ORACLE
OPTIMIZED SOLUTIONS

An Oracle White Paper
February 2013

Enable Data Collaboration Through Powerful Storage Efficiencies

Oracle Optimized Solution for Tiered Storage Infrastructure

ORACLE

Executive Overview

Rapid growth in digital data and tighter compliance requirements for data retention are driving many organizations to consider tiered storage solutions. Tiered storage is seen as a means to help control storage costs while enhancing data protection.

Exponential growth in storage requirements is making disk-only solutions impractical from a cost perspective. IT budgets are simply not growing enough to meet the today's increasing storage capacity and performance requirements with disk-based solutions. At the same time, offline tape archives are not a viable solution because compliance requirements and business usage requirements mean that data must be accessible on demand.

A tiered storage infrastructure blends the best of both worlds by taking advantage of the inherent cost savings and increased protection available through tape media while also delivering the necessary performance to keep up with user and application workloads. Balancing performance against cost requires a well-planned design that can meet the immediate performance and capacity requirements without extra costs and also provide a growth path that makes it easy to expand capacity and performance as needed.

As the leading provider of tape library systems and the number one enterprise database vendor, Oracle is in a unique position to leverage its knowledge and expertise to meet today's demanding requirements for data archives. The Oracle Optimized Solution for Tiered Storage Infrastructure offers the flexibility to scale both capacity and performance independently while also keeping storage costs low and providing reliable data protection over many years. The solution takes advantage of the product integrations inherent in Oracle's optimized hardware and software stack and also provides recommended configurations that have been further optimized to deliver specific performance levels for defined small, medium, and large workloads.

The Oracle Optimized Solution for Tiered Storage Infrastructure offers the following primary business advantages:

- **2.8x lower total cost of ownership (TCO) than EMC and 72 percent lower acquisition cost:**
  A comparison of acquisition costs and TCO for an EMC configuration that delivers comparable capacity shows that Oracle offers 2.8x better TCO than EMC over a five-year period and 72 percent lower initial acquisition cost. This significant cost difference is largely because EMC's tiered storage offering is 100 percent based on disk drives rather than taking advantage of tape media. In addition to lower acquisition costs, Oracle's solution also helps reduce deployment costs and ongoing operational costs. Pretested configurations can eliminate up to seven weeks of testing during solution design and deployment. Operational costs are also reduced due to the inherent space and power advantages of tape solutions.

- **Massive scalability for both capacity and performance:**
  Oracle's solution offers modular growth so that capacity and performance can be scaled independently with known results. Oracle engineers performed a series of performance tests to characterize the I/O workloads that can be achieved by each of three configurations for small, medium, and large deployments. The architecture of the system allows organizations to scale any configuration by adding additional disk drives, flash drives, or tape drives. Both the disk subsystem and the tape subsystem offer modular expansion so that the right capacity and performance can be easily achieved at any storage tier.

- **Reduced risk:**
  Three pretested configurations of this solution provide hardware and software components that are known to work together and offer a proven growth path for performance and capacity. The performance characterization results virtually eliminate the risk of purchasing an improperly sized configuration. In addition, there are numerous technical features in Oracle’s storage product portfolio to help protect against data loss. For example, the remote archive maintained by StorageTek Storage Archive Manager (StorageTek SAM) helps ensure data protection and continued data availability in the event of a site failure.

Data Trends and the Need for Tiered Storage

In its 2012 annual digital universe study, International Data Corporation (IDC) made a projection that the total amount of digital information will reach 40,000 exabytes by 2020. This represents an astounding 50x growth from the beginning of 2010 to the end of 2020. The study also points out that the proportion of data that requires protection is growing faster than the digital universe itself. IDC estimates that approximately one third of the data currently in the digital universe requires some type of security protection to help ensure privacy, to adhere to regulations, or to prevent digital snooping or theft. By 2020, more than 40 percent of this data is expected to require protection measures.[1]

In addition to these data growth and data protection trends, compliance regulations and the increased value of data for business intelligence are making data retention and access a greater requirement. Compliance regulations often dictate that specific types of data must be accessible for years, or even indefinitely. In some cases, businesses are keeping data because compliance rules are not clear enough. They might store all their data for long periods just to be sure that all retention requirements are covered.

In addition, more companies are finding value in being able to analyze historical data, thus making it appealing to save operational data and use it for business analytics for several years. Keeping the data in an easily accessible archive can also increase sharing for collaboration purposes so that data owners can cooperate with other users throughout the organization or even share beyond the organizational walls.

