



Technical Report

Iron Mountain NearPoint and NetApp Joint Solution

An Introduction to E-Mail Archiving, Disaster Recovery, and Cost-Effective Storage Using Iron Mountain NearPoint

Nathan Walker, NetApp Technical Marketing

Natalie Longhini, Iron Mountain Global Alliance Manager

June 2010 | TR-3825

MICROSOFT EXCHANGE JOURNALING USING NETAPP AND IRON MOUNTAIN

Iron Mountain NearPoint™ and NetApp® SnapManager® for Exchange enable enterprises to archive, manage, and retain long-term e-mail in a cost-effective manner. The e-mail journal archive solution in this paper describes an architecture that you can use to help your company manage Exchange e-mail for regulatory and litigation purposes. When used with NetApp storage solutions, NearPoint quickly captures, indexes, and archives messages to enable search and retrieval.

TABLE OF CONTENTS

1	OVERVIEW	3
2	JOINT SOLUTION OVERVIEW	3
3	IRON MOUNTAIN ARCHIVING SOLUTION	4
4	NETAPP STORAGE SOLUTIONS	4
5	NETAPP FLEXVOL	5
6	NETAPP SNAPLOCK.....	5
7	NETAPP DEDUPLICATION FOR FAS	5
8	SNAPMANAGER FOR MICROSOFT SQL SERVER	5
9	SNAPMANAGER FOR MICROSOFT EXCHANGE	6
10	SNAPMANAGER FOR EXCHANGE INTEGRATION WITH NEARPOINT	6
11	DISASTER RECOVERY INTEGRATION	8
12	KEY BENEFITS OF THE NETAPP AND IRON MOUNTAIN JOINT SOLUTIONS	9
13	CONCLUSION	9
14	REFERENCES	10

1 OVERVIEW

All organizations today deal with the challenge of managing growing volumes of e-mail. Between the four year period from 2008 to 2012 the average number of messages a typical corporate user receives in a day is expected to grow from 126 messages to 233 messages, an increase of 85%.

Correspondingly, the average size of these messages is also increasing. Many of these are either business records or valuable corporate knowledge that must be retained for a certain period of time. Keeping a large volume of e-mail leads to higher storage costs and headaches for IT staff in the form of longer backup windows and complex disaster recovery.

Osterman Research points out that, on average, message stores are growing by 35% annually. Also consider the e-discovery concerns driven by the 2006 amendments to the Federal Rules of Civil Procedure (FRCP). These rules cover the discovery of electronically stored information. FRCP rule 26 elevates the need for companies to begin implementing archiving solutions. This new legislation specifies that e-mail is discoverable and calls for timeliness of discovery. As archiving grows into a mission-critical initiative for organizations of all sizes, the need increases to leverage a flexible, cost-effective storage infrastructure to back up archiving systems and to store an ever-increasing amount of information.

The partnership between NetApp and Iron Mountain brings together best-in-class archival and information management solutions for user-generated content like e-mail, instant messaging, and free-range files. Customers enjoy high performance, simple manageability, and low operating costs by combining leading technology from NetApp and Iron Mountain in a single integrated solution. Companies using Iron Mountain NearPoint on NetApp storage solutions can expect to reduce their e-mail storage requirements by 25% or more, while satisfying regulatory and legal requirements. Now you can streamline your operations and minimize your risk with solutions from Iron Mountain and NetApp.

2 JOINT SOLUTION OVERVIEW

Data archiving is a concept that extends storage into different classes of service to provide a foundation for business units to assign value to their data. Iron Mountain NearPoint captures the richness of Exchange data and fits perfectly with the NetApp Unified Storage Architecture for block storage and NAS storage, replication and snapshot efficiency, allocation and provisioning, and backup and data deduplication. NetApp provides archive storage and continuous data protection with the SnapManager product family. Iron Mountain's archiving solution delivers full mailbox capture, storage lifecycle management, retention, and e-discovery. With the increased emphasis on e-discovery preparedness, organizations are storing a growing amount of critical data for extended periods of time. By leveraging NetApp storage solutions, Iron Mountain is able to deliver greater innovation for archiving and discovery by creating a simpler, more efficient solution.

