An Oracle White Paper
August 2011

Creating a Hierarchical Database Backup System using Oracle RMAN and Oracle SAM QFS with the Sun ZFS Storage Appliance

Introduction

In a perfect world, all backups would reside on enterprise-level storage and be available for recovery within seconds. However, for most enterprises, the cost to procure, manage, and maintain the required storage would devastate their IT budget.

As mission-critical databases grow to gigabytes – and even terabytes, enterprises find themselves challenged to provide storage with the capacity to quickly and cost-effectively absorb and retain backups consistent with retention policies. Administrators do not want to use prime disk space for data that is no longer relevant, but, at the same time, need to ensure the availability of valuable data even as it ages. A solution is needed that supports the backup of large amounts of data while automatically demoting older data to higher capacity/more cost effective storage.

This paper describes how to implement such a hierarchical database backup solution using the Oracle Recovery Manager (Oracle RMAN), the SUN Storage Archive Manager/Quick File System (SAM QFS) software, Sun ZFS Storage Appliances, and Oracle's StorageTek tape drives and libraries. This hierarchical database backup system is an ideal solution for customers with three or more mission-critical Oracle databases, 40 TB or more aggregate backup data, and a requirement to keep backups on tape for more than 180 days.

Overview of the Hierarchical Database Backup Solution

The hierarchical database backup system described in this paper is implemented using Oracle RMAN recovery management tools, a SAM QFS backup file system, one or more Sun ZFS Storage Appliances, and StorageTek tape drives and libraries. This solution automatically and transparently places data on the most cost-effective storage, based on specified retention and retrieval policies, to help align storage and archiving costs with business priorities and ease the management burden. Features of each of these components are described in *Appendix B. Solution Component Features*.

For this solution, Oracle RMAN is configured to place backups on a SAM QFS backup file system. The SAM QFS file system provides what appears to applications to be an "infinite disk" repository. A first-tier disk cache provides up to 4 PB of high-resilience, high-performance, relatively low capacity primary storage. Hundreds of PBs of additional storage is provided on higher capacity, lower-resilience, lower-performance storage tiers including disk archive and tape media.

SAM is configured to keep critical backups on Tier 1 storage to allow quick restoration. As a backup image ages, the relevance of the data typically diminishes, so older backups are moved to lower tiers.

As the data ages further, it is eventually available only on tape, typically in a traditional disk-to-tape backup system. Applications are unaware of the physical location of the data.

To reduce the complexity and costs of provisioning the tiers of a hierarchical system with multiple forms of back-end storage, the Sun ZFS Storage Appliance provides a single point management interface from which to configure each tier. SAM QFS then dynamically utilizes the storage tiers based on the organization’s data retention and retrieval policy. When a backup residing in a lower tier needs to be accessed, the SAM QFS software stages the backup images and associated files to Tier 1 storage for rapid database recovery.

Oracle RMAN automatically executes routine maintenance tasks before and after a backup or recovery operation. A multi-tiered storage configuration such as this reduces time, complexity and staffing costs for backup and recovery operations.

Physical Architecture of a Hierarchical Database Backup System

One goal of a successful backup system is to minimize the additional load or contention for resources on the backup storage system. Minimizing this additional load keeps the resources available for the Oracle Database, reducing service time and enhancing throughput.

To facilitate achieving this goal, it is recommended that a separate server be deployed to run as a SAM server that is connected to back-end storage comprised of a SUN ZFS Storage Appliance and a StorageTek tape library. The type of connectivity used between the SAM server, the SUN ZFS Storage Appliance, and the StorageTek tape library depends on the available infrastructure and on performance requirements (see the section *Performance Considerations* below). A suggested configuration is shown in Figure 1. (Figure 1 does not show the administration network to which all the components are connected.)

