary in order to achieve and maintain the desired level of optimization.
- Dynamically adjusts to avoid taking too many system resources while reallocating the volume or file.
- Running a LUN reallocation scan run before the creation of a Snapshot copy optimizes both the LUN and the Snapshot copy.

Snap Delta

Snapshot space is an important factor in effective planning for volume sizing, including the fractional-reserve value for LUN space guarantees. Snapshot space is determined by the number of Snapshot copies to be kept, the frequency of the Snapshot copies, and the amount of data that changes from one Snapshot copy to the next. To aid in planning, Data ONTAP 7G introduces the `snap delta` command, which displays the amount of data that has changed between Snapshot copies.

Snap Reclaimable

When Snapshot copies are taken frequently and large amounts of data change between copies, the amount of disk space used by the volume can grow quickly. Data ONTAP 7G introduces the `snap reclaimable` command. By running the command on one or more Snapshot copies on a volume (or multiple volumes), the storage administrator can determine the amount of space that can be recovered if the selected copies are deleted. This helps avoid or postpone the need to add additional storage resources.

Terminology Changes from Previous Releases

To support the ability to create flexible volumes and FlexClone volumes, as well as the ability to split clones, the structure of certain LUN commands has been modified to match those that now exist at a volume level. The entity that was referred to in previous releases as a writable Snapshot copy or backing store LUN is now referred to as a LUN clone. What used to be called a LUN clone is now referred to as a split LUN clone. This change removes any possible confusion when cloning LUNs and cloning volumes.

iSCSI Enhancements

Data ONTAP 7G provides enhancements for iSCSI support, including multiple iSCSI connections per session and better error handling and recovery. These features are available starting with Data ONTAP release 7.1.

An iSCSI session between an initiator and the storage system can use as many as 16 TCP/IP connections. Multiple iSCSI connections per session improves high availability and allows for bandwidth aggregation. Part of the multiple iSCSI connection feature are sophisticated error recovery techniques that reduce the impact of error conditions on applications.

Data ONTAP 7G now supports iSCSI error recovery levels 1 and 2. This improves data transmission error handling and recovery with initiators that also support this feature.

IMPACTS OF DATA ONTAP 7G FOR SNAPMIRROR

SnapMirror is a software product that allows a dataset to be replicated between NetApp controllers and NearStore® systems over a network, typically for backup or disaster-recovery purposes. SnapMirror is enhanced by the introduction of FlexVol and FlexClone technology, described earlier in this document, and

by the introduction of synchronous and semi-synchronous modes.

Synchronous and Semi-Synchronous Modes

NetApp recently added a synchronous mode to SnapMirror software, which sends updates from the source to the destination as they occur, rather than according to a predetermined schedule. This guarantees that data written on the source system is protected on the destination even if the entire source system fails. A semi-synchronous mode is also provided. It minimizes data loss in a disaster while also minimizing the extent to which replication impacts the performance of the source system.

No additional license fees need to be paid to use this feature; all that is required is appropriate hardware, the correct version of Data ONTAP software, and a normal SnapMirror license for each storage system.

Unlike asynchronous mode, which can replicate either volumes or quota trees, synchronous and semi-synchronous modes work only with volumes.

The synchronous and semi-synchronous modes of SnapMirror are discussed briefly in this document. Complete details can be found in the paper TR3324: “Synchronous SnapMirror Design & Implementation Guide.”

Flexible Volumes with SnapMirror

The size of a flexible volume can be changed dynamically. Additionally, it can act as a hard quota for a group or project assigned to it. In each volume, user- and group-level quotas as well as qtrees can be used to obtain finer granularity of quota management.

In Data ONTAP 7G, an aggregate is equivalent to a traditional volume. There is no aggregate SnapMirror. The following variations exist:

- Traditional volume SnapMirror (synchronous, semi-synchronous, and asynchronous)
- Flexible volume SnapMirror (synchronous, semi-synchronous, and asynchronous)
- qtree SnapMirror (asynchronous only)

In addition, there is AggrCopy for flexible volumes and VolCopy for traditional volumes.

The Disk Geometry Issue

In earlier releases of Data ONTAP, disk geometry was a concern for storage administrators who used SnapMirror. If the sizes of the source disk were not the same as the sizes of the destination disks, problems occurred that resulted in some spindles not getting any data spread over them.

The Three Modes of SnapMirror

SnapMirror can be used in three different modes: asynchronous, synchronous, and semi-synchronous.

Asynchronous Mode

In asynchronous mode, SnapMirror performs incremental, block-based replication as frequently as once per minute. Please consult with your technical team for the best plan for your environment, or to find out whether synchronous SnapMirror is a better match. Performance impact on the source controller is minimal as long as the system is configured with sufficient CPU and disk I/O resources.

The first and most important step in asynchronous mode, involves the creation of a one-time, baseline transfer of the entire dataset. This is required before incremental updates can be performed. This operation proceeds as follows:

- The primary storage system takes a Snapshot copy (a read-only, point-in-time image of the file system).
- This Snapshot copy is called the baseline Snapshot copy.
- All data blocks referenced by this Snapshot copy and any previous Snapshot copies, are transferred and written to the secondary file system.
- After initialization is complete, the primary and secondary file systems will have at least one Snapshot copy in common.

After initialization, scheduled or manually triggered, updates can occur. Each update transfers only the new and changed blocks from the primary to the secondary file system. This operation proceeds as follows:

- The primary storage system takes a Snapshot copy.
- The new Snapshot copy is compared to the baseline Snapshot copy to determine which blocks have changed.
- The changed blocks are sent to the secondary and written to the file system.
- After the update is complete, both file systems have the new Snapshot copy, which becomes the baseline Snapshot copy for the next update.

Because asynchronous replication is periodic, SnapMirror is able to consolidate writes and conserve network bandwidth.

Synchronous Mode

In synchronous mode, SnapMirror immediately replicates all data written to the primary file system. This guarantees zero data loss in the event of a failure, but can have a significant performance impact. It is not necessary or appropriate for all applications.

The first step involved in synchronous replication, is a one-time, baseline transfer of the entire dataset, just as in asynchronous mode, as described above.

Once the baseline transfer has completed, SnapMirror can change to synchronous mode, as follows:

- Asynchronous updates occur, as described above, until the primary and secondary file systems are very close to being synchronized.
- NVLOG forwarding begins. This is a method for transferring updates as they occur.
- Consistency point (CP) synchronization begins. This is a method for ensuring that writes of data from memory to disk storage are synchronized on the primary and secondary systems.
- New writes from clients or hosts on the primary file system are blocked until acknowledgment of those writes has been received from the secondary system..
- One final update occurs using the same method as asynchronous updates, as described above.

