ORACLE
OPTIMIZED SOLUTIONS

Extreme Scalability and Flexibility to Securely Store and Access Digital Media Content

Oracle Optimized Solution for Digital Media

ORACLE TECHNICAL WHITE PAPER | JUNE 2015

ORACLE

Table of Contents

Introduction

Exponential growth in digital media and tight regulatory compliance requirements are making data retention and data access a great challenge. At the same time, companies in the media and entertainment markets are finding value in being able to immediately search for and access historical digital media. At the same time they are reporting on new events, there also is the demand to report on historical activity related to the current event. However, IT budgets are simply not growing fast enough to meet today's increasing storage capacity and performance requirements with only disk-based solutions. Offline tape backup is not a viable solution, since accessibility is key.

This paper focuses on how to deliver a scalable, flexible, and yet cost-effective solution using the architecture and best practices defined in Oracle Optimized Solution for Digital Media. This solution takes advantage of Oracle's broad portfolio of storage products, including intelligent flash storage systems, high-performance disk storage systems, and archive tape systems—with all the data managed by DIVArchive software. The solution also utilizes the compute power, security, and I/O features in Oracle's X86 servers to provide a robust platform for managing digital media content from very small implementations to very large implementations with petabytes of data. In addition, the DIVArchive open API enables integration with all major video server, workflow, automation, traffic, media asset management, editing, and publishing platforms. Third-party transcoders are also qualified and tested with the solution to give users easy access to information that is archived in a legacy format so that assets can be shared between workflow silos, or content can be automatically published in multiple formats.

Oracle Optimized Solution for Digital Media is also designed to greatly simplify deployment and management and to provide guidelines for component selection based on performance and capacity requirements. Additional information about the benefits of simplification and scalability for users as well as IT staff can be found in the companion solution brief, "Increasing the Value of Media Assets While Reducing Cost and Risk."

Solution Technical Objectives

Oracle Optimized Solution for Digital Media orchestrates powerful object storage archiving onto a tiered storage infrastructure consisting of high-performance network-attached storage systems and tape archive systems. All of the data on the storage infrastructure is managed by DIVArchive software. The solution is designed to accomplish a range of objectives:

- **Increase storage efficiency:** Organizations need to decrease overall storage costs over the lifetime of data. They need dynamic access to data from any storage tier, enabling valuable collaboration and reuse.
- **Manage digital media**: Explosive data growth is challenging organizations' abilities to cope and respond. This content is required to be ingested from multiple sources and be delivered on demand; sometimes only a partial file is required.
- **Automate data placement**: Organizational decision-making requires accelerating the valuable data discovery process for making both tactical and strategic decisions based on current and historical data. Collaboration is vital and users need to be able to reuse and share digital media more expediently by eliminating the time-consuming search and restore process from a backup.

- **Survive digital media transformations and technology change**: Digital media content, coming from multiple sources, is in multiple formats. It is also distributed in multiple formats and must be transcoded on the fly. Overtime, these formats age and must again be transcoded when accessed. Technology changes also require digital content to be migrated onto new storage technology.

Oracle Optimized Solution for Digital Media

Oracle Optimized Solution for Digital Media (shown in Figure 1 below) takes advantage of the robust capabilities of DIVArchive software for managing content on storage tiers. The solution includes the following major area components:

- DIVArchive software including DIVArchive Manager, DIVArchive Actor, DIVAnet, and DIVAdirector
- Oracle ZFS Storage Appliance with hybrid pools of hard disk drives (HDDs) and solid-state drives (SSDs)
- Oracle's StorageTek modular tape library systems with Oracle's StorageTek LTO 6 or StorageTek T10000D tape drives

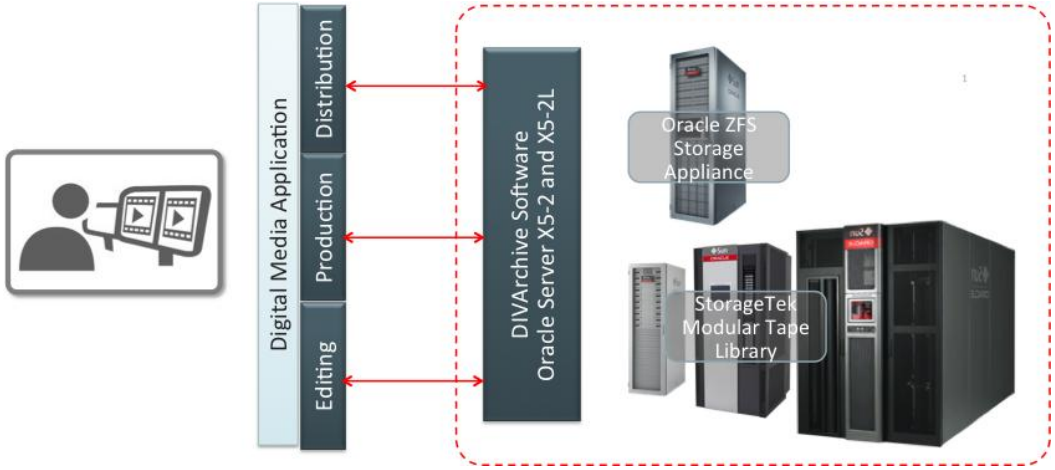

Figure 1: DIVArchive is tightly integrated with all major video server, workflow, automation, traffic, media asset management, editing, and publishing platforms and manages digital media stored on flash storage, primary drives, tape archive.

This solution provides a very scalable, flexible, and cost-effective storage platform to manage any step along the workflow of digital media. Such broad use demands an infrastructure that scales for both performance and capacity. The use of Oracle's disk storage systems and Oracle Server X5-2L delivers deployment flexibility and high performance. The StorageTek tape libraries provide long-term storage preservation enabling nondisruptive expansion.

