



Storage Best Practices and Resiliency Guide

Chris Lueth and Shawn Kung, Network Appliance, Inc.

July 2007 | TR 3437

TABLE OF CONTENTS

Abstract	3
1. Introduction	3
1.1. Scope	3
2. General Storage Best Practice Guidelines	4
2.1. Top Resiliency Best Practices	4
2.2. Storage Performance Guidelines.....	5
2.3. SATA Storage and Mixed FC and SATA Best Practices	6
3. Storage Resiliency Features and Options	7
3.1. Inherent Storage Resiliency Features.....	7
3.2. Guidelines That Increase Storage Resiliency	10
3.3. Options That Increase Storage Resiliency.....	13
4. Mapping NetApp Storage Resiliency Features and Options to Storage Tiers.....	17
4.1. Factors Considered in Establishing Storage Tiers.....	17
4.2. Compliance, Reference, and Archival Environments (Tier 4)	17
4.3. Corporate Data, Departmental Workgroups, and Distributed Remote Environments (Tier 3)...	19
4.4. Business-Critical Environments (Tier 2).....	20
4.5. Mission-Critical Environments (Tier 1)	21
5. Application Mapping to Tiered Storage	22
5.1. High-End Storage Applications.....	22
5.2. Upper Mainstream Storage Applications	22
5.3. Lower Mainstream Storage Applications	23
5.4. Near-Line Storage Applications	23
Appendix A: Additional Resources for Storage Resiliency Information.....	24
Appendix B: Revision History.....	26

ABSTRACT

Enterprise data typically has varying degrees of business value that in turn dictate what levels of storage resiliency are required to meet service-level agreements. NetApp offers a rich set of inherent features and available options to implement the appropriate storage resiliency tier for any class of enterprise data. Understanding these storage resiliency features and options, along with NetApp recommendations for each, provides a solid foundation to understand and implement a highly reliable storage environment.

1. INTRODUCTION

When used properly in production deployments, a core set of storage best practice guidelines that apply to all NetApp storage systems helps customers realize maximum data availability and integrity along with optimal performance. Following these core best practice guidelines often requires minimal additional costs in NetApp storage solutions and is easy to incorporate as part of the storage deployment. This report begins by covering general storage best practices, and the information provided should be used as key guidelines for building the optimal NetApp storage architecture to meet the specific needs of your IT environment.

Storage resiliency is a function of data availability and data integrity. Generally, data of higher value requires higher levels, or tiers, of storage resiliency than data of less value. NetApp storage solutions feature a rich set of built-in features and add-on options that provide customers with very fine granularity when selecting the storage resiliency tiers that best meet their business needs.

This paper addresses the following topics:

- Storage best practice recommendations
- Data ONTAP® storage resiliency features and options for improved fault tolerance
- Tiered storage: Mapping features, options, and configurations for different storage resiliency levels based on business requirements for stored data:
 - Tier 1: Mission-critical storage environments
 - Tier 2: Business-critical environments
 - Tier 3: Corporate data, work group, and distributed enterprise environments
 - Tier 4: Compliance, reference, and archival environments
- Application segmentation: General mapping of IT applications to tiered storage
- References to additional information

1.1. Scope

This report does not provide detailed technical information or deployment guidelines on the built-in features and add-on options that address storage resiliency. Instead, the appendix provides links to additional resources for those who desire more technical depth. As always, if there are specific questions not answered by this paper or the NetApp technical library, please contact your local NetApp sales or support team.

2. GENERAL STORAGE BEST PRACTICE GUIDELINES

2.1. Top Resiliency Best Practices

- Use RAID-DP™, the NetApp high-performance implementation of RAID 6, for better data protection.
- Use the default RAID group size when creating aggregates or traditional volumes.
- Allow Data ONTAP to select disks automatically when creating aggregates or volumes.
- Use the latest Data ONTAP general deployment (GD) release available on the NOW™ ([NetApp on the Web](#)) site.
- Use the latest storage controller, shelf, and disk firmware available on the [NOW site](#).
- Maintain at least one hot spare for each type of disk drive in the storage system:
 - Disk drive differences are FC or SATA, capacities, and rotational speed (RPM).
- Maintain two hot spares for each type of disk drive in the storage system to take advantage of Maintenance Center.
- Reserve space for at least one week of Snapshot™ backups.
- Do not put user data into the root volume.
- Replicate data with SnapMirror® or SnapVault® for disaster recovery (DR) protection.
- Replicate to remote locations to increase data protection levels.
- Use an active-active storage controller configuration (clustered failover) to eliminate single points of failure (SPOFs).
- Use Multipath HA with Active-Active storage configurations to improve overall system availability as well as promote higher performance consistency.
- Deploy SyncMirror® and RAID-DP for the highest level of storage resiliency.

2.2. Storage Performance Guidelines

- Use larger aggregates with flexible volumes to provide more disk drives to applications, especially ones that would otherwise not consume much storage space.
- If using traditional volumes, follow the same recommendation and use larger volumes with more spindles to avoid potential bottlenecks from insufficient spindle I/O.
- If using SnapMirror or SnapVault:
 - Stagger schedules to reduce impact.
 - Throttle bandwidth on transfers when on a LAN by using the `kbs` argument in the `snapmirror.conf` file (actual replication size using this option is kB/sec).
 - Schedule transfers during off-peak workload times when possible.
- Leave enough performance headroom under the planned peak workload on your storage controller to ensure adequate resources for routine activities like:
 - Snapshot copies
 - SnapMirror and SnapVault transfers
 - RAID reconstructions
 - RAID scrubbing
 - Backups to tape
- Use Multipath HA for Active-Active configurations to increase bandwidth to storage loops.
 - In addition to improving storage resiliency, the bandwidth of the second storage loop path is aggregated with the first loop and storage workload is load-balanced between the two data pathways.
- Using larger SATA disk drives such as the 500GB or 750GB size requires special performance considerations:
 - With these larger disks sizes, fewer of them are required to reach the maximum aggregate size of 16TB.
 - Because fewer data disks are available in an aggregate to handle the application workload, it could take multiple aggregates to reach the performance capabilities for a particular platform.
 - With all things equal, such as RAID group size and same storage workload, in the event of a drive failure, reconstruction of larger disks takes longer to complete.
 - The longer reconstruction time for larger disks is normal.
 - Longer reconstruction times for larger disks means that the slight performance impact associated with reconstruction simply lasts for a longer period of time.
 - Use RAID-DP for maximum data protection during the longer reconstruction.