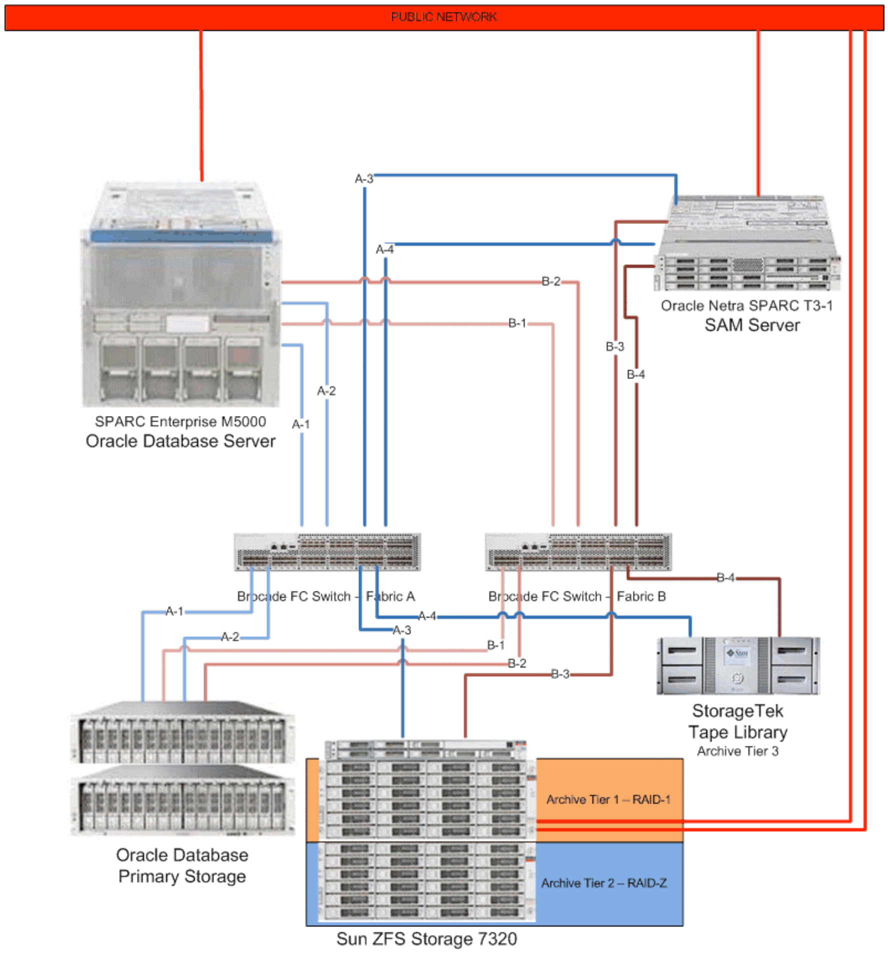

Figure 1. Physical architecture of a hierarchical database system

In the configuration in Figure 1, the SAM server is an Oracle Netra SPARC T3-1 server with two dual-channel Fibre Channel (FC) host bus adapters (HBAs). The HBAs are connected to existing FC switches in a dual-fabric configuration.

The back-end storage for the SAM server is provided by a FC-attached Sun ZFS Storage Appliance. In this case, a Sun ZFS Storage 7320 with two disk expansion trays is used to provide two archive tiers with different characteristics. Half the LUNS presented from the SUN ZFS Storage 7320 are masked to a channel on one of the FC HBAs in the Netra SPARC T3-1 server and half to a channel on the other FC HBA The SUN ZFS Storage Appliance is particularly well suited to meet a variety of storage requirements due to its flexible storage configuration, which allows different sets of IO requirements to be satisfied within a single appliance.

A StorageTek SL48 tape library, populated with four half-height LTO-4 tape drives, is also connected to each FC switch. The drives and robotics interface in this tape library are assigned to the same zone as the second channel on each of FC HBAs on the Netra SPARC T3-1 server.

Sizing of components, including the Sun ZFS Storage Appliance and StorageTek tape drives, is covered in the section *Sizing Oracle RMAN Backups* below.

The SAM QFS file systems maintained by the SAM server are presented to the Oracle Database server using Network File System (NFS) protocol. The Sun ZFS Storage Appliance LUNs can be accessed using iSCSI protocol when an existing network infrastructure is available. With respect to the SAM server, the only difference between an iSCSI infrastructure and a FC infrastructure is the way in which the devices are discovered initially. Several differences in performance between an iSCSI and a FC configuration are described in the section *Performance Considerations* below.

Hierarchical Database Backup System Logical Architecture

Figure 2 shows the logical architecture of the hierarchical database backup system including the dataflow paths between the logical components in the solution.

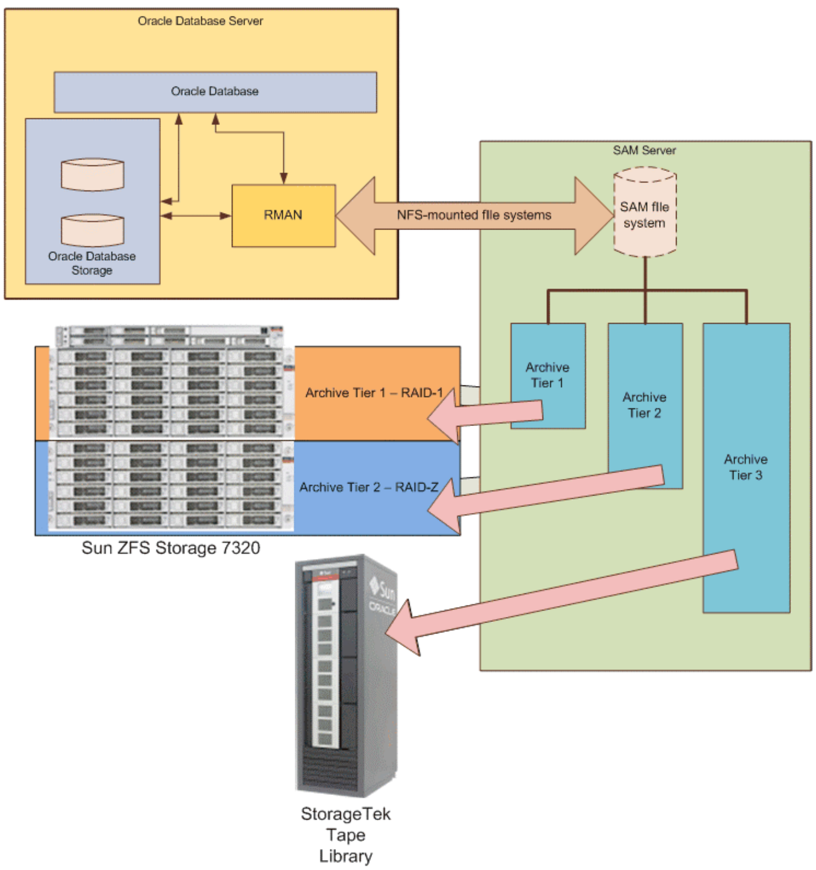

Figure 2. Logical architecture of the hierarchical database backup system

SAM Server Configuration

In this solution, two LUNs are presented by the Sun ZFS Storage Appliance – one configured as RAID-1 and one as RAID-Z. The RAID-1 LUN provides higher resilience and potential performance, while the RAID-Z LUN provides higher capacity but at the cost of reduced performance.

The StorageTek SL48 tape library contains four tape drives, which are split into two storage pools on the SAM server. The tape library can be shared with other applications, such as traditional backup applications, using the StorageTek Automated Cartridge System Library Software (ACSLS) Manager, which allows the tape population to be partitioned to create logical, application-specific libraries.

