



Using Oracle10g™ Database Automatic Storage Management with Network Appliance™ Storage

A Joint NetApp and Oracle White Paper

TR-3329 | June 2004

Ravi Animi, Jeff Kimmel, Network Appliance, Inc.

Paul Manning, Bill Bridge, Oracle Corporation

TECHNICAL REPORT

Network Appliance, a pioneer and industry leader in data storage technology, helps organizations understand and meet complex technical challenges with advanced storage solutions and global data management strategies.

Introduction 3
Best Practices Recommendations 5
 Common Recommendations 5
 Shared Disk Groups 6
 Separate Disk Groups 7
Conclusion 11
Appendix 1: Implementing ASM on NFS-Based Storage 12
Appendix 2: Using NetApp Snapshot Copies with ASM-Based Data 13
Appendix 3: Introduction to NetApp Technology and
NetApp Technical Reports 15

Using Oracle10g™ Database Automatic Storage Management with Network Appliance™ Storage

INTRODUCTION

The purpose of this joint white paper is to share best practices for use of the new Oracle Database 10g feature Automatic Storage Management (ASM) with Network Appliance (NetApp) Fibre Channel and iSCSI SAN storage products. Oracle and NetApp have a long-standing technology relationship that has helped many joint customers realize solutions that leverage technology from both companies. This white paper is another example of what this partnership can provide in terms of joint best practices. ASM provides a new way to manage the storage underlying the database. NetApp filers simplify the deployment and management of enterprise data. Both ASM and NetApp storage lower total cost of ownership, and they complement each other to offer even greater cost savings. This paper shares some of the considerations as to how the combined technology can best be configured.

ASM provides an integrated file system and volume manager for the database files, built into the Oracle Database kernel. With this capability, ASM provides an alternative to some of the third-party file-system and volume management solutions for database storage management tasks, such as creating/laying out databases and managing the use of disk space.

This paper seeks to educate the Oracle DBA on the combined benefits of using NetApp storage for Oracle Databases and database layered applications and includes best practices recommendations for deployment of ASM with NetApp networked storage solutions. Like Oracle on NetApp NAS storage, Oracle ASM on NetApp SAN and iSAN storage gives customers an alternative capability for volume management on the Oracle server host using familiar create/alter/drop SQL statements, simplifying the job of DBAs with regard to database storage provisioning.

The following topics summarize how features of ASM and NetApp storage complement and combine with one another:

I/O Load Balancing

Load balancing avoids performance bottlenecks by assuring that the I/O workload utilizes all available disk drive resources.

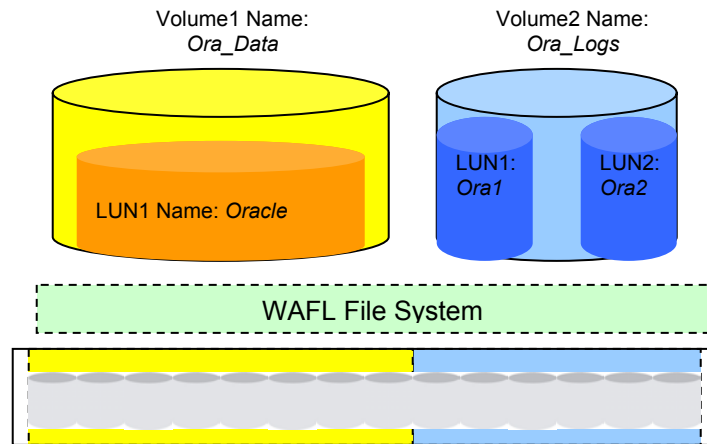
NetApp storage automatically load-balances I/O among all the disk drives in a WAFL® volume (see Figure 1). All LUNs and files placed within a single WAFL volume can be assured of utilizing all the volume's disk drives in a balanced fashion.

ASM provides load-balanced I/O across all LUNs in an ASM disk group by distributing the contents of each data file evenly across the entire pool of storage in the disk group based on a 1MB stripe size.

When used in combination, NetApp load balancing allows multiple LUNs and file-system data to share common disk drives, while reducing the number of LUNs per ASM disk group for improved manageability. ASM further allows load balancing across multiple WAFL volumes or filers. The WAFL file system and RAID4 are optimized for random small I/O. WAFL ensures the data is evenly spread across all the disk drives of a RAID group in a volume. This ensures optimal read/write performance in high-transaction database environments.

Figure 1) NetApp storage overview.

This figure shows NetApp WAFL file-system volumes, each residing on a distinct set of physical disk drives. The volumes contain files and LUNs, which are distributed across all the disk drives in the volume.



NetApp volumes can be configured from a pool of available physical disk drives—FC (filers) or ATA (NearStore® system)—in one or more software RAID groups in a RAID4 or RAID-DP configuration.

Adding/Removing Storage Capacity

Support for online addition and removal of storage capacity maximizes data availability while making necessary storage configuration changes (refer to Figure 1 for an illustration of these principles).

ASM allows online resize, addition, or removal of LUNs within each ASM disk group, automatically rebalancing data distribution for new configurations. Entire disk groups can be migrated from one set of LUNs to another (e.g., on different volumes or filers) by adding then removing LUNs from a disk group without database downtime.

NetApp storage allows online addition and replacement of disk drives within each WAFL volume and rebalances the load to utilize new disk drives as writes occur.

NetApp I/O rebalancing across physical disk drives only occurs when disk drives are added to expand a NetApp volume nondestructively. Adding LUNs to a volume does not cause any additional I/O rebalancing to occur as long as there is enough space in the volume to add a new LUN or grow an existing LUN. Please see TR 3001, TR 3002 and other introductory material listed in Appendix 3, which explain NetApp architecture basics. The key benefit to volume and LUN resizing is that it is a nondestructive process that is rapid and easy to accomplish. This allows you to grow your ASM disk group easily and rapidly with no downtime, thus allowing you to provision storage dynamically in sync with your database growth needs.