Iron Mountain NearPoint has high-performance storage requirements for both the SQL database and the archive content index, which is used for searching the archive. Because of their access patterns, both the database and the index are best stored on high-performance Fibre Channel (FC) or iSCSI attached block storage. The NetApp platforms provide excellent support for NearPoint's block-based storage requirements. In addition to the database and index storage, NearPoint uses NetApp storage for retaining the archived objects themselves on the same array or on a different NetApp storage device. The NetApp architecture supports a wide variety of storage configurations, including FC or SATA devices accessed through block or CIFS protocols. NetApp solutions are ideal to support all the storage requirements of NearPoint, giving administrators a full range of options for connecting the NearPoint application to a NetApp infrastructure.

NearPoint is a next-generation e-mail archiving solution, distinguished by its unique method of capturing e-mail content from Exchange. Unlike traditional e-mail archiving solutions that rely on MAPI or journaling, NearPoint captures e-mail directly from the Exchange stores and database transaction logs. The benefits of NearPoint's data capture technology are twofold. First, NearPoint is able to capture complete mailbox content and context, something that Exchange journaling and MAPI alone

cannot accomplish. Second, NearPoint puts no additional strain on the Exchange Server, while standard Exchange journaling can add as much as 35% overhead to the Exchange Server.

Iron Mountain calls its unique data capture technology Continuous Application Shadowing. Iron Mountain and NetApp have worked together to integrate Continuous Application Shadowing with NetApp SnapManager for Exchange. Iron Mountain calls this “shadowless” archiving, and NetApp is the first storage vendor to integrate with NearPoint in this manner. This approach saves archive disk space that would otherwise be necessary to copy the Exchange store files.

Combined, Iron Mountain NearPoint and NetApp offer a unique and innovative approach to e-mail archiving and archive data protection. NearPoint has been integrated to take advantage of the NetApp Unified Storage Architecture as well as SnapManager for Exchange to offer a uniquely comprehensive yet simple e-mail recovery and archive solution. Storage efficiency technologies like single instancing and block deduplication enhance the benefits of the combined solution.

3 IRON MOUNTAIN ARCHIVING SOLUTION

The Iron Mountain NearPoint software solution for e-mail archiving is based on standard Microsoft® technologies, and it runs on standard Intel® servers. NearPoint is an application-intelligent solution that provides deep integration with Microsoft Exchange and Microsoft Outlook® for seamless end-user access to messaging content. NearPoint is best known for its unique data capture technology and zero footprint on production Exchange Servers and desktops. It provides instant search and access to e-mail that has been protected, archived, and extended via a standard Microsoft Outlook or Outlook Web Access user interface.

- NearPoint performs continuous capture of Exchange data, metadata, and context, including all e-mail, folders, deletions, calendars, contacts, notes, tasks, and more.
- NearPoint minimizes the risks and costs associated with legal discovery through its comprehensive e-discovery solution.
- NearPoint optimizes Exchange mailbox storage by reducing the load that active mail puts on the Exchange environment. Offline e-mail stored in PST files can be easily migrated into the NearPoint archive.

Iron Mountain NearPoint minimizes risks associated with e-mail litigation readiness with a comprehensive e-discovery solution, going beyond basic e-mail to include every aspect of a mailbox—public folders, offline PST files, calendars, contacts, notes, and more. End users and auditors can quickly search complete archive information. For advanced legal discovery and workflow, the NearPoint eDiscovery option is available. NearPoint speeds legal discovery and increases search accuracy by providing a single location to search all current and historical e-mail information.

4 NETAPP STORAGE SOLUTIONS

NetApp storage system solutions offer compelling advantages to the data management scenario. The ability to provision storage with primary and archive workload characteristics on a single system offers simplified management and leverages or minimizes IT skill sets. Administrators are only required to manage, protect, and maintain a single system that is capable of providing multiple service levels. Recent regulatory requirements and IT operating models have driven many enterprises to pursue a more systematic and optimized archiving process. When new technologies are properly aligned with the business and legal requirements, companies can expect to reduce risks and costs while boosting IT efficiencies.

The use of NetApp's advanced storage management technologies is a significant benefit of the joint solution. These features are completely transparent and seamless to the Iron Mountain NearPoint archive because standard storage interfaces are used. You can leverage the powerful benefits of the NetApp data management tools without worrying about special on-site integration.

The four major NetApp products that are integrated with Iron Mountain NearPoint are NetApp SnapManager for Exchange for shadowless archiving, NetApp FlexVol® to create thin provisioned storage partitions, FAS deduplication to remove redundant data blocks on the storage system, and NetApp SnapLock® to protect archived data against incidental deletion—a requirement of the financial services industry.