Balancing Cost Versus Capacity, Performance, and Access

These trends imply that organizations must be prepared to expand storage capacity and also define long-term data protection and retention strategies. A key factor in addressing these requirements is optimizing the costs involved. This means identifying the lowest-cost solution that can meet core requirements and still offer scalability for future growth. In addition, the archive must deliver sufficient performance to make the data archive valuable to users.

Data protection, too, must be done in a cost-effective way. This includes ensuring proper access controls throughout the data lifecycle and properly disposing of data at the end of its lifecycle. Of the data that organizations maintain today, much of it is never used after 90 days. Yet many companies maintain a 100-year archive. One reason that many organizations don't really use their archived data is that they do not have the right tools for search and retrieval. If organizations solve the issue of how to provide better access to archived data in a cost-effective way, this can also increase the value of the data and, thus, contribute to the growth of the company.

[1] "The Digital Universe in 2020," sponsored by EMC, December 2012. Available at http://www.emc.com/leadership/digital-universe/index.htm.

Additional Factors that Affect Costs for Retaining Data

According to a study by Horison Information Strategies, a three-tiered storage implementation using tape for tier-3 archival data can provide significantly lower acquisition cost compared to an all-disk implementation. For a 100-TB archive, Horison estimated the acquisition cost of a three-tier solution to be 2.1x lower than a two-tier disk solution and 3.7x lower than a single-tier disk solution.[2]

With respect to archive costs, it is also important to consider operating or operational expenditure(s) or expenses (OpEx) costs, which include the cost of power and cooling as well as the cost to manage the storage. As a disk farm grows to petabytes of content, these operational costs grow with it. However, since operating cost for tape is significantly less than those for disk, a three-tier archive offers much better TCO than disk-only solutions. An independent study by The Clipper Group found that 12-year TCO for tape-based solutions was, on average, 15x less than the TCO for disk-based solutions. The study also found that disk solutions used 238 times more energy over the 12-year period. In fact, the energy costs alone for a disk-based solution approached the total cost for tape.[3]

Better Value with Oracle Optimized Solution for Tiered Storage Infrastructure

The Oracle Optimized Solution for Tiered Storage Infrastructure is designed to maximize the value of the data archive by improving access to data while simultaneously reducing storage costs. To help reduce the risk of data loss and to maintain data access in the event of a disaster, the solution offers both local and remote storage locations with the option to archive up to four copies of each data file.

Oracle Optimized Solution for Tiered Storage Infrastructure

The Oracle Optimized Solution for Tiered Storage Infrastructure is designed to provide cost-effective data archives while delivering reliable access to data with scalable performance and capacity. The solution is based on a combination of disk and tape storage systems. The tape archive is based on Oracle's StorageTek modular library systems and the disk storage can utilize either Oracle's Sun ZFS Storage Appliances or Oracle's Pillar Axiom 600 storage systems. StorageTek Storage Archive Manager software provides dynamic management of data placement so that the entire environment functions as a highly efficient tiered storage infrastructure.

The solution also provides high availability (HA) and reduced risk by taking advantage of Oracle Solaris Cluster HA technology and the superior data protection enabled by the tape technologies in StorageTek modular library systems.

---

[2] "Tiered Storage Takes Center Stage," Horison, Inc. Available at http://www.horison.com/OracleTieredStorageTakesCenterStage.pdf.

[3] "In Search of the Long-Term Archiving Solution—Tape Delivers Significant TCO Advantage over Disk," The Clipper Group, December 23, 2010. Available at http://www.oracle.com/us/corporate/analystreports/corporate/clipper-tape-storage-tco-322730.pdf

Figure 1 provides a high-level architecture diagram of the solution. The solution consists of three configuration options for small, medium, and large environments. The recommended products for each of these configurations are detailed in Table 1, and an overview of the hardware and software components follows immediately below Figure 1.