Data Protection and Availability

Enterprises require very high levels of protection and availability for their data, regardless of faults or component failures in storage subsystems.

NetApp storage implements RAID4 or RAID-DP protection on all WAFL volumes. RAID4 assures full data availability in spite of the loss of any given disk drive in a RAID group. RAID-DP implements double-parity protection to guard against the loss of any two disk drives. The WAFL write anywhere data layout assures excellent RAID write efficiency without incurring the storage costs implied by full mirroring. NetApp cluster failover protects against failure of other filer hardware components, such as within the filer head itself. Additionally, NetApp

mirroring solutions can be implemented for additional protection for business continuance and disaster recovery requirements.

ASM optionally implements two-way or three-way extent-based mirroring, offering an additional means for protection against storage failures. ASM also provides the ability to mirror between sets of ASM disks (implemented as LUNs on NetApp SAN storage), called failure groups, and can keep data available after a component failure causes all storage in one of those failure groups to become unavailable.

Combining the data protection technologies from NetApp and Oracle offers customers several ways to address their data protection requirements. The following best practices address how these might best be configured.

Shared Storage for RAC Clustered Nodes

Oracle10g Real Application Clusters (RAC) provides high availability and scalability for Oracle Databases implemented across clustered database servers. RAC requires that all servers have full, direct read/write access to the database files.

ASM implements simple and Oracle RAC-optimized cluster file system and volume manager functionality for shared storage.

NetApp networked storage also provides the shared access RAC requires. In SAN environments, ASM provides the cluster-capable file system and volume management for Oracle10g RAC.

ASM and NetApp SAN storage combine to create a simple and scalable storage facility for Oracle10g RAC. Together, NetApp and Oracle toolsets provide a rich assortment of data management options for use separately or in combination. Customer selection of techniques will depend on their unique requirements and the context of their enterprise's data infrastructure.

BEST PRACTICES RECOMMENDATIONS

This section discusses guidelines for configuring ASM disks, creating ASM disk groups, and creating databases within ASM, discussing some best practices along the way. Please refer to Figures 2 through 4 for further illustration of some of the best practices concepts and recommendations proposed.

Common Recommendations

The following recommendations apply regardless of which of the two high-level deployment models above is selected:

Use host file systems or NFS to store Oracle binaries and other non-ASM-enabled files. In Oracle10g RAC environments, NFS can be used to simplify deployment using a shared Oracle Home.

Use either Fibre Channel or iSCSI SAN protocols to access the storage for ASM disk groups.

Combine each filer's disk drives into few large WAFL volumes to minimize administrative overhead.

Make sure that the LUNs within each disk group are balanced in terms of throughput per capacity, so that the I/O throughput available to a disk group from each LUN is proportional to the LUN's size. In more concrete terms, all LUNs within the disk group should supply roughly equivalent I/O operations per second per gigabyte. Generally this will be best accomplished using disk drives with very similar capacity and performance properties throughout the underlying WAFL volume(s).

Configure several (e.g., four) LUNs per ASM disk group per WAFL volume. Configuring multiple LUNs can maximize I/O concurrency to the disk group, overcoming any per-LUN limits of the host driver stack (operating system, host bus adapter driver, and/or multipath driver). Even if initially unnecessary, this provisioning can avoid unneeded data movement that would occur if LUNs within the same volume were later added to the disk group.

Deployment Models

Either of two high-level deployment models may be appropriate for ASM on NetApp, depending on the intended use of data management facilities at the storage subsystem (filer) level:

If the ASM deployment will use filer-based features for data management (such as filer-based facilities to implement full database backup/restore, cloning, mirroring, and/or migration), then separate ASM disk groups should be established for each database deployed, and each disk group should be stored within a single WAFL volume. This deployment model is termed “**Separate Disk Groups**” below.

If filer-based data management will not be used, then it will be simpler to group multiple databases into shared disk groups. This deployment model is termed “**Shared Disk Groups**” below.