Once SnapMirror has determined that all data acknowledged by primary has been safely stored on the secondary, the system is in synchronous mode. At this point, the output of a SnapMirror status query will show the relationship is “In Sync.”

Note: If the environment is not able to maintain synchronous mode (because of networking or destination issues), SnapMirror will drop to asynchronous mode. When the connection is reestablished, the source controller will asynchronously replicate data to the destination once each minute, until synchronous replication is reestablished. Once this occurs, a message will be logged of the change of status (“into” or “out of” synchronous status). This “safety net” is known as failsafe synchronous.

Semi-Synchronous Mode

Semi-synchronous mode provides a middle ground that keeps the primary and secondary file systems more closely synchronized than asynchronous mode, but with less impact on application performance. Configuration of semi-synchronous mode is identical to configuration of synchronous mode, with the addition of an option that specifies how many writes can be outstanding (unacknowledged by the secondary system) before the primary system delays acknowledging writes from the clients.

Internally, semi-synchronous mode works identically to synchronous mode in most cases. The only difference lies in how quickly client writes are acknowledged; the replication methods used are the same.

FlexClone Impact on SnapMirror

Cloning offers a nearly instantaneous replica of a volume within the same aggregate. It also offers substantial space savings for work environments that require multiple copies of the same data (e.g., source trees, chip simulations, or database development) without causing any performance bottlenecks.

A FlexClone volume can be split from its parent to create a new standalone volume. It is also possible to create a writable volume from a read-only SnapMirror destination.

Cloning is available only with flexible volumes, not with traditional volumes. Cloning does not require any special hardware.

Deleting the Source in a SnapMirror Deployment

SnapMirror creates Snapshot copies at the beginning of an update and deletes copies at the end of an update. Customers often set their SnapMirror schedules such that many updates are initiated at the same time. As a result, SnapMirror creates many simultaneous Snapshot copies on the source volume, which can be wasteful of space. If the source is deleted, then the SnapMirror relationship will abort. The reason for this is that the source of the SnapMirror relationship is actually the Snapshot copy. This caveat applies to the volume where the Snapshot copy lives as well. If the administrator deletes the source, and therefore the Snapshot copy as well, SnapMirror must abort.

When a FlexClone is created, either the source volume, or the destination volume, it is bound to the source Snapshot copy. Even if the FlexClone volume was created remotely at the destination, deleting the source Snapshot copy on the source system will result in a failed SnapMirror update when SnapMirror attempts to relay the deletion to the destination of the SnapMirror relationship.

Implications for Synchronous and Asynchronous SnapMirror

With synchronous SnapMirror, a Snapshot copy is made on the destination volume every time a write is done on the source. The Snapshot copy may be deleted from the clone but not from the source volume while the SnapMirror relationship is “in sync.” Synchronous SnapMirror has a “hard lock,” whereas asynchronous SnapMirror has a “soft lock.” If the process falls out of synchronous mode, it will revert to asynchronous mode and become a “soft lock.”

Synchronous SnapMirror keeps the source and destination in sync as much as possible:

- If the NVLOG channel requests (per op) time out .
- If the CP on the source takes more than one minute.
- If network errors persist even after three retransmissions.
- If the source or destination fail to restart.
- If the network connection fails.

In such situations, synchronous SnapMirror completes an *asynchronous* update within one minute. It also turns on consistency point forwarding and NVLOG forwarding.

Soft Locks and Parent Volumes

When a replication-based backup is run, administrators sometimes need to update the secondary system from a source other than the most recent transfer. For example, if the secondary system is destroyed, the extra copy can be moved into its place. However, since the data on the backup secondary system is probably older than the original secondary system, the next backup update will need to be based on an older point in time. Primary systems cannot retain the state from each transfer, so secondary systems (and their backups) must be able to specify specific backup contexts that are requests. The soft lock is a method of specifying these requests. A soft lock is a request to a system to retain the context for rerunning a transfer. This type of lock request is not required by the system, and the locking mechanism does not prevent accidental or other types of deletion.

A FlexClone volume cannot be used as a SnapMirror destination. However, users can clone SnapMirror volumes.

FlexClone data resides in the parent Snapshot copy, so operations that would destroy the parent Snapshot copy are not allowed. The following operations are not allowed:

- Destroy parent volume (but it *can* be taken offline)
- Destroy parent shared Snapshot copy
- “SnapMirror initialize” over a parent volume
- “SnapRestore” a parent volume

- "SnapMirror resync" a parent volume

Splitting a Clone in a SnapMirror Deployment

Splitting a FlexClone volume from its parent removes the connection between the clone and its parent.

The administrator can split a FlexClone volume in a SnapMirror environment without concern, because it will not impact the SnapMirror transfer. This is because once the FlexClone volume is split, it becomes an independent entity. In fact, it is a good idea to split clones that have been used for an extended period of time, in order to avoid any impact on SnapMirror, especially if the source Snapshot copy could be deleted.

Impact on Synchronous SnapMirror in an FCP Environment

Before the release of Data ONTAP 7G, the requirements were as follows:

- The administrator must "SnapMirror Break" the destination in order to make the volume writable.
- The administrator must use Qtree SnapMirror (QSM) for LUN cloning.

With Data ONTAP 7G, the following applies:

- The administrator can create a clone volume that already attains writable attributes.
- The administrator can use either QSM (LUN cloning) or Volume SnapMirror (VSM) (FlexClone) provided they desire writable LUNs on the destination or the ability for clients to map and use LUNs on the destination.

Volume Capacity and SnapMirror

Before the release of Data ONTAP 7G, a SnapMirror destination volume's size had to be at least the same as the source volume's size before SnapMirror initialization and before an incremental transfer.

With ONTAP 7G, the source capacity must be less than or equal to the destination capacity when using flexible volumes. When the administrator performs a SnapMirror Break and the destination capacity is greater then the source capacity, the destination volume shrinks to match the capacity of the smaller source volume. This is a much more efficient usage of disk space, since it avoids consumption of unused space.

Guarantees in a SnapMirror Deployment

Guarantees determine how the aggregate preallocates space to the flexible volume. SnapMirror never enforces guarantees, regardless of how the source volume is set. As long as the destination volume is a SnapMirror destination (replica), the guarantee is volume-disabled. Subsequently, the guarantee mode is the same as the volume mode when the volume is broken off using SnapMirror Break.

Impact on SnapMirror when Overcommitting the Aggregate

When users require additional space, the administrator can increase the size of the aggregate by assigning additional disks to the aggregate. In a SnapMirror configuration, overcommitting the aggregate allows for a more efficient use of disk space on the destination. Only the data that is used on the SnapMirror source will be used by the FlexVol volume on the SnapMirror destination. If that SnapMirror destination is broken, the disk usage is deducted from the overall aggregate. Unless mirrors are broken, you can have many source volumes of varying sizes all mapped to flexible destination volumes.

