ORACLE
OPTIMIZED SOLUTIONS

Maximize Data Access Throughout Its Lifecycle With Extreme Scalability

The Oracle Optimized Solution for Lifecycle Content Management

ORACLE WHITE PAPER | FEBRUARY 2015

ORACLE

Table of Contents

Introduction

In today's competitive business environment, enterprisewide collaboration is one of the keys to success. As the lifeblood of collaboration, data has potentially infinite business value. To enable successful products and corporate strategy development, both private and shared content must be collected, easily accessed, collaborated upon, analyzed, evaluated, and reused in an endless cycle. Content management is more than merely collecting the data when it is created. To be useful, information must be easily searchable with reliable and dependable access, even over long retention periods, which in some cases could be forever. Delivering scalability and longevity to enterprise data requires a content management solution that is not only dynamic and sustainable but also cost effective and standards compliant.

Complicating matters, data can have vastly differing management requirements. Some data is transient, needing to exist for only seconds, minutes, or hours. Other types of data must be kept accessible for months, years, or even indefinitely. Investments in a content management solution that stores and accesses data successfully must also be flexible and scalable in order to nondisruptively grow with the organization and address technology changes through time. Ultimately, organizations need to manage unstructured content with the capability to search and access 100 percent of the content, no matter where it is stored, how old it is, or what its original format.

This paper describes the Oracle® Optimized Solution for Lifecycle Content Management, with a focus on how to greatly simplify infrastructure deployment and management. Guidelines are provided for component selection based on performance and capacity requirements for small, medium, and large alternative reference configurations. Suggestions are also provided for storage best practices to deliver a virtualized, scalable, flexible, and yet cost-effective solution that meets the most demanding content management needs. Additional information about the benefits of simplification and scalability for users of content management systems as well as IT organizations can be found in the companion business paper, "Maximize Data Access Through Extreme Scalability."

Solution Objectives

Oracle Optimized Solution for Lifecycle Content Management integrates Oracle WebCenter Content with a tiered storage infrastructure that is managed by Oracle Hierarchical Storage Manager (Oracle HSM). The solution is designed to accomplish the following objectives:

- *Increase information value* through dynamic access of data. Content-based collaboration now can occur both inside and outside the company enabling faster product time to market. Dynamic access also produces faster time-to-insight for better business decisions and for product validation, even for information that might be many years old.
- *Lower business risk* by reducing the complexity of managing tiered storage. Multiple copies of the content can be maintained in different physical locations providing assurance that any data needed in the future will be available. Content management infrastructure can grow easily with little or no downtime as capacity and access requirements grow.
- *Increase employee productivity* by enabling IT personal to spend less time testing and more time executing the corporate IT strategy. Users spend less time looking for data and more time analyzing data. Easier collaboration can also improve employee productivity.
- *Enable cost savings* by deploying multiple tiers of storage, including archival to tape, 100 percent of the content remains dynamically available on demand. In addition, pretested solutions and documented components result in much shorter proof of concept (POC) times that can speed up solution adoption.

Architecture Overview

Oracle WebCenter Content is the primary software application in the latest generation of the Oracle Optimized Solution for Lifecycle Content Management (Figure 1) and takes advantage of the perfect technology match of the Oracle HSM software application with Oracle's storage products. Oracle Optimized Solution for Lifecycle Content Management is applicable to many different use cases, so it demands an infrastructure that scales independently in terms of both performance and capacity. Oracle's SPARC T-Series servers running Oracle Solaris provides an integrated platform that delivers deployment flexibility and high performance with a powerful and consistent methodology for creating logically isolated virtualized server environments. This flexible allocation of server resources helps control costs for both application and hardware requirements. Oracle WebCenter Content runs on the virtualized SPARC T-Series server infrastructure, providing document and imaging management, web content management, digital asset management, and records and retention management. Oracle HSM also runs on the virtualized SPARC T-Series server infrastructure and is the critical component of the data storage because it provides automatic and dynamic access to content, while hiding the complexity of the storage devices and presenting a single file system to applications and users. Thus, Oracle HSM provides transparent access regardless of where data is stored, while delivering all the benefits of performance, scalability and flexibility across multiple storage tiers.

Oracle ZFS Storage Appliance or Oracle FS1-2 storage system both can provide highly scalable and flexible disk storage platforms for the solution. As with previous generations of this solution, Oracle's StorageTek modular library systems continue to provide long-term storage preservation enabling nondisruptive expansion.

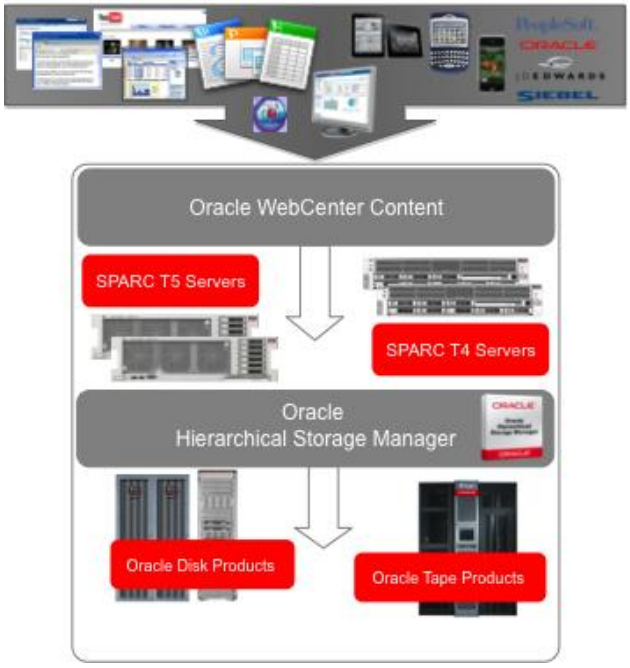

Figure 1. Oracle Optimized Solution for Lifecycle Content Management integrates Oracle WebCenter Content and Oracle HSM with Oracle's SPARC T-Series servers as well as Oracle's disk and tape storage products.

The solution architecture can be logically divided into the following categories:

- **Tiered storage.** Tiered storage is comprised of Oracle HSM together with Oracle's storage devices that provide storage areas for primary disk, disk archive, and tape archive.
- **Software application infrastructure**. Oracle WebCenter Content is the primary software application for the solution, providing organizations with a unified repository to house unstructured content and delivering it to business users in the proper format—all within context of familiar applications to fit the way they work.
- **Server infrastructure**. The content management applications are deployed on a modular server architecture based on the SPARC T-Series servers.

The following subsections provide an overview of these three major components of the architecture.

Tiered Storage

Tiered storage is a feature that accommodates diverse storage methods for different data storage needs and is implemented in Oracle's hardware products to provide various levels of storage access performance. Some content must be kept for long periods, while some use cases require fast ingest as well as fast access for recently ingested data. Oracle ZFS Storage Appliance and Oracle FS1-2 storage systems are used to hold the recently stored and most active data, yet they can also provide disk archiving in the same storage system. Oracle's StorageTek tape products are used for archival data that requires only occasional access, offering simple and cost-effective data protection through multiple on-tape copies.

Oracle Hierarchical Storage Manager

Oracle HSM is the critical software component of the tiered storage architecture because it ties together the different levels of storage hardware performance, known as storage tiers, by providing automatic and dynamic access to content from any storage device, while presenting a single file system to applications and users. Oracle HSM hides the complexity of the tiered storage infrastructure, and provides transparent access regardless of where data is

stored. Oracle HSM likewise helps simplify management, while providing all the benefits of scalability and flexibility across multiple storage tiers. Oracle HSM is a storage software application that runs on Oracle Solaris and requires Oracle Solaris Cluster for an active/passive high availability (HA) environment on two SPARC T-Series servers.

Table 1 provides an overview of the components of Oracle HSM and the specific version numbers used in the Oracle Optimized Solution for Lifecycle Content Management.

TABLE 1. ORACLE HSM SOFTWARE COMPONENTS

	SOFTWARE	RELEASE
ORACLE HSM SERVERS	Oracle Solaris	Oracle Solaris 11U1
	Oracle HSM	5.3-01
	Oracle Solaris Cluster	4.1

Oracle HSM accesses content from the primary disk cache based on preset policies, and creates copies on archive disk and/or tape devices. Oracle HSM then can dynamically allow access to the content from any device. For data protection, up to four copies can be made available, both locally and remotely, possibly eliminating backup requirements.

Within the Oracle HSM file system primary storage can scale to 4 petabytes (PB) in size. However, the capacity under Oracle HSM management can easily reach hundreds of petabytes through the use of tape media. The archiving file system policies within Oracle HSM automatically manage the lifecycle of the archive data.

- The *Archive Policy* transparently archives data from the disk cache to archive disk and/or tape without operator intervention. The Oracle HSM archiver uses policies based on file system characteristics, such as path name, wildcard, size, age, owner, group, or date to automatically manage the copies.
- The *Release Policy* automatically manages the disk cache and releases files that have been archived when the high capacity threshold is reached on the primary storage. The list of files eligible to be released is prioritized based on policies such as archive status, size, release status, and age.
- The *Stage Policy* automatically stages released files back to the disk cache or directly to the requesting application as files are accessed. Staging options include prestaging and bypassing the disk cache. Removable media access is optimized for mounting and positioning.
- The *Recycle Policy* repacks archive media onto new media in order to reclaim space. The recycling process can be used to migrate from older to newer technology.

Throughout a file's lifecycle, the relevant Oracle HSM metadata remains online and available to the content management application. All files appear to be directly located on the disk cache when they might actually exist only on tape. The result is cost-effective management and use of tiered storage while providing dynamic and immediate direct access to 100 percent of the data without operator intervention or human knowledge about where the data actually resides. Users then have access to data that might be many years old or might not have been accessed in many years.

More information on Oracle HSM can be found at the following URL:

http://www.oracle.com/us/products/servers-storage/storage/storage-software/storage-archive-manager/overview/index.html

Oracle ZFS Storage Appliance

The Oracle ZFS Storage Appliance family consists of enterprise tier 1 storage systems that radically simplify and accelerate storage management by providing simple installation, administration, and capacity scaling. Oracle ZFS Storage Appliance systems are available in a range of configurations to meet a variety of customer needs for capacity, price, and performance. Problem-solving capabilities are easily supported with the integrated DTrace Analytics feature. All configurations come bundled with the same software including data protocols, compression, and DTrace Analytics software for system troubleshooting and performance optimization. Oracle ZFS Storage ZS3-4 is ideal for large enterprises requiring both high-performance storage as well as high capacity for disk archive, with a dramatically easier and faster way to manage and scale. For smaller configurations, Oracle ZFS Storage ZS3-4 redefines midrange storage for enterprises, with simplified management, performance, efficiency, and seamless expansion to meet growth needs.

The Sun ZFS Storage ZS3-4 appliance scales up to 2.59 PB of raw storage capacity, and supports the hybrid storage pool architecture delivering the needed performance for ingesting data from sources such as Oracle WebCenter Content. The Oracle ZFS Storage ZS3-2 appliance provides a high-availability, entry-level cluster option with scalability up to 432 terabytes (TB) of raw capacity, and also supports hybrid storage pools. Both the Oracle ZFS ZS3-2 and ZS3-4 appliances also can be configured with up to 4 TB of read-optimized cache and optional write-optimized cache for enhanced application performance. The combination of read- and write-optimized solid state devices (SSDs) with high performance hard disk drives (HDDs) delivers optimal performance for primary storage for Oracle HSM while the high capacity HDDs meet the requirements of a disk archive (Figure 2). The hybrid storage pool architecture also provides both SAN and NAS interfaces, supporting the requirements of Oracle HSM as it manages multiple copies of the data in a tiered storage system. For even the highest performance tier in a tiered storage environment, the hybrid storage pool introduces an easy method of writing data to the correct high-speed storage without user or storage administrators having to make decisions. To deliver high performance, Oracle Solaris ZFS running on Oracle ZFS Storage Appliance seamlessly recognizes I/O patterns automatically and places data on the best storage media using hybrid storage pools.