Shared Disk Groups

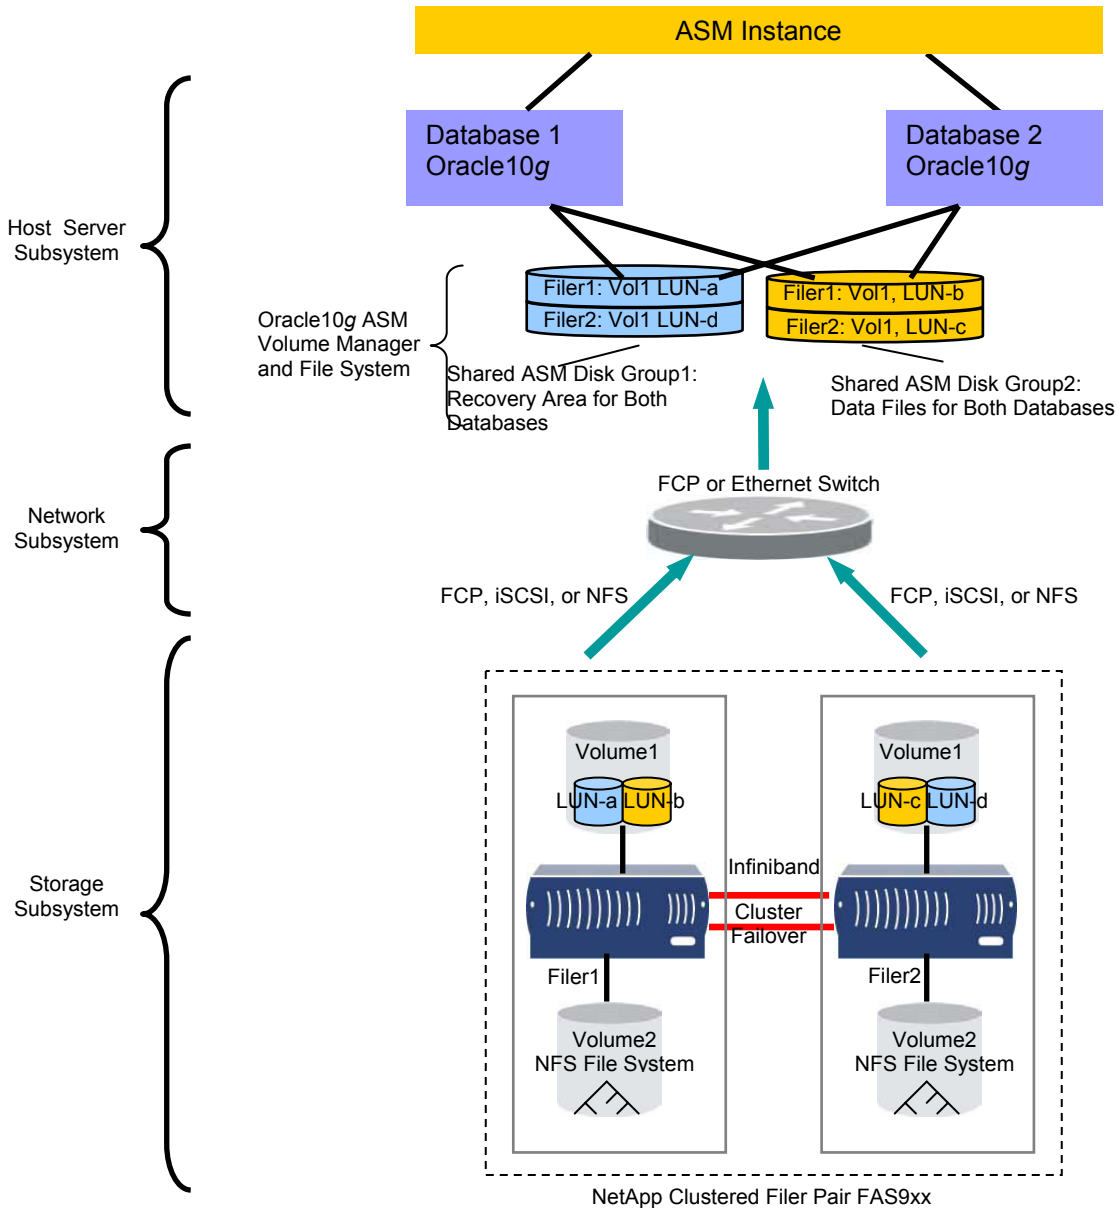
When all data management methods will be based exclusively on Oracle tools, the easiest deployment model will group multiple databases into shared disk groups that provide a common shared pool of storage for several databases. The following additional recommendations are specific to this model. See Figure 2 for an illustration of these recommendations:

Use two or more Oracle ASM disk groups to store the databases. The database work area should be stored in separate disk groups from their respective flashback recovery areas. For databases that may require independent control over the underlying storage resources (e.g., to assure quality of service or independent data migration capabilities), assignment into distinct ASM disk groups may be considered.

Each ASM disk group may include storage from one or more WAFL volumes or filers. Spreading a disk group across multiple WAFL volumes will distribute the disk group’s storage and throughput demands across volumes based on the combined size of the LUNs that each volume contributes to the disk group to improve performance. When adding disk drives to existing volumes for greater capacity or throughput, expand all the WAFL volumes and the disk group’s LUNs within them to maintain uniform throughput-per-capacity ratios and to minimize unneeded data movement due to load balancing.

If adding storage from a new filer within a disk group, configure the new filer’s LUN(s) with similar throughput-per-capacity properties to those already in the disk group. If the new filer contains multiple LUNs for host I/O concurrency, then all LUNs should be added to the disk group at once to limit ASM’s load-balancing movement to that required.

Figure 2) Multiple Oracle Databases running on shared ASM disk groups that are distributed across multiple NetApp filers.