To overcommit an aggregate, create flexible volumes with a guarantee of *none* or *file* so that the volume size is not limited by the aggregate size. The total size of the flexible volumes can be larger than the containing aggregate.

TAPE BACKUP AND RECOVERY WITH DATA ONTAP 7G

All Data ONTAP operations involving tape that operated on traditional volumes before Data ONTAP 7G now work on both traditional volumes and flexible volumes, but not on the aggregate container. Specific commands include `dump` and `restore` as well as `snapmirror store` and `snapmirror retrieve`. This is also true when these commands are issued via a third-party backup or management application through the NDMP protocol. Table 3 lists the options.

Command	Traditional volume	FlexVol	Aggregate
`dump`	Yes	Yes	No
`restore`	Yes	Yes	No
`snapmirror store`	Yes	Yes	No
`snapmirror retrieve`	Yes	Yes	No

Table 3) The effect of management commands on various types of volumes.

NetApp ensures both forward and backward compatibility of the `dump` and `restore` commands. This allows recovery of data from previous releases of Data ONTAP onto a system running Data ONTAP 7G or even recovery of data from a system running Data ONTAP 7G to a system running an earlier release. There is only one limit on this capability: if Fibre Channel or iSCSI LUNs are involved, data must be recovered to a release that supports those LUNs.

Since the `snapmirror store` and `snapmirror retrieve` commands use volume-based SnapMirror, it is not possible to retrieve data with these commands to an earlier major release of Data ONTAP. It is also not possible to retrieve data stored from a flexible volume onto a traditional volume, or vice versa. These commands can be used only to replicate data between two flexible volumes or two traditional volumes running the same major release of Data ONTAP.

Prior to the release of Data ONTAP 7G, users were restricted to a maximum of eight simultaneous `dump` and `restore` operations per NetApp storage system. This restriction also applied to operations submitted via NDMP from a third-party backup application. With Data ONTAP 7G, this limit has been increased to 16 simultaneous `dump` and `restore` operations. Individual NetApp platforms may have limits below 16 simultaneous operations based on design specifications.

The `dump` command operates slightly differently when the volume being backed up is the destination for a volume-based SnapMirror relationship. The `dump` command always operates on a Snapshot copy of the volume in question. If it is not explicitly pointed to a specific Snapshot copy, it will create a new Snapshot copy and operate on that copy. For SnapMirror destination volumes, it will either use the Snapshot copy specified, or, if no Snapshot copy was specified, it will use the most recent copy. This is because Snapshot copies cannot be created on a SnapMirror destination volume since it is read-only. This behavior has been in place since Data ONTAP 6.2.

A side effect of this behavior is that the Snapshot copy being dumped on the destination volume will become locked and thus cannot be deleted on the destination until the dump operation completes. A SnapMirror update that includes the deletion of this Snapshot copy from the source system will therefore not succeed until the dump completes. This can create a larger than normal lag between the source and destination systems.

With the introduction of FlexVol and FlexClone technology, it is now possible to work around this situation if the SnapMirror source and destination volumes are flexible volumes. Here's how. Before initiating the dump of the SnapMirror secondary flexible volume, create a FlexClone copy of the volume. Split the volume off immediately from the SnapMirror secondary volume, then perform the dump operation on the split volume as normal. This will allow the dump to operate completely independent from SnapMirror updates and eliminate any unnecessary lag between the source and destination volumes.

When properly implemented, the `dump` command works faster with flexible volumes. Since aggregates generally contain more disk drives than traditional volumes, and since flexible volumes utilize all of these drives, performance in general is improved. This benefit will be particularly noticeable for small flexible volumes and flexible volumes with large numbers of small files.

Tape SAN Support with Data ONTAP 7G

NetApp supports a wide variety of tape libraries attached via Fibre Channel and Gigabit Ethernet SAN configurations. For the most up-to-date information on supported configurations, please see the following Web site: http://www.netapp.com/osn/info/config.html. This site is constantly updated as new configurations are added. Data ONTAP 7G supports all configurations listed on this location.

ENHANCEMENTS IN DATA ONTAP 7G FOR V-SERIES APPLIANCES

Data ONTAP 7G has several enhancements that directly affect V-Series systems:

- Support for V-Series zone checksums
- Support for Fibre Channel attachment to servers
- Support for selected HP XP storage arrays
- Support for selected IBM DS4000 storage arrays
- A new V-Series appliance called GF980c
- Support for aggregate volumes

Checksums

Zone checksums have been introduced on V-Series products as a measure to improve the storage capacity available to the customer. Zone checksums on V-Series systems are a revival of a technology introduced in Data ONTAP 6.0.1. The current block checksum method of data validation was introduced as a continuation of the standard block checksum (BC) method used on regular controllers.

V-Series Block Checksum

With the introduction of the V-Series product family, the block checksum method was modified to accommodate the 512-byte sector presented by the storage subsystem. Every ninth sector is used to store checksum data for the previous eight sectors. Every 64th sector is unused to preserve the alignment and spacing of the WAFL volume block. The normal controller volume block is eight sectors on a 520-byte/sector drive. The volume block includes 4K for data and 64 bytes for a checksum. Since the checksum per block is 8 bytes in length (for a total of 64 bytes for 8 sectors), 448 bytes in the ninth sector is not used. In addition to the every-ninth-sector data correction, the unused 64th sector amounts to a 1.4% storage penalty of unavailable space due to alignment. This amount, factored in with the checksum data, represented 12.5% overhead for checksum processing, as illustrated below.

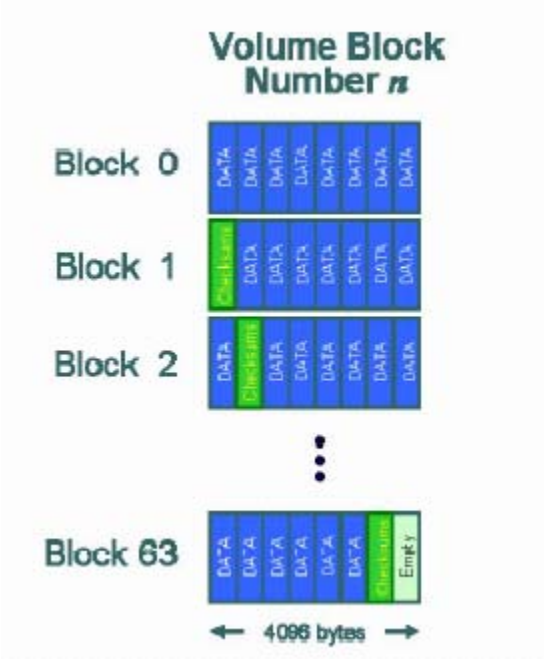

Figure 4) The effect of management commands on various types of volumes.