The RAID-1 LUN is used as a block device by the SAM QFS server to create a SAM file system that is mounted and exported as a shared file system to NFS clients on other servers. In the example SAM configuration shown in Table 1 below, this file system is called `tier1` and is mounted on the SAM server on `/tier1`.

The RAID-Z LUN is used to create a second file system called `tier2`, which is mounted on `/tier2` on the SAM server, but is not exported using NFS protocol. This file system may be either a SAM QFS file system or a ZFS file system.

The archive command file `archiver.cmd` must be modified to enable archiving on the `tier1` file system. If you mounted a SAM QFS file system on `/tier2`, you must also disable archiving on the `tier2` file system in the `archiver.cmd` file. Disabling archiving on the `tier2` file system is not needed if you mounted a ZFS file system on `/tier2`.

In the example in Table 1, a disk volume with the name `level2` has been created from the file system `/tier2`.

TABLE 1. SAM STORAGE CONFIGURATION

FILE SYSTEM OR POOL	MOUNT POINT OR LABEL	NOTES
`tier1`	`/tier1`	Provided by Sun ZFS Storage Appliance (RAID-1)
`tier2`	`/tier2`	Provided by Sun ZFS Storage Appliance (RAID-Z, no archiving)
`level2`	`/tier2`	Pool of disk volumes providing SAM with fast archive media
`tape-pool-1`	DG001-DG020	Tape storage pool storing first copy of the backup images
`tape-pool-2`	DG021-DG040	Tape storage pool storing second copy of the backup images

SAM is configured to wait at least 'n' minutes after a file has been modified ('n' is archive age of the file) before creating a copy of the file to ensure the file contents are complete. For this example, SAM has been configured to write three archive copies to the `tier1` file system at three different archive ages. These are:

- Copy taken after 5 minutes and written to disk volume `level2`.
- Copy taken after 8 hours and written to `tape-pool-1`.
- Copy taken after 10 hours and written to `tape-pool-2`.

SAM writes the three archival images into a single tar image and archives it to disk or tape.

The exported SAM QFS file system can be configured to release copies when space becomes an issue (highly recommended). Thus, correct sizing of the tier1 and tier2 file systems is an important consideration and is covered in the section *Sizing Oracle RMAN Backups* below.

Releasing a copy releases most of the space that was allocated to the file in the `tier1` file system, while a record is maintained indicating where the other copies of the file reside. The released space in the `tier1` file system is now available for newer backup images. However, from the user's perspective, the original file continues to appear with the expected file size in the output of a standard utility such as `ls(1)`.

If the file has not been released, I/O requests are satisfied from the `tier1` file system. If the file has been released, SAM will check each copy location in turn for the file. Thus, the ordering of the copies is significant. In this example, copy 1 is stored in the `tier2` file system. Archive images are stored as `tar(4)` images on the disk volume, so SAM determines which `tar` image holds the most recent copy of the file. This file is then copied out of the `tar` image into the `tier1` file system and requests are processed from `tier1`.

If copy 1 has expired or been un-archived (as may be the case for a very old file), SAM requests the latest version of the file from the appropriate tape in the library. Depending on the age of the file, a tape may need to be imported to the tape library to retrieve a file. After the file has been read from tape (in `tar(4)` format), the data is written to the `tier1` file system to satisfy the I/O request.

If the file is modified, it is treated like a new file and is again subject to the archive copy rules as shown in Figure 3.

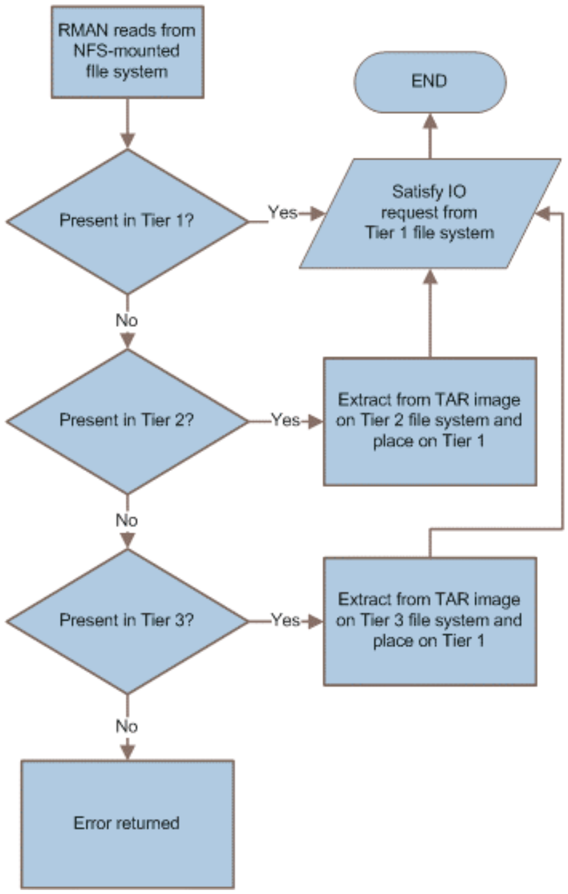

Figure 3. SAM I/O flow diagram of archive copy rules

Within the `tier1` file system, it is advisable to create a hierarchy of directories compliant with local policy to ensure some amount of segregation of the backup files. One method is to create a directory named after the database server and, in that directory, create directories corresponding to the Oracle System ID (SID) of each database on the database server as shown in Figure 4.

When files on more than one Oracle RMAN host are backed up to the SAM QFS solution, it may be advisable to provide multiple file systems, at least one for each Oracle RMAN host, to provide some level of segregation.