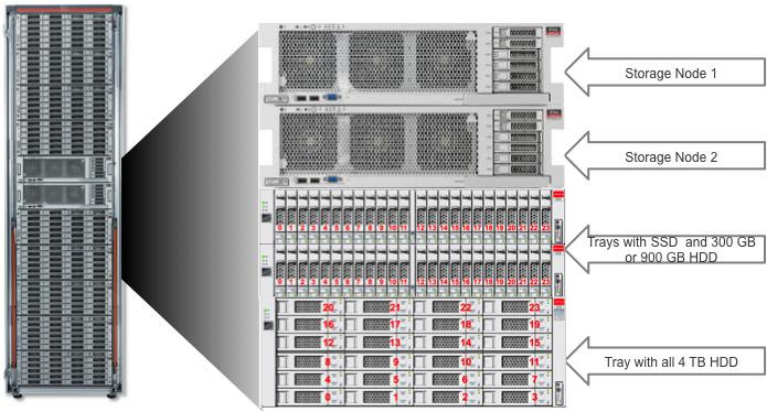

Figure 2. The Sun ZFS Storage ZS3-2 delivers the flexibility required for a tiered storage environment with two storage nodes, a tray with all HDDs, and a tray with a mix of SSDs and HDDs.

For example, Oracle Solaris ZFS transparently executes writes to low-latency SSDs so that writes can be acknowledged quickly, allowing the application to continue processing. Oracle Solaris ZFS then automatically flushes the data to hard disk drives as a background task. The dynamic use of the SSDs by Oracle ZFS Storage Appliance adds to the range of tiered storage that is managed by Oracle HSM.

Provisioning and management are dramatically simplified in Oracle ZFS Storage Appliance through the browser user interface (BUI) that takes the guesswork out of system installation, configuration, and tuning. The BUI provides an intuitive environment for administration tasks, visualizing concepts, and analyzing performance data. The management software is designed to be fully featured and functional on a variety of web browsers.

Unprecedented observability is provided with DTrace Analytics software, a feature of Oracle ZFS Storage Appliance. It provides the industry's only comprehensive and intuitive analytics environment. Administrators have all of the tools they need to quickly identify and diagnose system performance issues, perform capacity planning, and debug live storage and networking problems before they become challenging for the entire network. The real-time analysis and monitoring functionality utilizes built-in instrumentation to provide in-depth analysis of key storage subsystems, helping administrators quickly identify the source of performance bottlenecks Figure 3 shows an example of the type of data that can be visualized with DTrace Analytics. The top part of the screen shows the CPU utilization across four CPUs, and the bottom graph highlights NFS throughput in operations per second broken down by latency. The information depicted can be used to identify CPU or I/O bottlenecks across the storage systems providing the capability to drill down into the applications running on those systems for more information. The high level of observability provided by DTrace Analytics is an industry first and enables the identification of application bottlenecks that might previously have gone undetected.

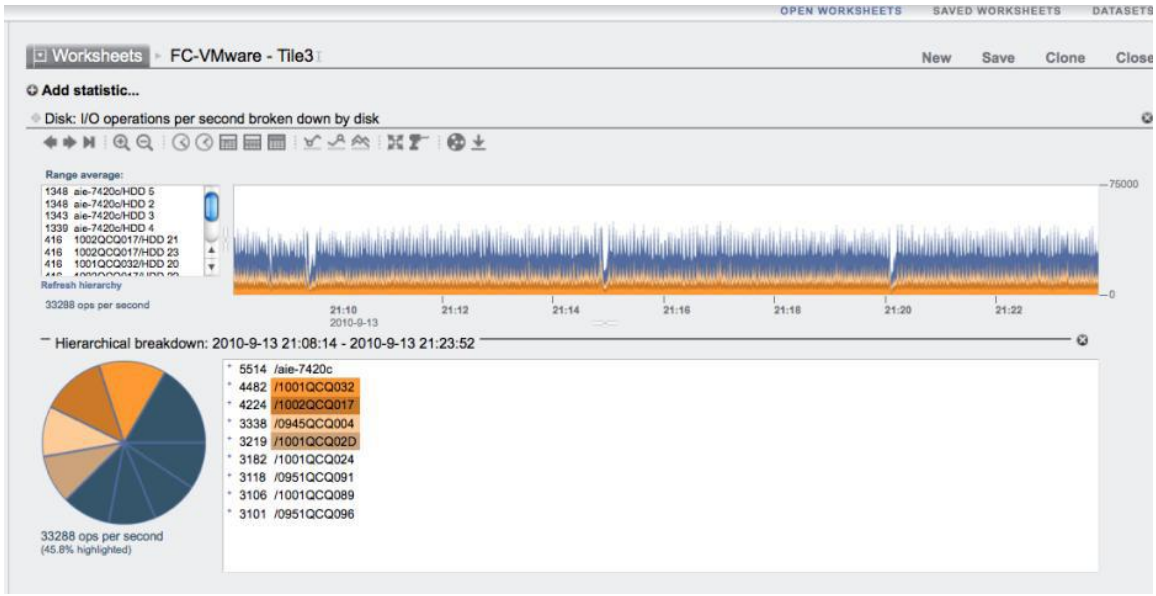

Figure 3. Oracle ZFS Storage Appliance browser user interface displays real-time analytics.

More information on Oracle ZFS Storage Appliance can be found at the following URL:
http://www.oracle.com/us/products/servers-storage/storage/nas/overview/index.html

Oracle FS1-2 Storage System

Oracle FS1-2 is a modular enterprise-class storage system that provides multiple levels of performance and capacity. The flexibility provided by these features makes the Oracle FS1-2 storage system a good fit for content

management, serving both very active content with possible high ingest rates, as well as content that is older or accessed less frequently, but still has high value.

In this optimized solution, three classes of storage can be implemented on separate logical unit numbers (LUNs) in the Oracle FS1-2 storage system. Metadata should be stored in the highest-priority storage for fastest access, while two additional storage classes can be provided for Oracle HSM primary storage and Oracle HSM archive storage. Figure 4 further defines these storage classes, which include:

» **Premium level**. Oracle HSM metadata, database metadata, and database redo logs, as shown in red.
» **High level.** For the Oracle HSM primary disk, as shown in green.
» **Medium, low, or archive level.** For Oracle HSM disk archive, as shown in blue.

The multiple levels of performance and capacity provides flexibility making the Oracle FS1-2 storage system an excellent storage solution with Oracle HSM for managing unstructured data. A single storage system with the three storage classes and five priority levels of service delivers all tiers of disk storage in a single storage system with a single management tool.

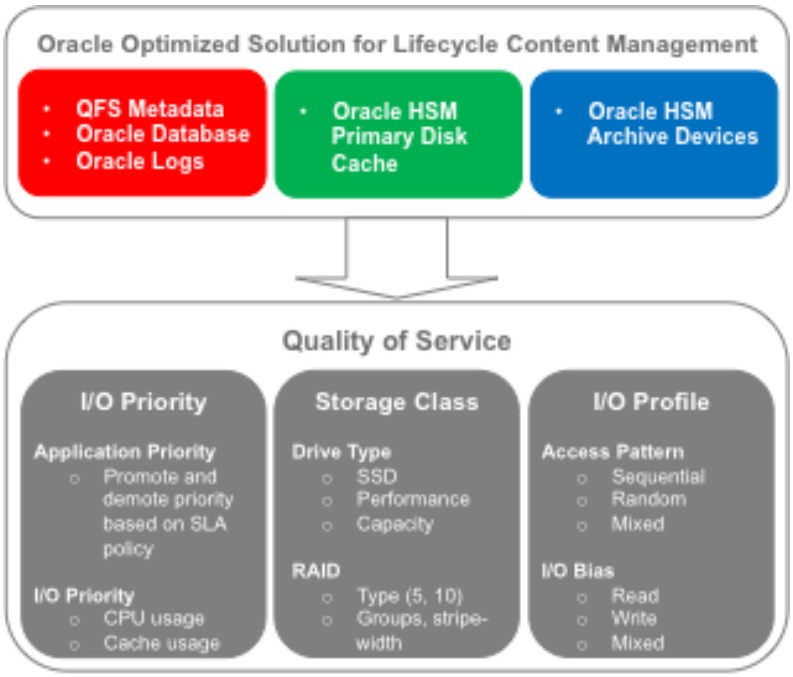

Figure 4. The combination of I/O prioritization, multiple storage classes, and I/O profiles each customized to specific requirements, results in the delivery of deterministic performance under any load condition.

Patented Oracle Flash Storage Systems Quality of Service technology is also represented in Figure 4, providing a significant differentiator over traditional controller-based disk storage. Oracle FS1-2 Quality of Service is delivered by prioritizing data access and ingest for different LUNs based on an assigned level of business priority. Advanced quality of service software manages system resources (CPU, cache, and capacity) to automate storage provisioning based on business priority.

Oracle FS1-2 Storage System Architecture

The Oracle FS1-2 storage system is designed to scale performance along with capacity. Unlike most storage systems, which have a fixed number of storage controllers (usually a maximum of two), the Oracle FS1-2 storage

system can be scaled in multiple dimensions by independently adding more storage controllers or more trays of disks and/or SSDs as needed. The Oracle FS1-2 storage system is built on three intelligent hardware assemblies. Oracle FS1-2 Pilot is the management platform, Controller is the storage controller, and DE is the drive enclosure shown in Figure 5.

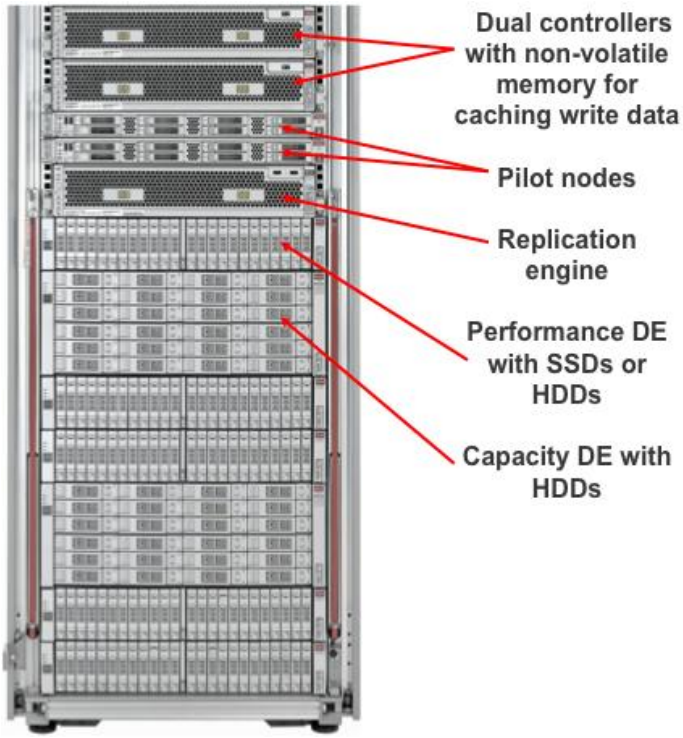

Figure 5. Oracle FS1-2 storage system components provide intelligence for considerable system flexibility.

Oracle FS1-2 Controller and Disk Enclosures can be flexibly combined to meet unique application performance and storage capacity requirements. The high flexibility combinations are especially valuable to the dynamics of a content management solution.