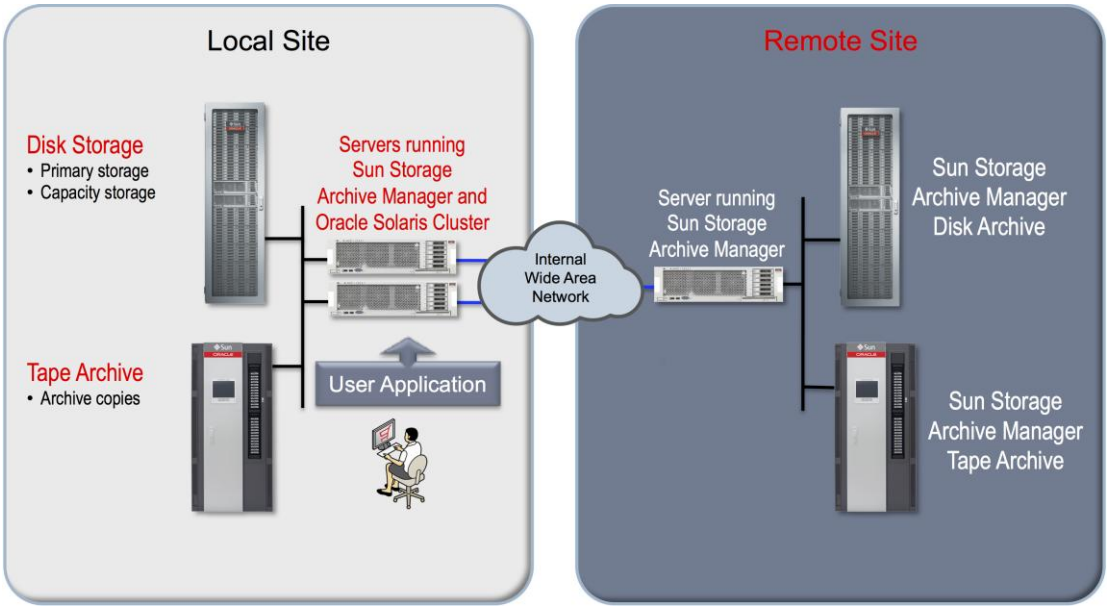

Figure 1. High-level architecture of Oracle Optimized Solution for Tiered Storage Infrastructure.

The following hardware and software components comprise the Oracle Optimized Solution for Tiered Storage Infrastructure:

- **Oracle's SPARC T-Series Servers**: Running the Oracle Solaris operating system, these servers provide a scalable and secure platform for running StorageTek Storage Archive Manager software. At the local site, two servers operate together in a cluster configuration running Oracle Solaris Cluster software.
- **Oracle's StorageTek Storage Archive Manager**: StorageTek Storage Archive Manager software provides a configurable file system with storage, archive management, and retrieval capabilities. The software delivers dynamic management of data placement on appropriate storage based on policies. It also automates the migration of data to new devices to provide continuity to the archive through data center transformations.
- **Oracle Solaris Cluster**: Oracle Solaris Cluster is a comprehensive HA and disaster recovery (DR) solution for Oracle's SPARC and x86 servers. It supports a broad spectrum of Oracle technologies and is the only HA solution that integrates with Oracle Solaris at the kernel level, enabling faster failure detection for servers and storage.

- **Oracle's Sun ZFS Storage Appliances**: The Sun ZFS Storage Appliance family delivers enterprise-class network attached storage (NAS) and block-attached storage area network (SAN) capabilities with leading Oracle integration, simplicity, efficiency, performance, and low TCO. The appliances utilize flash-enabled Hybrid Storage Pools to accelerate I/O response time by automatically placing active data in flash drives that offer very low latency. Sun ZFS Storage Appliances have demonstrated clear leadership in price/performance in audited SPC-1 and SPC-2 benchmark results[4].
- **Oracle's Pillar Axiom 600 storage systems**: Pillar Axiom 600 storage systems are based on a highly scalable architecture that can scale performance and capacity independently. Unlike traditional storage arrays that are limited to two storage controllers, the Pillar Axiom 600 system scales up to eight storage controllers. It also supports up to 128 RAID controllers in a single system, whereas traditional SAN architectures are often limited to two RAID controllers. Pillar Axiom's unique quality of service (QoS) capability enables the storage system to align storage resources as needed to provide guaranteed service levels for critical tasks requested from StorageTek Storage Archive Manager software.
- **Oracle's StorageTek tape solutions**: StorageTek modular library systems provide cost-effective long-term storage and data protection as part of the tiered storage architecture. Their performance and capacity is enhanced by the StorageTek T10000C tape drive, which delivers leading throughput and capacity, storing up to 5.5 TB of uncompressed data in a single tape cartridge. These tape drive efficiencies help improve I/O throughput while also enabling a high-density, low-cost tape archive. In addition, the Data Integrity Validation (DIV) feature of the StorageTek T10000C tape drive validates the data on a regular schedule and ensures that the data that was written is the same data that is read back.

[4] Sun ZFS Storage 7420c appliance SPC-1 results as of October 3, 2011 available at http://www.storageperformance.org/results/benchmark_results_spc1/#spc1. Sun ZFS Storage 7420 appliance SPC-2 results as of April 12, 2012 available at http://www.storageperformance.org/results/benchmark_results_spc2.