To comply with the least-access (or least-privilege) principle, user access to the NFS shares should be restricted to the appropriate Oracle RMAN servers. The `ro` parameter of the `share(1M)` command can be used along with a qualified list of hosts to ensure that backup images in a directory are accessed only from these named hosts. Permissions on the directories should be set such that minimum required access is granted, so that they are accessible only by the user that creates the backups (typically, the `oracle` user).

Figure 4 shows an example configuration for a multiple-server model.

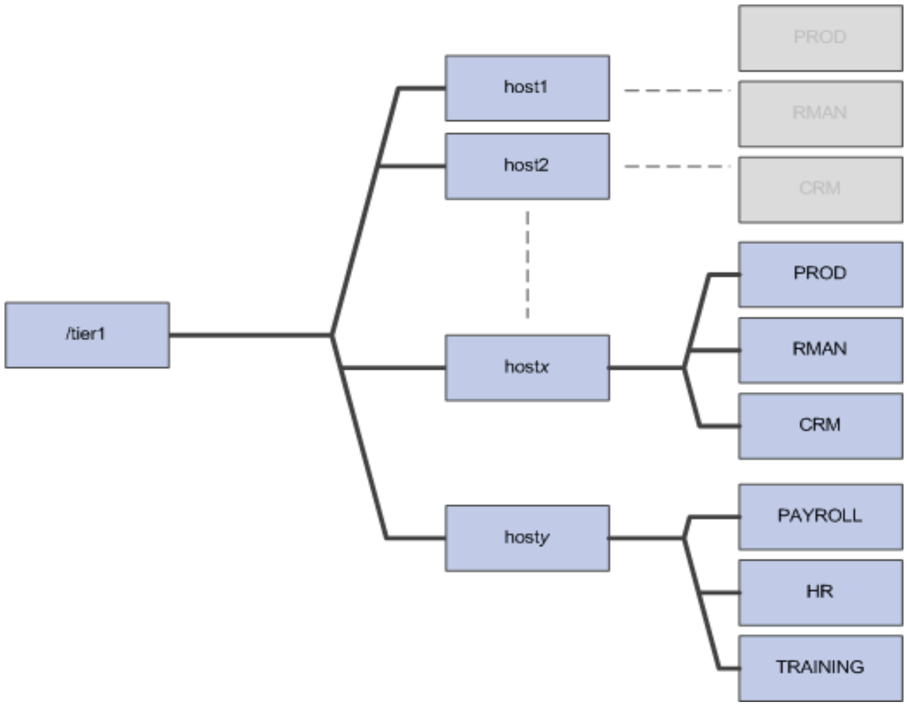

Figure 4. Example directory hierarchy for a backup system serving multiple hosts

Configuring the Oracle RMAN Server

Each directory corresponding to an Oracle RMAN server is mounted by the server under `/tier1` via NFS. Since these NFS-mounted file systems are to be used for Oracle RMAN backups, the mount options must be configured so that Oracle RMAN is able to use them efficiently. Table 2 lists the minimum options required for NFS mounts.

NOTE: It may be desirable to implement an Oracle Direct NFS-optimized NFS client configuration by replacing the standard Oracle Disk Manager (ODM) library with one that supports the Direct NFS Client. (See Hodak, W. & Closson, K. (July 2007) in *Appendix C. References* for a description of Direct NFS.)

TABLE 2. NFS MOUNT OPTIONS MANDATORY FOR ORACLE RMAN

PARAMETER	VALUE	DESCRIPTION
`rw`	–	Mount read-write.
`bg`	–	Retry the mount request in background if no response from the NFS server.
`intr`	–	Allow interrupts to kill a process that is hung while waiting for a response from a hard-mounted file system.
`rsize`	`32768`	Read blocks of 32K.
`wsize`	`32768`	Write blocks of 32K.

Table 3 describes a highly recommended NFS mount option for Oracle RMAN.

TABLE 3. NFS MOUNT OPTION RECOMMENDED FOR ORACLE RMAN

PARAMETER	VALUE	DESCRIPTION
`noatime`	–	Do not update the v-node 'access time' parameter when the file is read.

Oracle RMAN Configuration

Oracle RMAN must be configured to write database backup images to the directory mounted from the SAM server. For this example, it is assumed that Oracle RMAN runs on the Oracle Database server, although this is not required.

Assuming the file system exported by the SAM server is mounted on `/rmanbackup`, the parameters listed in Table 4 should be set to the values shown (highly recommended). With the parameters set to the values shown in the table, scheduled backups can be configured in the Oracle Enterprise Manager to use the SAM QFS server configuration.

TABLE 4. ORACLE RMAN SYSTEM CONFIGURATION VALUES

PARAMETER	VALUE	DESCRIPTION
CHANNEL DEVICE TYPE DISK FORMAT	/rmanbackup/%d/%U	Append SID as directory and standard Oracle RMAN unique filename.
DEFAULT DEVICE TYPE	DISK	By default, perform backups performed to DISK and placed them in the area specified by the parameter CHANNEL DEVICE TYPE DISK FORMAT.
CONTROL FILE AUTOBACKUP FORMAT FOR DEVICE TYPE DISK	/rmanbackup/%d/%F	Append SID as directory and standard RMAN control file name.
CONTROLFILE AUTOBACKUP	ON	Enable automatic backup of control files after backups.

Performance Considerations

A number of variables can affect the performance of Oracle RMAN backups to the SAM QFS server. Inadequate backup performance can directly affect database performance by causing contention for database back-end storage and CPU time.