The architecture can be logically divided into the following categories:

- **Data management:** DIVArchive software is a powerful policy engine that enables organizations to automatically move data to the appropriate storage tier based on access requirements.
- **Server infrastructure**: Oracle Server X5-2 and Oracle Server X5-2L servers provide the platform that runs the DIVArchive software.
- **Tiered storage**: Tiered storage includes a range of network-connected storage devices consisting of Oracle ZFS Storage Appliance and StorageTek tape libraries.

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

Data Management: Oracle DIVArchive Software High-Level Overview

The following is a high-level overview of a subset of the software that manages digital media in a very efficient and scalable solution:

- **DIVArchive Manager:** This main component in a DIVArchive system controls all archive functions. All user requests from the API come through DIVArchive Manager, which then instructs the DIVArchive Actor to run the actual data movement from user into the archive and back. Metadata is maintained in Oracle Database and includes all source information and archive location information, as well as an open format 4,000-character field for comments.
- **DIVArchive Actor**: This data mover carries out the various actions between the user environement and the storage devices. This function also can provide optimized file transfer or transcoding, and for every read and/or write, a checksum is calculated using SHA1 or MD5 to assure the data is unchanged throughout the workflow process. DIVArchive Actor takes all instructions from DIVArchive Manager.
- **DIVAnet**: This global connector links multiple DIVArchive systems together to support a variety of workflows for large globally networked systems that enables content sharing across the complete system. Moving content seamlessly amount these sites enables disaster recovery strategies that support business continuity. Under DIVAnet, each DIVArchive site operates independently with no dependency on other sites. DIVAnet assures that data is always available and will instruct from which location to retrieve the content.
- **DIVAdirector:** This software is directly integrated with DIVArchive to ensure content is easily searchable using simple descriptive asset-level metadata and thumbnails. This secure function allows the user to create frame-accurate shot lists to assist in further editing. Once the project is completed, content can be delivered to any DIVArchive-managed workflow, including online publishing through DIVApublish or integrated to other editing systems.
- **Archive eXchange Format (AXF):** This open standard for file archiving, preservation, and exchange includes features necessary for true preservation operations as defined in the Open Archival Information System (OAIS) model. AXF is a file system within an object that provides the ability to archive one file to a whole collection of related files in a single container or object. This concept is similar to the tape archive (TAR) format, only with many advanced features. This AXF container includes the metadata that describes the content, thus preserving not just the content but also the description and provenance of that content. Included in the AXF object is the individual file checksum, which ensures validity and is used for comparison on every read and write. If the metadata in the database is destroyed, it can be extracted from these objects and reloaded into a database. AXF media and objects can be moved between systems that comprehend AXF in a manner very similar to that offered by TAR.

  To summarize the AXF format of the container, each object contains:

  - Object Header: includes the object's universally unique identifier or unique material identifier, creation date, object provenance and file permissions, and paths
  - Generic Metadata Package(s): optional to store object specific metadata supplied by the application self-contained and can be structured, unstructured, open, vendor specific, binary, or XML
  - File Payload: the actual content
  - File Footer: contains the exact size of the file and each file checksum

  The paper, "An Overview of Archive eXchange Format, an Open Standard for File Archiving, Preservation, and Exchange" provides additional information about AXF.
- **Additional Features:** Other functions and features of the DIVArchive software are available within the workflow. The infrastructure described in this paper supports additional functions that include:

- The application filtering feature of Oracle DIVArchive: This is a security feature that allows multiple customers, both external and internal, to utilize the same archive components, such as tape and tape drives, as well as the DIVArchive software, yet protect the content, isolating it to predefined users.
- The protect feature of Oracle DIVArchive: This feature monitors and tracks resources, providing administrators with analytical reports in order to be proactive about failure and system monitoring and deliver on the SLAs.
- The multiple restore feature of Oracle DIVArchive: This feature allows for a single read of content from tape and distribution to multiple destinations or video servers. A use case for this is when there is a primary video server and multiple backup servers. The content will be read once from tape and distributed to all video servers.
- The partial file restore feature of Oracle DIVArchive: This time- and resource-saving feature allows a user to select a partial file based on timecode in and timecode out, transferring only the part of the file required by the user. The use case is when a user needs only 10 minutes in a one-hour news clip.
- Transcoding using third party transcoder applications: Oracle DIVArchive can automatically transcode content inline during archive or restore as part of the overall workflow. This is a great feature that's relevant when formats change and older content needs to be converted to new formats. Another use case is distribution of multiple formats starting with a mobile device and distribution to the large, high-resolution iMax theater.
- The paper, “Unique Features in the Oracle DIVArchive Solution” provides more details on these features and additional valuable features.

This link provides complete details about the DIVArchive software and all of the features: https://www-sites.oracle.com/content-storage-management/solutions/diva-content-storage/divarchive/index.html

Server Infrastructure: Oracle Server X5-2 and Oracle Server X5-2L

The server infrastructure deployed for Oracle Optimized Solution for Digital Media is based on x86 series servers from Oracle. The DIVArchive applications are both throughput- and CPU-demanding; therefore, the flexibility and scalability of Oracle Server X5-2 and Oracle Server X5-2L make them a perfect fit for this solution.

For performance, the servers provide up to 2x Intel® Xeon® processor CPUs in 8-, 10-, 12-, or 18-core options and 24x DDR4-2133 dual inline memory modules (DIMMs), allowing memory to scale up to 768 GB.

For configurations that use onboard disk, both servers support drive slots for up to eight 2.5-inch SAS-3 drives. For configurations in which additional onboard disk capacity makes sense, Oracle Server X5-2L can scale to a 12x 3.5-inch drive cage that supports 4 TB capacity disks or a 24x 2.5-inch drive cage that has several performance drive options.