2.3. SATA Storage and Mixed FC and SATA Best Practices

- General information:
 - FC and SATA disks must be in different shelves and on separate loops.
 - On a standalone storage controller, FC and SATA disks can be connected to the same controller.
 - With active-active configurations, FC and SATA disks can be attached to both controllers or separated on different controllers.
 - Aggregates or traditional volumes cannot span both FC and SATA disks.
- Typical mixed storage deployments target FC disks (10,000 or 15,000 RPM) for high-performance applications and SATA disks (7,200 RPM) for applications that are less sensitive to higher latencies.
- With Data ONTAP 7.2 and later, consider using FlexShare™ in mixed storage deployments to prioritize workloads running on the FC disk drives and to isolate high-performance applications from possible performance issues due to slower SATA disks. See section 6.2 in the [FlexShare Design and Implementation Guide](#) for more information about using FlexShare for mixed storage deployments
- With Data ONTAP 7.1 and earlier, consider using time-scheduled workloads with mixed storage to eliminate possible performance issues from slower SATA disks:
 - During peak load periods requiring the high throughput and IOPs of FC disks, smaller workloads should be generated on the SATA disks.
 - In off-peak periods when performance requirements are reduced, workloads can be spread across FC and SATA disks as needed.
 - Periodic SnapMirror replications from FC to SATA disks during peak workload periods, followed by backup of data on SATA after peak periods, is an example of a time-scheduled workload.
- Perform an extensive application workload analysis to determine IOPs demand and latency requirements before deploying primary applications on SATA storage.
- SATA is not the proper storage solution if the primary application requires low latency or a high number of IOPs, especially with a read-intensive workload.
- When running primary applications on SATA disks, it is important to use large aggregates with flexible volumes so that more disk spindles are available to the applications' workload:
 - Small volumes with three to seven SATA disks are likely to cause storage performance bottlenecks for primary applications.
- When using the larger capacity SATA disk drives such as the 750GB size, it could take several aggregates to reach the performance capabilities for a given platform.
 - It takes fewer large disk drives to reach the maximum aggregate size of 16TB.
 - Fewer data disks in an aggregate could limit the application workload that an aggregate can provide.

3. STORAGE RESILIENCY FEATURES AND OPTIONS

NetApp offers a rich set of features and options that allow customers to meet a wide range of storage resiliency business requirements. As mentioned earlier, the value of data drives the level of storage resiliency required in a particular storage deployment. This section describes a set of best-practice resiliency features and configuration options that provide the building blocks for tiered resiliency.

3.1. Inherent Storage Resiliency Features

RAID checksums. When Data ONTAP writes data to disk, each 4kB block generates a checksum that is stored as part of the block's metadata. When data is later read from disk, a new checksum is calculated and compared to the original checksum from the block's metadata. If the read data generates a different checksum, the requested data is recreated from parity and provided to the client. In addition, the data from parity is rewritten to the original 4kB block, then read back to verify its accuracy.

- Self-healing storage activity.
- Protects against data path and disk drive errors.
- This protection is always on and is an integral part of Data ONTAP.

RAID scrubs. The RAID checksum approach described earlier ensures that data being read from disk is accurate before it is served to requesting clients. In other words, frequently read data tends to benefit the most from RAID checksums. However, not all data is frequently read from disk. Examples are files in home directories, which tend to be accessed less frequently as they age, or archived files that are rarely accessed over the course of their retention life. RAID scrubs are a configurable NetApp feature that traverses the storage to read each 4kB block, which triggers the RAID checksum protection discussed earlier. Regardless of how seldom or how often data is accessed, proactive RAID scrubs ensure that media errors occurring over time do not affect the integrity of stored data.

- Proactive self-healing storage activity.
- Detects and corrects disk checksum and parity errors.
- NetApp recommends using the default setting for RAID scrubs to run once a week.
 - RAID scrubs can be scheduled to run on a different day than the default setting.

Background media scans. Similar to RAID scrubs, a background media scan is a process to detect media errors and take corrective actions. RAID scrubs are a weekly process that runs for several hours; background media scans are a resident process that runs in the background only when the controller is not busy. This approach eliminates any performance impact on the storage controller when it is under operational workloads. Background media scans can be thought of as a supplement to RAID scrubs.

- Self-healing storage activity running in background.
- Protects against disk media errors.
- NetApp recommends using the default setting to enable background media scans.

Rapid RAID recovery. When a disk fails, the Data ONTAP RAID implementation automatically begins recreating data contained on the affected disk from parity and storing the information on a hot standby disk. However, if Data ONTAP or a disk generates a predictive failure event, rapid RAID recovery tries to read as much information from the affected disk as possible rather than recreating all of its data from parity. Data read from the disk that passes the RAID checksum process discussed earlier is stored on the hot standby disk. Missing or corrupted data from the affected disk is recreated from parity and stored on the hot standby disk as needed. NetApp implemented rapid RAID recovery because copying good data from the affected disk is much faster with less impact on the CPU than recreating all of its data from parity.

- Self-healing storage activity.
- Faster RAID reconstructions with less performance impact on CPU.
- Greatly reduces exposure to uncorrectable disk media errors during RAID reconstruction.
- NetApp recommends using the default setting to enable disk copy.

ATA disk offline. ATA drives have been known to attempt heroic error recovery when reading or writing blocks in a bad media patch. The duration of such recovery can range from several seconds to several minutes. This design works well for desktop usage, where the inability to read a block can be fatal. However, in server environments, where many data access protocols (for example, CIFS, FC, and NFS) rely on lease timeouts for maintaining session state, long I/O response times can lead to unwanted connection terminations. With NetApp storage using either RAID4 or RAID-DP, Data ONTAP can quickly recreate data from parity in the event that an ATA disk goes into an error recovery state, also known as an ATA spasm. In the event of an ATA spasm, Data ONTAP ceases all I/O operations to the disk, marks the affected disk as offline, and serves reads from parity while queuing writes until the disk recovers.

- Self-healing storage activity.
- Ensures data availability during ATA disk spasms.
- This protection is always on.