5 NETAPP FLEXVOL

With FlexVol volumes, administrators can respond to changing storage needs fast, with little effort and no disruption. FlexVol storage-virtualization technology allows administrators to effectively share large pools of storage between the Iron Mountain NearPoint archive and other applications without having to predict future requirements or overallocate storage because predicting usage is not possible. FlexVol pools physical storage, so it is possible to create and resize virtual volumes as archive needs change. By using FlexVol, administrators can improve the utilization of existing storage and save the expense of acquiring more disk space. I/O performance and reduced bottlenecks are achieved by distributing volumes across all available disk drives.

6 NETAPP SNAPLOCK

Iron Mountain NearPoint is able to retain messages for varying lengths of time, and can even prevent manual deletion as required. Data immutability is a critical piece of a compliance solution. NetApp SnapLock is a flexible data permanence solution that enables you to meet the strictest data retention regulations or internal IT governance policies. SnapLock aids in compliance with records retention regulations that require e-mails, documents, audit information, and other data to be retained in an unalterable state until the end of the specified retention period. SnapLock is especially beneficial in retrieving unregulated but crucial reference data that has not been changed or deleted but must be accessed quickly. For technical information about SnapLock, see TR-3263: “WORM Storage on Magnetic Disks Using SnapLock Compliance and SnapLock Enterprise.”

7 NETAPP DEDUPLICATION FOR FAS

NetApp deduplication is a core component of Data ONTAP®. NetApp has the first deduplication capability that can be used broadly across many applications, including primary data, backup data, and archival data managed by NearPoint. When combined with NearPoint single instance storage, data deduplication enables compelling storage efficiencies. NetApp is the only vendor with a WORM-compliant deduplication storage system. Data ONTAP 7.3.1 and later feature the ability to deduplicate a SnapLock FlexVol volume in both compliance and enterprise modes. NetApp offers the only suitable platform for strict adherence to SEC 17a-4, NASD 3110, DOD 5015, Sarbanes-Oxley, and HIPAA requirements. For more information about deduplication, see TR-3505: “NetApp Deduplication for FAS and V-Series Deployment and Implementation Guide.”

8 SNAPMANAGER FOR MICROSOFT SQL SERVER

NetApp and Iron Mountain recommend SnapManager for Microsoft SQL Server® to protect the NearPoint databases. SnapManager for Microsoft SQL Server is tightly integrated with Microsoft technologies to help you streamline database storage management while simplifying storage layout planning, backup, and restore operations for SQL Server databases. SnapManager for Microsoft SQL Server can save you time and money with space-efficient backup capabilities and automated data management processes. With SnapManager, you can dramatically reduce SQL Server data recovery times from hours to minutes, making it one of the fastest backup and recovery solutions available. You can also use SnapManager for SQL Server with its PowerShell™ cmdlets to automate backup, recovery, and database cloning.

SNAPMANAGER FOR MICROSOFT SQL SERVER BENEFITS

Achieve near-instantaneous backup and fast restore of entire SQL Server databases and full text indexes by using NetApp Snapshot™ technology:

- Reduce storage costs with space-saving backup technologies
- Increase productivity by automating routine database tasks
- Protect more data and increase your backup frequency without affecting performance
- Move databases to SAN and IP SAN environments with easy migration wizards
- Restore a failed database of any size to full production in minutes
- Easy-to-use, intuitive graphical user interface
- Rich backup scheduling and reporting
- Integration with NetApp SnapMirror® for wide area data replication

For a detailed discussion of the capabilities and best practices for SnapManager for Microsoft SQL Server, see TR-3768.

9 SNAPMANAGER FOR MICROSOFT EXCHANGE

NetApp and Iron Mountain recommend using SnapManager for Microsoft Exchange to protect and manage your Microsoft Exchange environments. With its Microsoft Management Console GUI and PowerShell cmdlets, you can centrally manage and automate the complex, manual, and time-consuming processes associated with the backup, recovery, and verification of Exchange Server databases.

