An Oracle White Paper
July 2010

Optimizing Information with a Tiered Storage Architecture

ORACLE

Disclaimer

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.

Executive Overview

As you address data management needs, whether to make information available for long periods of time, meet legal or corporate regulations or to keep information safe, secure and accessible, a solution is needed that is not only dynamic and sustainable but also cost-effective. Oracle's approach to the use of tiered storage provides a solution that optimizes performance, cost and capacity across the breadth of storage media available today.

According to Fred Moore, President of Horison Information Strategies, "Storing information is one thing. Accessing information is everything." Tiered storage that is dynamically managed provides the most cost effective yet high performance method of storing information for both its most aggressive access requirements and its longest point of retention, which could be forever.

Introduction

This paper is an introduction to the concept of tiered storage, the business needs that drive it and the value of implementing a tiered storage strategy. This strategy helps realize the value of information, drives down costs, lowers risks, improves accessibility, ensures the correct information is saved and protected and even takes advantage of technology changes into the future.

Implementing a tiered storage strategy by utilizing the features of Oracle's Storage Archive Manager (SAM) to dynamically manage different types of storage are highlighted in the context of meeting the business requirements of storing information.

Business Needs for Implementing Tiered Storage Solutions

Information Is An Asset

Successful businesses consider information to be a strategic company asset. Examples of information include invoices, user guides, design documents, scientific research data, video surveillance, healthcare records, MRI images, weather data, digital images of ancient books, and contracts. Access to information, no matter how old, helps businesses gain a competitive edge by making intelligent decisions, improving support services and delivering products at an increased pace. Collaboration of information outside the line of business as well as outside the company also adds value to the

information; further improving business decisions and data analysis. For a shirt manufacturing company, an example would be: "should I make and ship long sleeve sweat shirts for early July because last year we ran out?". Looking at weather data from last year, the forecast for this year and current long-sleeve shirt run rates, derives the best answer. The conclusion will be the optimal number of shirts required to manufacture and deliver to optimize revenue and reduce waste. This strategy of collaboration of data in a timely manner results in success across the complete supply chain.

Legal and Corporate Regulations

The growth of data and ability to store and access information opens a business to laws and regulations governing the security and mandated retention of both business and client information. Some data must be kept for a very specific length of time and all copies destroyed exactly when that time is over. Some data may be put on legal hold and be part of litigation activities that could go on for many years, forcing those data to be kept long after its normal retention. Other data, such as digital archives and weather data, has 'forever' retention. While similarly, healthcare data must be kept longer than the life of the patient. Data is an asset and should be managed within the law or it could become a significant legal and financial liability.

Huge and Growing Capacity Requirements

The huge capacity requirements resulting from the vast use and access of digital content leads to the need to store and manage all of this information as cost effectively as possible. It must be kept accessible now and in the future and therefore be moved through storage and format technology changes. In a business, people come and go, hardware changes, software is upgraded but the data lives on. Moving that information through all of this change dynamically and without interruption brings additional challenges.

The above factors, - information value, regulations and capacity requirements - mandate a controlled, intelligent, dynamic approach to the use of tiered storage.

Data Characteristics When Using Tiered Storage

The characteristics of data determine if it can be classified for tiered storage. This applies to all industries and includes business data as well as industry data. File sizes can be billions of small files or millions of large Gigabyte or even Terabyte size files, both resulting in the need for managing petabytes of storage. Transactional or frequently changing data is typically a small component of the overall storage requirement, generally less than 20%, and takes advantage of tiered storage by leveraging high speed flash or solid state drives (SSDs) and high speed disk storage. However, this paper will focus on the vast majority of a company's information that is less transactional in nature and needs to be securely stored and accessed in a more cost effective manner. Some of this information, such as MRI Images, scans of book pages and video is never changed throughout its lifecycle, starting from its point of creation. Other information, such as current orders or inventory reports, frequently change during its workflow process and is then 'closed', given a status of static, and can then be archived.

Tiered Storage Architecture

The use of tiered storage started in Mainframe datacenters that were faced with the same problem at a different scale at the time. Even in the days when a megabyte of data was large, the ability to store less active information on more cost effective storage was a requirement. Then, there were two types of storage media, disk and tape. Today, there are numerous types of storage media such as SSD, several types of disk storage with varying I/O characteristics and different types of tape drives and libraries with varying features such as encryption and compression. Each type of storage is assigned to a storage tier.