A primary goal of a backup system is to meet stated business targets for recovery point objective (RPO) and recovery time objective (RTO) metrics. These metrics have a direct impact on how fast and how frequently backups must be created. To keep these metrics in the desired range, several performance areas can be examined:

- NFS traffic between the Oracle RMAN server and the SAM QFS server must be directed over a high-speed, low-contention network to allow the backup to be transferred as efficiently as possible. A 1GbE or 10GbE network should be used to provide sufficient bandwidth for transferring the images.
- When the SAM QFS server communicates with the Sun ZFS Storage Appliance using iSCSI protocol for LUN access, the network used for backups should be of similar bandwidth to provide the maximum possible bandwidth for traffic. If the SAM QFS server accesses the Sun ZFS Storage Appliance over a FC connection, zoning should be deployed with single-target/single-initiator zones and minimum access authorized.
- The back-end storage for the SAM QFS server provided by the Sun ZFS Storage Appliance must be configured with an appropriate block size. The block size should be large enough to minimize the number of requests that a single logical I/O request produces while, at the same time, small enough to keep to a minimum the amount of unnecessary data transferred, such as when the size of an I/O request is much smaller than the underlying storage block size.

Because RMAN requires 32 KB blocks for NFS reads and writes, it makes sense to keep the block size at 32 KB if block alignment can be guaranteed. If block alignment cannot be guaranteed, an underlying storage block size of 64 KB or 128 KB is more appropriate. A block size of 128 KB allows database backups to be written in sequential streams and supports the read performance required for efficient database restores.

Sizing Oracle RMAN Backups

Sizing of Oracle RMAN backup images is dependent on whether compression is enabled, the compressibility of the data involved, and whether the backups are incremental at various levels or are full backups. Including archived logs will also affect the size of the backup image, and the number of archive logs backed up will depend on the volatility of the database.

The Oracle Enterprise Manager provides historical backup reports, such as the example shown in Figure 5. These backup reports provide a means to identify trends and project future backup image sizes. SAM QFS is a scalable file system, so increasing capacity is a non-disruptive action if trends indicate more capacity is needed.

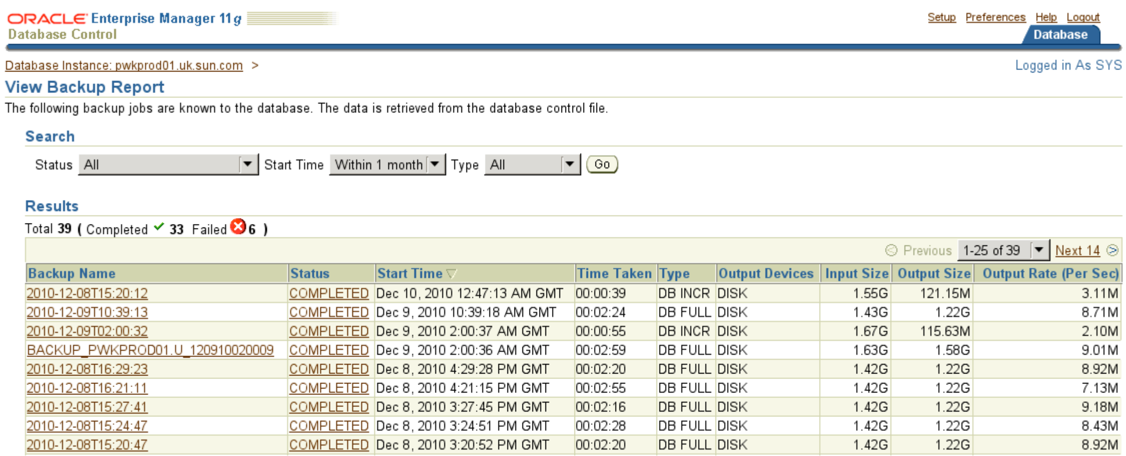

Backup Name	Status	Start Time	Time Taken	Type	Output Devices	Input Size	Output Size	Output Rate (Per Sec)
2010-12-08T15:20:12	COMPLETED	Dec 10, 2010 12:47:13 AM GMT	00:00:39	DB INCR	DISK	1.55G	121.15M	3.11M
2010-12-09T10:39:13	COMPLETED	Dec 9, 2010 10:39:18 AM GMT	00:02:24	DB FULL	DISK	1.43G	1.22G	8.71M
2010-12-09T02:00:32	COMPLETED	Dec 9, 2010 2:00:37 AM GMT	00:00:55	DB INCR	DISK	1.67G	115.63M	2.10M
BACKUP_PWKPROD01.U_120910020009	COMPLETED	Dec 9, 2010 2:00:36 AM GMT	00:02:59	DB FULL	DISK	1.63G	1.58G	9.01M
2010-12-08T16:29:23	COMPLETED	Dec 8, 2010 4:29:28 PM GMT	00:02:20	DB FULL	DISK	1.42G	1.22G	8.92M
2010-12-08T16:21:11	COMPLETED	Dec 8, 2010 4:21:15 PM GMT	00:02:55	DB FULL	DISK	1.42G	1.22G	7.13M
2010-12-08T15:27:41	COMPLETED	Dec 8, 2010 3:27:45 PM GMT	00:02:16	DB FULL	DISK	1.42G	1.22G	9.18M
2010-12-08T15:24:47	COMPLETED	Dec 8, 2010 3:24:51 PM GMT	00:02:28	DB FULL	DISK	1.42G	1.22G	8.43M
2010-12-08T15:20:47	COMPLETED	Dec 8, 2010 3:20:52 PM GMT	00:02:20	DB FULL	DISK	1.42G	1.22G	8.92M

Figure 5. Oracle Enterprise Manager backup reports

In the example backup report shown in Figure 5, the backup that was initiated at 10:39 a.m. December 9, 2010, took 2 minutes 24 seconds to complete, and backed up 1.22 GB in that time.