For connections to archive devices such as the StorageTek T10000 tape drives, Oracle Server X5-2 has 4x PCIe 3.0 slots (2x 16-lane, 2x 8-lane PCIe Gen 3), and Oracle Server X5-2 has 6x PCIe 3.0 slots (2x 16-lane, 4x 8-lane PCIe Gen 3) to support the 16 Gb Fibre Channel HBA interface cards.

Connections to the user’s application servers for accessing and writing digital content is through an IP-based protocol such as FTP, SMB, or NFS. Oracle Server X5-2 and Oracle Server X5-2L both have 4x Ethernet 10 GBase-T ports included. Additional ports can be added through a PCI-based interface card for additional copper ports or for optical ports.

The solution, Oracle Optimized Solution for Digital Media focuses on the use of the servers in the archive part of the digital media lifecycle; however, these servers also meet the demands of the active part of the lifecycle that includes production, postproduction, distribution, and all other activity in the active workflow. Using this server to support software outside this solution, extends the value of a single vendor and proven solution to include the complete lifecycle workflow and into the

active part of the lifecycle. A complete, powerful infrastructure is now available to support all steps in the workflow of processing digital media from camera to archive.

Tiered Storage Options

Tiered storage is a critical infrastructure as content must be kept for long periods, yet some use cases require fast ingest as well as fast access for recently ingested data. Oracle has two tier 1 storage products that meet the requirements of a tiered storage environment. DIVArchive software works best with shared storage; therefore, Oracle ZFS Storage Appliance is well suited to meet the requirements of Oracle Optimized Solution for Digital Media. In addition, Oracle offers the most complete line of tape systems for low-cost, high-reliability archiving.

» **Oracle ZFS Storage Appliance:** Oracle ZFS Storage Appliance is the disk storage for both the recently saved data and the most active data, such as data in the production workflow being rendered, no matter how old the data. This storage supports mixed combinations of high-performance SSD caching and various-speed, larger-capacity HDD storage for DIVArchive software and Oracle Database. Oracle ZFS Storage ZS3-2 entry-level engineered storage system delivers extreme efficiency and reduces cost, complexity, and risk while meeting high ingest and access requirements, thereby providing primary storage as well as disk archive for small configurations. Oracle ZFS Storage ZS4-4 is a larger-scale engineered NAS storage system that delivers reduced complexity and risk for enterprise customers demanding high-performance storage with extreme efficiency and low TCO for the largest and most demanding workloads.

  This solution focuses on the infrastructure for the archiving of digital media; however, in the workflow outside of the archive process, there is great demand for large storage farms with high performance and availability. These rendering storage farms are used to support the production, postproduction, and distribution workflows that are part of the complete content lifecycle of digital media. Oracle ZFS Storage ZS4-4 is a proven storage solution to meet these demands of extreme performance and high availability at a low cost. This storage option extends the single-vendor, proven solution outside the archive workflow of the digital media lifecycle and into the active workflow of the lifecycle.

  More information on Oracle ZFS Storage Appliance can be found at the following website: http://www.oracle.com/storage/nas/index.html.

» **Oracle's StorageTek tape and library systems.** StorageTek tape and library systems help organizations maximize secure data access, manage complexity, and control costs. Tape is used for archival of data and provides efficient access to all of the data, regardless of where it is stored or the retention period, while also providing data protection through multiple copies. The StorageTek T10000D Data Integrity Validation feature is based on ANSI standard cyclic redundancy checks (CRCs). This capability provides additional data protection by validating that what was sent to tape is what was actually written. If inactivity or environmental factors have deteriorated the media, DIVArchive is notified and a new archive copy is created from an alternate copy. More information on StorageTek tape libraries and tape drives can be found at the following website: http://www.oracle.com/goto/tape.

Security for Oracle Optimized Solution for Digital Media

In a 2015 survey by EiQ Networks of 145 CIOs and top IT professionals, 90 percent indicate concern they will face a data breach while only 15 percent report they are "well prepared."[1] Oracle designs, develops, and delivers security in every layer within the solution. Oracle also coengineers security across all layers. Solutions that include products that are developed, tested, and proven to work together offer a unique advantage in the ability to protect their data over solutions that utilize a mixed-vendor environment.

There are many facets to security in the digital world. The following list includes a few points of security that protect digital media content from unwanted or illegal access and accidental or on-purpose destruction and provide the highest level of availability:

1 Source is EiQ Networks annual Security Monitoring, and SIEM study for 2015. http://www.eiqnetworks.com/abouteiqnetworks/news/pressrelease/2014/'Data-Breach'-Top-Concern-of-IT-Security-Pros-for-2015.php

- DIVArchive Application Filtering: The access gateway feature within the Oracle DIVArchive software, either running on the same server as the DIVArchive Manager or on a separate server, acts as a manager proxy. This allows several applications to share the archive resources such as a tape library and tape drives while keeping their respective content private. Each application is only allowed to handle the content to which it has access. Additional information can be found in the Oracle paper, “Oracle DIVArchive Application Filtering.”
- DIVAnet: This option provides disaster recovery on a global scale. Content is replicated across independent, active DIVArchive installations, locally and remotely, based on securely defined policies and access controls within each site. Additional information can be found in the white paper, “Protect Your Content with a Distributed Disaster Recovery Solution.”
- Data encryption on tape: Oracle Key Manager provides a comprehensive platform for key management of encrypted data. This provides FIPS 140-2 Level 3 compliance and secure key protection throughout key lifecycles, with a dedicated key management and delivery network. Additional information can be found at “Oracle Key Manager 3.”
- DIVArchive Storage Plans: This defines the policy that applies to where an object is to be archived and how many copies are created, with a default of two. It provides a level of security through assurance that all copies are always in place. If a copy is accidently removed, the storage plan automatically creates a new copy.
- End-to-end encryption: Oracle ZSF Storage Appliance has optional inline data services that include highly secure, easy-to-implement two-level AES 256/192/128/-bit granular data encryption at the project, share, or LUN level and key management flexibility for data breach protection and security,
- Oracle Database: The DIVArchive metadata is stored in an Oracle Database that has had security options built in for many years and contines to include new security enhancements with every release. It stays ahead of the latest ‘hacking’ methods. These features enable Oracle Database to comply with privacy and regulatory mandates. Although the digital media itself is not stored in Oracle Database, this level of security protects the content description and location in the archive, which prevents unwanted and illegal accesses. The solution architecture is configured to run active/passive on two servers in the small, medium, and large configurations, with the passive server configured as failover to assure access to content.