Separate Disk Groups

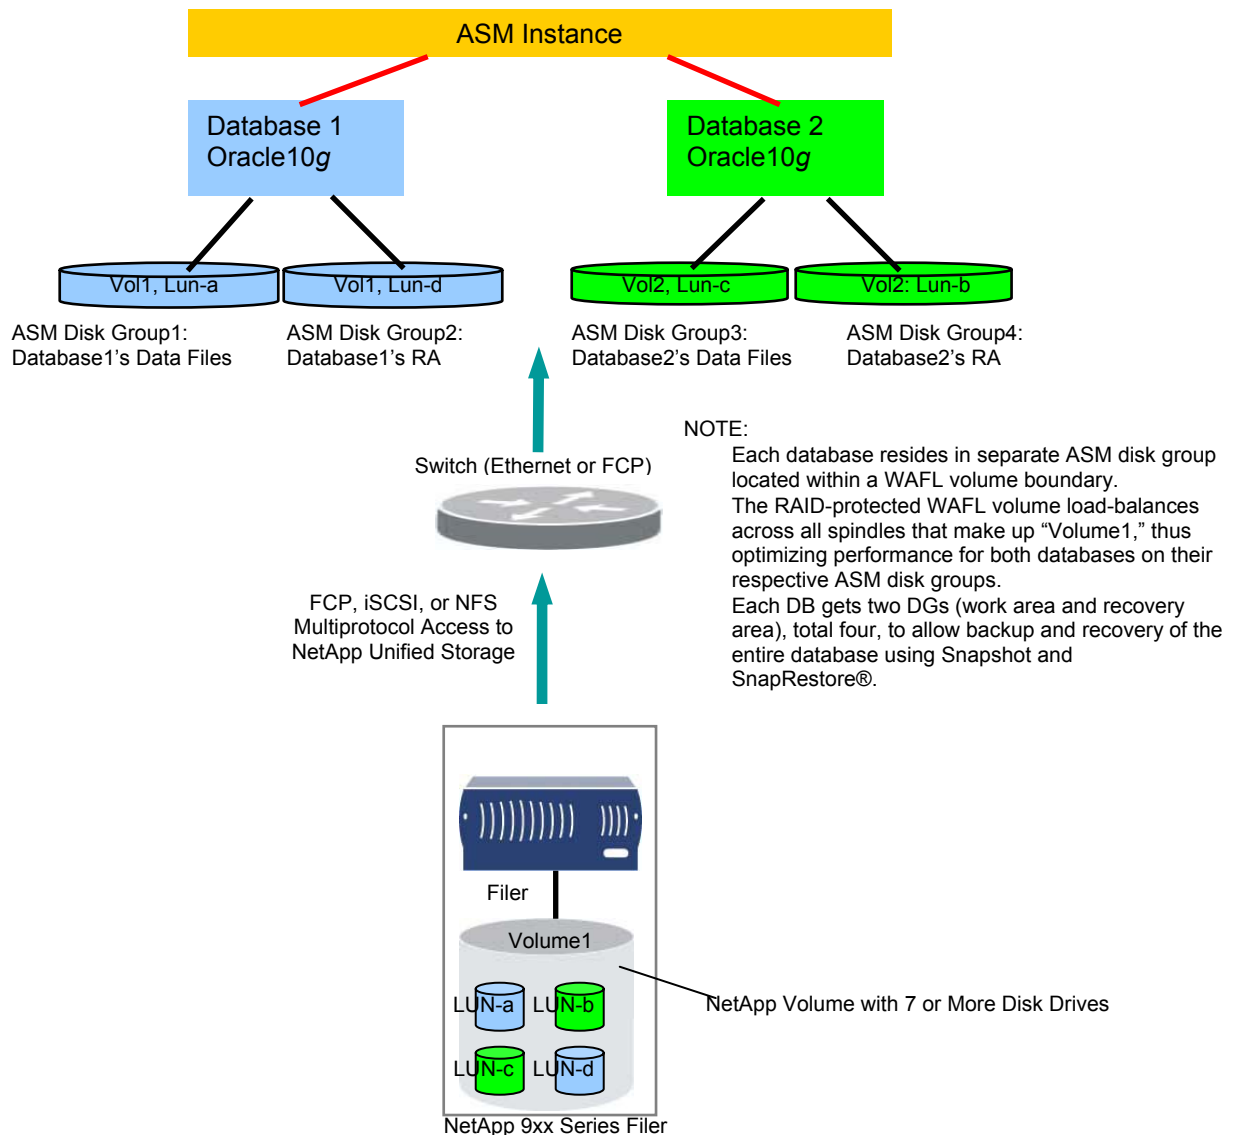
In order to leverage the filer’s ability to implement data management, each database should be stored in distinct ASM disk groups with each ASM disk group stored in a single WAFL volume. If that is not possible, an ASM disk group that spans WAFL volumes requires that the disk group be unmounted when creating Snapshot™ copies to assure ASM metadata consistency among the Snapshot copies. See Figure 3 for an illustration of these recommendations.

Use two Oracle ASM disk groups per database, one containing database work area data files and the other containing the flash recovery area.

Place all LUNs within a given disk group in a separate WAFL qtrees (quota tree directory) to maximize data management flexibility.

When adding disk drives for greater capacity or throughput, expand the WAFL volume and then resize the LUNs on the filer, in the host operating system, and finally in ASM to expand the ASM disk groups. Perform all the ASM add disk and drop disk commands after you have set the power limit parameter to 0 in order to pause the rebalance process until you are ready to increase the power limit and have the rebalance done in one pass.

Figure 3) Multiple Oracle Databases running on separate ASM disk groups that map to separate NetApp WAFL file-system volumes.



Failure Groups and Redundancy

Generally, ASM's external redundancy option should be configured for ASM disk groups on NetApp storage. This choice minimizes the amount of storage hardware required to realize data protection for a given level of database capacity and throughput. With external redundancy, the NetApp LUNs that comprise the ASM disk groups can be RAID parity protected. Standard RAID4 parity protection allows for continued operation with a single disk drive failure in a RAID group. NetApp

double-parity (RAID-DP option) protection allows for continued operation even after a simultaneous double disk drive failure in a RAID group, dramatically reducing the risk of data loss due to simultaneous double disk drive failure.

The NetApp cluster failover (CFO) option should generally be configured to assure availability in case of unscheduled filer downtime. As an alternative to using CFO, ASM disk groups could be configured with two-way mirroring (i.e., without the external redundancy option) to provide this assurance, for a cost of up to double the required disk drive hardware. If ASM mirroring is used to protect against filer failure, all LUNs in a disk group hosted by one filer should be placed within the same failure group.