As this was a *full* backup, it can be assumed that the database is very small or relatively unpopulated at the moment. The incremental backup that was initiated nearly 14 hours later, at 12:47 a.m. December 10, 2010, only backed up 121.15 MB indicating the database is not particularly dynamic.

The sizes of the backups previous to the last full backup were fairly static. However, if the capacity of the backup file system is restricted to this size, no additional space will be available if the database grows or data volatility increases.

For this example, let's consider a scenario in which an import of data into the database is planned that will triple the size of the database and that the data will be fairly static once it is imported. Taking this scenario into consideration, we'll assume that, after the data is imported, the norm for full backups will be 8 GB, and a full backup, rather than an incremental, will be run daily.

The next step is to determine how long a backup will stay relevant and need to be available for fast restoration and recovery of the database. For example, if the Service Level Agreement (SLA) states that it must be possible to recover within minutes any data from any backup created in the last 14 days, the `tier1` file system must have at least enough capacity to hold 14 days of 8-GB full backups, or be at least 112 GB in size. To continually store metadata as well, the minimum capacity needs to be at least 128 GB.

To further support SLA requirements, the `tier2` file system will hold more than the 14 days of backups. For this example, we want to be able to hold 56 days worth of 8 GB backups. The `tier2` file system will hold `tar` images of the backups, so minimal space is needed for metadata. Thus the required capacity of the `tier2` file system is approximately 448 GB. To determine the size of the tape library, consider that the larger the internal cartridge population and the capacity of the tapes themselves, the less often physical attendance by an administrator will be needed for importing and exporting tapes.

The storage sizing requirements for this example are summarized in Table 5.

TABLE 5. SIZING REQUIREMENTS FOR SAM QFS EXAMPLE

FILE SYSTEM / DEVICE	USE	SIZE	DESCRIPTION
`/tier1`	NFS-exported Tier 1 archive file system	128 GB	RAID-1
`/tier2`	Tier 2 archive file system	448 GB	RAID-Z
Tape library	Dual image copies made to tape	48 carts	StorageTek SL48 tape library

When sizing for the Sun ZFS Storage Appliance SSD devices, consider using SSD log devices to enable fast low-latency NFS metadata operations. At least one SSD log device should be provided for each concurrent database backup.

NOTE: When SSD log devices are used, to optimize write throughput during the backup process, set the `logbias` property to `throughput` (the default is `latency`). The `throughput` setting causes SSD devices to be bypassed when data is written, ensuring that the large sequential writes associated with creating a backup are made directly to the disk subsystem. Set the `logbias` parameter back to `latency` when the backup is complete.

Additional Considerations

A few additional considerations when running a SAM QFS-based backup system are described below:

- When the Oracle RMAN command `CROSSCHECK BACKUP` is run, Oracle RMAN will attempt to access and read from every available backup image. This can be problematic when an attempt is made to access a file in a backup that is registered as still available but no longer resides on disk-based archives. SAM will attempt to stage the backup image from a lower tier, which can potentially take an extended time due to the necessity of loading, reading, and unloading of tapes to verify availability of each backup image.

  Using the modifier `BETWEEN` *date1* `AND` *date2* with the `CROSSCHECK BACKUP` command restricts the RMAN search to backups created between the specified dates, thus limiting the number of files staged back from tape.

- When an existing campus or wide-area SAN is deployed, it may be advisable to locate the tape library at a site remote from the SAM and Oracle Database servers. Thus, copies of data will be stored at multiple sites providing resilience in terms of distance as well as redundancy.

- Once the configuration described in this paper has been implemented, it may be used to store application data as well as RMAN backups. The ideal storage access profile for application data is to provide access to files in disk storage while they are being heavily modified (usually for a relatively short period) and then migrate them to tape storage once they are no longer being frequently modified or access is predominantly read-only.

  To address the requirements of an individual application, the archive age for the relevant data should be set in the RMAN retention policy to accommodate the period of volatility expected for the application, as the archive age determines how long a file is retained on a particular tier. The retention policy of SAM QFS should be configured so that backup images are kept at each tier for at least as long as the retention time set by the RMAN retention policy to ensure that RMAN backups of active application files aren't removed from higher tiers prematurely.

Conclusion

The solution described in this paper addresses the needs of customers with three or more mission-critical Oracle databases, 40 TB or more aggregate backup data, and a requirement to keep backups on tape for more than 180 days. It combines Oracle RMAN recovery management tools, a SAM QFS backup file system, one or more Sun ZFS Storage Appliances used to provision multiple hierarchical storage tiers, and StorageTek tape drives for lower cost, longer term storage. A Sun ZFS Storage Appliance provides all benefits of unified storage in an easy-to-use appliance package with a powerful, intuitive browser user interface and business analytics functionality that allows administrators to quickly diagnose and resolve performance issues across the storage environment.

Deployment and management of backup storage is simplified while costs are reduced through benefits such as:

- Ease of management of tiered data protection environments utilizing integrated storage management tools.
- Automated administration of backup maintenance that reduces complexity and the number of personnel involved in any major event requiring recovery of a database.
- Ability to assign control and management of database backup, restoration, and recovery tasks to a database administrator (DBA), resulting in a faster, simpler recovery process and eliminating latencies introduced when inter-team communications are required.

With the exploding growth of data, backup and recovery of Oracle Databases in an enterprise environment requires an effective and efficient data backup solution. Oracle provides an automated, hierarchically-tiered architecture that ensures the performance and availability characteristics of the backup media on which each Oracle Database backup is stored complements the value of the backup.