SNAPMANAGER FOR MICROSOFT EXCHANGE BENEFITS

- Reduce backup times to seconds and restore times to just minutes.
- Add storage capacity and expand volumes without taking the Exchange Server or the NetApp storage system offline.
- Streamline management by automating common tasks so that administrators spend less time on maintenance.
- Deploy into the existing infrastructure.
- Integrate SnapManager with native Microsoft technology and frameworks.
- Integration with SnapMirror (synchronous and asynchronous) and SnapVault®.
- Provide disaster recovery and business continuity.
- Provide role-based access control.

For a detailed discussion of the capabilities and best practices for SnapManager for Microsoft Exchange, see TR-3598.

10 SNAPMANAGER FOR EXCHANGE INTEGRATION WITH NEARPOINT

NearPoint's Continuous Application Shadowing data capture technology is designed to perform off-host and is uniquely suited to integration with SnapManager for Exchange. When used with SnapManager for Exchange, NearPoint uses a Snapshot copy of the Exchange database to populate the archive with content created after the previous archiving task. This enables better storage utilization and a high-performance function for Exchange archiving. NetApp deduplication can then be applied to the NearPoint archive volume to maximize storage efficiency.

The simplest use case for the NearPoint integration with SnapManager for Exchange is when a local e-mail archive is populated from a Snapshot copy of the production Exchange databases. Here are the basic steps of the integrated archiving:

1. Schedule regular backups with SnapManager for Exchange.
2. SnapManager for Exchange notifies NearPoint that a Snapshot copy is complete and ready for processing.
3. NearPoint mounts the Snapshot copy and extracts archive e-mail from the copy.
4. When archiving is complete, NearPoint unmounts the Snapshot copy.

A common deployment, shown in Figure 1, illustrates how NearPoint is configured to use the same storage array as Exchange. NearPoint coordinates with SnapManager for Exchange to create a Snapshot copy of the Exchange database and uses this virtual image as the source for database extraction. NearPoint processes the database extracted by SnapManager for Exchange by creating a full text index of the newly archived data and storing the metadata in SQL. The index is used for rapid content searches by end users or discovery administrators. The Snapshot copy of the Exchange database is taken on the NetApp storage system, which means that there is negligible performance impact on the production Exchange Server operations. For a detailed discussion of the server impact of SnapManager for Exchange, see TR-3598.