TABLE 1. COMPONENT OPTIONS FOR SMALL, MEDIUM, AND LARGE CONFIGURATIONS

COMPONENT	SMALL CONFIGURATION	MEDIUM CONFIGURATION	LARGE CONFIGURATION
		LOCAL SITE	
Primary disk storage (Select either Sun ZFS Storage Appliance or Pillar Axiom)	• Sun ZFS Storage 7320 with 120 TB of high-capacity disk for primary storage and disk archive and 292 GB write-optimized cache for primary storage	• Sun ZFS Storage 7420 cluster with 26.4 TB disk and 292 GB write-optimized cache - OR - • Pillar Axiom 600 with 28-TB disk	• Sun ZFS Storage 7420 cluster with 52.4 TB disk and 292 GB write-optimized cache -OR - • Pillar Axiom 600 with 57 TB disk and 1.5-TB solid-state disk (SSD)
Capacity disk storage (Select either Sun ZFS Storage Appliance or Pillar Axiom)	• Sun ZFS Storage 7320 with a total of 120 TB of high-capacity disk for archive storage and primary storage	• Sun ZFS Storage 7420 cluster with 144 TB High-capacity disk - OR - • Pillar Axiom 600 with 104 TB high-capacity disk	• Sun ZFS Storage 7420 cluster with 288-TB high-capacity disk - OR - • Pillar Axiom 600 with 208-TB high-capacity disk
Tape archive	Oracle's StorageTek SL150 tape library with 6 LTO tape drives for 90 TB; scales to 450 TB total capacity	StorageTek SL3000 modular library system with 6 StorageTek T10000C tape drives and 4.8 PB; scales to 15 PB total capacity	StorageTek SL8500 modular library system with 10 StorageTek T10000C tape drives and 16 PB; scales to 500 PB total capacity
Storage archive server(s) infrastructure	• 2 x SPARC T4-1 servers • StorageTek Storage Archive Manager license for 4 cores • Oracle Solaris Cluster for 4 cores	• 2 x SPARC T4-2 servers • StorageTek Storage Archive Manager license for 4 cores • Oracle Solaris Cluster for 4 cores • 2 x 10-GbE switch	• 2 x SPARC T4-2 servers • StorageTek Storage Archive Manager license for 8 cores • Oracle Solaris Cluster for 8 cores • 2 x 10-GbE switch
		REMOTE SITE	
Primary disk storage	• Sun ZFS Storage 7320 with 36-TB SAS II high-capacity disk and 1-TB write-optimized cache for primary storage	• Sun ZFS Storage 7420 cluster with 26.4-TB disk and 2-TB write-optimized cache - OR - • Pillar Axiom 600 with 28-TB disk	• Sun ZFS Storage 7420 cluster with 52.4-TB disk and 4-TB write-optimized cache - OR - • Pillar Axiom 600 with 57-TB disk and 1.5-TB SSD disk
Tape archive	StorageTek SL150 tape library with 6 StorageTek LTO tape drives and 90 TB; scales to 450 TB total capacity	StorageTek SL3000 modular library system with 6 T10KC tape drives and 4.8 PB; scales to 15 PB total capacity	StorageTek SL8500 modular library system with 10 T10KC tape drives and 16 PB; scales to 500 PB total capacity

TABLE 1. COMPONENT OPTIONS FOR SMALL, MEDIUM, AND LARGE CONFIGURATIONS

COMPONENT	SMALL CONFIGURATION	MEDIUM CONFIGURATION	LARGE CONFIGURATION
Storage archive server(s) infrastructure	• 2 x SPARC T4-1 servers • StorageTek Storage Archive Manager license for 4 cores • Oracle Solaris Cluster for 4 cores	• 2 x SPARC T4-2 servers • StorageTek Storage Archive Manager license for 4 cores • Oracle Solaris Cluster for 4 cores • 2 x 10-GbE switch	• 2 x SPARC T4-2 servers • StorageTek Storage Archive Manager license for 8 cores • Oracle Solaris Cluster for 8 cores • 2 x 10-GbE switch

2.8x Better TCO than EMC and 72 Percent Lower Acquisition Cost

Figure 2 graphically illustrates the savings available with this solution compared to a similar configuration based on EMC VMAX for primary storage and on Data Domain systems for archive storage. The five-year TCO is nearly 2.8x better for the solution than with EMC, and the initial investment saves 72 percent.[5]