More information on the Oracle FS1-2 storage system can be found at this URL:

https://www.oracle.com/storage/san/fs1/index.html

Oracle's StorageTek Modular Library Systems

Oracle's StorageTek tape libraries, tape virtualization, tape drives, tape media, and tape device software contribute to byte-stream preservation in a content management solution. Tape provides the highest level of availability at the lowest cost and content can be stored for years in multiple locations without the use of power to keep a disk device spinning.

These systems also provide for local and remote copies of data while the content is still present on spinning media, delivering data protection and potentially eliminating the requirement for a backup utility. The StorageTek modular library systems that are proposed in the small, medium, and two large content management reference configurations scale from 30 to 100,000 slots and meet all capacity, archive, and access requirements.

Table 2 provides a comparison of the features of Oracle's StorageTek tape libraries recommended for Oracle Optimized Solution for Lifecycle Content Management.

TABLE 2. ORACLE'S STORAGETEK TAPE LIBRARIES PROVIDE A RANGE OF CAPABILITIES

	StorageTek SL150 Modular Tape Library	StorageTek SL3000 Modular Library System	StorageTek SL8500 Modular Library System
			
NUMBER OF CARTRIDGE SLOTS	30 to 300	200 to 5,925	1,450 to 100,000
STORAGETEK T10000D CAPACITY	n/a	1,700 to 50,364 TB	12,325 to 850,000 TB
STORAGETEK LTO 6 CAPACITY	75 to 750 TB	500 to 14,812 TB	3625 to 250,000 TB
MAXIMUM NUMBER OF TAPE DRIVES	20	56	640
MAXIMUM NATIVE THROUGHPUT (TB/HR)	11.5	48.4	552.9
SUPPORTED TAPE DRIVES	HP LTO 6/5	StorageTek T10000D, StorageTek T10000B, StorageTek T10000A, StorageTek T9840C, StorageTek T9840D, and HP and IBM LTO 6/5/4/3/	StorageTek T10000D, StorageTek T10000B, StorageTek T10000A, StorageTek T9840C, StorageTek T9840D, and HP and IBM LTO 6/5/4/3/2
NUMBER OF PHYSICAL PARTITIONS	2	8	8
REDUNDANT COMPONENTS	Control path, fans, power	Robotics, electronics, control path CAPS, fans, power	Robotics, electronics, control path CAPS, fans, power
REDUNDENT HOT-SWAPPABLE COMPONENTS?	yes	yes	yes
ANY CARTRIDGE ANY SLOT?	yes	yes	yes

Table 3 provides a description of the features of the two tape drives tested and recommended for Oracle Optimized Solution for Lifecycle Content Management.

TABLE 3. STORAGETEK TAPE DRIVES AT A GLANCE

	STORAGETEK T10000D TAPE DRIVE	STORAGETEK LTO 6
MEDIA CAPACITY	8.5 TB	2.5 TB
THROUGHPUT	252 Mb/sec	160 Mb/sec
NUMBER OF GENERATIONS MEDIA SUPPORT	3	3
DATA INTEGRITY VALIDATION	Yes[1]	Yes[2]

Selection of a library and tape drives for a content management solution is made through the review of requirements that include the following:

- Retention period of the content, which contributes to capacity requirements
- Number of copies on tape, which contributes to capacity requirements (a minimum of two copies is recommended)
- Current content capacity to be archived
- Daily content capacity to be archived, which contributes to performance requirements
- Estimated yearly content growth
- Other applications that will share the library
- Estimated activity of staging from tape to disk
- Requirements for data integrity validation

More information on Oracle's StorageTek modular library systems and tape drives can be found at the following URL: http://www.oracle.com/goto/tape

Software Application Infrastructure

The primary software application in this solution is Oracle WebCenter Content. The following subsections describe Oracle WebCenter Content and the additional software and management components used in this optimized solution. All software components run on SPARC T-Series servers.

Oracle WebCenter Content

Oracle WebCenter Content is designed to scale horizontally, making it a good fit for the SPARC T-Series servers. Oracle WebCenter Content has many different software components that can be installed all on a single server or across multiple servers to achieve greater throughput. Table 4 provides an overview of the components of Oracle WebCenter Content and the specific version numbers used in this optimized solution.

TABLE 4. SOFTWARE COMPONENTS

SOFTWARE	RELEASE

1 StorageTek T10000D uses CRC as defined in ANSI X3.139 and is calculated in the chip for a no-impact to performance

2 LTO uses the Reed Solomon CRC calculation in software at an 88 percent decrease in performance and is supported in IBM Tivoli.

ORACLE WEBCENTER CONTENT SOFTWARE SERVERS	Oracle Real Application Clusters (Oracle RAC)	11.2.0.1.0
	Oracle WebLogic Suite	11*g* (10.3.5)
	Oracle WebCenter Content	11.1.1.8.0

As the workload increases and more demand is placed on ingest and access of data, solution performance can be increased using two methods. Due to the power of the Oracle's SPARC T-Series processors and use of Oracle VM Server for SPARC, additional cores can be added to the virtual server used by Oracle WebCenter Content. If additional cores are not available, additional physical servers can be added to the configuration and Oracle WebCenter Content then distribute the users across the multiple servers through load balancing technology, while all servers have access to the same content over the storage network. In addition to scaling ingest and access performance, the storage infrastructure scales in capacity through the use of tiered storage, allowing users to have access to any data, no matter how old it is or where it is stored.

The following major features in Oracle WebCenter Content are important for content management:

» Capture and Imaging leverages Oracle WebCenter Capture for image capture and Oracle WebCenter Forms Recognition for intelligent data capture, and it provides annotation and markup of images, automates routing and approvals, and provides a scalable repository supporting enterprise-wide applications.

» Records Management enables both electronic and physical records to be managed. Records management policies and practices can be applied on both records and non-record content, as well as on content in remote repositories, such as file systems, content management systems, and e-mail archives, which enables holds to be placed from a single console.

» Content Publishing enables documents and web pages to be published as a website and enables organizations to maintain accurate, timely, and current Web content with consistent branding and presentation. Role-based tools are provided to ensure high productivity and fast updates to critical content.

» Digital Asset Management provides industry-leading digital asset management for storing, finding, and accessing brand assets. Oracle's digital asset management functionality allows users to automate the creation of thumbnails and renditions, convert high-resolution assets into web-friendly formats, and storyboard video to help users select key scenes.

» Oracle Content Services provides common set of services and functionality to the independent content management systems listed above. The common features include file plan management or warehouse management for digital and physical records, robust transformation for video files or for digital assets, WYSIWYG editors, layouts, and templates, and dynamic and static publishing models for websites. Inbound Refinery includes features such as taking Adobe Photoshop files and transforming them into different formats, resolutions, and sizes. Inbound Refinery features can evolve enterprisewide and organizational content into the future as new format technology is developed.

» The newest features in Oracle WebCenter Content includes the support of the Integrating the Healthcare Enterprise (IHE) standards-based specification for sharing patient electronic health records across many systems. This interface is commonly known as *cross-enterprise document sharing*, or XDS. Oracle WebCenter Content supports the role of XDS Document Repository (XDS.b and XDS-I.b). The XDS Document Registry communicates with the Oracle Master Person Index in order to access all known identifiers of a patient. The result is a common healthcare repository for ingest and access, enabling the secure sharing of documents between any health care enterprise, including a private physician office, a clinic or a large acute care in-patient facility. Oracle WebCenter Content also supports other document access protocols that allow clinical applications to share data with nonclinical applications in the appropriate setting and access controls.

» This opens Oracle WebCenter Content to the healthcare industry, meeting new requirements to improve clinical workflow and productivity for safer, collaborative care decisions. Leveraging Oracle WebCenter Content to manage the consolidated ingest and retrieval of all information and yet abstract the storage infrastructure from

Oracle WebCenter Content results in a cost-effective solution to handle the complete lifecycle of patient records that may have a retention period of 50 or more years.

More information on Oracle WebCenter Content can be found at the following URL:

http://www.oracle.com/us/products/middleware/webcenter/content/overview/index.html

Management Applications

The following management applications are deployed separately from Oracle WebCenter Content and Oracle HSM, but they reside in the same infrastructure.

- Oracle's StorageTek Automated Cartridge System Library Software manager is used to configure and control the tape drives.
- Oracle Enterprise Manager Ops Center is used to load, patch, and control the SPARC T-Series servers.

Server Infrastructure

The server infrastructure for the solution is based on a modular architecture. SPARC T-Series servers provide scalability and flexibility for the overall solution used to deploy Oracle WebCenter Content, Oracle Database, and Oracle HSM. SPARC T-Series processors are the industry's most highly integrated system-on-a-chip, supplying the most high-performance threads of any multicore processor available and integrating all key system functions. When scaling for the medium and large configurations, multiple servers are used to run multiple instances of the content management application. As workload increases, additional servers can be added, utilizing load balancing technology and the ability of the application to federate across servers for access to the content.

Oracle's SPARC T4 Servers

The SPARC T4 core architecture was initially designed with features that are usually associated with superscalar designs. These include out-of-order instruction execution, sophisticated branch prediction, prefetching of instructions and data, deeper pipelines, and multiple levels of caches, with support for a 2 GB page size, and multiple instruction dispatching. These innovations resulted in an improvement in single-thread execution, networking, and overall throughput performance for demanding applications such as Oracle Optimized Solution for Lifecycle Content Management.

The SPARC T4 processor provides eight cores, with each core supporting up to 8 threads (64 threads per processor). In addition, each core provides two integer execution pipelines, so a single SPARC T4 core is capable of executing two threads at a time. Oracle's SPARC T4-1 server is a very cost-effective solution for entry-level content management requirements and contains one 2.85 GHz SPARC T4 processor with up to 512 GB of RAM and six PCIe expansion slots. Optional 8 Gb Fibre Channel HCAs are typically configured for redundant Fibre Channel storage access requirements.

More information on Oracle's SPARC T4 servers can be found at this URL:

http://www.oracle.com/us/products/servers-storage/servers/sparc-enterprise/t-series/overview/index.html

Oracle's SPARC T5 Servers

SPARC T5 servers expand the SPARC portfolio with best-in class systems that are engineered and optimized to accelerate Oracle software and business-critical applications with extreme performance, mission-critical reliability and scalability. SPARC T5 servers improve the price/performance leadership established by their SPARC T4-based predecessors, while complementing and extending the optimization for performance, scale and availability. The SPARC T5 CPU delivers the best-in-class single thread performance, with 16 cores per CPU, and it delivers up to a remarkable 2.5 fold system throughput over the previous generation. In addition, the SPARC T5 servers provide three times faster security with CPU integrated zero overhead encryption accelerators and also provide two times the I/O bandwidth by utilizing PCIe 3.0 to deliver higher application performance.

Oracle's SPARC T5 servers provide an innovative and efficient 8-way design that allows for near-linear scaling. The most demanding cloud and enterprise applications like Oracle WebCenter Content can benefit from intelligent, real-time scaling with virtualization. Oracle's SPARC T5 servers can scale from one to eight socket systems with binary compatibility and scalable architecture, for seamless application performance uplift. PCIe 3.0 is new on Oracle's SPARC T5-series servers and provides for higher speed networking which aids in removing I/O bottlenecks in highly virtualized environments

Reliability, availability, and serviceability (RAS) features are built into the SPARC T5 server workload, including built-in redundancy for resilient mission-critical operation. SPARC T5 servers provide hot-swap components, including the CPUs advanced power management, which enable CPUs to gracefully degrade systemwide resources for greater availability. Automated software lifecycle management is built in with Oracle Enterprise Manager at no extra cost.

The SPARC T-Series systems lend themselves to very high throughput demands as well as CPU-intensive applications such as Oracle WebCenter Content. Oracle's SPARC T5-2 and SPARC T5-4 servers represent scalable building blocks that can be configured into several reference configurations for deploying lifecycle content management as follows:.