Maintenance Center. Maintenance Center (MC) software is the newest addition to the NetApp suite of proactive, self-healing storage resiliency tools. MC provides configurable in-place disk drive diagnostics to determine the health of suspect disk drives. If Data ONTAP disk health monitoring determines that a disk drive has surpassed an error threshold, RapidRAID Recovery is initiated to a hot spare. Afterward, the suspect disk can be placed into MC, where it undergoes a series of diagnostic tests. If diagnostic testing shows disk health to be normal and the error condition to be an anomaly, then the disk is returned to the spares pool. If diagnostics do not show normal disk health, or, by default, if the disk is in MC for a second time, then MC flags the disk as broken and an RMA process is initiated. Consisting of Storage Health Monitor (SHM), NetApp Health Triggers (NHT), and NetApp Drive Self Tests (NDST) software, Maintenance Center promotes drive self-healing and preventive or corrective maintenance while reducing customer impact from the RMA processing of disk drives that are in normal working order, resulting in lower drive maintenance costs.

- Self-healing storage activity.
- Comprehensive proactive drive diagnostics.
- Maintenance Center protection is on by default.
- A minimum of two hot spare disks for each disk type is required before a suspect disk can be placed in MC to undergo diagnostic testing
- NetApp recommends providing storage with enough hot spares to utilize the MC self-healing functionality.

Disk firmware nondisruptive upgrades (NDUs). NDU disk firmware upgrades for either RAID-DP groups or RAID4 groups with SyncMirror can run automatically as a background task, upgrading a single disk at a time with minimal performance impact and resulting in zero downtime. Provided that all aggregates and volumes, including the root volume, are either RAID-DP or RAID4 with SyncMirror, to begin the automatic disk firmware upgrade in Data ONTAP 7.0.1 or later, copy the new disk firmware to the `/etc/disk_fw` directory in the root volume. A process in Data ONTAP checks every few minutes for the presence of a new disk firmware level in this directory. If one is detected, Data ONTAP automatically initiates the nondisruptive background task to upgrade the disk firmware. Conversely, if RAID4 without SyncMirror is used for any aggregate or volume, including the root volume, then the disk firmware upgrade process takes place during the next storage controller boot up and results in an extended boot-up period. The reason is that aggregates remain offline until the disk firmware upgrade process has completed and the reboot process can continue.

- The NDU process improves both storage and system availability without service disruptions.
- This feature is on by default.
- NetApp recommends using RAID-DP for all aggregates and volumes, including the root volume, to utilize NDU disk firmware upgrades.
- It is possible to convert RAID4 aggregates and volumes to RAID-DP for the duration of the NDU disk firmware upgrade and back to RAID4 afterward.
- NetApp recommends using the latest disk firmware available on the NOW site to increase resiliency levels.

Storage shelf firmware nondisruptive upgrades (NDU). NDU shelf firmware upgrades are available for FC disk-based shelves equipped with ESH (electronically switched hub), ESH2, and ESH4 shelf modules. The shelf firmware is upgraded automatically when upgrading Data ONTAP, but otherwise requires manual steps to initiate the process.

Note: Shelves that contain SATA disks drives do not offer NDU shelf firmware upgrades and require a service outage of approximately 45 to 70 seconds per shelf module.

- Improves both storage and system availability.
- This feature is always on.
- NetApp recommends using the latest shelf firmware available on the NOW site to increase resiliency levels.

Topology sort. This feature automatically spreads disks over loops and shelves when creating or expanding aggregates and traditional volumes. The topology sort algorithm in Data ONTAP is optimized to select disks to improve performance and fault tolerance. Performance improvements can be realized by spreading storage workload over numerous loops. Similarly, fault tolerance is improved by spreading disks over numerous shelves and loops to reduce the likelihood of storage failures resulting in the possibility of RAID group reconstructions being initiated.

- Improves both performance and data availability.
- This feature is always on but can be overridden by manually selecting disks when creating or expanding volumes or aggregates.
- NetApp recommends allowing topology sort to automatically select the disk when provisioning storage.

3.2. Guidelines That Increase Storage Resiliency

RAID-DP. RAID-DP, the NetApp high-performance implementation of RAID 6, is double-parity RAID that adds a second parity stripe to dramatically increase data availability. With RAID-DP, aggregates and volumes can withstand up to two failed disks in a RAID group, or the more common event of one failed disk followed by an uncorrectable bit read error from the disk drive. RAID4 (single-parity stripe) cannot protect against either of these scenarios. If a single disk in a RAID4 group fails, Data ONTAP is able to detect and regenerate data from parity. However, as with any single-parity RAID type, during the period in which RAID4 is reconstructing data, it is vulnerable to unrecoverable media errors. Although RAID4 can still detect the bit error, it cannot recreate the data. RAID-DP software protects the system against these types of failure for improved data availability.

From a value proposition standpoint, RAID-DP offers RAID1 levels of data reliability, at RAID4 prices.

- Protects against double disk failures or uncorrectable bit read errors from the disk while in reconstruction mode.
- No impact to required capacity because RAID-DP groups can be double the size of RAID4 groups.
- RAID-DP groups are used by default when creating aggregates for both FC and SATA disks.
- RAID-DP enables NDU disk firmware upgrades that result in zero downtime.
 - All aggregates and volumes, including the root volume, must be RAID-DP to take advantage of NDU disk firmware upgrades.
- NetApp recommends always using RAID-DP because of its higher reliability with negligible performance costs.
- NetApp recommends using the default RAID group sizes when using RAID-DP.
- For even higher resiliency and data reliability, consider using smaller RAID-DP group sizes.

Snapshot reserves. Snapshot copies are backups of how a volume looks at a particular point in time. The unique NetApp approach allows Snapshot copies to work almost instantaneously with very little impact on storage capacity. Later, a storage administrator or even an end user can recover data from the desired Snapshot copy. By default, Snapshot reserve space is set to 20% of the total volume capacity and also by default keeps a minimum number of weekly, daily, and hourly Snapshot copies online and available. To meet business needs, a customer might increase the number of Snapshot copies kept online, the frequency of Snapshot copies, or both.

- Snapshot copies provide a natively integrated and easy-to-use data protection utility that helps reduce impact on storage administrators by enabling end users to recover their own data.
- Protects against inadvertent file modification or deletion by making point-in-time backups available.
- NetApp recommends leaving Snapshot reserve and frequency at the default values or increasing the frequency and scheduling to meet business requirements.
- Rapidly changing data increases the sizes of Snapshot copies, resulting in more impact on reserve space.
- More frequent Snapshot copies also have more impact on reserve space.
- For either of these conditions, be sure to keep ample spare capacity in the Snapshot reserve.

Local backups. NetApp provides two bundled backup solutions that do not require add-on licenses or third-party applications. The `dump` command and `ndmpcopy` are available to replicate data to tape drives or to other storage, respectively. Both commands are easy to use and include incremental functionality that backs up only files that have changed, reducing impact on both storage and tape backup libraries.