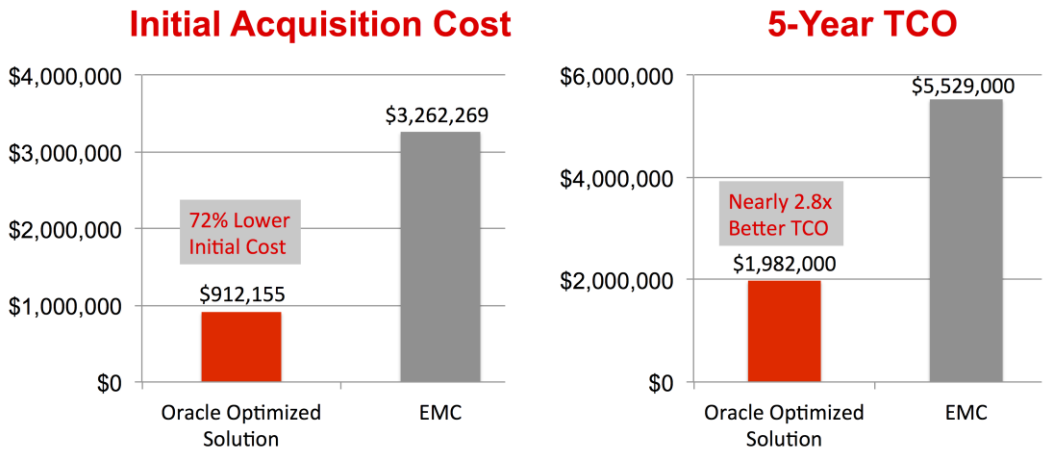

Figure 2. Cost savings with Oracle Optimized Solution for Tiered Storage Infrastructure.

Even though the Oracle solution has a dramatically lower acquisition cost, it provides 38 percent greater capacity than the EMC solution. Table 2 shows the raw storage capacity of the Oracle and EMC configurations used in the price comparison.

TABLE 2. USABLE CAPACITY OF ORACLE AND EMC CONFIGURATIONS USED IN PRICE COMPARISON

ORACLE		EMC	
COMPONENT	**CAPACITY**	**COMPONENT**	**CAPACITY**
High-performance disk in Sun ZFS Storage 7420 appliance	12 TB	High-performance disk in EMC VMAX SE	9 TB
High-capacity disk in Sun ZFS Storage 7420 appliance	72 TB	High-capacity disk in EMC VMAX SE	30 TB
Tape capacity in StorageTek SL3000 modular library system	500 TB (100 cartridges)	High-capacity disk configured as virtual tape in EMC Data Domain DD890	384 TB
TOTAL CAPACITY	**584 TB**		**423 TB**

[5] EMC pricing information sourced from December 1, 2012 EMC master price list available at http://www.emc.com/collateral/emcwsca/master-price-list.pdf

Table 3 summarizes an analysis of the annual costs associated with the EMC configuration compared to the annual costs for the Oracle solution. The total cost of ownership over five years is computed by summing these undiscounted annual costs over the five-year period. The major factors considered in the analysis of ongoing annual costs include costs for power and cooling, floor space, hardware service, and license fees for software or hardware components.

TABLE 3. TCO COMPARISON OVER FIVE YEARS

INVESTMENTS	INITIAL	YEAR 1	YEAR 2	YEAR 3	YEAR 4	YEAR 5	TOTAL
EMC solution undiscounted cash flows (US$K)	$3,262	$270	$321	$583	$427	$665	$5,529
Oracle solution undiscounted cash flows (US$K)	$912	$144	$158	$308	$204	$256	$1,982
Annual savings (US$K)	$2,350	$126	$163	$275	$223	$409	$3,547
EMC 5-year TCO	US$5,529K						
Oracle 5–year TCO	US$1,982K						
TCO savings with Oracle	72% (or 2.8X lower than EMC)						

Details about the cost comparison analysis, including specific configurations compared, can be found in Appendix A.

Reduce Deployment Time by up to Seven Weeks to Save Additional Costs

The Oracle Optimized Solution for Tiered Storage Infrastructure has been tested and optimized to define the best configuration options and parameters for each of the small, medium, and large configurations. Testing performed by Oracle engineers has enabled Oracle to document predefined configurations that virtually eliminate the need for onsite testing and optimization. This up-front work performed by Oracle enables organizations to save up to seven weeks during onsite deployment and setup.

The solution also includes an implementation guide that provides detailed instructions for how to configure each product so that the entire solution can achieve the required performance, protection, and access. The implementation guide also refers readers to additional product documentation and product installation guides to help simplify the process of finding the information needed during installation and setup.

Oracle Technologies that Enable High Performance at Lower Cost

Archival solutions require that data be accessible to users on demand, but cost is often a more important factor than performance. What is needed is a solution that can deliver good performance at low cost.