» SPARC T5-2 server with two 3.6 GHz SPARC T5 processors, up to 512 GB of memory, four 10 Gb Ethernet ports, and 16 PCIe 3.0 expansion slots.
» SPARC T5-4 server with four 3.6 GHz SPARC T5 processors, up to two TB of memory, four 10 Gb Ethernet ports, and 16 PCIe 3.0 expansion slots.
» Optional 8 Gb Fibre Channel HCAs are typically configured into the PCIe slots of both servers for redundant Fibre Channel storage access requirements.

More information on the SPARC T5 servers can be found at this URL

http://www.oracle.com/us/products/servers-storage/servers/sparc-enterprise/t-series/overview/index.html

Understanding Performance and Capacity

The previous sections described the optimal components that are included in the proven Oracle Optimized Solution for Lifecycle Content Management. These are best-of-breed products that work well together and have been tested with known results for performance, scalability, and flexibility. The following subsections include the general architecture of the solution followed by the configuration of the component building blocks and performance benchmarks that become the basis for the reference configurations that follow in the next section.

Scalability and flexibility of the reference configurations are achieved by consolidating the applications required for lifecycle content management onto the highly expandable SPARC T-Series servers using integrated, no-cost virtualization through Oracle VM Server for SPARC. Sun Logical Domains (LDOMs) offers a powerful and consistent methodology for creating logically isolated virtualized server environments across the entire SPARC T-Series server range. The solution's applications run on this infrastructure and the strength of Oracle's approach is that the multithreaded layers of the architecture can be allocated core resources as needed, with near-optimal performance. LDOMs flexible allocation of core resources helps control license and server costs for applications that are based on the number of cores by allocating only the resources that are needed to meet current requirements while reserving the headroom for future scalability. Resources then are available to run additional applications in the same physical server. For high availability, SPARC T-Series servers are deployed in pairs in the reference configurations. This pairing also delivers very high ingest rates and database performance by federating Oracle WebCenter Content and Oracle Database across the available cores on more than one server.

Both performance and capacity must be taken into consideration when selecting the hardware and software configuration for the solution to meet content management requirements. Ingest performance is important for use cases that have high ingest requirements. Capacity is important for use cases that may have a smaller daily ingest requirement but have a very long retention period (often forever) for potentially millions of files. All use cases have specific requirements in terms of both performance and capacity. The goal of providing the reference configurations that follow is to offer guidance for selecting a configuration that most closely meets current requirements and provides the ability to scale to easily meet future requirements.

Small, Medium, and Large Configuration Overview

Figure 6 shows the general overall architecture of Oracle Optimized Solution for Lifecycle Content Management.

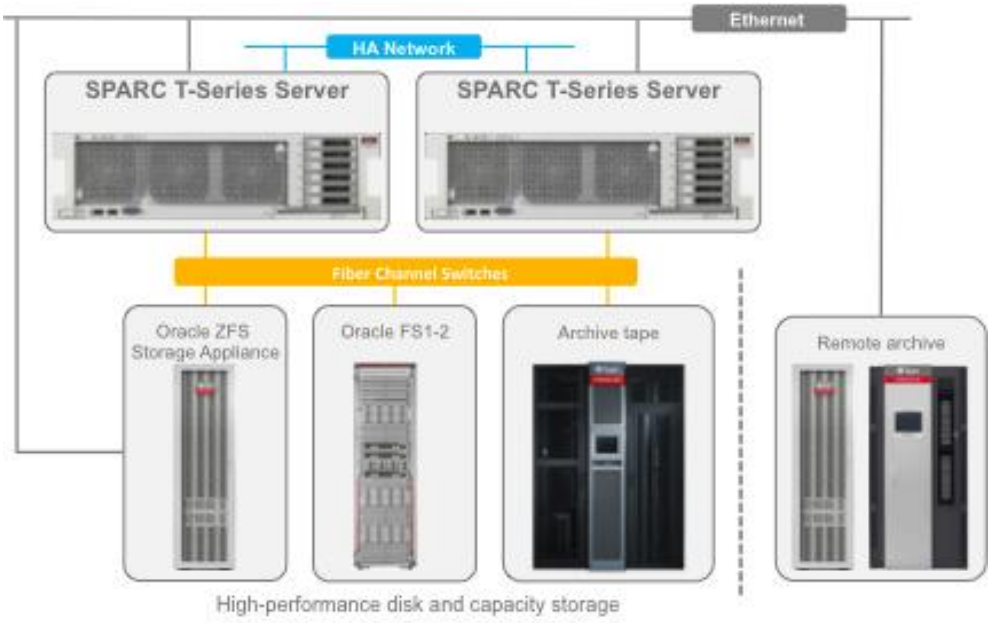

Figure 6. General architecture for Oracle Optimized Solution for Lifecycle Content Management.

Oracle Optimized Solution for Lifecycle Content Management provides capacity and performance comparisons for small, medium, and large configurations. The suggested configurations provide guidelines for initial expectations and show the scalability of the solution as business needs grow and data capacity increases. Oracle WebCenter Content is horizontally scalable, and delivers higher ingest throughput by adding additional cores, processors, or larger servers. Oracle HSM, Oracle ZFS Storage Appliance, and the Oracle FS1-2 storage system all have the ability to migrate data to new technology nondisruptively. Moreover, Oracle's tape system strategy is to provide read access for three generations of tape media from the newest tape drives. This degree of backward compatibility gives businesses many years of data access without forcing migration of the content, along with the ability to take advantage of the latest tape drive technology and media for new archives. As tape media becomes increasingly dense and appealing, Oracle HSM provides the tools to migrate the content to this new technology. Table 5 shows the configuration of the individual building blocks for the suggested reference architectures that follow.

TABLE 5. SERVER, STORAGE, AND TAPE SYSTEMS WITH SOFTWARE COMPONENTS FOR HIGHLIGHTED CONFIGURATIONS

	SERVER VIRTUALIZATION				DISK STORAGE OPTIONS (ORACLE ZFS STORAGE APPLIANCE OR ORACLE FS1-2)							TAPE SYSTEMS	
Configuration	Server	Oracle HSM Cores	Oracle Database Cores	Oracle WebCenter Content and Oracle WebLogic Cores	Oracle ZFS Storage Appliance				Oracle FS1-2 Storage System (option to Oracle ZFS Storage Appliance)			Tape Library	Tape Drives
					Oracle ZFS Storage	SSD	Perf. Drive Trays	Cap. Drive Trays	SSD DE	Perf. Drive DE	Cap. Drive DE		
Small	SPARC T4-1	4	Up to 4		Sun ZFS 7320	2		2				StorageTek SL150	LTO
	SPARC T4-1	4	Up to 4										
	SPARC T4-1			Up to 8									
	SPARC T4-1			Up to 8									
Medium	SPARC T5-2	4	Up to 12	Up to 16	Sun ZFS ZS3-2	4	2	2		4	4	StorageTek SL3000	StorageTek T10000D
	SPARC T5-2	4	Up to 12	Up to 16									
Large Option 1	SPARC T5-4	8	Up to 24	Up to 32	Sun ZFS ZS3-4	8	4	4	2	8	8	StorageTek SL8500	StorageTek T10000D
	SPARC T5-4	8	Up to 24	Up to 32									
Large Option 2	SPARC T5-2	8	Up to 24		Sun ZFS ZS3-4	8	4	4	2	8	8	StorageTek SL8500	StorageTek T10000D
	SPARC T5-2	8	Up to 24										
	SPARC T5-2			Up to 32									
	SPARC T5-2			Up to 32									

Performance and Capacity Scalability

While scaling capacity is clearly a requirement through time, it is vital that any solution also be able to scale in performance to support growth in ingest rates performance as well. Oracle Optimized Solution for Lifecycle Content Management provides the capability for both the scalability of ingest rate and capacity growth over time, as described in the following subsections.

Ingest Performance Scalability

Oracle WebCenter Content is particularly scalable horizontally, and delivers further ingest throughput with the addition of cores or processors. Oracle WebCenter Content servers can also be added to achieve additional throughput scalability.

To provide a basis for sizing of the reference configurations that follow, Oracle conducted data ingest testing for Oracle WebCenter Content. The graph in Figure 7 shows the ingest rates reported from the Faban test (www.faban.org) for the SPARC T4 and the SPARC T5 single processor configurations when writing 100 K byte files using 10 Gigabit (10 GbE) and the NFS protocol to access the Oracle HSM file system. As shown, even single-processor SPARC T4 and T5 servers provide a scalable range of ingest rates.

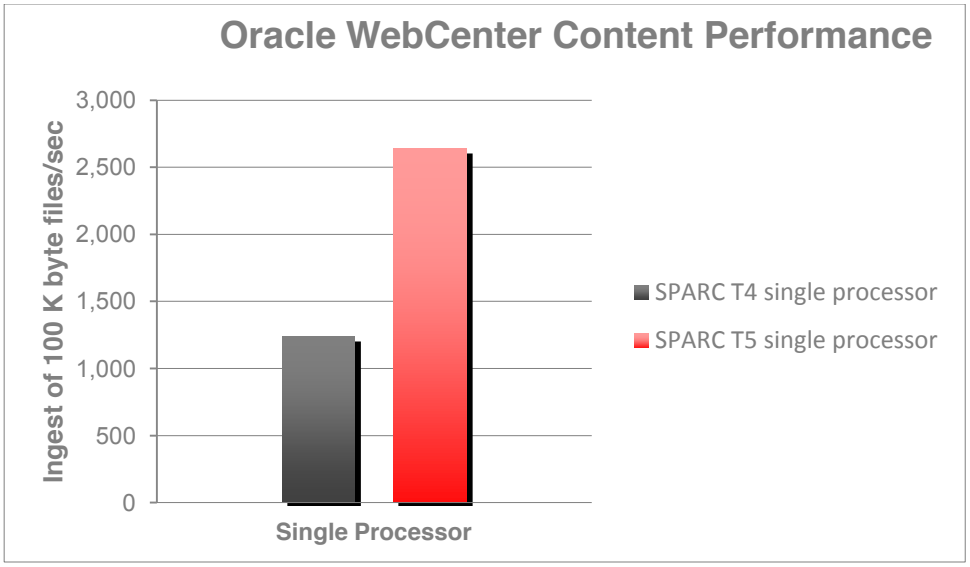

Figure 7. Single-processor SPARC T4 and T5 servers provide scalable ingest rates of 100 K byte files when running Oracle WebCenter Content.

This range is even more pronounced when additional cores, processors, or systems are added to run Oracle WebCenter Content. The graph in Figure 8 shows the ingest rates reported from the Faban test as series of runs on increasing numbers of SPARC T5 processor cores. The Fabian tests were run as above by writing 100 K byte files using 10 Gigabit (10 GbE) and the NFS protocol to access the Oracle HSM file system. Figure 8 demonstrates that a significant scalability of ingest rates is easily achieved, by just adding more SPARC T5 cores.