- Both `dump` and `ndmpcopy` are available in Data ONTAP.
- Tape backups generated by `dump` can be stored off site, while `ndmpcopy` can replicate to NetApp storage across LANs or WANs without the need for tape library overhead.
- Both backup utilities increase data protection.
- NetApp recommends using data backup as a key foundation piece of enterprise data protection.

Root volumes. The root volume can exist either as the traditional standalone two- or three-disk volume or as a FlexVol® volume that is part of a larger hosting aggregate. Both approaches are supported, and each approach provides a slight benefit over the other. Smaller standalone root volumes offer fault isolation from general application storage, whereas flexible volumes have less overall storage utilization impact because two disks are not being dedicated to the root volume and because of its small storage requirements. However, if a FlexVol volume is used for the root volume, file system consistency checks or recovery operations can take longer to finish than with the standard two- or three-disk traditional root volume. The reason is that with FlexVol recoveries, these types of commands work at the aggregate level, meaning that all containing disks are targeted during the operation. One way to mitigate this effect is to use a smaller aggregate with only a few disks to house the FlexVol volume that contains the root volume.

In practice, having the root volume on a FlexVol volume makes a bigger difference with smaller capacity storage systems compared to very large ones, where two dedicated disks for the root volume have little impact. Regardless of whether traditional or flexible volumes are used for the root volume, the same general guideline to not store application data in the root volume still applies.

- Root volumes can use either flexible or traditional volumes.
 - If using a FlexVol volume for root, consider allocating 100GB capacity for activities such as storage of system files.
 - If the root volume is a FlexVol volume, more concise guidelines on recommended minimum size by platform are available in Section 4, [Understanding the Root Volume](#), of the Data ONTAP System Administration Guide on
- For higher resiliency, use a separate two-disk root volume, and for maximum resiliency implement SyncMirror on the traditional root volume or its containing aggregate.
- For small storage systems where cost concerns outweigh resiliency, a flexible root volume on a regular aggregate may be more appropriate.
 - If available capacity is an issue, migrate the root volume to a FlexVol volume to free up a two-disk traditional root volume.
- NetApp recommends using RAID-DP for the root volume, whether it is standalone or a FlexVol volume, in order to take advantage of NDU disk firmware upgrade capabilities.
- NetApp recommends that customers not put user data in the root volume.
 - SnapRestore® recoveries that are typical in environments with user data are not recommended on the root volume.

Hot-standby and spare disks. With NetApp self-healing RAID software, disk failures automatically trigger parity reconstructions of affected data onto a hot standby disk. The one caveat is that a hot spare disk must be available to Data ONTAP for this self-healing process to begin. Therefore, at a minimum, resiliency planning should include keeping at least one hot spare disk for each type of disk drive present in the storage system. For storage controllers with large capacities, or to take advantage of Maintenance Center, NetApp recommends having at least two hot spare disks available for each disk drive type. Maintaining on-site spare disks can also be incorporated into overall storage resiliency planning. This approach allows rapid replacement of failed disks that are in turn refreshed through regular NetApp RMA procedures.

- Providing adequate hot spare disks is a critical component of storage capacity planning.
- Adequate hot spare disks increase storage resiliency and data availability.
- For active-active configurations, hot spares must be owned by the right controller, and with SyncMirror, hot spares must be in the right pool.
- NetApp recommends having at least two hot standby disks for each disk drive type in the system.
 - Two hot spares per disk type are required to utilize Maintenance Center diagnostics.

Cabling best practices. General cabling best practices are usually enforced as part of technology infrastructure deployments. Avoiding extremely long data cable runs, not bending data cables at right angles, and avoiding close proximity to EMF sources such as fluorescent lighting are examples of cabling best practices. For highest resiliency levels, additional cabling best practices should be employed when deploying and maintaining NetApp storage. Verify that cables are securely connected with a gentle tug during initial configuration, and when possible route dual-path cabling to different HBAs on standalone storage controllers.

- General industry standard cabling best practices should be included as part of NetApp storage deployments.
- Cabling best practices help increase overall storage resiliency and data availability.
- NetApp recommends that in addition to industry cabling standards, ensure that cables and SFPs are securely connected and that dual-path cables to shelves connect to different HBAs in the storage controller.

Maximize the number of loops to increase bandwidth for reconstructions. If a disk failure occurs, loop traffic generated by the reconstruction operation competes with existing production workload and can result in slower performance for each. One way to minimize the impact of disk reconstruction on loop performance is to employ more loops than would normally be required to support the amount of storage capacity. With RAID group disks spread across additional loops, any reconstruction activity is also spread across the extra loops, resulting in faster reconstruction times and shorter performance impact on application-generated workloads (assuming that the storage processor CPU/memory does not initially bottleneck the reconstructions). A second way to limit the performance impact of reconstructions is to change the Data ONTAP option setting `raid.reconstruc.perf_impact` to low. This approach reduces performance impact on primary applications, but the reconstruction process takes longer.

- Maximizing the number of loops allows faster reconstructions after a disk failure.
- Shorter reconstruction periods improve storage resiliency and data reliability by reducing the chance of a media error during reconstruction.
- NetApp recommends maximizing the number of loops for highest levels of storage resiliency.
 - When maximizing the number of loops, spread shelves horizontally across the loops and HBAs.

3.3. OPTIONS THAT INCREASE STORAGE RESILIENCY

SnapMirror and SnapVault. SnapMirror and SnapVault are data replication products that improve data protection by automatically maintaining duplicate copies of data either locally or remotely. After the initial baseline data transfer, SnapMirror and SnapVault replicate only changed blocks from the primary storage controller to minimize performance impact on storage and bandwidth impact on the network. Because only changed blocks are replicated and bandwidth impact is limited to the rate of data change on the primary storage controller, both SnapMirror and SnapVault are excellent choices to replicate data over generally slow WANs to increase data protection options. Each replication option is highly configurable to meet business requirements. SnapMirror can be configured to replicate data in asynchronous mode, semisynchronous mode, and full synchronous mode. SnapVault replicates NetApp Snapshot copies, and the frequency of the Snapshot replication process can be configured during initial SnapVault configuration or changed as needed afterward.

- SnapMirror and SnapVault provide automated native data replication functionality.
- Replicating data to local secondary storage allows for faster backup and recovery times versus tape backups.
- Data replication improves data protection and is a critical component of disaster recovery deployments.
- NetApp recommends replicating data with either option to increase data protection.