Sizing the Solution

Two main points must be taken into consideration when selecting a hardware configuration for Oracle Optimized Solution for Digital Media: ingest/access performance and archive capacity. The first point, ingest/access performance, is calculated by looking at all activity from the applications that are putting content into the archive. Users should calculate the total ingest rate per day and determine how many devices and servers are required to meet that ingest rate. The requested number of accesses is added to the total requirement. The second point, capacity, is a simple calculation of ingest rate/day times number of days to be retained. This is frequently a forever retention period so a good practice is to keep enough archive capacity for one year and scale the capacity each year. Tape archive storage scales nondisruptively.

Oracle performs testing to provide guidance for selecting a configuration that most closely meets current requirements. This testing has resulted in the following configuration categories that match a range of ingest performance and capacity requirements. The goal is to architect a configuration that can archive two copies, requiring two reads from disk, of the content and never gets behind and runs out of disk space. The ingest and write to disk never can be faster than the archive activity. Scalability of a solution is focused on accomplishing this smooth movement of digital media through the system, leaving enough capacity and performance to access the digital content to copy to a requesting user application environment.

Capacity Considerations

When capacity is calculated based on ingest rates running 24/7, 365 days per year, it quickly grows beyond the typical projection based on 50 percent growth per year. Table 1 represents capacity of content with a retention of seven years based on proven continuous ingest rates for 24 hours per day, seven days per week, ingesting at the maximum rate to the given number of tape drives. Remember the ingest rate is half as fast as the archive rate because two copies are being archived. Following this “rule of thumb,” the solution stores content on disk generally for two days and 200 percent on tape.

The ‘normal’ processes have peaks and quiet times. It's necessary to understand the peaks and configure for performance and to understand the quiet times and configure for capacity.

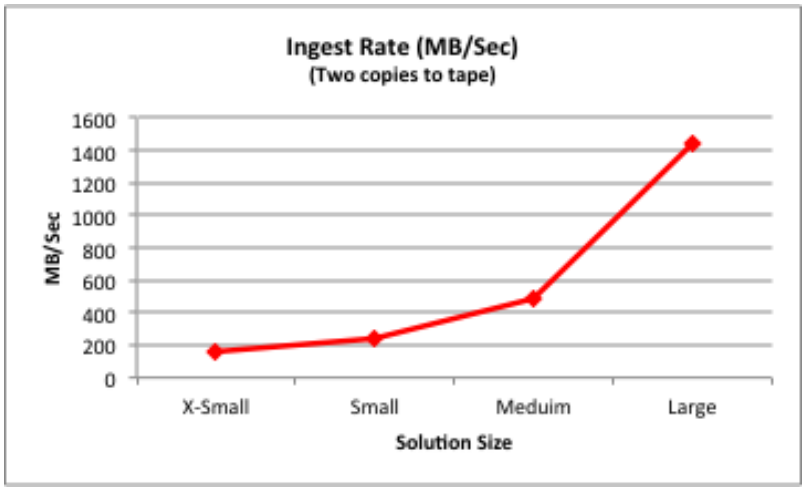

Figure 2. Capacity for one year based on ingesting digital media 365 days/year, 24 hours/day.

Realistically, the maximum ingest rates used for calculations, which are represented in Figure 2, will be experienced only at peak times, not 365 days/year; however, they must be processed with little or no impact to users during those peak times. The total ingest capacity is more likely a much smaller percentage of the totals shown. To achieve a solution that meets both performance and capacity requirements, the x-small, small, medium, or large configuration should be selected based on the peak ingest and access rates required. Total capacity should be selected based on actual expected growth and retention time. The scalability and flexibility of the solution enables the infrastructure to grow to meet new requirements.

All requirements must be taken into consideration when selecting an x-small, small, medium, or large configuration, with performance and capacity being the primary factor. The Oracle Optimized Solutions team has proven that the components work together, and the test results provide guidelines for size selection. With DIVArchive, Oracle ZFS Storage Appliance, and StorageTek tape systems providing tools for migrating data nondisruptively, it is possible to begin with one configuration size and easily and confidently move to the next size.

Performance Considerations

DIVArchive Actor functions as the data mover among the user application, storage infrastructure, and the archive devices. DIVArchive Manager functions as the traffic cop, directing to DIVArchive Actor the digital media to store or access, from where and to where. As digital media is ingested, it is immediately written to tape with the best practices and default of creating two copies on tape. For performance, it is a balancing act to write the digital media to disk and then read it twice in order to write two copies to tape. The bottleneck is not writing the original content to disk and then reading it twice for archive. The bottleneck is creating the checksum and comparing it to the known checksum for every read and write. This is a key and differentiating feature of DIVArchive that ensures the digital media ingested today is the same digital media accessed 10, 15, or 20 years from its first write. This is a multithreaded operation, and therefore, more cores can improve performance.