ASM Disks

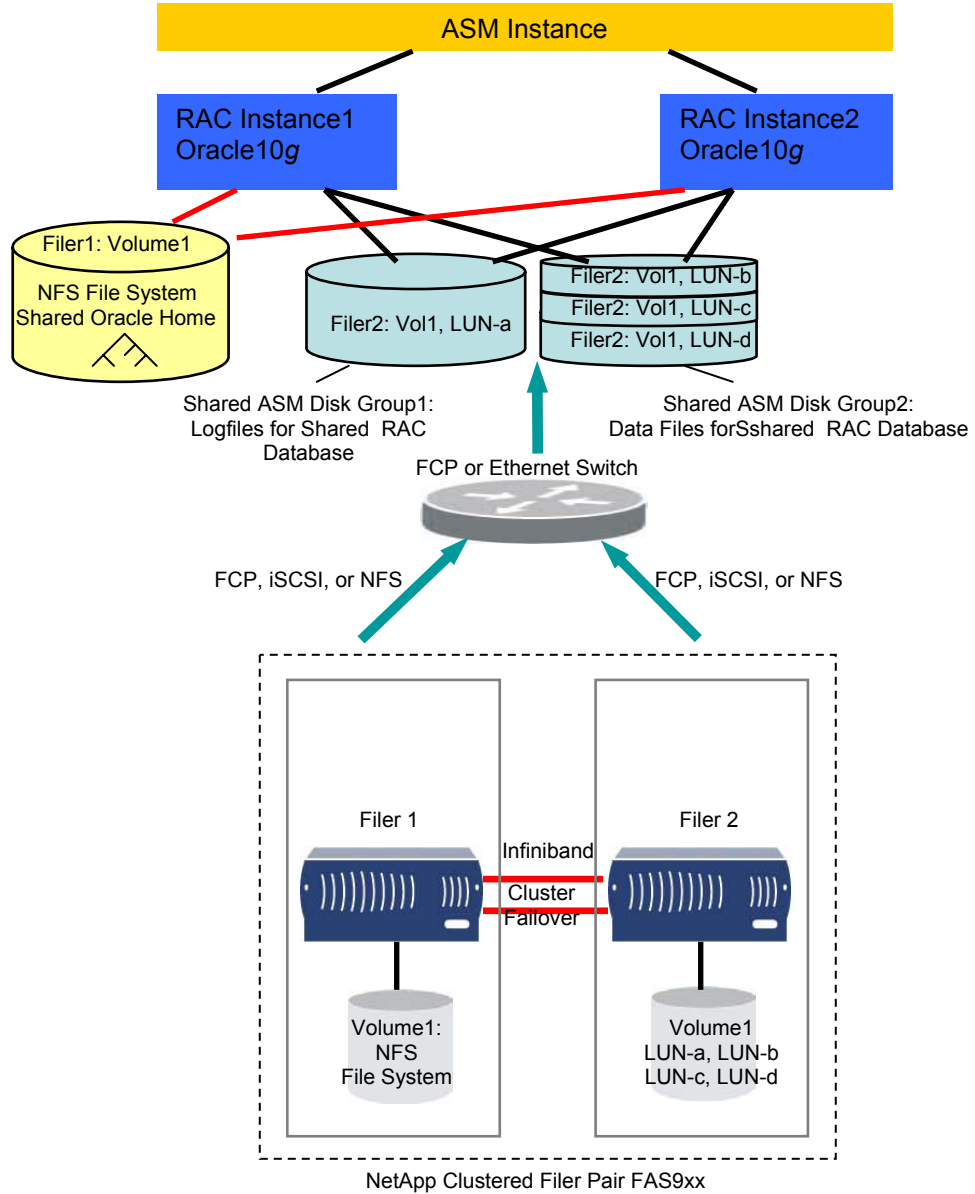
NetApp LUNs have the same semantics on the host side as any other LUNs, and as such they are incorporated into ASM disk groups as documented in the 10g ASM documentation. The best practice recommendation is to use LUNs of the same size and performance capacity when configuring an ASM disk group. In order to achieve good load balancing for a disk group stored on multiple WAFL volumes, the total capacity from each WAFL volume should be proportional to the throughput available to the disk group from that volume.

Running 10g RAC on ASM Disk Groups Derived from NetApp

When Oracle10g RAC is deployed with ASM on storage derived from NetApp, we recommend placing the shared Oracle Home on a NetApp NFS volume for ease of manageability. In order to install and enable cluster synchronization services in Oracle10g one needs to follow the directions in the product documentation. Alternatively, some clustered host file systems can support a shared Oracle Home on SAN storage, or each node can have its own Oracle Home on a private host file system. See Figure 4 below.

With the release of Oracle10g RAC, Oracle has made it extremely easy, as compared to Oracle9i™ RAC, to share a single Oracle Home on a NAS volume among multiple nodes of an Oracle10g RAC cluster. The ability to share the Oracle Home binaries and configuration files between multiple RAC instances adds tremendous benefits to manageability and administration of a multinode Oracle10g RAC cluster. The shared instance specific admin and configuration files (TNSnames.ora, initSID.ora, etc.) have separate parameter entries for each cluster node. When deploying ASM on NetApp SAN storage, it is generally preferable to implement the Oracle10g RAC CRS files (voting disk and OCR device) on SAN storage as well; consult the configuration details used for the applicable Oracle10g RAC validation.

Figure 4) Multiple Oracle10g RAC instances running on shared ASM disk groups.



CONCLUSION

Automatic Storage Management simplifies the configuration and management of storage management for the Oracle10g database data files. ASM complements NetApp Fibre Channel and iSCSI SAN storage products, and they can be combined in ways that deliver the full value of both sets of technologies. Although several possible combinations of redundancy protection and data management options exist, the partnership between NetApp and Oracle provides best practices for how these can be combined to meet the business requirements of our joint customers. Depending on the high-availability and performance requirements, there are some powerful combinations of our joint technology that meet those demands. Although the increased number of possible combinations may seem complex at first, this paper sets the criteria in place to help choose the most cost-effective solution that fits the specific requirements. It is another example of how the partnership between Oracle and NetApp delivers unique value for our joint customer base. Together we are helping reduce the complexity and cost of achieving high-performance and high-availability business systems with a reduced management overhead.

Summary of Best Practices

- 1) ASM is the recommended file system and volume management solution in a NetApp SAN/iSAN environment for Oracle Databases.
- 2) Configure ASM disk groups with external redundancy for use with filers.
- 3) Use NetApp cluster failover to assure no single point of failure.
- 4) Use ASM for management of Oracle Database files.
- 5) Use the NetApp clustered file system for shared Oracle Home, alert logs, and trace files.
- 6) Build ASM disk groups with LUNs of similar performance/capacity properties.
- 7) Distributing I/O across the front end with ASM and striping within the filer are complementary and can further simplify the combined configuration versus using either approach in isolation.