Local SyncMirror. Local SyncMirror provides synchronous mirroring between two different volumes or aggregates on the same storage controller so that a duplicate copy of data exists, resulting in higher storage resiliency and data availability. Software disk ownership (SANOWN) or system configuration ownership is the mechanism to assign disks in the local SyncMirror configuration. (**Note:** FAS270, FAS3040, FAS3070, and FAS6000 have SANOWN turned on by default.) Although SnapMirror in synchronous mode does provide a similar capability, it is generally used for replicating data between geographic locations to improve disaster recovery options in the event of a data center outage. At the local level, SyncMirror provides storage resiliency capabilities that SnapMirror or even active-active configurations by themselves do not. The additional resiliency features that SyncMirror offers over SnapMirror in synchronous mode are protection against system downtime due to shelf failure, triple disk failure for RAID-DP groups, and Fibre Channel loop failure.

- When used in active-active configurations, SyncMirror provides the highest resiliency levels in NetApp storage for a local data center.
- Highest levels of storage resiliency ensure continuous data availability within a data center.
- NetApp recommends SyncMirror and active-active configurations for very high levels of storage resiliency.

FC disk shelf modules. In order of introduction, the loop resiliency circuit (LRC) shelf module appears first in NetApp storage, followed by the electronically switched hub (ESH) modules, ESH, ESH2, and ESH4. All four shelf modules provide FC-AL connectivity to FC disk-based storage, but they have different storage resiliency levels. The newer ESH, ESH2, and ESH4 modules introduce the ability for NetApp storage controllers to detect and isolate disk drives that can disrupt the FC-AL operations. This functionality, which does not exist with older LRC technology, improves resiliency by offering protection against a rogue disk bringing down an entire loop. For best protection from all rogue disk activities, use the ESH2 module or, better still, the ESH4 module with the latest firmware release. Similar protection with the two older shelf modules can be achieved by using the latest disk and module firmware. To obtain the latest firmware levels, use the [Disk Drive and Firmware Matrix](#) on the NOW site.

The new ESH4 module and DS14mk4 shelf combination offers the best storage resiliency levels, because both include the latest firmware and technology enhancements while providing the performance benefits of 4Gb/sec speed on each storage loop. The ESH4 has the same intelligent fault isolation capabilities as its ESH and ESH2 predecessors, but when used in combination with the DS14mk4 shelf it provides new features such as enhanced diagnostics and drive power cycling. Individual disk drive power cycling, supported starting with Data ONTAP 7.2.1, is as a new tool in the NetApp storage error recovery process. In Data ONTAP 7.2.1 and 7.2.2, this feature is off by default but can be enabled for customers via the `options disk.powercycle.enable` command.

All NetApp Fibre Channel shelves include fully redundant ESH4, ESH2, ESH, and LRC shelf modules. In addition, each shelf contains dual power supply units. This fully redundant configuration, along with the best practice of multipath to different storage processor HBAs, ensures a highly tolerant subsystem with minimal risk of SPOF.

- ESH, ESH2, and ESH4 modules provide the ability to detect and isolate disk drives that might disrupt FC-AL operations (for example, FC-AL protocol violations).
- ESH modules help increase storage resiliency compared to previous LRC modules, with ESH4 offering the latest technology for maximum performance and resiliency.
- To increase storage resiliency, keep all shelves, shelf modules, and disk firmware levels current.
- NetApp recommends using ESH4 modules and DS14mk4 shelves for best storage resiliency.

Active-active configurations. Active-active configurations, also known as clustered failover, eliminate the storage controller as a single point of failure. In this configuration, each storage controller has its own dedicated pool of disk drives and handles all I/O operations during normal operation. Each clustered storage controller pair is connected to its partner's disk drives as well and maintains a heartbeat status of its partner. If a heartbeat check reveals that a paired partner is down, the remaining controller initiates a takeover operation of the failed storage controller's disks and handles all I/O requests until the down controller can be brought back online.

- Active-active configurations prevent a storage controller from becoming a single point of failure.
- Active-active configurations increase storage resiliency levels.
- NetApp recommends using active-active configurations for an environment's highest data availability.

Multipath HA storage configuration. Multipath HA storage configuration further enhances the resiliency of active-active controller configurations. Although cluster failover software provides high availability by providing fault tolerance in the event of controller failure, storage triggered events often result in unneeded failovers or prevent successful takeovers. Multipath HA storage enhances storage resiliency by reducing unnecessary takeover by a partner node due to a storage fault, improving overall system availability and promoting higher performance consistency.

- Multipath HA provides added protection against various storage faults, including:
 - HBA or port failure
 - Controller-to-shelf cable failure
 - Shelf module failure
 - Dual intershell cable failure
 - Secondary path failure
- Multipath HA helps provide consistent performance in active-active configurations by providing larger aggregate storage loop bandwidth.
- NetApp recommends Multipath HA for all active-active controllers to increase storage resiliency levels.

MultiStore®. MultiStore is a NetApp option that allows a storage controller to be divided into many “virtualized” storage systems. With MultiStore, each virtual storage system is completely segregated from every other one, and this functionality provides extremely granular options for storage provisioning. The MultiStore functionality can also provide a simple and economical disaster recovery solution. MultiStore supports dynamic assignment and reassignment of storage and network resources. Therefore, if a virtual storage system is replicated to a destination virtual storage system on a different storage controller with SnapMirror, MultiStore can be failed over from the original virtual storage system to the replicated one without any changes or impact on clients.

- Among other storage virtualization capabilities, MultiStore provides simple and economical automatic failover disaster recovery functionality.
- Disaster recovery solutions improve storage resiliency by providing higher data availability.
- NetApp recommends that disaster recovery solutions be implemented for data that is of high value to the enterprise.

MetroCluster with SyncMirror. Although both Stretch MetroCluster and storage controllers in active-active configurations are supported up to 500 meters apart with 2Gb/sec connectivity, or 270 meters with 4Gb/sec speeds, this generally means that all controllers are in the same data center. Fabric MetroCluster allows the active-active configuration to be spread across data centers up to 100 kilometers apart. In the event of an outage at one data center, the second data center can assume all affected storage operations that were lost with the original data center. SyncMirror is required as part of MetroCluster to ensure that an identical copy of the data exists in the second data center in case the original data center is lost.

- Fabric MetroCluster along with SyncMirror extends active-active clustering across data centers up to 100 kilometers apart.
- Fabric MetroCluster and SyncMirror provide the highest level of storage resiliency across a local region.
- Highest levels of regional storage resiliency ensure continuous data availability in a particular geography.
- NetApp recommends that MetroCluster be used with SyncMirror for the highest level of storage resiliency.