V-Series Zone Checksum

The V-Series zone checksum method is the logical progression of checksum methods for the V-Series product line. Like the original zone checksum method, the V-Series zone checksum method works with 512-byte blocks, which is what the storage subsystem serves up.

A 4096-byte (4KB) block of checksums is appended to every 63 4KB blocks of data written. The checksum block is read every time one of the 63 data blocks is read or modified. The checksum block is updated whenever a data block is modified. Data and checksum reads and writes require separate I/O commands, which increases disk utilization and can affect performance.

Reverting to any release earlier than Data ONTAP 7G is not possible with zoned checksums.

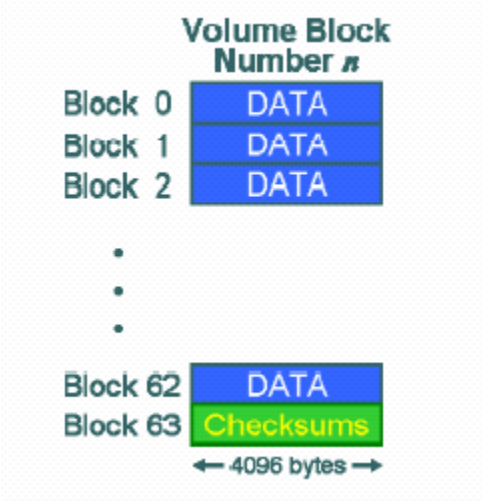

Figure 5) The effect of management commands on various types of volumes.

Zone and Block Checksum Caveats

As with the controller zone and block checksum schemes, the V-Series zone and block checksum schemes are not interchangeable. A V-Series aggregate volume that has a RAID checksum type of zone can use LUNs that have been defined as zone or block checksum. However, a V-Series aggregate volume that has a RAID checksum type of block can use only LUNs that have been defined as block checksum. It is recommended that the LUN be reformatted to the same checksum type as the aggregate volume in which it will be placed.

The command that is used to define the LUN as zone or block is the `disk assign` command. This command is normally used to assign LUNs to a V-Series system. The default assignment is block. To make the LUN a zoned LUN, the `-c` switch is used with the zoned parameter. This command can be used singly or with the all parameter for bulk assignments.

For example, to create a single disk from a LUN with zone checksum enabled, issue the following command:

```
disk assign myLUN.L125 -c zoned
```

The `-c` block switch can also be used to convert a LUN back to block format. Before the command can be issued, it must be unassigned first using the `disk assign` command:

```
disk assign myLUN.L125 -s unowned -f
```

Once the LUN has been unassigned, it can be reassigned using the `-c` block switch. The result will be a reformatting of the LUN and a loss of data.

Upgrading

The V-Series zone and block checksum are not interchangeable formats and would require a volume migration tool, to move the contents of this volume, to another volume.

Performance Trade-offs

The V-Series block checksum structure formatting economizes on I/O at the expense of capacity. Here's a summary of the results for the standard benchmarks. On the whole, block checksum performs between 0% and 35% better than zone checksum in terms of throughput.

System File Server (SFS)

Block = 10.4k ops at 2.8MB per second average response time vs. Zone = 10.0k ops at 3.5MB per second average response time. Block squeezes out 2% more throughput and achieves lower response times.

4KB random writes

Block = 5.8k ops. Zone = 5.7k ops. Again, block gets 2% more ops.

4KB random reads

Block = 4.25k ops at 83% CPU. Zone = 3.16k ops at 74% CPU. Block achieves 35% more ops, and this seems to be disk-limited in both cases. Block incurs much less NFS response time at higher throughput.

Aggregate NFS reads

Block = 103 MB per second. Zone = 105 MB per second. Interestingly, block is actually slightly lower in performance here but also has a lower NFS response time.

Aggregate NFS writes

Block = 43.3 MB per second. Zone = 44 MB per second. Again, block is slightly slower in this case.

In summary, the CPU overhead for zone versus block is mostly within 2% except for random reads, where it jumps to 20%. The biggest impact appears in the disk utilization and disk response times. For random read-intensive workloads, such as databases, zone may not be worth the ~12.5% space savings.

FCP Block Access to Hosts Now Supported

Fibre Channel Protocol (FCP) attachment to hosts is supported in Data ONTAP 7G. The same FCP block access that a regular FAS system uses is now available on the V-Series product line, meaning the complete suite of block access protocols is now available. Customers can now take advantage of NetApp "Snap" technologies for applications that require FCP block access to data even if they already own HDS, HP, or IBM storage.

Selected HP XP Storage Arrays Now Supported

The current models of the HP XP storage line are now supported. These arrays are licensed by HP from Hitachi Japan. The HP models XP48, XP512, XP128, and XP1024 are the equivalent to the HDS Lightning series models 9910, 9960, 9970v, and 9980v respectively. The door panels are different between the two model lines.

The HP arrays are differentiated from the HDS arrays by the HP StorageWorks software that is used to control and configure the HP XP models. Some of the terminology differs slightly between the two vendors, but most of the nomenclature is the same. Documented restrictions for the V-Series products attached to HDS storage arrays will be the same for the HP XP storage arrays. The microcode levels for the HP XP arrays use a different number scheme than HDS. See the "gcontroller Support Matrix for HP StorageWorks Disk Arrays" for specifics. The Integration Guide for the HP Storage Subsystem will document the equivalent microcode levels for the HP XP arrays.

Selected IBM DS4000 Storage Arrays Now Supported

IBM introduced the Fibre Array Storage Technology (FAStT) storage array as a low-cost storage subsystem to support the open systems community. IBM recently changed this product line, designating it the DS4000 line. (See the table below for details.) DS4000 array products are "SAN only," meaning that they can serve block-based data only on Fibre Channel network infrastructures. For NAS access to data stored on a DS4000 array, a gateway product such as a NetApp V-Series system must be deployed.

V-Series connectivity to the newly certified DS4000 arrays is similar to the connectivity to the ESS storage arrays. The DS4000 arrays can be connected to the V-Series system through an approved Fibre Channel switch, or they can be directly connected. No modifications, other than the new Data ONTAP 7G upgrade, are required on V-Series appliances to use these new arrays. The DS4000 models currently supported by V-Series product family include the DS4300 with the persistent reservation feature, the DS4300 Turbo, the DS4400, and the DS4500. Only Fibre Channel expansion units are supported on these models.