With this storage complexity and increase in capacity requirements comes the need for an intelligent way to dynamically manage where the information is stored as cost effectively as possible. The cost is not just the CAPEX of the storage but also the ability to manage the movement of information through the storage tiers and to new storage technology with as little human intervention as possible to keep the OPEX costs down as well. Ultimately, without dynamic storage management, the different tiers of storage become disparate storage silos and detract from the end-goals of efficiency, cost effectiveness, shared access and security.

Data management is composed of 4 basic layers. Each layer provides its own independent contribution to the whole. For a sustainable strategy, each of the layers can change and the data will remain available and accessible. Figure 1 describes these four basic layers.

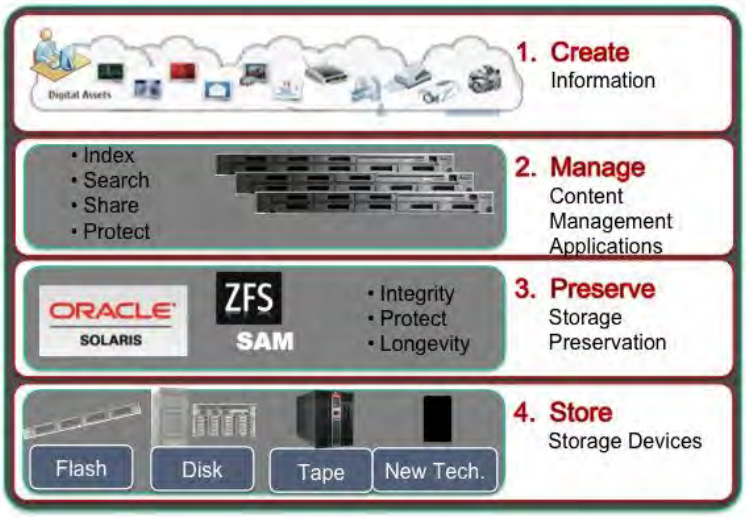

Figure 1: Data Management Architecture

Layer 1 is the source of information. It can be a desktop and a person typing and saving files or a machine that generates information at a rate of 1 terabyte a day or more.

Layer 2 is managing the information from a content point of view. Building metadata that describes the content and makes it searchable and sharable.

Layer 3 is the storage preservation layer, providing the required data availability and sustainability. Key features include integrity checking, preemptive error detection and correction, maintaining multiple copies based upon policy.

Layer 4 is the physical storage device(s) where the information is stored. This includes: high speed disks for metadata and active content, cost or capacity optimized lower performing disks for less active content, and tape for long term retention, least active content and often to provide multiple copies as required.

SAM Architecture and Functionality

Included in the Storage Preservation Layer, the Sun Storage Archive Manager (SAM) and QFS file system. These components work together as a hierarchical storage management file system to dynamically manage the data on the different tiers of storage based on policies. Files are archived for long-term, compliant retention by copying them from the online file system to disk and/or tape archive media using the gnu tar format. Up to four archive copies can be written to disk or tape archival media, including remote sites. These copies can be processed in parallel or at different times, according to policy rules. The archiving file system works seamlessly with automated tape libraries.

The file system is designed to provide high performance in a shared disk, heterogeneous cluster environment. Each node in the cluster has direct access to the shared disk, providing concurrent, low latency, and high bandwidth I/O for applications.

A single SAM QFS file system can scale to 4PB in size, limited only by the number of devices available to a system. SAM can then manage many 100s of PBs of data through the use of tiered storage.

The SAM file system automatically manages the life cycle of the data through five distinct Functions:

- **Archiver**

  Transparently and dynamically copy data from the primary disk location to other disk or tape devices without operator intervention. The primary disk location, often called disk cache, is the SAM file system. The target of the copies can be a file system, an NFS mounted file system such as Oracle's Unified Storage System Sun Storage 7000, a remotely connected SAM-FS server or it can be a tape system. This function uses policies based on path name, wildcard, size, age, owner, group, or date and can copy data to up to 4 different devices, including remote sites.

- **Releaser**

Automatically, based on policy, manage the release of files from the primary disk location after the files have been copied. The list of eligible files is prioritized based on policies such as archive status, size, release status and age.

- **Stager**

  Automatically restore a released file back to disk cache when an application or user accesses the file. Staging options include pre-staged and bypassing the primary disk cache. Removable media (tape) access is optimized for mounting and positioning.

- **Recycler**

  Recycling repacks archive media to reclaim space. Reclaimed space is based on retention policies. The recycling process can be used to migrate from older storage technologies to newer technologies.

- **Access**

  When an application or user accesses a data file, the data is automatically and transparently retrieved from the media on which it is located. Although there could be a slight latency when retrieving from removable media, the data is automatically retrieved with no human intervention required.