SnapValidator®. For Oracle® deployments, SnapValidator can be used to provide an additional layer of integrity checking between the application and NetApp storage. SnapValidator allows Oracle to create checksums on data transmitted to NetApp storage for writes to disk and to include the checksum as part of the transmission. When Data ONTAP receives the data, it generates a new checksum and compares it to the one generated by Oracle. If the two match, the Oracle write is acknowledged and the data is written to disk. As part of writing data to disk, the inherent features of Data ONTAP, such as RAID checksums, are engaged and continue to guarantee data integrity going forward. If the checksums do not match, the write is aborted, no data corruption occurs, and an alert is generated so that corrective action can be taken.

- SnapValidator provides an additional layer of data integrity checking for Oracle deployments.
- SnapValidator is tightly integrated with the Oracle Database architecture and complies with the Oracle HARD initiative.
- Combined with other NetApp storage resiliency capabilities, SnapValidator provides the highest levels of storage resiliency for Oracle deployments.
- NetApp recommends using the appropriate resiliency features and SnapValidator for all Oracle deployments.

Disk-to-disk-to-tape backup. Staging primary storage backups to secondary disk-based near-line storage generally allows shorter backup windows than going directly to tape. Disk-based backups are even more beneficial if data recovery is required. Data that resides on disks is online and readily available to facilitate rapid restoration. Data backed up to tape is offline, in contrast, and probably not even on site at the data center. Both of these circumstances increase the time and logistics involved in data recovery. However, even with the benefits of disk-based secondary storage solutions, tape archiving can still be a critical component of an enterprise's data protection strategy. Allowing the secondary storage to keep data available online for rapid recovery and to handle the overhead of writing to tape media improves overall storage resiliency while minimizing performance impact on primary storage.

- Disk-to-disk-to-tape backup provides an enterprise-class multistaged and multitiered storage solution for data protection, disaster recovery, and data archiving.
- Combined with other NetApp resiliency features and options, disk-to-disk-to-tape backup solutions provide the highest levels of storage resiliency.
- NetApp recommends using disk-to-disk-to-tape backups as part of a highly resilient storage environment for enterprise data.

4. MAPPING NETAPP STORAGE RESILIENCY FEATURES AND OPTIONS TO STORAGE TIERS

Having established the various features and options for deploying highly resilient NetApp storage environments, the focus shifts to mapping them to the different tiers, or service levels, that are typically a part of enterprise storage deployments. Probably the easiest way to determine the tier for stored data is to consider its value to the enterprise. Data of the highest value dictates a very high tier of service levels, which in turn requires highly resilient storage deployments. As data moves down the value curve, the appropriate resiliency options to meet the corresponding storage tier can be selected for the deployment. However, even for data with the lowest service levels, such as archiving or remote office, there are inherent features and cost-effective options to increase storage resiliency capabilities.

4.1. Factors Considered in Establishing Storage Tiers

Although the value of stored data drives its service level, in terms of storage, two primary considerations establish the correct storage tier level. The first is how long it takes to resume operations in the event of an outage, and the other is what amount of data loss is acceptable. Given that reducing the impact of either consideration adds to the overall storage cost, establishing the business value of each data class will probably involve some tradeoffs with availability and an acceptable amount of lost data.

Mission-critical data in enterprise data centers typically demands no unplanned outages and zero data loss, and falls into tier 1 classes of storage. Departmental storage requirements typically have lower service levels and can withstand longer unplanned downtime and minutes of lost data. Service levels for small remote offices without shared storage can be such that as much as a day's worth of data is lost and outages are resolved in hours instead of minutes.

4.2. Compliance, Reference, and Archival Environments (Tier 4)

Compliance, reference, and archival environments typically target near-line storage on lower cost media such as SATA disks. In general, these environments can tolerate longer downtime, because the data is infrequently accessed for reads and the archival application can queue up hours of writes if the storage is unavailable. Acceptable loss of data on storage can vary from hours for reference data to zero data loss for compliance data.

- Resiliency targets:
 - **Recovery point objective (RPO):** From minutes up to one day of lost data acceptable; no acceptable lost data for compliance archiving.
 - **Recovery time objective (RTO):** Recovery can take minutes to hours.

Table 1 shows inherent storage resiliency features for tier 4, along with NetApp recommendations. Table 2 shows guidelines for tier 4 storage resiliency. Table 3 shows options for tier 4 resiliency.

Table 1) Features for tier 4 storage resiliency.

Inherent Storage Resiliency Feature	NetApp Recommendation
RAID Checksums	Always on
Background RAID Scrubs	Leave in default on state
Background Media Scans	Leave in default on state
Rapid RAID Recovery	Always on
Lost Writes Protection	Leave in default on state
ATA Disk Offline	Always on
Disk Firmware Nondisruptive Upgrades	Always on
Disk Topology Sort	Allow Data ONTAP to automatically select disks when creating or adding capacity to new volumes and aggregates

Table 2) Guidelines for tier 4 storage resiliency.

Storage Resiliency Guideline	NetApp Recommendation
Data ONTAP and Storage Firmware	Always use the latest general deployment (GD) release and the latest storage firmware available on the NOW site .
Maintain Snapshot Reserve	Leave the default 20% Snapshot reserve and increase if necessary. If Snapshot reserve shows 100% utilization, either increase the reserve or decrease Snapshot frequency.
Use RAID-DP	Use RAID-DP for aggregates and traditional volumes.
Root Volume	Do not put user data in the root volume. For higher resiliency levels, use a separate two-disk traditional root volume.
Local Backups	Perform local backups either to secondary storage with SnapMirror or SnapVault or to tape drives with applications such as Symantec® NetBackup™.
Adequate Disk Spares	Factor in adequate hot standby disks when provisioning storage. Maintain at least one spare for each type of disk drive in the storage system.
Cabling Best Practices	Follow general and NetApp cabling best practice guidelines.

Table 3) Options for tier 4 storage resiliency.

Storage Resiliency Option	NetApp Recommendation
SnapMirror or SnapVault Data Replication	NetApp recommends replicating data either locally or to a remote storage controller to increase data protection and disaster recovery capabilities.