The following two graphs in Figure 3 indicate the maximum throughput for each configuration. The requirements used for configuration guidelines are for the peak time performance.

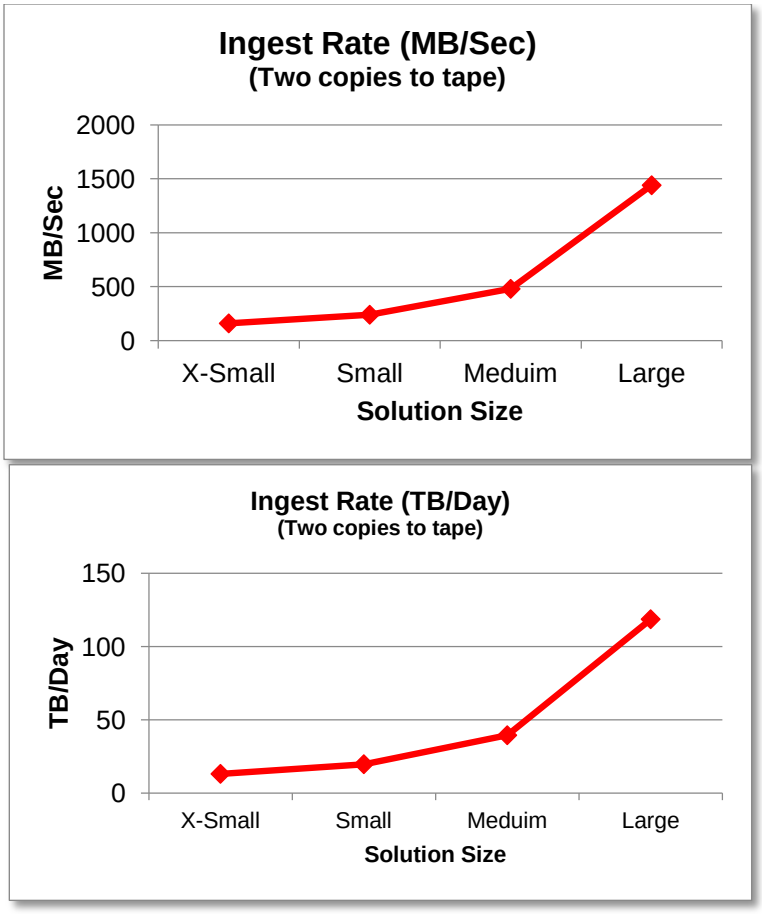

Figure 3. The maximum ingest rates for each solution size are represented in MB/sec and TB/sec.

Storage Hardware Considerations

Oracle ZFS Storage Appliance systems are configured with tiers of solid-state storage including large DRAM pools—up to 2 TB in dual-controller configurations—and both read and write cache areas use flash memory. The Oracle ZFS Storage ZS3 appliances are available with single- or dual-controller options in two basic models:

- The Oracle ZFS Storage ZS3-2 appliance is an entry-level engineered storage system for smaller deployments equipped with up to 15 TB of cache (2 TB DRAM, 13 TB flash) and as much as 1.5 PB of capacity per cluster.
- The Oracle ZFS Storage ZS3-4 appliance is an engineered storage system for larger deployments, with up to 25 TB of cache (2 TB DRAM, 23 TB flash), and it scales up to 3.5 PB of raw uncompressed capacity per cluster.

The highly scalable StorageTek tape libraries ensure data availability in heterogeneous tape storage environments of any size, including those that comprise Oracle Applications, Microsoft Windows servers, mainframes, and supercomputers. The StorageTek tape libraries that are proposed in the x-small, small, medium, and large tiered-storage solutions scale from 30 to 100,000 slots (75 TB to 850,000 TB), meeting virtually all capacity, archive, and access requirements.

The following sections provide both capacity comparisons for x-small, small, medium, and large for Oracle Optimized Solution for Digital Media. This discussion can help organizations select the appropriately sized configuration to meet initial capacity requirements, and the information shows the nondisruptive scalability of the solution as business needs grow and data capacity requirement increases.

Scalability

Oracle Optimized Solution for Digital Media can scale capacity and performance independently or together. When the recommended maximum number of tape drives attached to the DIVArchive Actor, the software easily scales horizontally. Additional DIVArchvie Actors with additional tape drives attached will scale to the new performance requirements. The addition of slots to the library will easily increase the capacity in the archive.

Configurations

The following configurations are based on testing and evaluating the best configuration for the size. Ingest and restore rates are included; therefore, with knowledge of the throughput profile at peak and slow times, users can select the correct configuration. These start at the x-small size, providing small organizations with a cost-effective solution that gives them an enterprise-class archive solution. This scales up to the large configuration for the customers that process large amounts of data 365 days/year, 24 hours/day.

X-Small Configuration

Figure 4 below shows the x-small reference configuration with its starting storage capacity for tiered storage. All DIVArchive software is running in the same server. There is no failover server in this configuration, and the storage is based on onboard disks in a 24-drive cage configuration of Oracle Server X5-2L.