APPENDIX 1: IMPLEMENTING ASM ON NFS-BASED STORAGE

ASM can be configured with NFS-based files as ASM disks. For optimal performance this is not considered best practice, but it will work. When using ASM on top of NFS, the NFS files are used as ASM disks, so the provisioning and data management model is similar to that for ASM on SAN storage, while NFS is used to perform I/O.

There may be circumstances where ASM on NFS is applicable, whether for uniformity of operating procedures or to satisfy product, host OS, and/or other environmental constraints. The following points discuss such configurations further.

NFS-based storage from NetApp and ASM provides some common functions. Oracle DBAs can directly employ the NFS-based file system functionality to implement simple, scalable, RAC-ready file storage. An Oracle DBA can directly provision storage from a NetApp filer using NFS storage and the familiar create/alter/drop SQL statements. ASM on NFS similarly implements scalable, RAC-ready storage and further extends scalability at the high end, but provisioning additional filer storage to Oracle files within Oracle will require some additional steps when using ASM on NFS storage.

When implementing ASM on NFS, one should combine the guidelines for deploying Oracle directly on NFS with those above for implementing ASM on NetApp SAN storage:

- Use NFS mount options as recommended for non-ASM Oracle files on NAS. See the applicable NetApp technical reports (see Appendix 3) for guidance on configuring specific NFS clients for use with Oracle. (Specifically, refer to TR 3183, “Using the Linux® NFS Client with Network Appliance Filers,” and the relevant NFS configuration sections in TR 3189, “Oracle9i Real Application Cluster Release 2—Installation with a Filer in Red Hat Advanced Server Linux Environment.”)

- Create NFS-based files according to the guidelines for LUNs discussed in the “ASM Disk Groups” best practices above. Either the “shared disk groups” or “separate disk groups” model may be followed as appropriate. Multiple files and/or mount points from a single WAFL volume can be used to overcome I/O concurrency limits, just as multiple LUNs were discussed for this purpose in the SAN best practices above.

- Use mkfile to size the files to be incorporated into the ASM disk groups (as if they were LUNs). If later resizing is required, establish up-to-date backups, then alter the files’ sizes as required without truncating or overwriting any live data; the tools needed to do this vary among host environments.

- Set the permissions on the files to allow "oracle:dba" read/write access.

- Use an ASM disk string that allows the ASM instance to find the file, for example, /nfs/<host>/oradata/<db_name>/*. Symbolic links can be used to make such a disk string work for files spread across multiple NFS mount points.

APPENDIX 2: USING NETAPP SNAPSHOT COPIES WITH ASM-BASED DATA

Backup and recovery methods are of course critical elements of any Oracle production deployment. Refer to Figures 5a and 5b for an illustration of the concepts presented in this section.

NetApp Snapshot copies implement inexpensive point-in-time copies, which can be used to accelerate media recovery, making them a very popular feature for Oracle environments. Snapshot copies also provide the basis for various other backup/recovery methods and can also be used to clone databases nearly instantaneously. In order to effectively apply Snapshot copies with ASM-based Oracle data, it is important to focus on the ASM disk group as the unit of Snapshot copy-based recovery, copying, or cloning.

ASM also works with the Oracle full suite of data management tools, including RMAN, Data Guard, Flash recovery, and Flashback queries. These technologies can be used when databases span multiple WAFL volumes or more than one filer.

RMAN is the only interface able to take hot and cold backups of Oracle Databases on ASM disk groups, and, more importantly, RMAN is the only interface for single file restore capability from a backup set. However, you can also do hot and cold backups of Oracle using the NetApp volume-level point-in-time Snapshot copy and SnapRestore to back up and restore entire LUN(s) and hence entire ASM disk group(s) to a point in time. In order to preserve ASM disk group integrity, all disks within the ASM disk group must be snapped atomically with the current NetApp technology; this practically implies following the “separate disk groups” guidelines discussed in the best practices section earlier, including placement of all LUNs for a given disk group into a single WAFL volume. ASM maintains all disk group metadata within the disk group itself, so a disk group can be restored to a previous state by dismounting, restoring, and remounting the disk group. A disk group’s metadata will be fully consistent for any atomic point-in-time copy of the disk group, so there is no need for flushing and freezing the metadata as is required prior to creating Snapshot copies of some traditional host file systems. All that is required for doing a hot backup is to put the individual tablespaces of the database in online or hot backup mode and then create a Snapshot copy of the volume containing the LUN(s).

Figure 5a) Overview—Backing up an Oracle10g Database on ASM with NetApp storage.