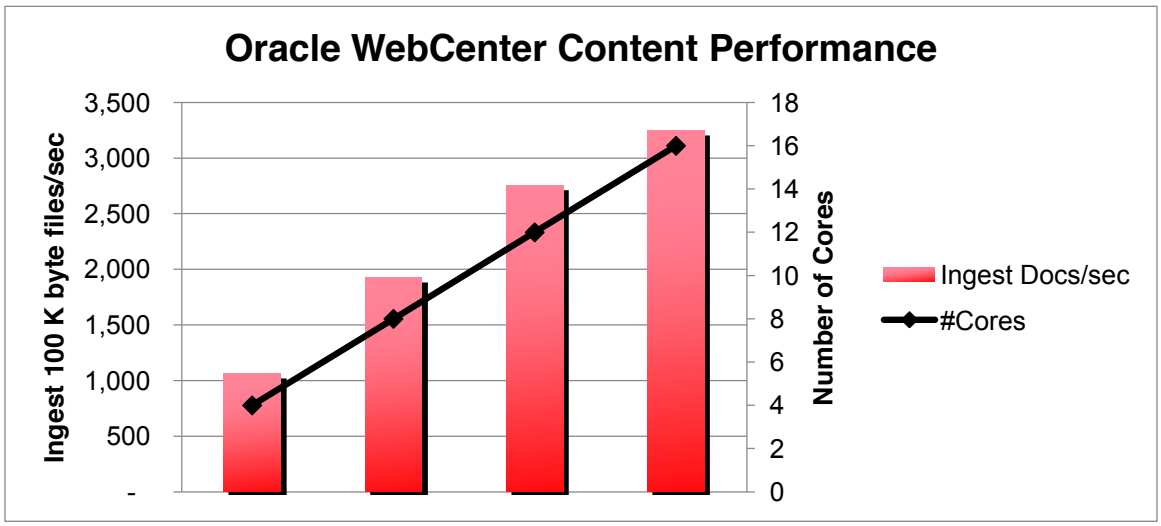

Figure 8. Increasing the number of SPARC T5 cores provides scalable ingest rates of 100 K byte files when running Oracle WebCenter Content.

Evaluating Capacity Requirements Through Time

As a part of Oracle's testing, a moderate baseline ingest range rate was established. While this is far below the maximum capability of Oracle WebCenter Content, this configuration has the potential to generate very large capacity repositories based on ingest rate capabilities. Table 6 shows the capacity requirements for the baseline test configuration assuming a continuous ingest rate for 24 hours per day, seven days per week.

TABLE 6. ONE-YEAR CAPACITY REQUIREMENTS BASED ON INGEST RATE AT 24/7/365

Number of 100 K Byte Records Ingested Per Sec	Ingest Capacity Per Sec (GB)	Ingest Capacity Per Day (TB)	Ingest Capacity Per Year (PB)
3,280	3.2	270	96

Realistically, the ingest rates noted in Table 1 are experienced only at peak times. However, they must be processed without a negative impact to the users. The total ingest capacity is more likely a much smaller percentage of the total shown above. Selection of system size based on peak ingest rates, along with total capacity, expected growth, and the retention time, can be used to determine overall capacity requirements.

As an example of capacity requirements, Figure 9 indicates the capacity requirements when starting with 200 TB.

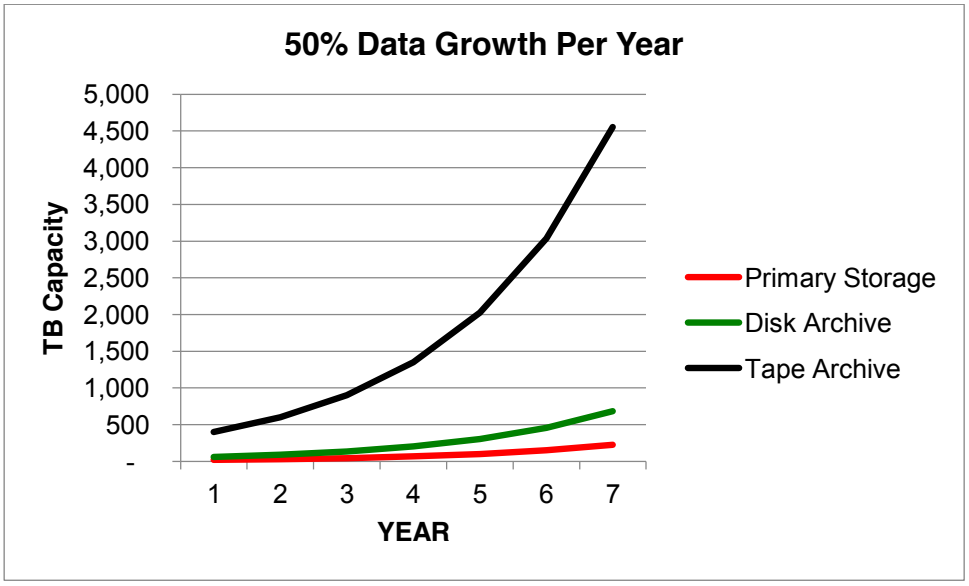

Figure 9. Total capacity increase over seven years based on 50 percent growth per year.

The chart in Figure 9 starts with year one capacity of 200 TB and shows an expected growth of 50 percent per year up to seven years for each of the different storage tiers. The chart uses the following assumptions: 10 percent for primary disk, 30 percent for disk archive, and 200 percent for tape archive. Even though dramatic growth in overall capacity is shown, the primary disk storage tier, which is the most expensive tier, remains relatively small, thus keeping costs low. The most cost-effective storage—tape—carries the largest capacity.

Figure 10 shows the capacity scalability of Oracle Optimized Solution for Lifecycle Content Management, comparing the ranges of overall tape storage capacity available for the small, medium, and large (option 1 and 2) reference configurations.

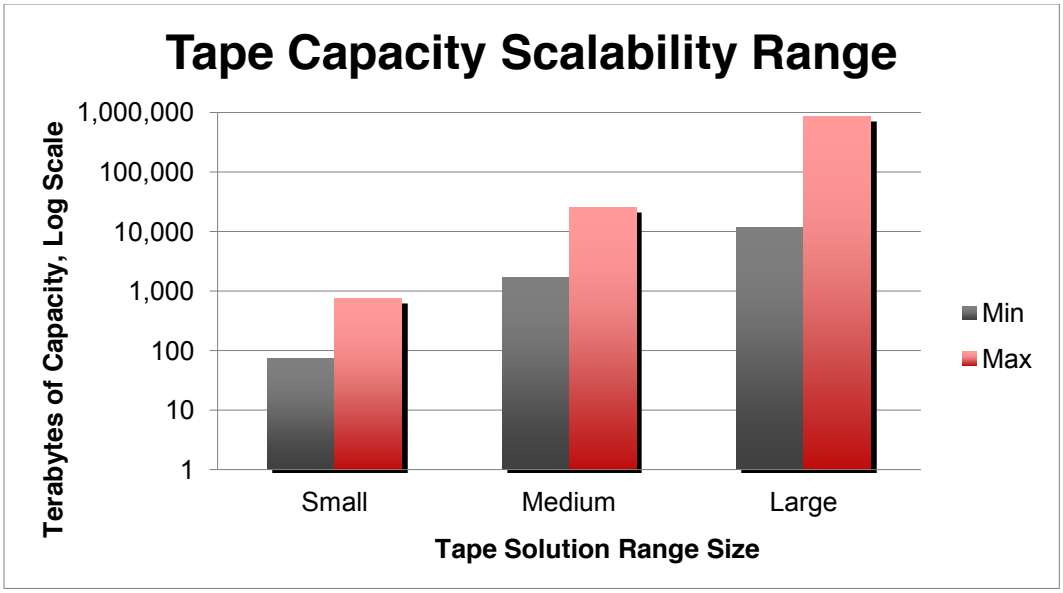

Figure 10. Tape scalability of the solution using the capacity ranges for small, medium, and large (option 1 and 2) StorageTek library system configurations.

Businesses can only estimate capacity, ingest, and access requirements for the future when implementing a content management system. Oracle Optimized Solution for Lifecycle Content Management scales nondisruptively in capacity and performance, meeting requirements during unexpected (but welcomed) growth of the company. Growth in data flow and capacity is easily and quickly handled by Oracle Optimized Solution for Lifecycle Content Management, leaving administrators and users free to focus on important activities other than worrying about where the data is and whether it is safe.

All requirements must be taken into consideration when selecting a small, medium, or large solution with performance and capacity being the primary place to start. The Oracle Optimized Solutions team has proven that the components work together and have test results to provide guidelines for size selection. With Oracle WebCenter Content, Oracle HSM, Oracle ZFS Storage Appliance or Oracle FS1-2, and StorageTek tape libraries providing tools for migrating nondisruptively, it is possible to begin with one solution size and easily and confidently upgrade to the next size.

Small, Medium, and Large Reference Configurations

As mentioned previously, both performance and capacity must be taken into consideration when selecting the hardware and software configuration for the solution to meet content management requirements. This section describes four performance and capacity configuration levels that are each designed to be used as a guideline to meet expectations for sizing various use cases that most closely meet current requirements and also provide the ability to flexibly scale to meet future requirements. The applications required for lifecycle content management are easily consolidated through virtualization and onto the highly expandable SPARC T-Series to provide the most affordable initial implementation and future incremental scaling.

Small Configuration

The small configuration consolidates lifecycle content management applications onto four SPARC T4-1 servers with integrated virtualization with up to two LDOMs. The configuration described below is flexible as it can start with a nominal smaller configuration and easily scale to meet additional performance or capacity requirements in the future. Figure 11 depicts the configuration for the small reference architecture.

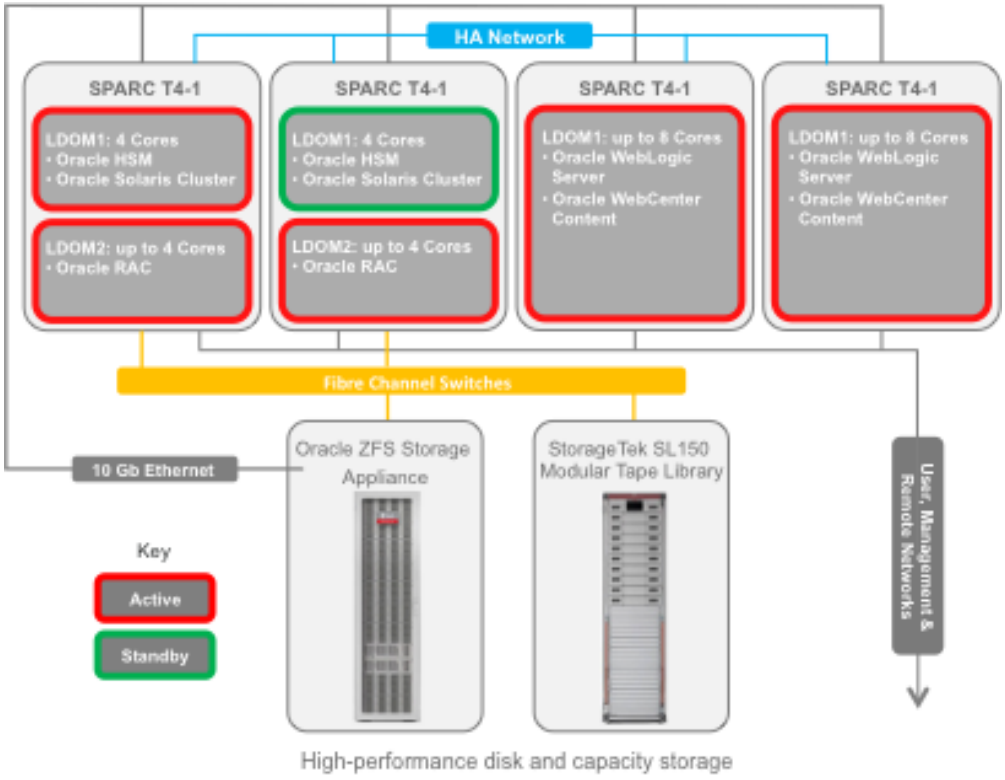

Figure 11. The small reference configuration uses four SPARC T4-1 servers, Sun ZFS Storage ZS3-2 arrays, and a StorageTek SL150 modular tape library with StorageTek LTO 5 tape drives.