The functions listed above enable SAM to provide both data protection functionality as well as an archive functionality. The result is you can combine data protection, archive and storage tiering in a common architecture.

SAM Metadata

SAM uses existing, standard, posix compliant filesystem metadata with an extended set of SAM metadata that includes the actual location(s) of the file as well as additional SAM attributes. This extended SAM metadata can be stored on high speed, primary storage, possibly even flash, since this metadata requires only a small capacity. Since it is also stored in the file, it can be rebuilt should it be inadvertently destroyed. Until a user or application accesses the file that the SAM metadata is managing, the file remains in its policy driven storage location. Applications can access the SAM API to retrieve the status and location of the file content; therefore, if there is an expected latency to access a file, an appropriate message can be presented to the user while the file is being staged.

SAM running on multiple storage clients can copy and stage files to a centralized tape library, locally or in remote sites, facilitating automated management of data for disaster recovery. Communication between the remote clients and sever uses TCP/IP, offering secure and flexible methods for data movement between sites. Figure 1 shows a high level view of a configuration for a SAM file system and its related copies.

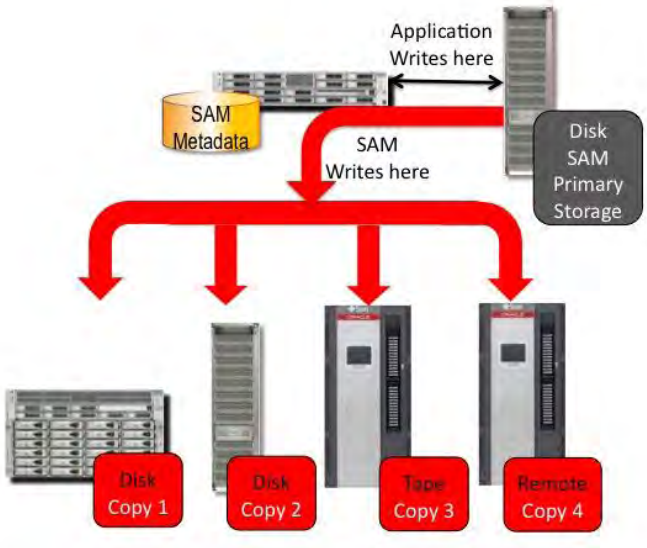

Figure 2: File Archive Options

The use of tiered storage, dynamically managed by SAM, allows businesses to turn their data into an asset, meet compliance regulations, seamlessly leverage advances in storage media and store a growing amount of data at a sustainable cost.

Use Cases: Putting SAM to Work for Applications

Using tiered storage with Content Management and Archive applications makes sense on several different levels. Businesses are now able to get to know their data, matching data attributes, which determine data value, with the most cost-effective technologies in the storage hierarchy. The SAM architecture is an ideal data mover for managing data in tiered storage. Up to four copies of the data can be stored to disk, tape and remote sites. Since this data often has to be retained for periods of time longer then the life of the media, migration to new media and new devices is necessary. SAM can accomplish this dynamically, overtime, without disruption to the applications.

Use cases for Today's Tiered Storage Infrastructure

Tiered storage managed by Sun Storage Archive Manager can be utilized as the storage infrastructure for several different use cases.

- Use case one is the integration of SAM with Oracle Universal Content Manager (UCM). The primary file store is Oracle SecureFiles with UCM using its records management features to move records into a SAM file system for further archiving through tiers including tape. UCM is the primary manager of all of the content.
- Use case two is to archive Oracle Recovery Manager (RMAN) backup disk images to tiered storage making it feasible to maintain years of backups on tape while RMAN only has to manage a posix file system disk image, not knowing it is on tape.

- Use case three is to move files to tiered storage using an archive application. As in use case one, many years of data can be maintained with online access while utilizing lower cost tape as its storage media. All the metadata and full text indexing remains searchable and accessible in a database while the content is safely stored on disk, tape and/or at remote sites.

SAM Use Case 1: Oracle Universal Content Management

Including SAM as an integral component of the storage infrastructure, the Oracle Universal Content Management application can dynamically manage unstructured content on multiple tiers of storage keeping active data in Oracle SecureFiles. The content metadata is stored in an Oracle database, which includes the content file store location. UCM provides the policies that can move the data from file store 1 to file store 2, which is a SAM file system on disk. SAM will then further archive the content to disk, tape and/or a remote site.

The following diagram describes the data flow of information through the architecture of this solution.