4.3. Corporate Data, Departmental Workgroups, and Distributed Remote Environments (Tier 3)

As with the previous near-line storage deployments, corporate data (which includes workgroups and distributed enterprise) environments typically can tolerate longer unplanned outages and some data loss. With these lower service levels, these types of deployments can leverage the reduced cost associated with lower tiers of storage. Two storage resiliency options available for primary storage generally used for this type of deployment are an active-active configuration and the ESH, ESH2, or ESH4 shelf module. Otherwise, all inherent features and guidelines covered in the previous section should also be used for this storage tier.

- Resiliency targets:
 - **Recovery point objective (RPO):** Data loss of minutes to hours for local workgroups, possibly up to one day for remote offices.
 - **Recovery time objective (RTO):** Recovery can take minutes to hours.

Table 4 shows tier 3 inherent storage resiliency options, along with NetApp recommendations.

Table 4) Options for tier 3 storage resiliency.

Storage Resiliency Option	NetApp Recommendation
Basic Foundation	All features and guidelines covered in section 4.2 should also be used for tier 3 storage deployments.
Active-Active Configuration	(Also known as clustered failover.) Use this option to prevent the storage controller from becoming a single point of failure (SPOF).
Multipath HA Storage Configuration	Use this option in conjunction with active-active configuration to improve overall system availability and promote higher performance consistency.
ESH, ESH2, or ESH4 Shelf Modules	Use ESH, ESH2, or ESH4 shelf modules for tier 3 deployments on FC disks.

4.4. Business-Critical Environments (Tier 2)

A corporation's business-critical environment typically processes high-value data, and loss or unavailability of data for more than a few minutes is very disruptive. This higher value data class requires storage resiliency levels provided by tier 2 NetApp storage. This section builds on the storage foundations established in the earlier tiers and introduces guidelines and available options that enable the higher resiliency levels associated with this tier of storage.

- Resiliency targets:
 - **Recovery point objective (RPO):** From zero lost data up to minutes of lost data.
 - **Recovery time objective (RTO):** Recovery can take minutes.

Table 5 shows tier 2 inherent storage resiliency options, along with NetApp recommendations.

Table 5) Options for tier 2 storage resiliency.

Storage Resiliency Option	NetApp Recommendation
Basic Foundation	Features, guidelines, and options covered in sections 4.2 and 4.3 should also be used for tier 2 storage deployments.
Local SyncMirror	To protect business-critical data from any local storage-related issues, use Local SyncMirror in conjunction with active-active configuration clusters.
Remote SnapMirror and SnapVault or MetroCluster	Replicate high-value data to remote locations to improve data protection and disaster recovery. For fast, long-distance replication that doesn't require full synchronization, SnapMirror or SnapVault is ideal. For fast, synchronous replication within 100 km, Fabric MetroCluster is ideal.
MultiStore Disaster Recovery (SnapMirror)	For a simple and economical failover solution, consider using MultiStore as a disaster recovery option for SnapMirror.
ESH2 or ESH4 Shelf Modules	Use dual ESH2 or ESH4 storage modules for tier 2 deployments on FC disks.
Maximize Number of Loops	Increase the number of loops for faster reconstruction times and reduced risks during reconstruction.
Smaller RAID Groups	Use smaller RAID groups for faster reconstructions and reduced risks during reconstruction.

4.5. Mission-Critical Environments (Tier 1)

Mission-critical data has the highest corporate value and by necessity requires the highest tier of storage resiliency. For data that requires this tier of storage, no amount of lost data is acceptable, and only occasional minutes of unplanned downtime can be withstood in order to meet service levels. NetApp tier 1 storage resiliency uses the same inherent features, guidelines, and options covered in earlier tiers as basic building blocks. There are also additional NetApp options and guidelines that are core components to highly resilient storage environments.

- Resiliency targets:
 - **Recovery point objective (RPO):** Zero data loss.
 - **Recovery time objective (RTO):** Zero downtime to recovery in minutes.

Table 6 shows tier 1 inherent storage resiliency options, along with NetApp recommendations.

Table 6) Options for tier 1 storage resiliency.

Storage Resiliency Option	NetApp Recommendation
Basic Foundation	F, guidelines, and options covered in sections 4.2, 4.3, and 4.4 should also be used for tier 1 storage deployments.
MetroCluster plus SnapMirror	Use MetroCluster (with SyncMirror) for metropolitan-sized area disaster protection. In addition, there should be a third remote site that receives mirrored data via SnapMirror. This multiple-hop model ensures true DR protection against local and regional site failures due to disaster or other broad continuity failures.
SnapMirror (Multisites)	An alternative to MetroCluster plus SnapMirror that still offers disaster protection across larger geographies is SnapMirror across three or more sites. If data does not require full synchronization, SnapMirror (Async) can be used to mirror from primary to secondary to tertiary sites. For regional synchronization, MetroCluster is recommended to provide fast, fully synchronous two-site mirroring with fast site failover.
SnapValidator	For Oracle deployments, use SnapValidator to ensure an additional layer of data integrity checking.
Disk-to-Disk-to-Tape Backups	To reduce impact on primary storage and increase service levels of data protection, NetApp recommends implementation of staged backups to secondary storage. Once on secondary storage, data is backed up to tape and sent off site, possibly in encrypted format, for robust data protection.

5. APPLICATION MAPPING TO TIERED STORAGE

After establishing the fundamental building blocks of NetApp storage resiliency and implementing the appropriate resiliency level based on storage tiers, the next topic of focus is mapping various business applications to the required storage tier. In general, corporations affix varying levels of value on data, depending on how much unplanned downtime or potentially lost data is allowable before there is significant impact on the business. The guidelines that drive the value of data to an enterprise and the underlying storage resiliency requirements are generally consistent across enterprises and provide the basis for the application mapping recommendations in this section.