The DS4000 devices are licensed from Engenio. Engenio also licenses the same storage and controller modules to other storage vendors. Engenio's other major OEM customers include StorageTek (the D-Series array family), SGI (the InfiniteStorage TP9xxx-series array family), and NCR Teradata (the Quad Modular array product). Some Fujitsu ETERNUS arrays are also based on Engenio array controller bases. The DS4000 line of storage arrays is similar in appearance and architecture to the NetApp FAS product line. A head unit controls the storage, which attaches to the controller head in shelf units called storage expansion units. The storage expansion units are either Fibre Channel (EXP700) or serial ATA (EXP100).

The DS4000 line of storage arrays uses a different storage management tool than the ESS storage subsystem arrays. The DS4000 uses the DS4000 Storage Manager to manage the creation, manipulation, and deletion of flexible volumes on the DS4000 arrays, among other functions. The Storage Manager can be accessed in two ways: in-band and out-of-band. The in-band agent resides on the host and accesses the DS4000 array through the Fibre Channel I/O path. The out-of-band agent communicates with the DS4000 array over an Ethernet connection. Its GUI is completely different from the ESS GUI.

Because the DS4000 family of storage arrays can use either Fibre Channel or serial ATA expansion units, administrators must carefully select applications to run on a V-Series system backed by DS4000 storage. The EXP100 serial ATA expansion units have larger capacity and slower performance than the Fibre Channel expansion units. A DS4000 using EXP100s will have similar characteristics to the NetApp NearStore product line, and customers might deploy these units in similar situations. Be sure to ask what type of expansion unit the controller is serving. The current release of Data ONTAP does not support the EXP100 serial ATA expansion units attached to a DS4000 controller.

One of the unique features of the DS4000 product line is the ability to create storage partitions. A storage partition is similar, yet different, from NetApp V-Series volumes. A storage partition will allow for the segmentation of DS4000 storage. Each partition can be accessed only by the host computers or host groups that have been assigned to it. For clustered V-Series systems this will "zone" the DS4000 storage so that only V-Series appliances assigned to the storage partition will be able to access the LUNs associated with the storage.

As with the IBM ESS Shark arrays, the DS4000 arrays are firmware-sensitive. The current minimum level of DS4000 firmware is 8.4.

Mapping of FAStT Names to DS4000 Series Names

On September 7, 2004, IBM changed the name of the FAStT product line to the DS4000 product line. The naming convention did not follow a one-for-one name change. Table 4 below maps the old product models to the new product models.

IBM nomenclature for DS4000 Series.	
Name Prior to September 7, 2004	**New Name as of September 7, 2004**
IBM TotalStorage FAStT Storage Server	IBM TotalStorage DS4000
FAStT	DS4000
FAStT Family	DS4000 Mid-range Disk Systems
FAStT Storage Manager Vx.y (e.g. FSM V9.10)	DS4000 Storage Manager Vx.y (e.g. SM V9.10)
FAStT100	DS4100
FAStT600	DS4300

FAStT600 with Turbo feature	DS4300 Turbo
FAStT700	DS4400
FAStT900	DS4500
EXP100	DS4000 EXP100
EXP700	DS4000 EXP700
FAStT FlashCopy	FlashCopy for DS4000
FAStT Remote Mirror (RVM)	Enhanced Remote Mirroring for DS4000

Table 4) IBM nomenclature for old and new product lines.

DS4300

The DS4300 is similar to the NetApp FAS270 in that it has 14 Fibre Channel disk drives within the processor enclosure. Additional storage can be dynamically added by connecting DS4000 EXP700 Fibre Channel disk enclosures. Upgrading to FAStT firmware version 8.4 on the DS4300 will enable the persistent reservation feature.

DS4300 Turbo

The DS4300 Turbo is an up-rate model of the DS4300. The Turbo model has more cache, storage, and connectivity options. The DS4300 Turbo also includes the persistent reservation feature. The DS4300 can attach the DS4000 EXP100 Serial ATA disk expansion units or DS4000 EXP700 Fibre Channel disk expansion units. Although the DS4300 Turbo can attach the SATA EXP100 expansion unit, this configuration is not supported by the V-Series product line at this time.

DS4400

The DS4400 is one of two models in the FAStT line that is exclusively Fibre Channel. The DS4400 is similar to the NetApp FAS940. It supports only the DS4000 EXP700 expansion unit.

DS4500

Similar to the NetApp FAS940, the DS4500 is the top-of-the-line model in the IBM DS4000 line. It supports both the DS4000 EXP700 and DS4000 EXP100 expansion units. Although the DS4500 can attach the SATA EXP100 expansion unit, this configuration is not supported by the V-Series product line. Currently the DS4500 will support either the Fibre Channel EXP700 or the Serial ATA EXP100, but not both at the same time.

New V-Series Model

A new V-Series system, the GF980, is introduced with the release of Data ONTAP 7G. The GF980 will also be available as a cluster model called GF980c. This will align the V-Series product line with the FAServer® product line.

New Aggregate and FlexVol Support

Both aggregates and FlexVol volumes are supported on the V-Series product line. There is no difference in the command structure or usage between a regular controller and a V-Series system. A V-Series system provisioned with LUNs defined on an IBM storage array will be able to create and manipulate aggregates and flexible volumes just like a regular controller.

Storage administrators can provision the V-Series system with LUNs that are correctly sized for the array performance in the SAN. Similarly, a FlexVol volume can be created and sized to match the application requirements. This allows a storage administrator to create a generic LUN on the storage array that will be usable for all applications utilizing the V-Series product line. This greatly simplifies the management and

provisioning of the array storage.

ROLE-BASED ACCESS CONTROLS

Role-Based Access Controls (RBACs) determine how users and administrators can use a particular computing environment.

Most organizations have multiple system administrators, some of whom require more privileges than others. By selectively granting or revoking privileges for each user, you can customize the degree of access that each administrator has to the system.

RBACs allow you to define sets of capabilities (roles) that apply to one or more users. Users are assigned to groups based on their job functions, and each group is granted a set of roles to perform those functions. Using this method, the only configuration required for an individual administrator is to ensure that the administrator is a member of the appropriate group or groups.

How RBACs Work in Data ONTAP

While the overall concept of RBACs is applicable to a wide range of operating systems and applications, the details of how RBACs are implemented vary depending on the OS or application in use. This section describes the specific terminology and architecture used in Data ONTAP. It is important to understand these concepts and definitions before configuring RBACs in Data ONTAP, especially if you have experience with RBACs implementations in other software, because the terminology and architecture are not necessarily the same.

Definitions

Users

A *user* is defined as an account that is authenticated on the NetApp system.

A domain user is defined as a nonlocal user who belongs to a Microsoft® Windows® domain and is authenticated by the domain.

Both users and domain users represent individual, authenticated individuals. While it is possible to define a user or domain user that represents a piece of software, or one that is shared among multiple people, this is not a common scenario and is thus not discussed extensively here.