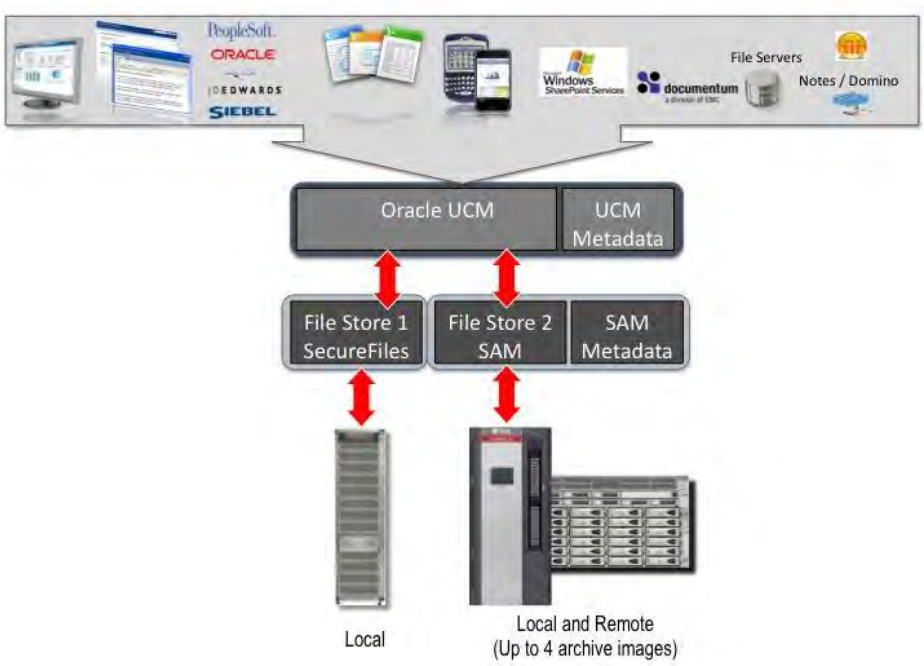

Figure 3: Use Case 1: UCM Enabling SAM Integration with SecureFiles

SAM Use Case 2: Archive RMAN Backup Images from Disk to Tape

Oracle RMAN is Oracle's utility for managing the backup and, more importantly, the recovery of the database. It eliminates operational complexity while providing superior database performance and availability. Oracle RMAN determines the most efficient method of executing the requested backup,

restore, or recovery operation and then executes the operation in concert with the Oracle Database server. It provides capabilities such as incremental backup and corrupt block detection while taking care of backup, block media recovery, and other tasks to ease the backup-and-recovery administration of an Oracle Database instance.

When used in concert with Oracle Sun Storage Archive Manager as the destination device for the RMAN Backup related files, SAM will dynamically archive all related RMAN images to tape, with no operator or DBA intervention. Policies are defined when SAM is installed as the RMAN BU device and from that point on, will automatically archive to disk and/or tape. SAM will also stage RMAN files back to disk without operator or DBA intervention when recovery is necessary. Although the Database Server has the ability to age out older RMAN related files, it will now be possible to keep many more versions available on a smaller disk space reducing the overall cost of backup. The DBA will never again receive a 'Recovery Area Alert Handling' message further reducing management costs. As new Database Servers are added, the shared RMAN BU space is available for use further reducing the storage costs though consolidation.

In the following figure, the dataflow diagram shows all files related to the backup are written to the SAM disk cache.