Medium Configuration

The medium configuration is an example of consolidating lifecycle content management applications onto just two SPARC T5-2 servers with integrated no-cost virtualization deploying three LDOMs on each server. The medium configuration displayed in Figure 12 is flexible as it can begin deployment with a smaller consolidated system by virtualizing the applications and easily scale to meet additional performance or capacity requirements in the future.

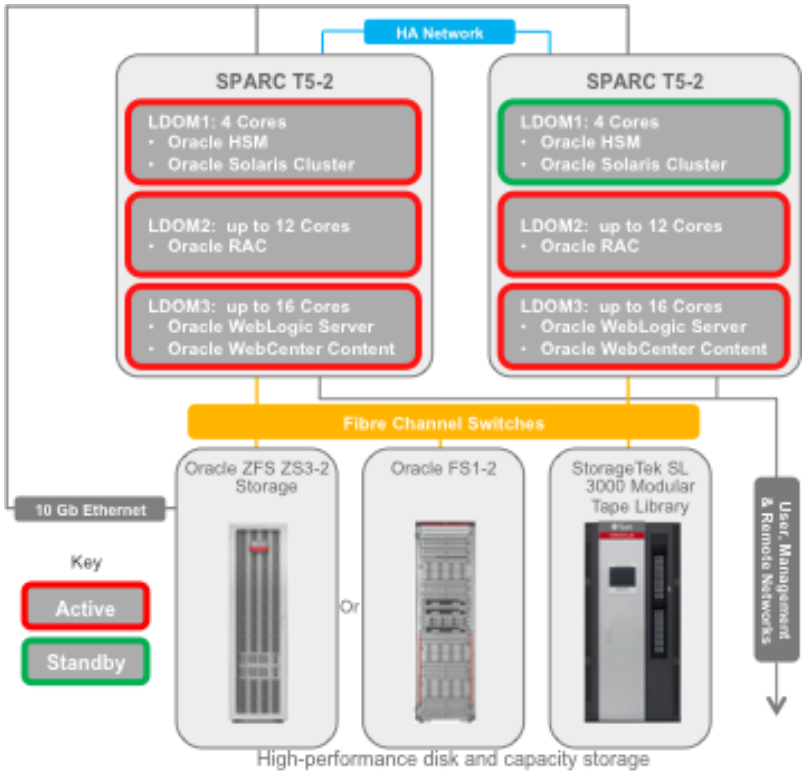

Figure 12. The medium reference configuration uses two SPARC T5-2 servers, a Oracle ZFS Storage ZS3-2 storage appliance or the optional Oracle FS1-2 storage system, and a StorageTek SL3000 modular library system with StorageTek T10000D tape drives.

Large Configuration (Option 1)

In the first large configuration, lifecycle content management applications are consolidated onto SPARC T5-4 servers with integrated virtualization deployed three LDOMs on each server. For high availability two SPARC T5-4 servers are configured to deliver very high ingest rates by federating Oracle WebCenter Content across both servers. Three logical domains are configured on each SPARC T5-4 server. The configuration is flexible as it can begin deployment with a smaller consolidated system by virtualizing the applications onto fewer cores and easily scale to meet additional performance or capacity requirements in the future. Figure 13 depicts the first large configuration.

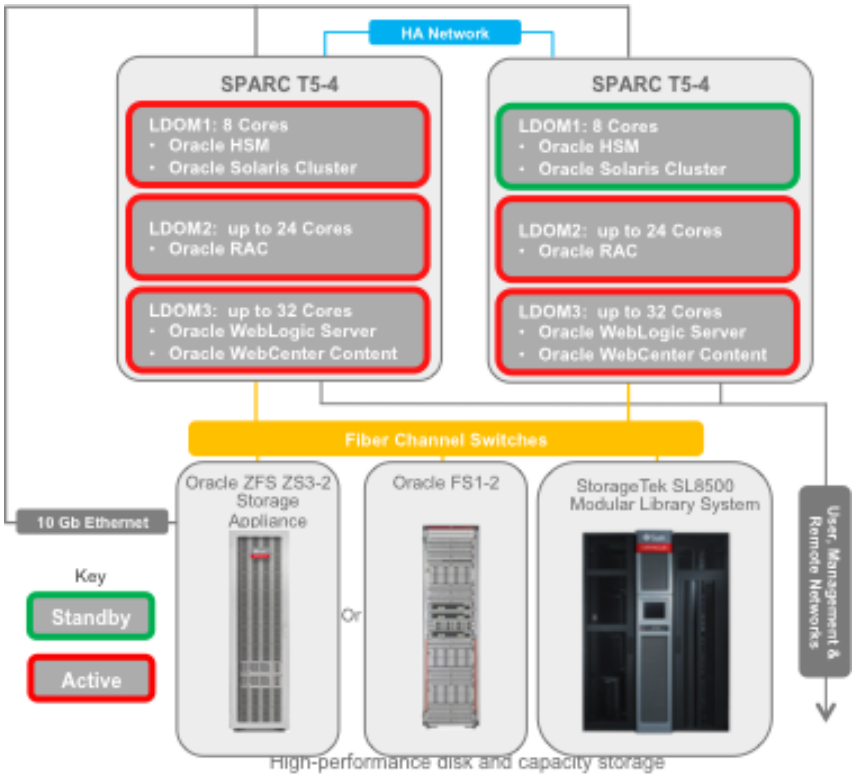

Figure 13.The first large reference configuration uses two SPARC T5-4 servers, Sun ZFS Storage ZS3-2 appliance or Oracle FS1-2 storage system and a StorageTek SL8500 modular library system with StorageTek T10000D tape drives.

Large Configuration (Option 2)

The second large configuration demonstrates another approach by consolidating lifecycle content management applications onto four SPARC T5-2 servers with integrated no-cost virtualization. For high availability two SPARC T5-2 servers are deployed for running Oracle RAC and Oracle HSM. Oracle WebCenter Content is federated across another pair of SPARC T5-2 servers and delivers very high ingest rates. The described configuration is flexible as initial implementation can start with fewer cores by virtualizing applications and easily scale to meet additional performance or capacity requirements in the future. Figure 14 depicts the second large configuration reference architecture.

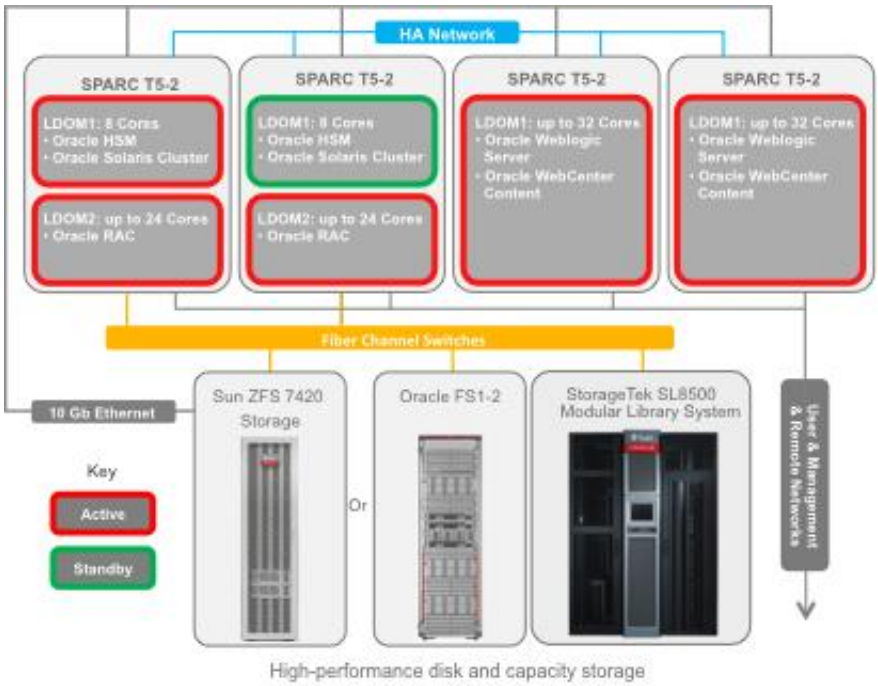

Figure 14. The second large reference configuration uses four SPARC T5-2 servers, a Sun ZFS Storage ZS3-2 appliance or an optional Oracle FS1-2 storage system, and a StorageTek SL8500 modular library system with StorageTek T10000D tape drives.

Scalability and Flexibility of the Reference Configurations

The components in Oracle Optimized Solution for Lifecycle Content Management reference configurations all have the capability to easily scale in overall performance and provide increased capacity to meet growing future demands. The SPARC T-Series with Oracle VM Server for SPARC virtualization features the capability to provide applications with exactly the compute resources that are needed. Adding additional cores to Oracle WebCenter Content increases the ability to ingest and access data as demands increase. Oracle ZFS Storage Appliance and Oracle FS1-2 storage system both have the capability to grow disk capacity and performance, and Oracle's StorageTek modular library systems all have the capability add additional tape slots increasing capacity and tape drives to increase performance. Oracle HSM has the capability to nondisruptively add both disk and tape storage capacity, to an existing file system, while maintaining access to 100 percent of the content.

The following summarizes how to grow the four configurations:

» As the user count increases or as new applications are brought online, Oracle WebCenter Content can be used as a tool to manage and store content. In order to manage this increase in demand, load balancers and federation by Oracle WebCenter Content provides the capability to add available cores in existing servers or by adding new servers to the existing configuration, resulting in a positive impact to the users. Oracle's SPARC T-Series servers provide an exceptional growth path to scale flexibility with performance growth requirements and include the SPARC T4-1 server, SPARC T5-2 server, SPARC T5-4 server and SPARC T5-8 server.

» As capacity demands increase, disk storage can be added and Oracle HSM can nondisruptively add capacity to an existing file system through the 'grow' command. Hybrid storage pools (Zpools) utilizing SSDs are the primary storage for Oracle HSM and Zpools without SSD are used for the disk archive. Capacity easily scales up for either or both of the Oracle HSM disk cache and disk archive. As new storage is added, new LUNs are created and mapped to the Oracle HSM server and can be dynamically added to the Oracle HSM file system or to the disk archive file system. For scalability, the Sun ZFS Storage ZS3-2 disks and trays in the small configuration are

compatible with the Sun ZFS Storage ZS3-4 trays in the medium and large configurations, providing an easy upgrade path by swapping in the existing trays with the data intact. Capacity of the Sun ZFS Storage ZS3-4 appliance can be increased nondisruptively, as required, to a maximum of 2.59 PB of raw disk storage. Capacity of the Oracle FS1-2 can be increased by adding units of performance disk enclosures or capacity disk enclosures to grow the primary and the disk archive storage capacity for Oracle HSM while also increasing performance. Total capacity of the Oracle FS1-2 can be increased nondisruptively, as required.

» Increasing tape capacity is accomplished nondisruptively by adding slots and/or drives to the library. If changing from one library to another, the tape media are removed from the old library and inserted into the new library. Performance also can be increased by adding tape drives to open bays or upgrading to faster tape drives. The StorageTek SL150 modular tape library as configured for the small reference configuration is expandable up to 300 slots for a total capacity of 750 TB. The StorageTek SL3000 modular library system, as configured for the medium reference configuration, can scale up to just under 6,000 slots with a total capacity of 14.8 PB, and the StorageTek SL8500 modular library system, as configured for the large configurations, can scale up to 554.5 PB with 100,880 activated slots.

Configuring Storage and Protecting Data

This section provides best practice guidelines for configuring Oracle ZFS Storage Appliance and Oracle FS1-2 storage system recommended in Oracle Optimized Solution for Lifecycle Content Management. In addition, a subsection is included on protecting data with Oracle's Data Integrity Validation feature.