- Ingest rate: 160 MB/sec, creating two tape copies
- Restore rate: 320 MB/sec
- Server: Oracle Server X5-2L with 24-disk drive cage
- Tape: StorageTek SL150 modular tape library
    - 2 StorageTek LTO tape drives
    - 30 tape slots; 78 TB
- Disk: onboard disk storage
    - Twenty-four 1.2 TB disks with the option of RAID
- Scalability:
    - Disk capacity: limited to the available slots in the server, or the small configuration can be referenced
    - Tape capacity: scale the slots to 300 for a total capacity of 780 TB
    - Server performance: scale horizontally, adding another server—the small configuration can be referenced

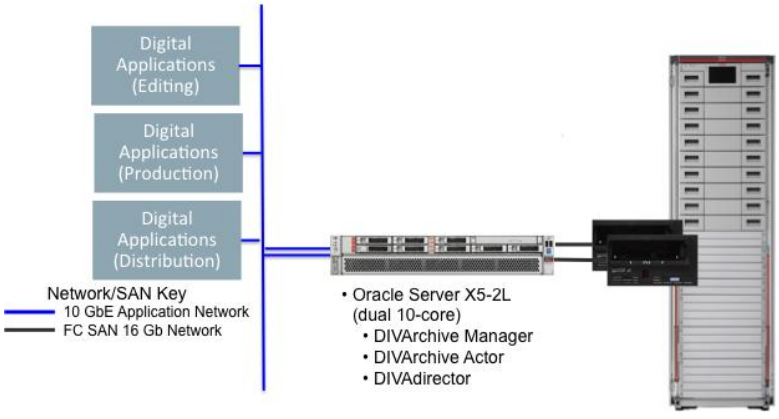

Figure 4. Above is the x-small reference configuration with the StorageTek SL150 modular tape library.

Small Configuration

Figure 5 below shows the small reference configuration with its starting storage capacity for tiered storage. DIVArchive Manager is running on Oracle Server X5-2 server 1 and DIVArchive Actor is running on Oracle Server X5-2 server 2. Each server is configured as the failover server to the other to provide a highly available configuration.

- Ingest rate: 240 MB/sec, creating two tape copies
- Restore rate: 480 MB/sec
- Servers:
    - Oracle Server X5-2 and DIVArchive Manager; failover for DIVArchive Actor
    - Oracle Server X5-2 and DIVArchive Actor; failover for DIVArchive Manager
- Tape: StorageTek SL3000 modular library system
    - 2 StorageTek T10000D tape drives
    - 1,245 tape slots; 10,582 TB
- Disk: Oracle ZFS Storage ZS3-2
    - 1 capacity disk enclosure with 20 x 4 TB disks and 2 SSD write accelerators
- Scalability of small configuration
    - Disk capacity: scale to 1.6 PB
    - Tape:
        - Capacity: scale the slots to 5,925 for a total maximum capacity of 50,360 TB
        - Performance: scale to four tape drives, two each server
    - Server performance: scale horizontally, adding another server—the medium configuration can be referenced

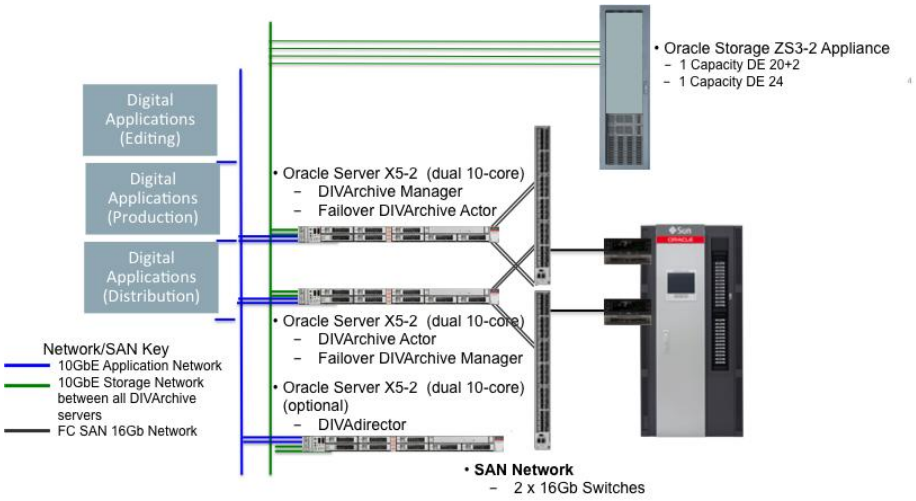

Figure 5. Above is the small reference configuration with the Oracle ZFS Storage ZS3-2 appliance and StorageTek SL3000 modular library system.

Medium Configuration

Figure 6 below shows the medium reference configuration with its starting storage capacity for tiered storage. DIVAarchive manager, including Oracle Database, is running on Oracle Server X5-2. DIVArchive Actor is running on the two Oracle Server X5-2L servers providing highly available data archive functions. The DIVArchive manager will fail over to one of the Oracle Server X5-2L servers, providing highly available data management functions.

- Ingest rate: 480 MB/sec, creating two tape copies
- Restore rate: 960 MB/sec

- Server:
  - Oracle Server X5-2 and DIVArchive Manager
  - 2x Oracle Server X5-2L and DIVArchive Actors
- Tape: StorageTek SL3000 modular library system
  - 4 StorageTek T10000D tape drives; two attached to each Oracle Server X5-2L server
  - 1,660 tape slots; 14 PB
- Disk: Oracle ZFS Storage ZS3-2
  - 1 capacity disk enclosure with 20 x 4 TB disks and 2 SSD write accelerators
  - 1 capacity disk enclosure with 24 x 4 TB disks
- Scalability of medium configuration
  - Disk capacity: scale to 1.6 PB
  - Tape
    - Capacity: scale the slots to 5,925 for a total maximum capacity of 50,360 TB
    - Performance: scale to six tape drives, three each server
  - Server performance: scale horizontally, adding additional servers—the large configuration can be referenced