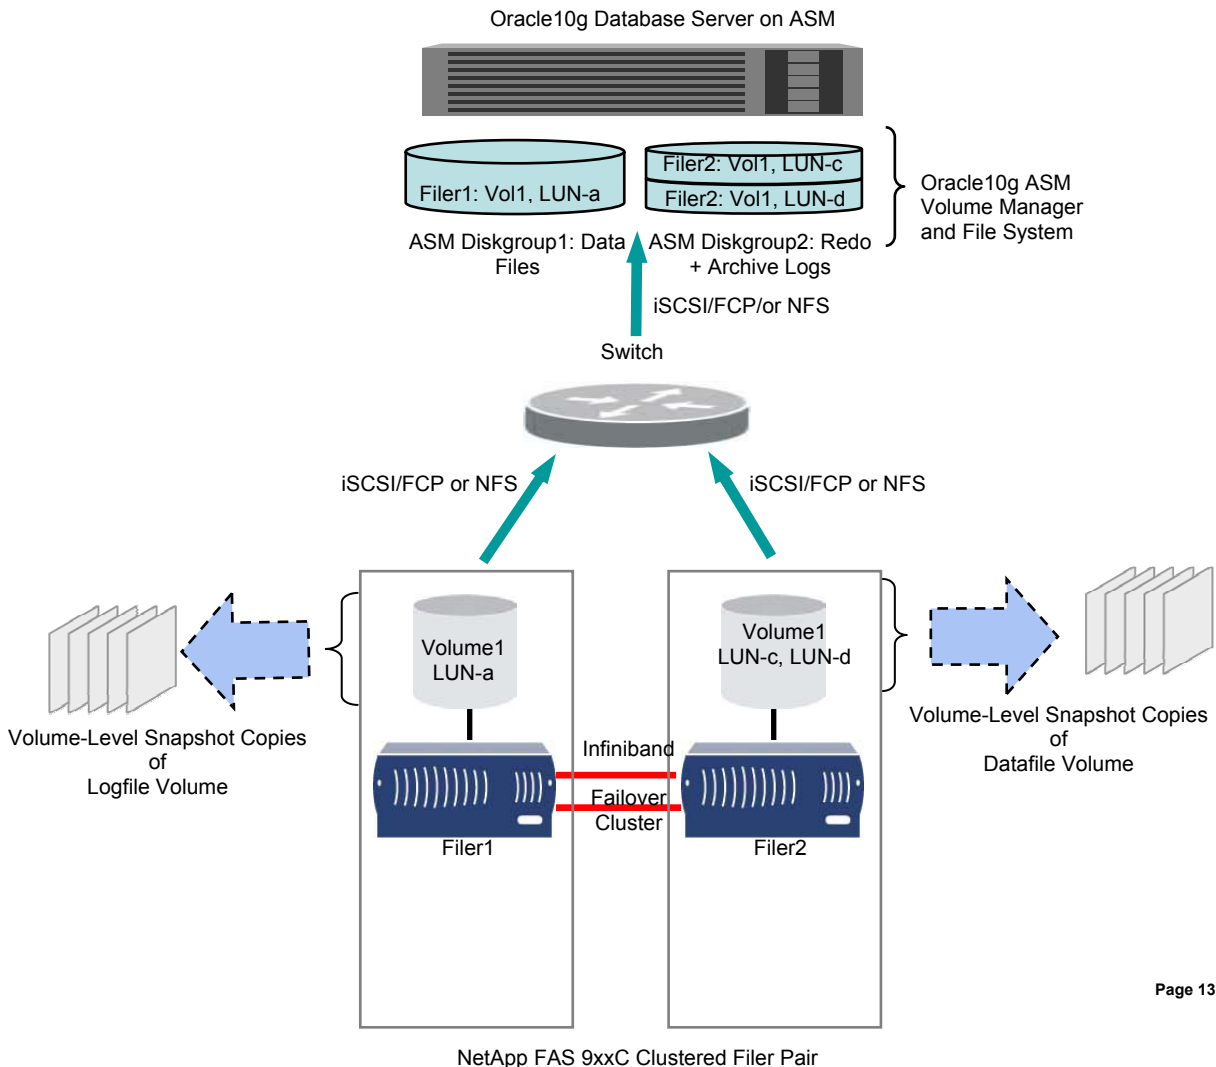
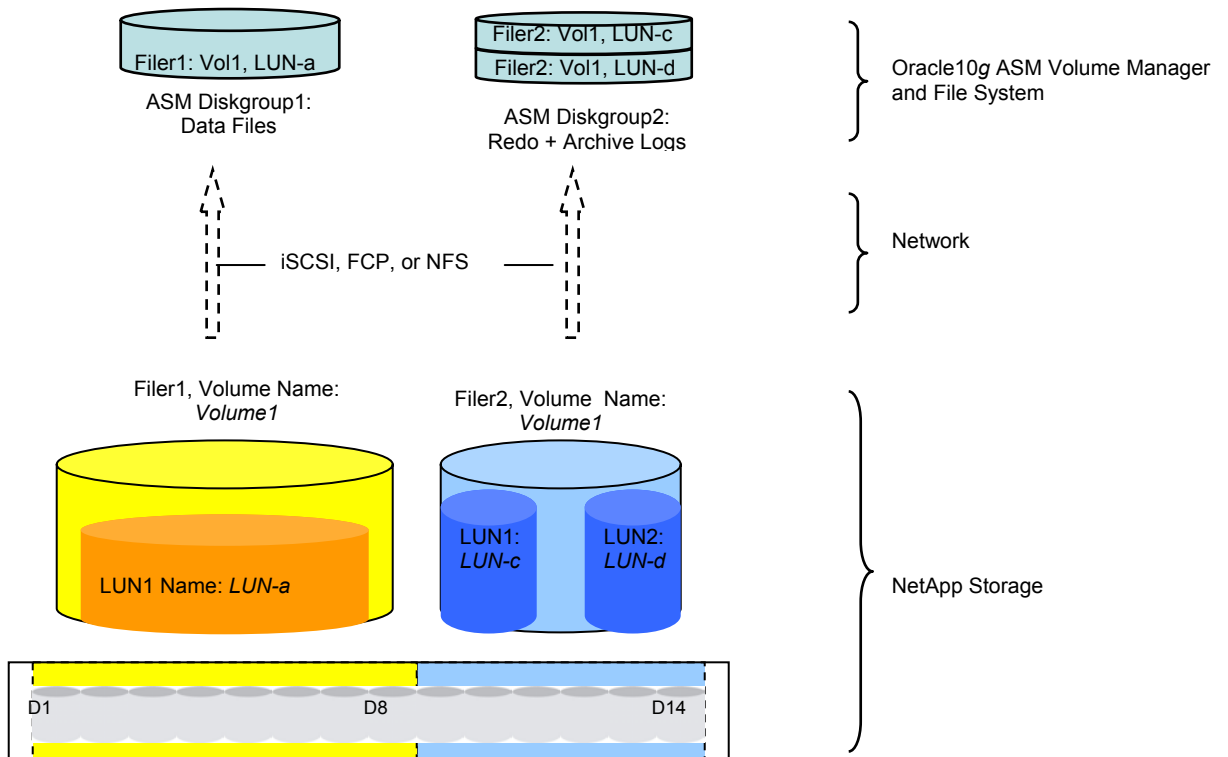


Figure 5b) Backing up an Oracle Database on ASM with NetApp storage.