Appendix A. Reference Scripts and Configurations

Some sample scripts and configurations are included below for reference.

Sample SAM QFS Server mcf(4) File

The Master Configuration File (MCF) file listed below provides a sample configuration in which two StorageTek tape libraries are configured, one local to the SAM QFS server and one on a remote site. The remote tape library is accessed over a Fibre Channel network and the local tape library using a SCSI HBA.

```
# Equipment         Eq     Eq      Family    Device    Additional
# Identifier        Ord    Type    Set       State     Parameters
# -----------       ---    ----    ------    ------    ----------
#SL48:
/dev/samst/c8t500110A0008BD82Au1  10 h4  SL48    on    /SAM-QFS/catalog/SL48_cat
/dev/rmt/1cbn   11      li      SL48    on
/dev/rmt/0cbn   12      li      SL48    on
#
#
#L25:
/dev/samst/c4t0u0       20      al      L25     on     /SAM-QFS/catalog/L25_cat
/dev/rmt/3cbn   21      li      L25     on
/dev/rmt/2cbn   22      li      L25     on
#
#tier1:
tier1   30      ms      tier1   on
/dev/dsk/c1t600144F0F05E906C00004CCFF7810007d0s0        31      md      tier1   on
#tier2:
#
tier2   40      ms      tier2   on
/dev/dsk/c1t600144F0F05E906C00004CCFFAF9000Bd0s0        41      md      tier2   on
```

Sample archiver.cmd(4) File

The `archiver.cmd` file defines the retention policy for each of the tiers defined in the MCF file. It also defines the tape labels that will be used for each tape copy set and a number of global configuration parameters.

```
#
#       archiver.cmd
#
#       Generated by config api Wed Nov 24 11:12:55 2010
#
#       Global Directives
interval = 60
logfile = /var/adm/samfs-archivelog
```

```
examine = noscan
bufsize = li 512
bufsize = dk 512
archmax = li 68719476736
archmax = dk 64G
#
#
#       File System Directives
#
fs = tier1
        1 5m 13w
        2 8h 26w
        3 10h 26w

examine = noscan
logfile = /var/adm/samfs-archivelog
interval = 4m
#
#
#        Copy Parameters Directives
#
params
tier1.1 -unarchage access
tier1.1 -recycle_hwm 90
tier1.1 -sort path
tier1.1 -startage 5m
tier1.2 -drives 2
tier1.2 -unarchage access
tier1.2 -sort path
tier1.2 -startage 8h
tier1.3 -drives 2
tier1.3 -unarchage access
tier1.3 -sort none
tier1.3 -startage 10h
tier1.3 –drives 2
endparams
#
#
#        VSN Pool Directives
#
vsnpools
diskpool dk level*
L25pool li HWR*
SL48pool li [D|S]
endvsnpools
#
#
#       VSN Directives
#
```

```
vsns
tier1.1 dk .
tier1.2 li -pool SL48pool
tier1.3 li -pool L25pool
endvsns
```

Sample SAM Server *dfstab(4)* File

The `dfstab` file content below demonstrates how access to the RMAN hosts is restricted by the policy of "least access" security. The `dfstab` file resides on the SAM-QFS server.

```
share –F nfs –d "Prestwick RMAN" –o rw=prestwick /tier1/rman/prestwick
share –F nfs –d "Troon RMAN" –o rw=troon /tier1/rman/troon
share –F nfs –d "Lavemill RMAN" –o rw=lavemill /tier1/rman/lavemill
```

Sample RMAN Server *vfstab(4)* File Entry

The following extract from an RMAN server `vfstab` file shows the method for mounting SAM-QFS file systems with the options required for RMAN operation.

```
#device         device         mount          FS      fsck    mount   mount
#to mount       to fsck        point          type    pass    at boot options
#
cis:/tier1/rman/prestwick -    /rmanbackup nfs – yes \
   hard,bg,rsize=32768,wsize=32768,nointr
```

Sample RMAN Incremental Backup Script

The following script is a simple RMAN backup script intended to be run daily to provide an incremental backup.

```
Daily Script:
run {
allocate channel oem_disk_backup device type disk;
recover copy of database with tag 'ORA_OEM_LEVEL_0';
backup incremental level 1 cumulative  copies=1 for recover of copy with tag
'ORA_OEM_LEVEL_0' database;
}
```

Appendix B. Solution Component Features

The primary components of the hierarchical database backup system solution described in this paper are the Oracle RMAN recovery management tools, a SAM QFS backup file system, and one or more Sun ZFS Storage Appliances, along with StorageTek tape drives. Features and capabilities of several of these components are described below.

Oracle Recovery Manager

The Oracle Recovery Manager (Oracle RMAN) is a command-line and enterprise manager-based tool that serves as the Oracle method for efficiently backing up and recovering Oracle databases. It is tightly integrated into the Oracle Database, providing block-level corruption detection during backup and restore.

Oracle RMAN optimizes performance and space consumption during backup using file multiplexing and compression. Oracle RMAN offers:

- **Native integration with the server** – Oracle RMAN is aware of the various data structures that the database uses, so the appropriate backup and recovery operations can be performed correctly. During backup and recovery operations, Oracle RMAN validates all blocks for corruptions. If corruption is found, Oracle RMAN can quickly and easily recover the corrupted blocks. Oracle RMAN backups can be performed while the database is online. Additional redo operations are not generated during the backup, as is the case for operating system backups when the database is in hot-backup mode.
- **Incremental backups** – Oracle RMAN optimizes incremental backup performance by only backing up blocks that have changed since the last full backup. This results in faster backup performance and smaller backup sizes.
- **Incrementally updated backups** – Incremental backups can be rolled into data file image copies in place, thus eliminating the need to apply incremental backups on recovery. The resultant image copies can be utilized as a newer full backup. This reduces overall recovery time, in addition to reducing the need to make full backups.
- **User-defined scripts** – Common Oracle RMAN tasks can be scripted to automate backup and restore tasks. Oracle RMAN stores and manages these scripts in a central repository, so that they can be applied to more than one database. In a large environment with multiple databases, this provides for ease of use, uniformity, consistency, and accuracy of backup and recovery operations.
- **Comprehensive reporting capabilities** – The Oracle RMAN repository contains detailed records about all aspects of backup and recovery, such as names, dates, sizes, and locations of backups. Oracle RMAN also keeps track of which backups are obsolete and which are required. In the event of data loss, being able to quickly and correctly determine which backups are needed for recovery is critical.