Both users and domain users are assumed to be authorized system administrators. Normal, nonadministrative users who access files on the system via CIFS or NFS, or who use client systems that mount LUNs via FCP or iSCSI, are not discussed in this document. They are unable to log into or manage a Data ONTAP system unless they have been specifically defined as either a user or a domain user via the `useradmin` command, as discussed below.

Groups

A *group* is defined as a collection of users and/or domain users. Groups may be assigned one or more roles.

Groups defined within Data ONTAP are separate from the groups defined in other contexts, such as a Microsoft Active Directory server. This is true even if the groups within Data ONTAP have the same names as groups defined elsewhere within your environment.

Roles

A *role* is defined as a named set of capabilities. Data ONTAP comes with several predefined roles, and users can create additional roles or modify the existing ones.

Capabilities

A *capability* is defined as the privilege granted to a role to execute commands or take other specified actions.

Data ONTAP uses four different types of capabilities.

1. **Login Rights:** These capabilities have names that begin with “login-” and are used to control which access methods an administrator is permitted to use for managing the system.

2. **CLI Rights:** These capabilities have names that begin with “cli-” and are used to control which commands the administrator can use within the Data ONTAP command-line interface.

3. **API Rights:** These capabilities have names that begin with “api-” and are used to control which application programming interface commands can be used. API commands are usually executed by programs rather than directly by administrators. However, it is possible to restrict a specific program to certain APIs by creating a special user account for it, or to have a program authenticate the administrator who is using the program and be limited by that administrator’s roles.

4. **Security Rights:** These capabilities have names that begin with “security-” and are used to allow certain users to execute advanced commands or change passwords for other users.

Putting It All Together

Users are members of groups; groups have one or more roles; and each role grants a set of capabilities. All configuration for RBACs occurs via the `useradmin` command provided by Data ONTAP. For example, users are added or modified with the `useradmin user add` or `useradmin user modify` commands.

Detailed documentation on how to use the `useradmin` command to define users, domain users, groups, and roles is provided in the “Data ONTAP System Administration Guide.”

Benefits of RBACs

Data ONTAP RBACs have the flexibility to meet the needs of almost any IT environment. How they are used will depend largely on local security policies and organizational structure. The following are just a few examples of how RBAC might be used to enhance security and manageability in an enterprise IT environment.

CIFS File Services in a DMZ Environment

Some demilitarized zone (DMZ) environments need to provide CIFS file access services without the benefit of a Windows domain controller or Active Directory server. Using RBACs, customers can set up local user accounts on the storage system for access via CIFS. They can also configure roles and group membership to restrict administrative login to authorized administrative users.

Separate CIFS, NFS, and Storage Administration Groups

Some IT organizations have separate groups of administrators for Windows, UNIX®, and storage administration needs. Using RBAC, it is possible to provide the Windows administrators with access to CIFS configuration functions, provide the UNIX administrators with NFS configuration functions, and provide the storage administrators with volume, qtree, and other storage management functions. This prevents the Windows and UNIX administrators from mistakenly changing the configuration for protocols that are managed by another group, and allows the storage administrators to focus on allocating storage resources while delegating Windows- and UNIX-specific tasks to the appropriate server administrators.

Separate NAS (CIFS and NFS) and SAN (FCP and iSCSI) Administration Groups

Some IT organizations may wish to delegate responsibility based on data class rather than operating system types. For example, RBACs could be used to delegate NAS protocol administration (both NFS and CIFS) to one set of administrators, and SAN protocol administration (both FCP and iSCSI) to another set of administrators. Similarly, an IT organization might wish to set up a network administration group that is permitted to manage DNS, network interface addresses, and other such details but without any protocol-level capabilities.

Performance Monitoring Pseudo-User for Scripts

In some environments, it is necessary to periodically run custom scripts on a UNIX or Windows administration host. This entails logging into the storage system and collecting performance or usage data. RBACs can be used to create a special user account for this purpose, ensuring that the performance collection script has the access level it needs to obtain the required data, and no more. In this way, it is possible to prevent the data collection procedure from executing potentially damaging commands.

Snapshot Copy Creation for Database Backups

Database environments frequently need to coordinate database activities with Snapshot copy creation on the storage system. For example, during a backup operation, the database is generally placed into a “hot backup” or “write suspend” mode. After a Snapshot copy is taken, the database must be placed back into normal operation. This process is generally managed from the database host, and is usually performed by a Database Administrator (DBA) or by a script that the DBA writes and maintains. RBACs allow the storage administrator to grant the DBA enough access to log in and create Snapshot copies as required by this process while preventing the DBA from accessing other functions, such as network configuration, security settings, or storage allocation.

REGULATORY COMPLIANCE

Data ONTAP 7G includes a number of new features related to regulatory compliance. These features are designed to address regulatory concerns about data permanence, security, and confidentiality.

Secure ComplianceClock™

Data ONTAP includes a secure timebase for regulatory compliance features called ComplianceClock. ComplianceClock is a tamperproof method for ensuring that data cannot be deleted before its retention period has expired by manipulating the system clock.

SnapLock

Starting with Data ONTAP 7.1, both Compliance and Enterprise editions of SnapLock can be licensed on the same storage system and infinite retention periods can be specified.

LockVault

LockVault is a product designed for retaining large amounts of unstructured data such as documents, project files, and home directories. LockVault is built upon the proven SnapLock and SnapVault products. With LockVault, retention periods are set on the Snapshot copy created automatically after a SnapVault transfer takes place. LockVault integrates with Open Systems SnapVault (OSSV) as well, creating a compliance solution for open systems without compliant storage.

MIGRATING FROM TRADITIONAL VOLUMES TO FLEXVOL VOLUMES

Because of the underlying differences in the way data is stored on traditional volumes versus new FlexVol volumes, it is necessary to migrate the data from existing volumes to FlexVol volumes rather than converting the data in place. This will require enough spare capacity to continue to store the data being migrated as well as enough capacity to create the new aggregates.

Migration Strategy

One high-level migration strategy involves creating separate aggregates for applications that generate sequential or random I/O access patterns. Use the workload descriptions found in the “FlexVol Guide” and “Aggregate Performance Considerations Guide” to determine which type of application is which.

The overall migration strategy involves consolidating existing traditional volumes onto one or more aggregates that contain the majority of disks attached to the storage appliance. If many small traditional volumes are currently in use, consider migrating their data to several flexible volumes on a large aggregate.

After the data is migrated, each aggregate will probably hold many FlexVol volumes. The maximum number of volumes allowed in Data ONTAP 7G is as follows:

- Up to 100 aggregates and traditional volumes are allowed on the same storage appliance.
- Up to 200 traditional and FlexVol volumes are allowed on the same storage appliance.

Using a few large aggregates versus many smaller aggregates or traditional volumes will result in better capacity utilization because there will be fewer overall parity disks.