5.1. High-End Storage Applications

Business-critical applications that require the highest levels of storage resiliency tend to involve the processing and analysis of databases. These types of database applications are typically the lifeblood of enterprise transactions and business planning activities, with large user communities both inside and outside the organization. The value of data in this environment warrants the highest levels of infrastructure fault tolerance to ensure continuous data availability, absolute data integrity, and overall quality of service. NetApp recommends that tier 1 storage resiliency solutions be used when deploying the following types of mission-critical applications:

- Online transaction processing
- Batch transaction processing
- Enterprise resource planning
- Data warehouse analysis
- Customer relationship management

5.2. Upper Mainstream Storage Applications

Mainstream applications typically process data classes with values that approach tier 1 storage resiliency requirements to data class values that may tolerate short outages or some number of minutes of lost data. For these types of applications, the level of storage resiliency can be tuned to meet the service levels based on data value versus storage infrastructure cost. NetApp recommends that the features, guidelines, and options covered for tier 2 storage resiliency be used for the following types of applications:

- Data warehousing
- Application and software development
- E-mail
- Web serving
- Networking

5.3. Lower Mainstream Storage Applications

Another class of mainstream applications that process slightly lower value data might be tolerant of minutes to hours of unplanned outages or to longer point-in-time data recovery. These lower end mainstream applications, like the higher end ones just covered, can have the correct levels of storage resiliency implemented to meet their service levels versus storage infrastructure costs. NetApp recommends that the features, guidelines, and options covered for tier 3 storage resiliency be used for the following types of applications:

- Scientific and engineering computation
- Workgroup collaboration
- File and print services

5.4. Near-Line Storage Applications

In general, near-line applications are more tolerant of unplanned data outages because the data is infrequently accessed after archiving. The amount of acceptable data loss in near-line deployments varies depending on the class of data. For lost backup data that still resides on primary storage, the business impact can be very small. Conversely, the loss of compliance data could affect business severely, including the possibility of significant fines. In either case, the potential amount of lost data from a catastrophe can be tuned to meet business requirements by implementing a data recovery point via SnapMirror or SnapVault replication solutions from NetApp. The following types of applications are good candidates for NearStore® deployments:

- Backup and recovery
- Compliance
- Archiving
- Reference data

6. CONCLUSION

Enterprises typically have data classes that vary in degrees of value, which in turn drive availability and protection requirements in the data storage system. NetApp storage includes a rich set of inherent features and available options that allow businesses to implement storage resiliency tiers that enable enterprises to meet availability and protection requirements for any class of data. Understanding NetApp features, available options, and recommendations helps customers during the storage design cycle and in selecting the appropriate storage resiliency level. For additional information about topics covered in this paper, or to determine what storage solution implementations can meet or exceed your business data requirements, please contact your NetApp sales team.

APPENDIX A: ADDITIONAL RESOURCES FOR STORAGE RESILIENCY INFORMATION

For further information on features and options covered in this paper, use the following links. **Note:** Some of these links require access to the NOW site.

Continuous Media Scrubs

[Continuous Media Scans Knowledgebase Article](#)

Rapid RAID Recovery

[Rapid Recovery with NetApp Storage Solutions Data Protection Strategies for NetApp Enterprise Strategy Group Report](#)

Disk Firmware NDU

[Data ONTAP Disk Firmware Upgrade Guide](#)

Shelf Firmware Upgrade Guide Including NDU

[Data ONTAP Shelf Firmware Upgrade Guide](#)

Latest Data ONTAP GD Release and Firmware

[Data ONTAP Upgrade Guide](#)

Snapshot Reserve

[Datasheet File System Design for an NFS File Server Storage Controller](#)

RAID-DP

[Double Parity RAID for Enhanced Data Protection Best Practices for Volume and RAID Group Configuration on NearStore R200](#)

Root Volume and Aggregate Configuration

[Data ONTAP Storage Management Guide Data ONTAP System Administration Guide Best Practices for Volume and RAID Group Configuration on NearStore R200 How to Move or Rename the Root Volume on a Storage System](#)

Local Backups and SnapVault

[SnapVault Best Practices Guide Data Protection Technical Reports Data Protection Solutions Overview SnapVault Deployment and Configuration Enabling Rapid Recovery with SnapVault](#)

Use Default RAID-DP Group Sizes

[Double Parity RAID for Enhanced Data Protection Best Practices for Volume and RAID Group Configuration on NearStore R200](#)

Active-Active Configurations (Clustered Failover)

[Active/Active Controller Configuration Overview and Best Practices Guidelines Clustered Failover High-Availability Solution](#)

ESH Storage Controllers

[Electronically Switched Architecture for NetApp Enterprise Storage Configurations](#)

Cabling Best Practices

[Fibre Channel Cabling Design and Management](#)

Local SyncMirror

[SyncMirror Software Rapid Recovery with NetApp Storage Solutions SyncMirror Management Adding a Plex](#)

Remote SnapMirror

[Data Protection and Recovery for NAS over TCP/IP Networks](#)

Smaller RAID-DP Group Sizes

[Double Parity RAID for Enhanced Data Protection Best Practices for Volume and RAID Group Configuration on NearStore R200](#)

MultiStore Disaster Recovery

[Storage Virtualization and DR Using MultiStore
NetApp MultiStore and SnapMover®
MultiStore Management Guide
Data Protection Technical Reports
Data Protection for Disaster Recovery
Distributed Storage for a Scalable Grid Architecture](#)

Maximum Number of Loops

[Fibre Channel Cabling Design and Management](#)

SnapMirror Synchronous

[Using Synchronous SnapMirror for Disaster Protection with Block-Access Protocols
Synchronous SnapMirror Design and Implementation Guide
Data ONTAP Data Protection Guide: Synchronous SnapMirror](#)

MetroCluster with SyncMirror

[MetroCluster Design and Implementation Guide
Storage Consolidation for Database Environments
Enterprise Strategy Group Management Brief](#)

Remote Disk-to-Disk-to-Tape Backup

[Data Protection for Backup and Recovery
Data Protection for Disaster Recovery
Open Storage Networking
Best Practices Guide for Data Protection with Controllers Running FCP
Rapid Recovery with NetApp Storage Solutions
Data Protection Tape Backup and Recovery Guide
Gigabit Network Shared Tape Backup Solution](#)

SnapValidator

[Ensuring Oracle Data Integrity With SnapValidator
SnapValidator Software
Using SnapValidator
SnapValidator Press Release](#)

APPENDIX B: REVISION HISTORY

Date	Name	Description
7/10/2007	Chris Lueth	Update
7/20/2006	Kiran Sreenivasamurthy	Update
7/7/2006	Shawn Kung	Update
3/1/2006	Chris Lueth	Update
12/10/2005	Chris Lueth	Creation



www.netapp.com

© 2007 Network Appliance, Inc. All rights reserved. Specifications subject to change without notice. NetApp, the Network Appliance logo, Data ONTAP, FlexVol, MultiStore, NearStore, SnapMirror, SnapMover, SnapRestore, SnapValidator, SnapVault, and SyncMirror are registered trademarks and Network Appliance, FlexShare, NOW, RAID-DP, and Snapshot are trademarks of Network Appliance, Inc. in the U.S. and other countries. Oracle is a registered trademark of Oracle Corporation. Symantec is a registered trademark and NetBackup is a trademark of Symantec Corporation or its affiliates in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.