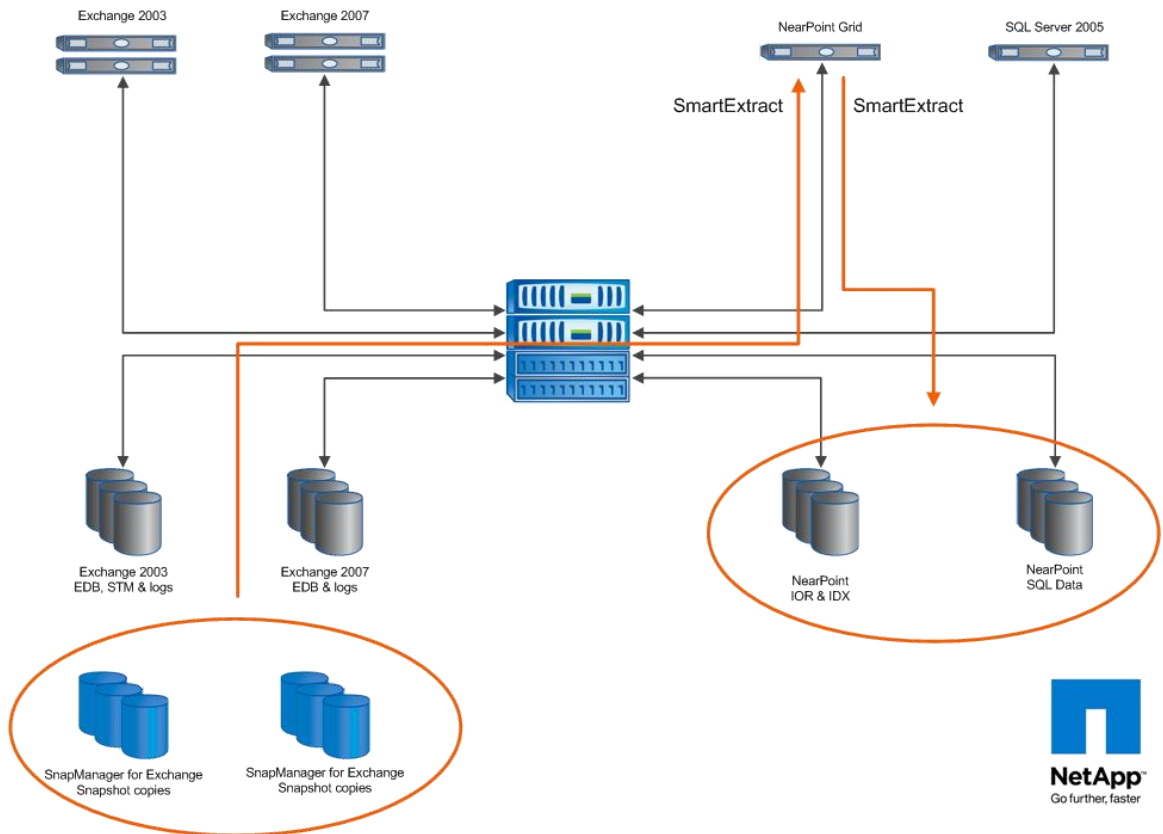


Figure 1) Typical Iron Mountain NearPoint architecture on NetApp storage solutions.

The joint solution delivers optimizations for customers who deploy NetApp storage and NetApp manageability software with NearPoint. Customers benefit from a reduction in archive storage and lower performance requirements for both servers and storage. NearPoint management is integrated into a single user interface. Administration is simplified by the coordinated integration of the archive application and the storage software; therefore only one tool is needed for recovery purposes.

Storage optimization comes from leveraging Snapshot copies of the Exchange database on the NetApp storage platform. The production Exchange environment has a disaster recovery copy made possible by SnapManager for Exchange. Iron Mountain NearPoint uses that copy as the source for the archiving; no other copy is required. This reduces storage requirements, and it also greatly enhances the disaster recovery and data protection processes.

An additional benefit is lower overall processing and transactional requirements. One of the heaviest CPU loads generated by NearPoint is the merging of log files into the NearPoint shadow database. This load is completely eliminated by the integrated solution. Customers see better performance from existing deployments, and future deployments will be satisfied with more cost-effective solutions.

11 DISASTER RECOVERY INTEGRATION

For single-site archive installations, the previous sections describe the compelling advantages of combining NearPoint with the features of SnapManager for Exchange. For many enterprises, having remote services availability for disaster tolerance is a basic requirement. To enable remote disaster recovery capability for the NearPoint archive, Iron Mountain has integrated NearPoint with several software options available to NetApp users. By integrating with SnapMirror and SnapManager for SQL Server, NearPoint is able to create an archive that brings all the advantages of Iron Mountain NearPoint to a disaster recovery capable storage platform for the complete messaging environment.

As in the single-site use cases, SnapManager for Exchange creates a copy of the Exchange database. In this scenario, the Snapshot copy of Exchange is sent to a remote NetApp storage system. SnapManager for Exchange uses this remote copy to provide backup and recovery of the Exchange environment, and Single Mailbox Restore is used to service individual mailbox requests. SnapManager for SQL Server is used by NearPoint to create local copies of the SQL databases, and SnapMirror is used to replicate the Snapshot copies of the index and the NearPoint archive to the remote destination. In this use case, full-site failover is supported for both the Exchange and archive environments, managed through one simplified user interface.