Oracle disk and tape storage systems enable superior performance at a lower price point than many competing systems. Sun ZFS Storage Appliances are consistently ranked among the highest performing storage solutions for a variety of workloads, and Pillar Axiom 600 systems are the only storage systems on the market that can enable up to 80 percent storage efficiency with no decrease in I/O performance. StorageTek modular library systems also enable high throughput at low cost, especially when taking advantage of the leading performance and capacity of the StorageTek T10000 tape drive.

Some of the key technologies that enable Oracle to deliver high performance at low cost include the following:

- Large memory and cache sizes in Sun ZFS Storage Appliances and Pillar Axiom 600 storage systems enable these storage systems to serve many requests from cache.
- Hybrid Storage Pools enable Sun ZFS Storage Appliances to automatically place data in memory, flash, or hard disk drives to optimize performance and cost.
- The Pillar Axiom 600 storage system supports prioritized storage I/O through its patented QoS technology, which enables deterministic performance even in consolidated storage environments.
- StorageTek modular library systems are a fraction of the cost of disk-based virtual tape solutions while offering superior data protection and enough performance to meet the needs of most archive applications.
- High-capacity StorageTek T10000C tape drives and media help reduce costs because fewer cartridges, drives, and libraries are required to meet the same performance and capacity requirements.

Massive Scalability for Both Capacity and Performance

Oracle's solution is designed to offer extensibility so that organizations can start with a cost-effective configuration and then scale both capacity and performance as their needs dictate. Of particular value is the fact that performance can be scaled independently of capacity.

Scale Performance Throughput to Nearly 3.8 GB/sec

Extensive testing has been performed by Oracle engineers to characterize the performance of archives that utilize StorageTek Storage Archive Manager software with StorageTek modular library systems and Oracle disk storage offerings. Testing has shown that read and ingest rates of nearly 3.8 GB/sec can be achieved for data that is being sent to or retrieved from the tape archive.

Based on the tested throughput rates for various configurations of StorageTek SL3000 and SL8500 modular library systems, Oracle engineers have designed a small, medium, and large configuration of the Oracle Optimized Solution for Tiered Storage Infrastructure. Table 4 shows the expected performance range for each of the three configurations. Each configuration has been optimized to achieve the best possible I/O throughput for the components that are used in the configuration.

Table 4 represents the estimated performance of the baseline configuration as recommended by Oracle. Addition of disk drives and/or FC paths will scale throughput between the configurations and addition of disk drives, tape drives and tape cartridges scales the capacity between configurations.

TABLE 4. PERFORMANCE THROUGHPUT RANGE FOR LOCAL SITE CONFIGURATIONS

PERFORMANCE GROWTH PATH FOR LOCAL SITE	SMALL CONFIGURATION	MEDIUM CONFIGURATION	LARGE CONFIGURATION
Expected ingest rates (baseline configuration to max)	Up to 2 GB/sec	Up to 2 GB/sec	Up to 3.8 GB/sec

Scale Capacity to More than 500 PB

The use of tape libraries in the Oracle Optimized Solution for Tiered Storage Infrastructure enables far greater capacity expansion than can be achieved with disk-only solutions. StorageTek modular library systems are known for their scalability, and the StorageTek SL8500 modular library system can be scaled up to 500 PB.

Table 5 shows the capacity ranges for each of the three configurations in the Oracle Optimized Solution for Tiered Storage Infrastructure. The first number in each capacity range represents the initial capacity of the baseline configuration as recommended by Oracle. The second number in the range represents the capacity of the configuration when the tape library has been scaled to its maximum capacity. The small configuration uses the StorageTek SL150 modular library system with LTO 5 drives. The medium configuration uses the StorageTek SL3000 modular library system and the large configuration uses the StorageTek SL8500 modular library system.

The disk capacity is a combination of high-performance disk for primary storage and high-capacity disk for disk archive. Therefore, the starting and ending capacities will be based on the requirements for ingest and access rates as well as the anticipated retention period. These capacities can be adjusted to meet those requirements. The ranges shown in Table 5 are, thus, an approximation of the expected capacities.