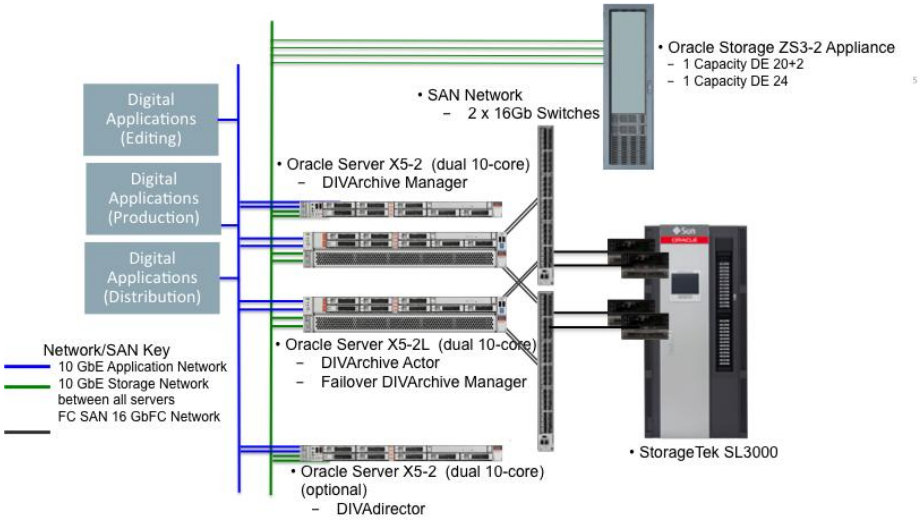

Figure 6. The medium reference configuration with the Oracle ZFS Storage ZS3-2 appliance and StorageTek SL3000 modular library system.

Large Configuration

Figure 7 below shows the large reference configuration with its starting storage capacity for tiered storage. DIVAarchive Manager, incuding Oracle Database, is running on Oracle Server X5-2 with a second Oracle Server X5-2 server as the failover server. DIVArchive Actor is running on the four Oracle Server X5-2L servers providing highly available data archive functions.

- Ingest rate: 1,440 MB/sec, creating two tape copies
- Restore rate: 2,880 MB/sec
- Server:
  - 2 Oracle Server X5-2 servers: one running DIVArchive Manager software and Oracle Database, and the second is a failover server

- 4 Oracle Server X5-2L servers running DIVArchive Actor software

- Tape: StorageTek SL3000 modular library system
  - 4 StorageTek T10000D tape drives, three attached to each Oracle Server X5-2L
  - 4,979 tape slots, 42 PB
- Disk: Oracle ZFS Storage ZS3-2
  - 1 capacity disk enclosure with 20 x 4 TB disks and 2 SSD write accelerators
  - 3 capacity disk enclosures with 24 x 4 TB disks
- Scalability of large configuration
  - Disk capacity: scale to 1.6 PB
  - Tape
    - Capacity: scale the slots to 100,000 for a total maximum capacity of 850,000
    - Performance: scale to 16 tape drives, 4 each server

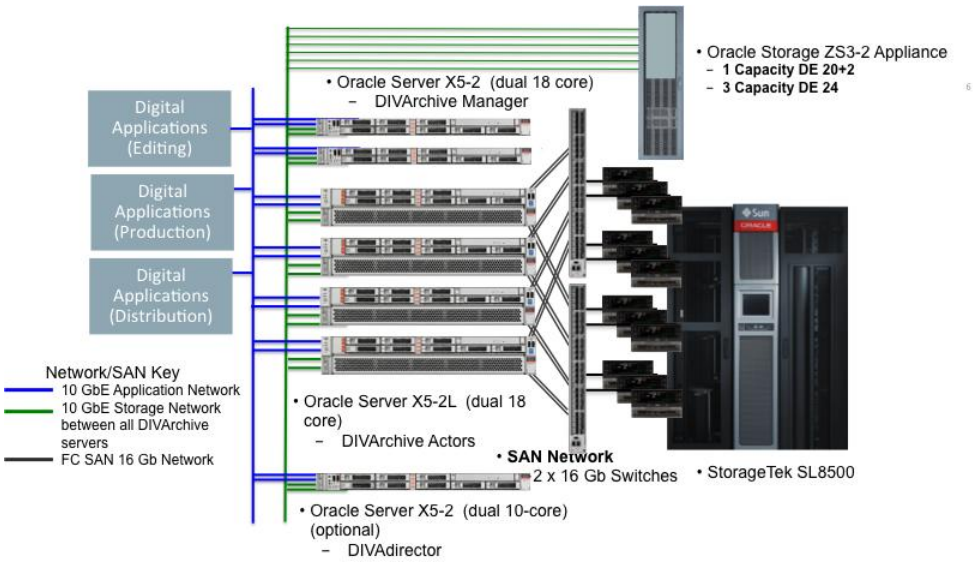

Figure 7. Above is the large reference configuration with Oracle Server X5-2, Oracle Server X5-2L, Oracle ZFS Storage ZS3-2, and the StorageTek SL8500 modular library system.

Summary Configuration

The following table is a summary of the capacity and expected performance of the different configurations described above.

Table 1. Capacity and performance of the four different configurations, each with a low end and high end of the configuration, scaling capacity and performance within the existing components.

Size		Expected Performance		Oracle Server		Tape System				Disk System		
Configuration Size		Ingest (MB Sec)	Restore (MB Sec)	Oracle Server X5-2	Oracle Server X5-2L	Library	Tape Drive	Qty	Slots	Disk Storage	Capacity Disk Enclosure	Disk Capacity (Usable)
X-Sm	Low	160	320	0	1	StorageTek SL150	LTO 6	2	30	In Server	24	28.8 Raw
X-Sm	High	240	480	0	1	StorageTek SL150	LTO 6	3	30	In Server	24	28.8 Raw
Small	Low	240	480	2	0	StorageTek SL3000	StorageTek T10000D	2	830	Oracle ZFS Storage ZS3-2	1	40
Small	High	360	480	2	0	StorageTek SL3000	StorageTek T10000D	3	1,245	Oracle ZFS Storage ZS3-2	2	59
Med	Low	480	960	1	2	StorageTek SL3000	StorageTek T10000D	4	1,660	Oracle ZFS Storage ZS3-2	3	79
Med	High	720	1,440	1	2	StorageTek SL3000	StorageTek T10000D	6	2,489	Oracle ZFS Storage ZS3-2	4	237
Large	Low	1,440	2,880	2	4	StorageTek SL8500	StorageTek T10000D	12	4,979	Oracle ZFS Storage ZS3-2	4	316
Large	High	1,920	3,840	2	4	StorageTek SL8500	StorageTek T10000D	16	6,639	Oracle ZFS Storage ZS3-2	6	560