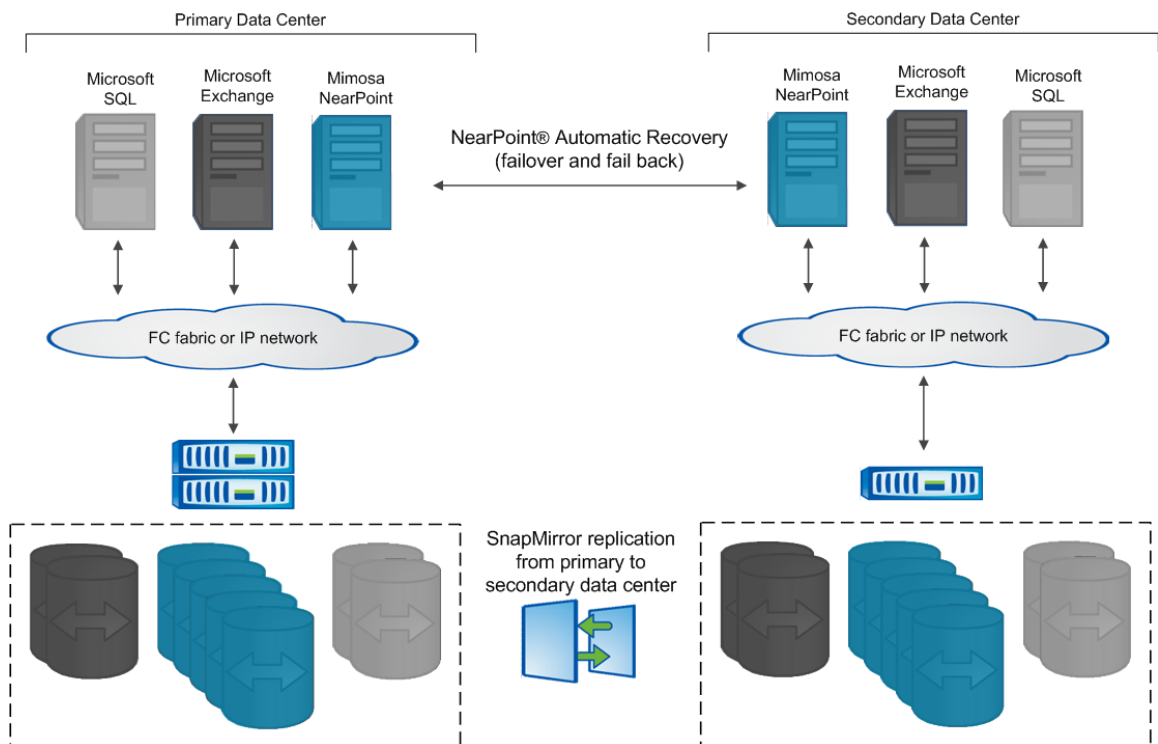


Figure 2) Iron Mountain NearPoint on NetApp storage to enable a disaster recovery configuration.

12 KEY BENEFITS OF THE NETAPP AND IRON MOUNTAIN JOINT SOLUTIONS

The NetApp storage solutions and Iron Mountain NearPoint can help you to manage your archives while spending less money. Together they create one of the industry's most compelling solutions for Exchange data protection and archiving. The combined solution from NetApp and Iron Mountain can meet the most challenging requirements for compliance, legal discovery, data recovery, and disaster protection in a single integrated solution. The benefits of an integrated solution are:

- Provides integration with SnapManager for Exchange to archive e-mail.
- Offers an integrated disaster recovery solution for e-mail.
- Reduces Exchange disaster recovery time and vulnerability.
- Reduces the cost of legal discovery with in-house review.
- Optimizes Exchange and e-mail archive storage while eliminating the need for PSTs.

13 CONCLUSION

Together, NetApp and Iron Mountain deliver valuable archiving and information management solutions for environments that use Microsoft Exchange. The joint offering provides valuable benefits in the areas of data protection, storage management, legal discovery, and compliance. The highlight of the joint solution is the integration of NetApp SnapManager for Exchange and Iron Mountain NearPoint. The NetApp value proposition of storage efficiency, performance, operational agility, and data protection strongly complements the unique data extraction capabilities offered by NearPoint. With NetApp and Iron Mountain, you achieve efficiency without compromising performance. We can help you increase cost savings AND improve business performance.

14 REFERENCES

1. Addressing Information Overload in Corporate Email: The Economics of User Attention
http://www.radicati.com/files/Addressing_Info_Overload.pdf
2. A Guide to Message Archiving http://www.ostermanresearch.com/whitepapers/or_or0108.pdf
3. <http://www.law.cornell.edu/rules/frcp/Rule26.htm>
4. http://www.uscourts.gov/rules/EDiscovery_w_Notes.pdf
5. <http://www.netapp.com/us/library/technical-reports/tr-3263.html>
6. <http://www.netapp.com/us/library/technical-reports/tr-3505.html>
7. <http://www.netapp.com/us/library/technical-reports/tr-3563.html>
8. <http://www.netapp.com/us/library/technical-reports/tr-3441.html>
9. <http://www.netapp.com/us/library/technical-reports/tr-3446.html>
10. <http://www.netapp.com/us/library/technical-reports/tr-3338.html>
11. <http://www.netapp.com/us/library/technical-reports/tr-3598.html>
12. <http://www.netapp.com/us/library/technical-reports/tr-3768.html>
13. <http://now.netapp.com/NOW/products/interoperability>

NetApp provides no representations or warranties regarding the accuracy, reliability or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