- **Centralized management with Enterprise Manager** – Oracle RMAN is integrated with Oracle Enterprise Manager, which allows the database administrator to manage all backup and recovery activities from a central Web-based console.
- **Scalable, tested, and proven deployments –** Oracle RMAN scales with the available hardware and performs operations in parallel to provide the best possible performance.

Oracle RMAN is Oracle's supported method for database back up to disk, a key component of an automatically tier-migrated storage solution.

Sun Quick File System and Storage Archive Manager

The Sun Quick File System (QFS) and Storage Archive Manager (SAM) software provide data classification, centralized metadata management, dynamic policy-based data placement, protection, migration, long-term retention, and recovery capabilities.

To help organizations effectively manage and utilize vast data repositories consistent with business requirements, SAM provides a powerful, easily managed, cost-effective way to access, retain, and protect business data over its entire lifecycle with continuous backup and fast recovery features to help enhance productivity and improve resource utilization.

The SAM Hierarchical Storage Management (HSM) feature provides the ability to control storage based on tiers defined by the system administrator. Copies of data are maintained in multiple tiers, and the logical migration of data through the tiers is managed based on policies. Storage tiers are defined by access performance, capacity, and resilience to component failure. Typically, Tier 1 comprises the highest performance and most resilient storage, but with a relatively low capacity. As the tier number increases, performance, and resilience decrease while capacity increases.

SAM is configured to copy data to different tiers based on file characteristics, which include file access time, creation time, owner, group, and name. Once a file ages to 'n' minutes (its archive age), a copy is made. Policies can be defined for multiple copies of the same file using different selection criteria.

A SAM QFS file system is built on underlying block storage. A highly suitable platform is comprised of one or more Sun ZFS Storage Appliances, which can be configured into multiple storage tiers.

Sun ZFS Storage Appliance

The Sun ZFS Storage Appliances are a family of unified storage systems that offer enterprise-class data services, massive scale, and industry-leading performance, while delivering significant cost savings. Provisioning and management are simplified through the easy-to-use browser interface to the underlying Oracle Solaris ZFS file system, which takes the guesswork out of system installation, configuration, and tuning. DTrace Analytics software provides a comprehensive, intuitive analytics environment using built-in instrumentation to provide in-depth analysis of key storage subsystems.

Administrators have all of the tools they need to quickly identify and diagnose system performance issues and debug live storage and networking problems before they become challenging for the entire network. SUN ZFS Storage Appliances also include the comprehensive self-healing capabilities of

Oracle's Fault Management Architecture (FMA). FMA automatically and silently detects and diagnoses underlying system problems and automatically responds by taking faulty components offline and notifying systems administrators.

To deliver high performance using cost-effective components, the Oracle Solaris ZFS file system is organized into Hybrid Storage Pools (HSPs) that enable seamless access to different types of media. Oracle Solaris ZFS is designed to automatically recognize different I/O patterns and place data in the best storage media for optimal performance, while helping lower power and cooling requirements.

The appliances leverage volume servers along with densely packaged high-capacity disk drives and can also be configured to utilize solid-state drives (SSDs). Leveraging an industry standard server in place of an expensive, proprietary disk controller reduces both cost and risk.

These components are integrated into an appliance framework that simplifies deployment and includes an integrated data services software stack at no extra cost. The rich set of data services includes compression, data deduplication, mirroring, snapshot, error correction, and system management features that help simplify ongoing data management.

See *Appendix C. References* for comparative study by the Edison Group that demonstrates the cost of ownership and the business value advantages of choosing Oracle's Sun ZFS Storage Appliance family of unified storage devices.

Appendix C. References

Oracle Database 11g Direct NFS Client by W. Hodak and K. Closson:
http://www.oracle.com/technetwork/articles/directnfsclient-11gr1-twp-129785.pdf

Oracle RMAN Disk-to-Disk Backup Methods Using Sun Storage Array, Oracle white paper:
http://www.oracle.com/technetwork/articles/systems-hardware-architecture/rman-disk-to-disk-backup-170902.pdf

Oracle Sun Storage Archive Manager web site:
http://www.oracle.com/us/products/servers-storage/storage/storage-software/031715.htm

StorageTek Tape Storage web site:
http://www.oracle.com/us/products/servers-storage/storage/tape-storage/index.html

Sun Unified Storage web site:
http://www.oracle.com/us/products/servers-storage/storage/unified-storage/index.html

Oracle ZFS Storage Appliance Comparative Management Costs Study by the Edison Group:
http://www.theedison.com/pdf/2010_Samples_Oracle_ZFS_Storage.pdf

ORACLE

Oracle is committed to developing practices and products that help protect the environment

Creating a Hierarchical Database Backup System using Oracle RMAN and Oracle SAM QFS with the Sun ZFS Storage Appliance
August 2011, Version 1.0
Author: Andrew Ness

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

oracle.com

Hardware and Software, Engineered to Work Together