TABLE 5. CAPACITY RANGE FOR LOCAL SITE CONFIGURATIONS

SCALABILITY GROWTH PATH FOR LOCAL SITE	SMALL CONFIGURATION	MEDIUM CONFIGURATION	LARGE CONFIGURATION
Disk storage capacity range (baseline configuration to max)	132 TB to 432 TB	170 TB to 340 TB	340 TB to 1.72 PB
Tape archive capacity range (baseline configuration to max)	45 TB to 450 TB	1.5 PB to 15 PB	5 PB to 500 PB

Scale Performance as the Archive Grows

In most cases, a data archive is constantly expanding to store more and more data each year. While the need for increased capacity is obvious, it's also important to consider that archive performance might need to scale as the archive grows over time. As the data archive get larger, its value to the user community often grows as well. This generally means there will be more requests to retrieve data from the archive, and it might also mean that higher ingest rates are needed because more data is being stored in the archive. To meet these growth requirements, organizations need an archive that can grow incrementally in terms of both performance and capacity.

Since the Oracle Optimized Solution for Tiered Storage Infrastructure has multiple configurations, one option is obviously to upgrade from a small or medium configuration to the next larger configuration. In many cases, however, this is not necessary. There are a number of ways to add performance to any of the solution configurations without going through the process of swapping out system components to upgrade to the next level. It is best to get assistance from an Oracle representative to identify how a configuration can be expanded to meet specific performance goals. Some of the most common ways to increase performance of Oracle's archive solutions include the following:

- Add more flash or high-performance disk drives to the Sun ZFS Storage Appliance.
- Add more DRAM or high-performance disk drives to the Pillar Axiom 600 storage system.
- Add more StorageTek T10000C tape drives to the StorageTek SL3000 or SL8500 modular library system.
- Although StorageTek Storage Archive Manager is a throughput application, if the utilities are demanding more processing power, additional CPU cores can be easily added to the server domain in which StorageTek Storage Archive Manager is running.

Scale Capacity Without Disrupting Access

Large archive libraries are often in near continuous use, so it's important not to disrupt access to the data for the inevitable capacity upgrades that will be needed. The Oracle Optimized Solution for Tiered Storage Infrastructure is simplified, because the system is designed for non-disruptive scalability so that continuous access to data can be achieved. Whether extending the capacity of an existing configuration of the solution or upgrading to the next larger configuration, the following scalability features inherent in the solution components make it possible to scale capacity without taking the archive offline:

- **Capacity on demand licensing.** Tape drives or slots can be added to StorageTek modular library systems while the library continues to operate. Both the StorageTek SL8500 and SL3000 modular library systems allow expansion frames for additional slots to be preinstalled with minimal up-front cost, and then this capacity can be licensed as needed with Oracle capacity-on-demand licensing.
- **Any cartridge any slot.** StorageTek tape libraries enable organizations to grow their library any way they need because multiple types of drives and cartridges can coexist within the same library. This makes it easy to migrate from LTO drives and cartridges, for example, to the StorageTek T10000C tape drive if additional performance or capacity is required.
- **Media reuse.** Oracle tape drives use the exact same media for two generations in a row, allowing customers to reuse media from an old drive when upgrading to newer tape drive technology. In addition to having backwards compatibility which enables the new drives to read old media, the new generation of drives can write to the media used by the previous generation of drives at the full new capacity. This also can help improve library density because writing to the old media at the new higher capacity will open up slots in the library.
- **4-PB file system for primary disk storage.** StorageTek Storage Archive Manager software enables non-disruptive growth of the primary disk storage file system by adding new LUNs without dismounting the file system. It supports a maximum file system capacity of 4 PB for primary storage. This means that additional drives can be added to Sun ZFS Storage Appliances or Pillar Axiom 600 storage systems without downtime to reconfigure the file system. Although this is a great feature, using tiered storage and using StorageTek Storage Archive Manager to dynamically copy and release data from primary storage keeps this higher-priced storage capacity at a minimum.

Reduced Risk

The components of this solution have been tested and proven to work together to provide scalable performance. Oracle's performance-characterization work greatly reduces the risk of purchasing an improperly sized configuration. Sizing guidelines and best practices for performance are provided in the related technical white paper for this solution. These guidelines provide the visibility that organizations need to buy a right-sized solution for their current requirements while knowing how the solution can scale to meet future needs.

Superior Data Protection

The solution is also designed to reduce the risk of data loss. Multiple copies of data are maintained in the tiered storage infrastructure to protect against data loss. In addition, the following advantages of the solution can help further reduce the risk of data loss:

- StorageTek Storage Archive Manager software takes advantage of the StorageTek T10000C Data Integrity Validation (DIV) feature to provide end-to-end protection of data during its movement and while the data is at rest. Storage protection technologies such as RAID, Error Correction Codes (ECC), and Cyclic Redundancy Codes (CRC) help ensure that data is not corrupted while contained within the storage device. However, these technologies do not protect data in transit nor enable data integrity to be verified later without sending the data back to an application.
- StorageTek SL8500 and SL3000 modular library systems have the industry's highest availability with hot-swappable and redundant robotics and library electronics.
- The StorageTek T10000C tape drive offers an outstanding uncorrectable bit error rate (UBER) of 1×10^{19}.
- Sun ZFS Storage Appliances were ranked #1 in overall enterprise and midrange NAS quality in the 2012 Quality Awards survey from *Storage* magazine.[6]
- Single vendor support for the entire solution means that one call to Oracle is all that is needed to identify and resolve issues. There is no need for internal IT staff to spend time isolating the issue to a specific system component before calling for support.