Configuring Oracle ZFS Storage for Oracle HSM

Unlike traditional file systems, which reside on single devices and thus require a volume manager to use more than one device, ZFS file systems are built on top of virtual storage pools called Zpools. A Zpool is constructed of virtual devices, which are themselves constructed of block devices. Zpools are a collection of storage devices that provides physical storage and data replication for ZFS datasets. Figure 15 shows how the storage pools can be configured for the primary storage for Oracle HSM.

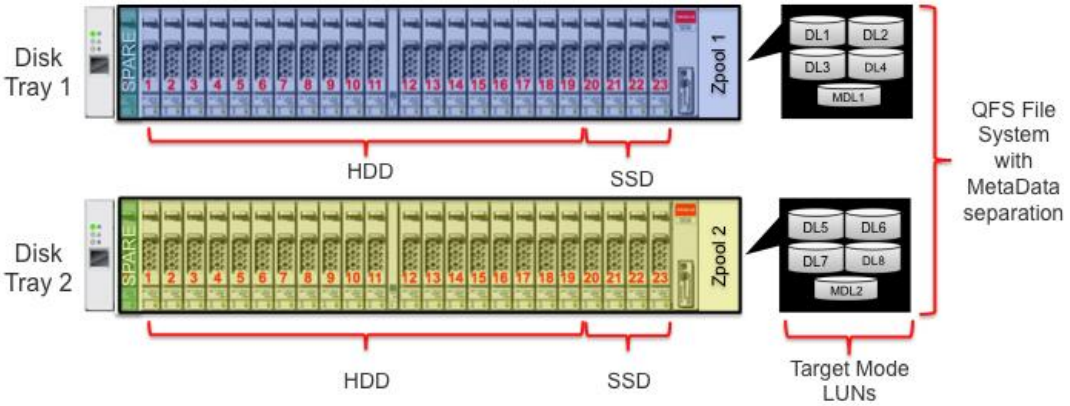

Figure 15. Oracle Storage Drive Enclosure DE2-24P with Zpools and LUN allocation for Oracle HSM primary storage.

As shown in the illustration in Figure 15, configuration settings for the content pools and for the metadata pool include.

» Two pools with all disks from each tray including the SSD
» Four data LUNs per pool
  » 128 KB record size
  » secondarycache=none

» One metadata LUN per pool
  » Metadata LUNs 16 KB record size
  » logbias=latency
  » secondarycache=all

Block storage is required for Oracle HSM primary storage; therefore, LUNs on Oracle ZFS Storage Appliance are created and mapped to the server. Oracle HSM then is used as a volume manager and selects small LUNs into a single small file system for the metadata and selects the larger LUNs into a single file system for the content.

Configuring the Oracle FS1-2 Storage System for Oracle HSM

Figure 16 describes the physical and logical configuration of the Oracle FS1-2 storage system. The logical components are created using the Oracle FS1-2 user interface. The Oracle HSM file system uses the option of metadata separation. The Oracle HSM metadata resides on one logical device, and the primary content resides on a different logical device. The separation of data types enables storing the metadata (which is small in size) on the highest-performing storage while storing the primary content (which represents the actual data) on the next highest performing storage.

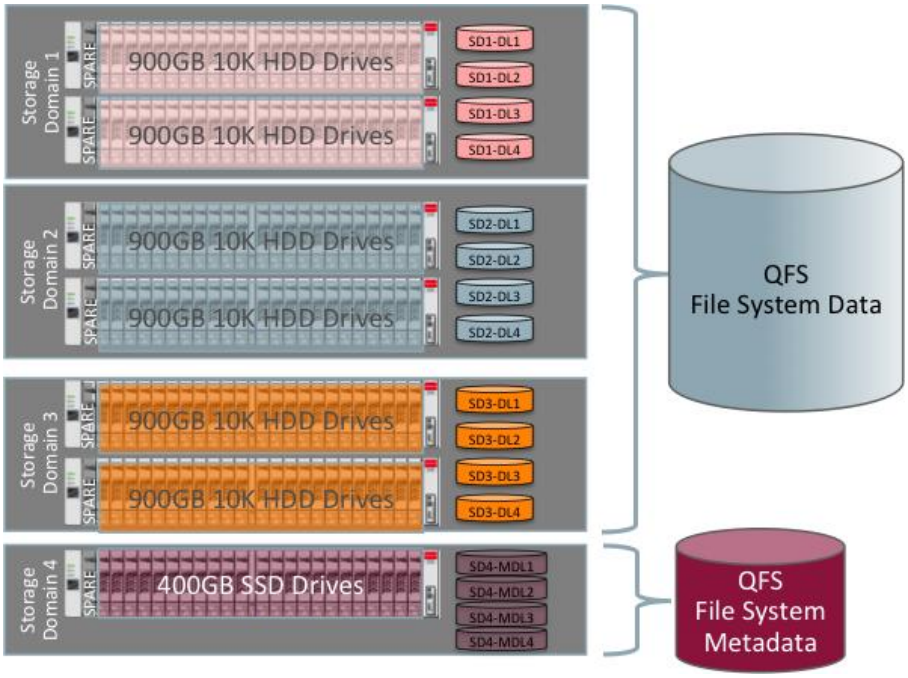

Figure 16. Recommended best practice physical and logical configuration of the Oracle FS1-2 storage system for best performance with Oracle HSM.

Figure 17 shows how the logical LUNs to be used for the Oracle HSM data file are mapped to the server though the controllers when two controllerss are available. For the best performance, metadata is stored at premium-priority, cached data is stored at high priority and archive data is stored at medium priority. All LUNs then are configured into the Oracle HSM file system and presented to the application as a single file system for access via NFS or CIFS or as a POSIX-compliant file system. The same logic is used for the metadata. If two LUNs are created on an SSD disk enclosure Oracle HSM can use both LUNs for the metadata. Oracle FS1-2 Quality of Service can ensure the appropriate priority for each record processed.

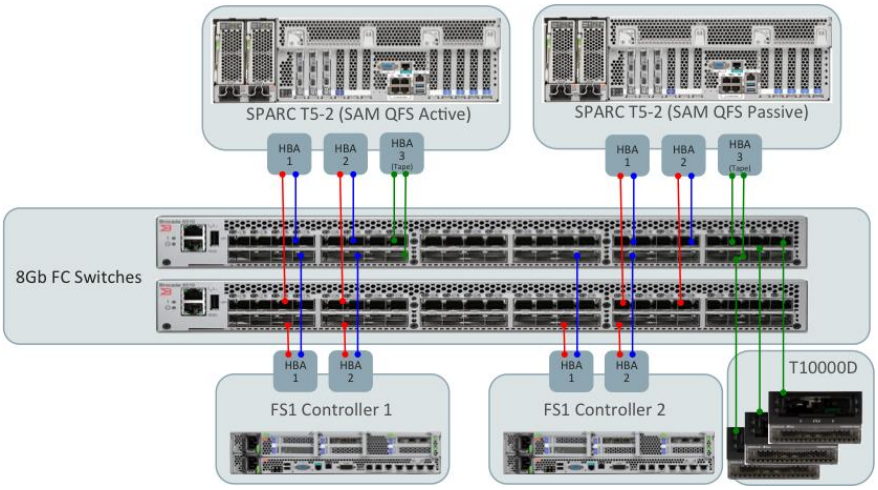

Figure 17. Mapping multiple LUNs from the Oracle FS1-2 storage domains involves utilizing all controllers paths to the Oracle HSM server.

Configuring LUNs on the Oracle FS1-2 Storage System

The previous section describes how and where the LUNs are created and how the Oracle HSM file system is created from those LUNs, and this section describes the best practices for configuring LUNs for Oracle HSM and the Oracle FS1-2 storage system.

- **Number of data LUNs per file system.** Within a storage domain of two disk enclosures, several LUNs are created and mapped to the server running Oracle HSM, and they utilize all of the controller paths on the back end between the disk controller and the controller HBAs. Oracle HSM then puts the LUNs back together into a single QFS file system.
- **Number of metadata LUNs.** Metadata performance has the greatest impact when using the Oracle HSM utilities. Testing has shown that the filesystem can easily grow to one billion while showing perfect linear scalability.

Protecting Data with Data Integrity Validation

It is critical to ensure that stored data is recorded accurately and just as important to be sure it remains unchanged through its retention time. In tape drives, data is protected with read-after-write verification as it is written, and Error Correction Code (ECC) is added to ensure data recovery once it is on the media. In addition, a typical tape drive adds Cyclic Redundancy Code (CRC) protection, as soon as a record is received. This practice ensures that the record is not corrupted while moving between internal memories. Unfortunately, these technologies do not protect data that is being moved outside the storage device. As a result, there is a chance for data corruption as it is migrated across the storage landscape. The Data Integrity Validation feature on Oracle's StorageTek T10000D tape drive takes this one step further by validating CRC checksums generated at the host. This integrity check for write, read, and validate has the highest level of importance when storing data that might be kept forever and yet has low access requirements.

Historically, data written to tape is verified after it is written or validated at the inefficient full file level. Data Integrity Validation starts this process at the server on a record level (default record size is 2 million bytes) and continues the CRC check throughout the write process until the data is written to tape media. The StorageTek T10000D validates the data at the drive to complete the process.

Oracle HSM is used to issue a validation request based on a policy. For instance, tape media that has been stored for more than a year in a library slot and has not been loaded in the tape drive is instructed by Oracle HSM to load

the tape media and verify all data is exactly what was written. It does this by calculating the CRC and comparing it with the CRC on the record.

Performance Implications of Data Integrity Validation

Both SPARC- and x86-based servers have the option to generate the required CRCs in hardware rather than software, requiring very little processor overhead. Test results from driving 10 StorageTek T10000D tape drives at optimal speed shows not only that Data Integrity Validation has essentially zero impact on performance but also has the scalability to add StorageTek T10000D tape drives.

Performance on the server during the validation-only step is not affected because the validation does not require data to be staged back to the server for the CRC check process. The validation is executed on the drive in the background, and the server is notified to take action only if an error is detected. This process ensures that all media, even if it contains dark archive files and is rarely, if ever, accessed, is loaded into the drive and read on a schedule, such as yearly or every six months. A migration to new media after many years of being stored in a library slot is not be the first time a file is read.

With vast amounts of critical data being stored digitally, it is essential that the content of the data remains unchanged during transfer from server to tape media and back as well as when stored for long periods of time. For legal and preservation purposes, the fixity of this data must be verified. Oracle's Data Integrity Validation on the StorageTek T10000D tape drive allows Oracle HSM to use CRC checksums to ensure the integrity of archived data is preserved.

Conclusion

Tiered storage is an ideal fit for enterprise content management and archive applications, but designing and configuring the tiered storage environment to optimize the content management solution can be a big undertaking. Oracle Optimized Solution for Lifecycle Content Management provides a detailed roadmap for how to configure and size the software applications, the hardware servers, and the tiered storage to meet specific performance and capacity goals. The proposed complete solutions employ all Oracle hardware and software components, and they have been tested together and proven to meet specific levels of performance.

All components in Optimized Solution for Lifecycle Content Management have the ability to easily scale performance and capacity to meet growing performance and capacity demands. Growing the storage capacity, both disk and tape, without downtime, maintains access to 100 percent of the content. Adding available cores or servers to run the application increases the ability to ingest and access data as demands increase. Oracle HSM dynamically keeps track of all content and creates policy-based copies to further enhance the solution by lowering cost and risk.