NetApp Snapshot and SnapRestore technology can be used for an Oracle ASM environment similar to how it is used for a non-ASM environment. The only constraint is that to back up your Oracle ASM disk groups using Snapshot copies, the entire ASM disk group must be contained within a NetApp WAFL file-system volume boundary.

APPENDIX 3: ADDITIONAL TECHNICAL REFERENCES

The following technical collateral has been selected as introductory material on NetApp technology. Additional sources for further reference are listed below.

INTRODUCTION TO NETAPP TECHNOLOGY

- TR 3001:** NetApp Architectural Overview: www.netapp.com/tech_library/3001.html
- TR 3002:** File System Design for NFS/NAS Appliances: www.netapp.com/tech_library/3002.html
- TR 3014:** Multiprotocol Data Access: NFS, CIFS, HTTP: www.netapp.com/tech_library/3014.html
- TR 3241:** iSCSI—Accelerating the Transition to Network Storage: www.netapp.com/tech_library/3241.html
- TR 3183:** Using the Linux NFS Client with Network Appliance Filers: www.netapp.com/tech_library/3183.html
- TR 3027:** Multiple Volumes and Multiple RAID Groups on NetApp Filers: www.netapp.com/tech_library/3027.html
- TR 3298:** Double-Parity RAID for Enhanced Data Protection with RAID-DP: www.netapp.com/tech_library/3298.html
- TR 3131:** Data Protection Solutions Overview: www.netapp.com/tech_library/3131.html
- TR 3194:** Rapid Recovery with Network Appliance Storage Solutions: www.netapp.com/tech_library/3194.html
- TR 3202:** Best Practices for Data Protection with Filers Running FCP: www.netapp.com/tech_library/3202.html
- TR 3127:** High-Availability Storage Appliance: www.netapp.com/tech_library/3127.html
- TR 3227:** NFS vs. SAN for Database Applications: www.netapp.com/tech_library/3227.html
- TR 3102:** Databases on NetApp Filers—A TCO Study: www.netapp.com/tech_library/ftp/3102.pdf
- TR 3189:** Oracle9i Real Application Cluster Release 2—Installation with a Filer in Red Hat Advanced Server Linux Environment: www.netapp.com/tech_library/3189.html
- TR 3294:** Dynamic File System Expansion on Solaris™: www.netapp.com/tech_library/3294.html

ORACLE TECHNICAL WHITE PAPERS ON ORACLE 10G AND AUTOMATIC STORAGE MANAGEMENT

Overview of Automatic Storage Management:

<http://otn.oracle.com/products/manageability/database/pdf/asmov.pdf>

Technical White Paper on Automatic Storage Management:

<http://otn.oracle.com/products/manageability/database/pdf/asmwp.pdf>

ASM – The new best practice:

http://otn.oracle.com/products/database/asm/pdf/40288_asm_fnl.pdf

Simplify Your Life – Automatic Storage Management:

http://otn.oracle.com/products/database/asm/pdf/40140_asm_fnl.pdf

Overview of Oracle 10g:

http://otn.oracle.com/products/database/oracle10g/pdf/oracle_database_10g_revolution.pdf

Description of new benefits of Oracle 10g:

http://otn.oracle.com/products/database/oracle10g/pdf/twp_general_10gdb_product_family_0104.pdf

Web Sites with more details about Oracle Database10g and ASM

<http://otn.oracle.com/products/database/asm/index.html>

<http://otn.oracle.com/tech/linux/asmlib/index.html>

<http://otn.oracle.com/tech/linux/asmlib/install.html>

<http://otn.oracle.com/products/database/oracle10g/index.html>

NETAPP AND ORACLE JOINT SOLUTIONS

In addition to the excellent technical reports and white papers on NetApp technology, we invite you to browse through additional technical and marketing collateral at the following URLs:

NetApp:

www.netapp.com/tech_library/

www.netapp.com/solutions/oracle-db.html

www.netapp.com/tech_library/index_db.html#to c

ORACLE

June 2004

Paul Manning, Bill Bridge

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
www.oracle.com

Oracle Corporation provides the software
that powers the internet.

Oracle is a registered trademark of Oracle Corporation. Various
product and service names referenced herein may be trademarks
of Oracle Corporation. All other product and service names
mentioned may be trademarks of their respective owners.

Copyright © 2004 Oracle Corporation
All rights reserved.



June 2004

Ravi Animi, Jeff Kimmel

Network Appliance Inc.
World Headquarters
495 E Java Dr.
Sunnyvale, CA 94089
U.S.A.

Worldwide Inquiries:
Phone: +1.408.822.6000
Fax: +1.408.822.6000
www.netapp.com

Network Appliance provides the hardware and software
that powers the data centers that run the enterprise

Specifications subject to change without notice. NetApp, the Network Appliance
logo, NearStore, SnapRestore, and WAFL are registered trademarks and Network
Appliance and Snapshot are trademarks of Network Appliance, Inc. in the U.S.
and other countries. Solaris is a trademark of Sun Microsystems, Inc. Linux is a
registered trademark of Linus Torvalds. All other brands or products are
trademarks or registered trademarks of their respective holders and should be
treated as such.

Copyright © 2004 Network Appliance, Inc. All rights reserved.
All rights reserved.



Network Appliance, Inc.
495 East Java Drive
Sunnyvale, CA 94089
www.netapp.com