Conclusion

Designed to address the challenges of rapid data growth and data management associated with active archiving of rich content, Oracle Optimized Solution for Digital Media automates data management processes to help organizations save time and money. The solution employs powerful, policy-based storage tiering and automated data movement to increase storage efficiency while reducing the risk of data loss and the risk of loss of access to data.

Oracle Optimized Solution for Digital Media optimizes storage efficiency by ensuring that data is always kept on the storage tier that is best for its preservation and for its availability. This automatic archiving of digital media into a secure and yet accessible archive increases value of the content by increasing revenue, simply by being the first to publish for collaboration or for repurpose.

Appendix I: Differentiate Between Oracle DIVArchive and Oracle Hierarchical Storage Manager

The software company Front Porch Digital, purchased by Oracle in the fall of 2014, has three main application areas:

- DIVASolutions, which protect, manage, store, and deliver rich media in an integrated and streamlined workflow
- LYNX, a media-grade private cloud solution enabling an integrated network for all digital content distributed worldwide
- SAMMA, an Emmy Award-winning migration process for converting content from analog to digital for production, monetization, or preservation

The Oracle Hierarchical Storage Manager (Oracle HSM) software has one major application area with many features and options. This is a very sophisticated hierarchical storage manager (HSM) that protects, manages, stores, and delivers unstructured data of any age from any storage location.

The following table lists a few of the major differences between DIVArchive software and Oracle HSM. The key in architecting the correct solution using the correct software is to understand the content that is to be archived and to understand the applications writing to the archive. The first question is, “Does the user application write to the DIVArchive API?” More than 150 major media and entertainment applications, from all areas of the business, support this API.

Table 2. Comparison of DIVArchive and Oracle HSM

Function	DIVArchive Software	Oracle HSM
Content Format	Object store	File system
File Access	Open standards AXF-based file format for asset protection, preservation, and accessibility. It is necessary for the application to support the industry-accepted API to write and access content. For access, content is not accessed directly by the end user application. DIVArchive puts the data into the user application storage space using standard IP protocols such as FTP, SMB, and NFS. Ingest also requires the DIVArchive software to extract the data from the end user application, generally using standard IP protocols such as FTP, SMB, and NFS.	Access through an open standard API such as NFS, CIFS, POSIX file system, HTTP, or OpenStack Swift. Files are accessed by the application directly from the Oracle HSM file system.
Metadata	DIVArchive metadata is stored in Oracle Database. This metadata consists of the object name and category, the source information, and the location in the archive. DIVAdirector metadata has additional content-based metadata that is extractable from the content or entered manually. This metadata is searchable.	Metadata is the file system characteristics stored in the POSIX file system metadata. Extended metadata is now available via the Oracle HSM API.

Function	DIVArchive Software	Oracle HSM
Archive Policy	Flexible archive policies, referred to as *storage plans*, are based on customer requirements. The recommended policy is to store content on disk and immediately create two tape copies. The content on disk will be removed once a high watermark is reached with file removal down to low watermark. The metadata remains in the database at all times.	Flexible archive policies are defined to create up to four copies on archive media potentially using all file system characteristics in defining the policy. The normal policy is to retain data on disk until the configurable high watermark is reached, which may be months or years based on access patterns.
Remote Archive	DIVAnet provides the function to asychronously place digital media into multiple DIVArchive sites globally, assuring content is always available. Content can be accessed from any site should the closest site fail.	Oracle HSM has a remote archive into which the metadata must also be copied and transferred in order to recover in the event of a total primary site failure.
Preservation	DIVArchive generates and checks a checksum upon every read and write to detect and correct corrupted content. Oracle's StorageTek Tape Analytics also can be used to check for ECC errors.	The Data Integrity Validation feature is tightly integrated with Oracle HSM, providing assurance that content will be readable when accessed and validation is run based on Oracle HSM policy.
File Transformation	The option exists for in-path content transcoding, processing, and rewrapping (on demand, archive, and restore). Multiple formats of transcoded images can be stored.	No file transformation.
Migration to New Technology	Migration tools exist that will migrate digital media objects from older storage technology to newer technology. For StorageTek tape storage, three generations of media can be read on the current release tape drive.	The recycle function is used to migrate data from old technology to new storage technology. For StorageTek tape, three generations of media can be read on the current release tape drive.
Partial File Retrieval	Using partial file retrieval via the DIVArchive API from tape based on timecode, only a subset of the file is retrieved and copied to the requesting application’s online storage.	If data copies are only on tape and have been released from Oracle HSM disk cache, the staging from tape will result in a complete file stage to Oracle HSM disk cache.

Function	DIVArchive Software	Oracle HSM
Storage Terminology Differences	• Online: data on the users application storage • Nearline: data on DIVArchive disk • Archived: data on DIVArchive tape • Offline: data on DIVArchive media outside the library	• Online: data on the user's application storage is the Oracle HSM file system • Online: data on Oracle HSM disk • Nearline: data on Oracle HSM tape • Offline: data on Oracle HSM media outside the library

Extreme Scalability and Flexibility to Securely Store and Access Digital Media Content
June 2015
Author: Donna Harland