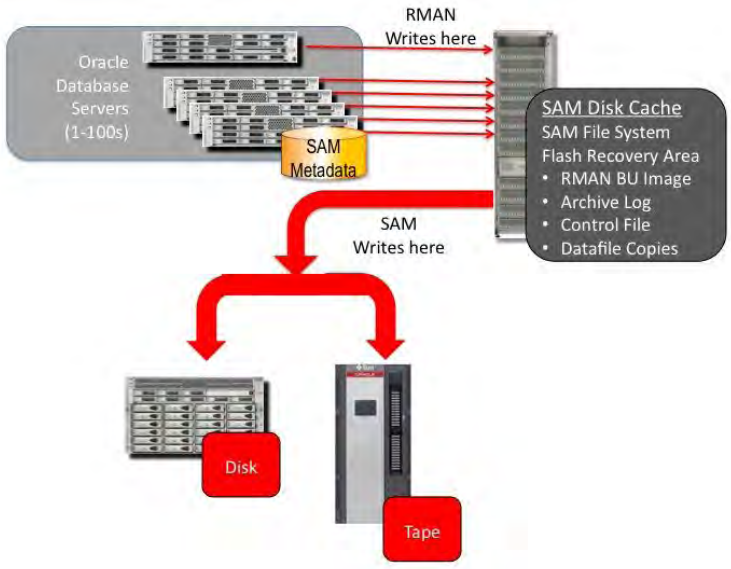

Figure 4: Use Case 2: RMAN, Flash Recovery Area, SAM and Tiered Storage

SAM Use Case 3: Archive Applications using Tiered Storage

A wide choice of Archive Applications have been qualified using tiered storage architectures that include Oracle's Sun Storage Arrays, Open Storage and Tape. The data is dynamically managed between tiers of storage based on the policies defined in SAM. The archive application metadata, used to store, search, access and link data is stored in an Oracle Database on high performing storage systems, often consisting of flash storage. Figure 5 is the basic architecture used by the majority of these archive applications.

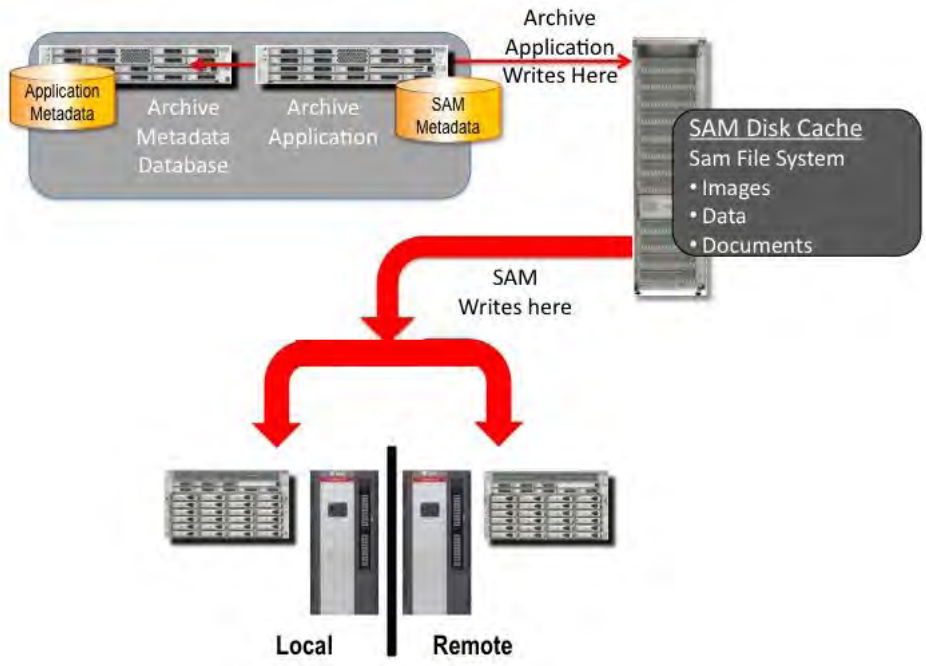

Figure 5: Use Case 3: Archive Application using SAM File System

Many of the software solutions in this use case are industry-specific. For example, the Scalable Online Archive and Repository (SOAR) solution bundle is integrated and delivered by ARROW ECS. This solution uniquely addresses enterprise IT and line of business managers' need for "appliance" –like storage repositories for ever-growing digital information. SOAR is delivered and supported through a network of Oracle's Sun value-add reseller partners.

The digital repository software (collectively referred to as Islandora) is an open source digital asset management framework with tools for loading data and providing access to that data via the web-based interface. Islandora combines the Drupal and DuraSpace Fedora Repository software applications, creating a robust solution for managing data-intensive digital repositories for Universities, Research Centers and Government Centers. Integration with tiered storage delivers the performance of high-end storage for ingest and search, scales capacity to petabytes and delivers advanced data protection, preservation and management provided by SAM.

Other industry specific applications are available, both open source with subscription services for support and commercial applications with standard warranties. Both offer a manageable environment for archive of up to many Petabytes. The key to lowering management costs of archive solutions is through easily managed tiered storage including low cost tape systems. This provides the archive applications with recommended configurations and standard access, which are seamlessly scalable at a low cost.

Summary

The integration of a content management solution that manages information based on its content and a solution that dynamically manages its location on tiered storage based on policy solves the problem of information growth as well as bringing value to that information that can directly affect the bottom line. Oracle brings unique value through fully integrated solutions that are standards-based hardware and software, high performance and highly scalable and provides support for the complete information lifecycle, taking information from creation to archive to destruction efficiently and effectively.

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

Optimizing Information with a Tiered Storage Architecture
July 2010 PN: A01-0003-10
Author: Donna Harland

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

SOFTWARE. HARDWARE. COMPLETE.