Additional information on migrating the root traditional volume to a FlexVol volume can be found in the “Storage Administrators’ Guide” at https://now.netapp.com/.

Qtree Considerations

For environments with a high number of qtrees on traditional volumes, these considerations apply:

- Data ONTAP 7G allows up to 4,995 qtrees per FlexVol volume or traditional volume.
- Migration of the entire traditional volume with all of its qtrees to a new FlexVol volume is generally the easiest approach, unless one of the following is true:
  - You wish to relocate some qtrees to another FlexVol volume.
  - You wish to convert some qtrees into a FlexVol volume.
- Migrating qtrees into their own individual FlexVol volumes buys extra functionality:
  - Snapshot copies are available at the FlexVol level.
  - FlexClone cloning is available at the FlexVol level.
- Migrating every qtree to its own FlexVol volume may not be possible if the following conditions exist:
  - Migrating each qtree to a FlexVol volume would exceed the 200-volume-per-storage appliance limitation.
  - SnapVault operates only at the qtree level, so migrating qtrees to FlexVol volumes may not be desirable.

Preliminary Activities

The performance of the storage environment can be improved and management activities can be reduced as part of the FlexVol migration. Start by documenting the existing volume and qtree configurations. Be aware of data layout issues in the traditional volumes, then evaluate the total capacity needed for the migration. As described below, there are ways to juggle capacity in environments where extra storage is not available.

Make sure you allow for the performance impact of the migration activities. For details regarding the potential performance improvements that can result from migrating traditional volumes to FlexVol volumes, please review the FlexVol and Aggregate Performance Considerations Guide.

NOTE: NetApp Professional Services can help with the data migration effort on a contractual basis and will provide the necessary capacity to complete the job.

Understanding the Migration Environment

Execute a `vol status and df –h` command on the controller from the command-line interface, or use Filerview® Volumes • Manage to see existing volumes and the capacity utilization of each. Then enter volume and volume utilization information in a spreadsheet to simplify the migration planning effort, as shown in Table 5. Run the command qtree from the CLI or use FilerView to determine existing qtrees and enter them into the spreadsheet.

Traditional Volume - Qtree Path	Total GB	GB Used	Capacity Utilization	Target Aggregate	FlexVol	Qtree Path
/vol/vol0	23	5	21%	aggr1	/vol/vol0	
/vol/homedir	95	58	61%	aggr1	/vol/homedirs	
/vol/homedirs/engineering				aggr1	/vol/engineering *	/vol/engineering/homedirs
/vol/homedirs/finance				aggr1	/vol/homedirs	/vol/homedirs/finance

/vol/homedirs/business				aggr1	/vol/homedirs	/vol/homedirs/business
/vol/mailserver	167	115	69%	aggr1	/vol/mailserver	
/vol/mailserver/engineerig				aggr1	/vol/engineering *	/vol/engineering/mailserver
/vol/mailserver/finance				aggr1	/vol/mailserver	/vol/mailserver/finance
/vol/mailserver/business				aggr1	/vol/mailserver	/vol/mailserver/business
Total Size (GB)	285	178		300 **		
* Engineering was broken out into its own FlexVol volume to take advantage of Snapshot functionality available at the FlexVol level.						
** The new aggregate will be 58% utilized after migration allowing adequate capacity for growth.						

Table 5) Planning for the migration effort.

Allocating the Necessary Capacity

If you have limited capacity to spare, then approach the migration in a stepwise fashion. For example, migrate some traditional volumes to a new FlexVol volume, then destroy the original volumes to free up capacity for additional migrations. If you are migrating qtrees, make sure all of them have been migrated before destroying the volume that is hosting them. Ensure the integrity of the migrated data before destroying the original traditional volume, and consider keeping the original traditional volume for an extended period until you are sure the data environment is intact. All of these guidelines are reiterated in detail in the “Migration Setup Tasks” section below.

Here are some additional suggestions for capacity planning:

- Consider backing up traditional volumes to tape or another NetApp storage device for ease of data recovery until Snapshot copies start to accumulate on the new FlexVol volume.
- Migrating the root volume to a FlexVol volume will free up two disks from the traditional root volume.
- If volumes are being replicated to provide extra copies of data (common in database environments), stop replicating data and use the disks instead for migration to FlexVol.
- Once data has been migrated to a FlexVol volume, a FlexClone volume can provide additional copies of data without requiring extra capacity.
- Using `ndmpcopy` or QSM, temporarily replicate data to a second NetApp appliance, or back up the data to tape and verify its integrity. Then destroy the original volume and create a new aggregate and FlexVol volume in its place. Then migrate data back to the new FlexVol volume.

Keep in mind that the data migration process will impact the performance of all storage appliances involved. If you are migrating data to a FlexVol volume on the same appliance, the performance impact will be even greater. If possible, plan to migrate data during nonpeak hours to minimize the impact to production systems.

Best Practices: For large data migrations that might overlap into business hours, consider qtree SnapMirror (QSM) and throttle the replication rate to reduce the performance impact.

Estimating the Time Required for Migration

Knowing how long it will take for a given set of data to migrate is a key part of planning a successful migration to flexible volumes. The results of several time trials are presented below. Trials were carried out

in several production NetApp storage environments. Environments covered include migration to a FlexVol volume on the same controller, migration to a FlexVol volume on a separate controller, and migration to a FlexVol volume on a NearStore appliance.

Migration Setup

As part of the planning for your data migration, it is a good idea to consider the activities that must be carried out immediately after the migration has completed. These are discussed under the heading “Post-Migration Activities,” following the discussions of the time trials. Postmigration activities can be performed relatively quickly, but until they are completed the FlexVol production environment is not ready for use. So be sure to include them in your migration planning.

In general, the migration setup should be a straightforward process. Careful attention to the tasks described here will maximize the benefits available with FlexVol technology. After the setup tasks have been completed you are ready to initiate the data migration from traditional volumes to flexible volumes.

Best Practices

Before migrating live data, rehearse some “practice” migrations with small data sets.

Migration Setup Tasks

1. Create the new aggregate for the flexible volumes. When creating aggregates, use the default setting of RAID-DP for RAID group type in Data ONTAP 7G and the default RAID group size for your appliance. Using the default RAID group size ensures a healthy balance between availability, useable capacity, and performance.

   RAID-DP is a unique double-parity RAID protection solution that can withstand two disk failures within the same RAID group at virtually no performance cost. (You can read more about RAID-DP at TR3298, in our Tech Library.)

   For FAS250 and FAS270 systems, the recommended configuration is one large aggregate consisting of RAID-DP or RAID groups of the default size. The root volume should be a FlexVol volume within the larger aggregate.