As businesses grow and content requirements increase, their goals related to a content management solution are to find the bottleneck in performance, understand the capacity growth trend, and grow the configuration to meet those demands before they happen. The following points summarize how to grow the four configurations:

- As the user count increases or new applications use Oracle WebCenter Content as a tool to manage and store content, it is possible to add servers easily to manage this increase in demand. Load balancers and federation by Oracle WebCenter Content servers make it possible to add servers resulting in a positive impact to the users.
- As capacity demands increase, disk storage can be added and Oracle HSM nondisruptively adds capacity to an existing file system through the 'grow' command.
- Growing disk from a small to a medium or large configuration using Oracle ZFS Storage Appliance requires only that the heads be swapped and the existing trays remain in place with the data intact.
- Increasing tape capacity is accomplished nondisruptively by adding slots and/or drives to the library. If changing from one library to another, the tape media are removed from the old library and inserted into the new library.

By following the guidelines and best practices outlined in this paper, organizations can save valuable time in deploying a content management solution while also reducing risk by implementing a single-vendor solution that addresses the full solution from application to disk. User productivity and agility is also achieved because the tiered storage solution provides dynamic and automatic access to 100 percent of the data.

References

For more information, visit the web resources listed in Table 7.

TABLE 7. WEB RESOURCES FOR FURTHER INFORMATION

WEB RESOURCE DESCRIPTION	WEB RESOURCE URL
• Oracle HSM configuration and documentation	http://www.oracle.com/us/products/servers-storage/storage/storage-software/storage-archive-manager/overview/index.html
• Oracle WebCenter Content documentation	http://www.oracle.com/technetwork/middleware/content-management/index-094708.html
• Oracle Solaris operating system	http://www.oracle.com/technetwork/server-storage/solaris/overview/index.html.
• Oracle Optimized Solution for Lifecycle Content Management	http://www.oracle.com/us/solutions/optimized-solutions-171609.html http://www.oracle.com/technetwork/server-storage/hardware-solutions/index.html

Appendix A: Small, Medium, and Large Reference Configuration Details

The hardware configuration for the small reference configuration contains four SPARC T4-1 servers. The hardware configuration diagram shown in Figure 11 includes the following:

- Two SPARC T4-1 servers for Oracle HSM and Oracle RAC in a high-availability configuration:
  - 64 GB DIMM for each server
  - One 8 Gb Fibre Channel HBA for each server
  - One optional 1Gb Ethernet quad port card
  - One optional 1Gb Ethernet quad port card (eight total with standard quad ports)
    - Two for storage network
    - Two for user and management network
    - One for Oracle Solaris Cluster
- Software
  - Oracle HSM running in four cores in one server (active/passive)
  - Oracle Solaris Cluster software license running in up to four cores in each server (active/active)
- Two SPARC T4-1 servers for Oracle WebCenter Content in a high-availability configuration:
  - 64 GB DIMM for each server
  - 1 Gb Ethernet
    - Two for storage network
    - One for HA network
    - Two for user and management network
  - Software
    - Oracle WebCenter Content running on up to eight cores in each server (active/active)
- One 1 Gb Ethernet Switch
- Two 8 Gb Fibre Channel SAN switches with 6 active ports per switch
- Tape storage
  - One StorageTek SL150 modular tape library (30 slots and one LTO 6 tape drive)
  - One 30-slot expansion shelf
  - Three additional StorageTek LTO 6 tape drives
- Disk storage
  - Sun ZFS ZS3-2 cluster
  - 48 GB DIMM per node
  - One 8 Gb Fibre Channel HBA (two ports) per node for target mode
  - Two disk trays
    - 40 three TB drives
    - Two SSD write flash accelerators

The configuration for the medium reference architecture deploys two SPARC T5-2 servers and provides a choice of either Oracle ZFS Storage Appliance or Oracle FS1-2. The hardware configuration diagram shown in Figure 12 includes the following:

- Two SPARC T5-2 Servers for Oracle HSM, Oracle RAC, and Oracle WebCenter Content
  - 256 GB DIMM for each server
  - Three 8 Gb Fibre Channel HBAs for each server (six ports)

- 1 Gb Ethernet for each server
  - One for failover heartbeat network
  - One for user and/or management network
- 10 Gb Ethernet for each server
  - Two for storage network
- Software
  - Oracle HSM software license running in 4 cores in one server (four total for active/passive)
  - Oracle Solaris Cluster software license for up to 12 cores in each server (24 total for active/active)
  - Oracle WebCenter Content for up to 16 cores in each server (up to 32 total for active/active)
- Two 10 Gb Ethernet Switches
- Two 8 Gb Fibre Channel SAN switches with 10 active ports each switch
- Tape storage
  - One StorageTek SL3000 modular library system
  - Six StorageTek T10000D tape drives
  - 900 activated slots
- Tape management
  - StorageTek Automated Cartridge System Library Software and server
  - StorageTek Tape Analytics software and server

The disk storage configuration provides a choice of either Oracle ZFS Storage Appliance or Oracle FS1-2.

- Disk storage (first option): Oracle ZFS Storage
  - Oracle ZFS ZS3-2 cluster
  - 256 GB DIMM per node
  - 10 Gb Ethernet per node
  - Two 8 Gb Fibre Channel dual port HBAs (four ports) per node for target mode
  - Two disk trays
    - Four SSD write flash accelerator and twenty 600 GB 15 K high-performance drives in one tray
    - 24 three TB capacity drives for the second tray
    - Additional trays can be added as capacity requires

- Disk storage (second option): Oracle FS1-2
  - One Controller
  - Four 8 Gb Fibre Channel Ports
  - Four instances of Oracle FS1-2 Disk Enclosures with 24 900 GB high performance drives in each enclosure
  - Four instances of Oracle FS1-2 Disk Enclosures with 24 4 TB capacity drives in each tray

The large reference configuration option 1 deploys two SPARC T5-4 servers and provides a choice of either Oracle ZFS Storage Appliance or Oracle FS1-2. The hardware configuration diagram shown in Figure 13 includes the following:

- Two SPARC T5-4 servers for Oracle HSM, Oracle RAC, and Oracle WebCenter Content in a high-availability configuration
  - 510 GB DIMM for each server
- Four 8 Gb Fibre Channel HBAs (eight ports) for each server

- 1 Gb Ethernet for each server
  - One for failover heartbeat network
  - Two for user/management network
- 10 Gb Ethernet ports for each server
  - Two for storage network
- Software
  - Oracle HSM software license for 8 cores in one server (8 total for active/passive)
  - Oracle Solaris Cluster software license for 24 cores in each server (up to 48 total for active/active)
  - Oracle WebCenter Content running on up to 32 cores on each servers (up to 64 total for active/active)
- Two 10 GbE switches
- Two 1 GbE switches
- Two SAN switches with 16 active ports per switch
- Tape storage
  - One StorageTek SL8500 modular library system
  - Twelve StorageTek T10000D tape drives
  - 1,750 activated slots
- Tape management
  - StorageTek Automated Cartridge System Library Software and server
  - StorageTek Tape Analytics software and server

The disk storage configuration provides a choice of either Oracle ZFS Storage Appliance or Oracle FS1-2.

- Disk storage (first option): Sun ZFS Storage
  - Oracle ZFS ZS3-2 cluster
  - 256 GB DIMM per node
  - Two 10 Gb Ethernet ports per node
  - Two 8 Gb Fibre Channel HBAs per node for target mode
  - Four disk trays
    - Four SSD write flash accelerators and twenty 600 GB 15 K high-performance drives in one tray
    - Twenty-four 3 TB capacity drives in each of three trays

- Disk storage (second option): Oracle FS1-2
  - Two instances of Oracle FS1-2 Controllers
  - Four 16 Gb Fibre Channel ports in each Controller
  - Eight instances of Oracle FS1-2 Disk Enclosures with 24 900 GB high-performance drives in each disk enclosure
  - Eight instances of Oracle FS1-2 Disk Enclosures with 24 4 TB capacity drives in each disk enclosure

The large configuration option 2 deploys four SPARC T5-2 servers and provides a choice of Oracle ZFS Storage Appliance or the Oracle FS1-2 storage system. The hardware configuration diagram shown in Figure 14 includes the following:

- Two SPARC T5-2 servers for Oracle HSM and Oracle RAC in a high-availability configuration
  - 256 GB DIMM for each server

- Six 8 Gb Fibre Channel HBA (twelve ports) for each server
- 10 Gb Ethernet
  - Two ports for storage network
- 1 Gb Ethernet
  - Two ports for user/management network
  - One port for failover network
- Software
  - Oracle HSM software license for eight cores on the first server (active/active)
  - Oracle RAC running on up to eight cores on each of the first two servers
  - Oracle Solaris Cluster license for up to eight cores on the first server

- Two SPARC T5-2 servers for Oracle WebCenter Content in a high-availability configuration
  - 256 GB DIMM for each server
  - 10 Gb Ethernet
    - Two ports for storage network
  - 1 Gb Ethernet
    - Two ports for user/management network
    - One port for failover network
  - Software
    - Oracle WebCenter Content running on up to 32 cores on each of the second two servers
- Two 10 Gb Ethernet switches
- Two Fibre Channel SAN switches
- Tape storage
    - One StorageTek SL8500 modular library system
    - Twelve StorageTek T10000D tape drives
    - 1,750 activated slots
- Tape management
  - StorageTek Automated Cartridge Library System Software and server
  - StorageTek Tape Analytics software and server

The disk storage configuration provides a choice of either Oracle ZFS Storage Appliance or Oracle FS1-2.

- Disk storage (first option): Sun ZFS Storage 7420 appliance
  - Sun ZFS ZS3-2 cluster
  - 256 GB DIMM per node
  - Two 10 Gb Ethernet ports per node
  - Two 8 Gb Fibre Channel HBAs per node for target mode
  - Four disk trays
    - Four SSD write flash accelerators and twenty 600 GB 15 K high-performance drives in one tray
    - Twenty-four 3 TB capacity drives in each of three trays

- Disk storage (second option): Oracle FS1-2
  - Two instances of Oracle FS1-2 Controllers
  - Four 8 Gb Fibre Channel ports in each Oracle FS1-2 Controller

- Eight instances of Oracle FS1-2 Disk Enclosures with 24 900 GB high-performance drives in each disk enclosure
- Eight instances of Oracle FS1-2 Disk Enclosure with 24 4 TB capacity drives in each disk enclosure

Appendix B: Testing Tools and Commands

Two test applications were used during this testing: File System Test (also called FS Test or fstest) and Faban.

The FS Test called Multiprocess File System Performance Test Tool (mpfstest), writes to an Oracle HSM file system. mpfstest aims to test a file system's basic functionality and I/O performance. Unlike fstest, which is a single process and single-file test tool, mpfstest is a multiprocess file system performance test tool that generates a multiprocess workload to measure and report a file system's read and write performance. The following areas are tested and measured by mpfstest:

- Multiprocess write performance test for files of fixed size
- Multiprocess write performance test for files of random sizes within a given range
- Multiprocess read performance test for files of any size
- The following information is reported:
  - Write rates
  - Read rates

Faban is an application that dives I/O through the Oracle WebCenter Content application, simulating user activity. For more information about Faban, refer to this URL: http://www.faban.org

The following commands dumps or restores Oracle HSM file control structure data

- The samfsdump command creates a dump file containing control structure information for each specified file.
- The samfsrestore command uses the contents of the dump file to restore control structures for all the files in the dumpfile or for each specified file.

For more information about samfsdump and samfsrestore, refer to this URL:
http://docs.oracle.com/cd/E22586_01/html/E22976/glaok.html

The following command checks and repairs a StorageTek QFS or Oracle HSM file system

- The samfsck command checks and optionally repairs a StorageTek QFS or Oracle HSM file system from the disk partitions that belong to family set name or a mount point.

For more information about samfsck, refer to this URL:
http://docs.oracle.com/cd/E22586_01/html/E22976/glalp.html#QFSRMglalp

Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

Extreme Scalability and Flexibility for Content Management Throughout Its Lifecycle
May 2013, Version 3.0
Author: Donna Harland