   A FAS2xx with all disks used in a RAID DP aggregate is more fault-tolerant than a FAS2xx with a single RAID 4 group and one hot spare.

   When creating the new aggregate, allow enough spare capacity to comfortably accommodate anticipated data growth rates in FlexVol volumes.

2. Create the FlexVol volumes and specify the aggregate name created in the above step in the `vol create` command. By default, `vol create` without the aggregate being specified will create a traditional volume. The FlexVol size must be a minimum of 20MB.

3. Enable `ndmpcopy` or QSM as required on the appliances involved.

4. If you are migrating a data set consisting of many files, make sure the FlexVol volume `maxfiles` and `maxdirsize` options are set appropriately. Use the `maxfiles` and `vol options` commands to view and change these two settings on volumes.

5. To migrate data from traditional to flexible volumes, use either `ndmpcopy`, or in some cases, qtree SnapMirror (QSM). `ndmpcopy` can be used for most data migrations, and no license is required.

   Use QSM if lowering (throttling) the data migration rate is important. This will minimize the performance impact to production environments during the migration process.

   When migrating data that may change during the transfer, use incremental replication with either `ndmpcopy` or QSM so that any data that changes after the initial baseline transfer will be replicated.

Sample Migration Time Trials

These conditions apply to all the time trials discussed below:

- The recommended `ndmpcopy` was used.
- There was no other load on any appliance during the time trials.

These are the characteristics of the source controller used for all migrations in time trials:

- F840
- Small file data set on 8-disk volume (7+1)
- Small file data set size (120GB)
- Large file data set on 5-disk volume (4+1)
- Large file data set size (60GB)

Best Practices

Run some practice migrations with small data sets to validate the migration environment.

Migrating data sets that consist of mostly small files will take longer than migrating data sets containing mostly larger files. For migration planning purposes, base your assessment of file size on the size of the majority of the files in the data set. If the sizes of the files in the data set are unknown, assume small file size when estimating your migration times.

Files less than a few megabytes are classified as small. Typical examples include documents contained in home directories such as word-processing documents, spreadsheets, and presentations, and personal e-mail folders.

Files larger than several megabytes in size are classified as a large. Typical examples include backup archives, image or streaming video files, and seismic data.

Migration on a Single Controller

Figure 6) Single controller Migration

- 11-disk destination volume (10+1)
- 120GB small file data set took 3 ½ hours
- 60GB large file data set took ½ hour

Migration to a Secondary controller over 10/100Mb/sec Network

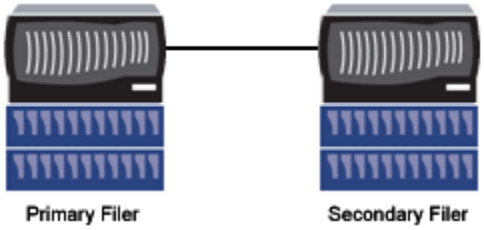

Figure 7A) Secondary controller Migration

- F840 secondary controller
- 14-disk destination volume (13+1)

- 120GB small file data set took 4 hours
- 60GB large file data set took 2 hours

Migration to a Secondary controller over 1000Mb/sec Gigabit Network

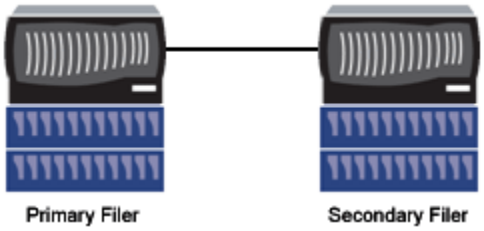

Figure 7B) Secondary controller Migration (GbE Network)

- F840 secondary controller
- 9-disk destination volume (8+1)
- 120GB small file data set took 3 hours
- 60GB large file data set took ½ hour

Migration to a Secondary NearStore Appliance over 10/100Mb/sec Network

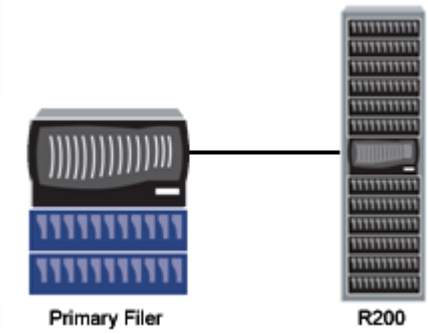

Figure 8) Secondary NearStore Migration

- R200 secondary
- 14-disk destination volume (12+2)
- 120GB small file data set took 3 ½ hours
- 60GB large file data set took 1 ½ hour

Postmigration Activities

These activities can be performed relatively quickly, but until they are completed the FlexVol production environment is not ready for use.

Validate Integrity of Migrated Data

Both `ndmpcopy` and QSM have been field-proven over many years to maintain replicated data integrity. If additional checking is desired, you should perform some statistical sampling data integrity checks with standard checksum tools such as cksum in UNIX or md5sum in the Microsoft environment. Either checksum tool can be incorporated into a script or batch file to make it easy to verify the migrated data.

Note: Snapshot copies are not migrated from the original volume to the new FlexVol volume with `ndmpcopy` or QSM. Automatic Snapshot copies created on the new FlexVol volume will begin with the default Snapshot

schedule setting in Data ONTAP.

Best Practices

- Since Snapshot copies are not migrated, manually generate a Snapshot copy of each new FlexVol volume immediately after migration or after validating the migrated data.
- Keep the original volume for as long as possible for access to its Snapshot copies and back up to tape, if possible, before destroying the original volume.
- Rehearse recovering files on the new FlexVol volume from Snapshot copies on the original volume or tape backups of volumes.

Validate Volume Settings

Review the following settings on the traditional volume from which data was migrated, and make sure the settings on new FlexVol volumes are set correctly:

- Snapshot schedule
- Security settings (Windows, UNIX, mixed)
- Quotas

To take advantage of FlexVol functionality, the settings will not always match those of the traditional volume from which the data was migrated.

Modify Mount Share Information

- Review RBAC (Role-Based Access Control) and security in the “Data ONTAP 7G Guide” to learn more about new options and functionality available for FlexVol access.
- Update NFS export settings to reflect the new FlexVol volume for accessing data.
- Update CIFS shares to reflect the new FlexVol volume for sharing data.
- Remove old NFS or CIFS settings for the original volume when it is destroyed.
- Update CIFS or NFS settings on servers to reflect new share information for each FlexVol volume.

Update Data Protection for New FlexVol Volume

- Update the SnapMirror configuration on the destination to reflect the new FlexVol data source, then reinitialize the mirror.
- Update the SnapVault configuration on the destination appliance to reflect the new FlexVol data source, then reinitialize the relationship.
- Update the NDMP backup application to reflect the new FlexVol data source.
- Update in-house scripts that may have generated backups to reflect the new FlexVol data source.