Conclusion

Many organizations are struggling to control costs as their digital archives continue to grow rapidly. The Oracle Optimized Solution for Tiered Storage Infrastructure provides the means to keep storage costs low while delivering the necessary performance and data protection to meet user demands and compliance requirements.

The primary advantages of the Oracle solution are as follows:

- **Cost savings.** StorageTek tape libraries and cost-effective disk storage solutions enable Oracle to offer 2.8x better TCO and 72 percent lower acquisition cost than EMC's virtual tape solution.
- **Massive scalability.** Modular growth is made possible by the ability to start small and then scale capacity and performance independently as needed.

[6] http://searchstoragechannel.techtarget.com/survey/NAS-products-Quality-Awards-Oracle-storage-systems-rated-highest

- **Reduced risk.** Oracle offers industry-leading data protection and availability features as well as single-vendor support for all the hardware and software components in the solution.

Learn More

For more information about Oracle Optimized Solutions visit http://www.oracle.com/optimizedsolutions or call 1-800-ORACLE1 to speak to a representative.

Table 6 also provides links to additional Web resources for the Oracle components that complete this solution.

TABLE 6. WEB RESOURCES FOR FURTHER INFORMATION

PRODUCT WEB PAGES	
StorageTek Storage Archive Manager software	http://www.oracle.com/us/products/servers-storage/storage/storage-software/storage-archive-manager/overview/
Sun ZFS Storage Appliances	http://www.oracle.com/us/products/servers-storage/storage/nas/
Pillar Axiom storage systems	http://www.oracle.com/us/products/servers-storage/storage/san/pillar/
StorageTek tape storage solutions	http://www.oracle.com/us/products/servers-storage/storage/tape-storage/overview/

Appendix A Financial Analysis for Cost Comparisons

This appendix provides detailed analysis for the EMC cost comparison outlined in the main body of the paper.

Initial Purchase Cost Comparison Detail

Table 7 shows a comparison of initial purchase costs (all in US$) for Oracle's tiered storage offering with a similarly configured solution using tiered storage from EMC. The Oracle purchase cost is 72% less than EMC. The specific Oracle and EMC product components used in the analysis are listed below.[7]

TABLE 7. COMPETITIVE COST COMPARISON AGAINST EMC CONFIGURATION

	ORACLE SOLUTION		EMC SOLUTION	
CATEGORY	**COMPONENTS**	**COST**	**COMPONENTS**	**COST**
Server hardware	• 2 x SPARC T4-2 servers • 2 x Sun Server X3-2 servers from Oracle	$89,552	• 2 x SPARC T4-2 servers	$69,522
Tiered storage hardware and media	• Sun ZFS Storage 7420 appliance • StorageTek SL3000 modular library system • StorageTek T10000C tape drives and media • Network connectivity • Total capacity = 584 TB	$645,098	• EMC Data Domain DD890 • EMC VMAX SE • Network connectivity • Total capacity = 423 TB	$2,324,027
Software and hardware licenses	• StorageTek Storage Archive Manager • Oracle's StorageTek Tape Analytics • Oracle Automated Cartridge System Library Software • Oracle Solaris Cluster, Enterprise Edition	$159,000	• EMC TimeFinder Clone • EMC Enginuity software • EMC Symmetrix Management • EMC VMAX FAST Base chassis license • EMC DD890 Expansion capacity • EMC Data Domain Boost	$751,901

[7] EMC pricing information sourced from December 1, 2012 EMC master price list available at http://www.emc.com/collateral/emcwsca/master-price-list.pdf

TABLE 7. COMPETITIVE COST COMPARISON AGAINST EMC CONFIGURATION

Integration and deployment	$18,505	$116,819
Total initial investment	**$912,155**	**$3,262,269**
Total savings on initial investment by choosing Oracle		**$2,350,114**
Percent savings on initial investment		**72%**

Even though the Oracle solution costs dramatically less, it includes 38 percent greater capacity than that of the EMC solution. The Oracle solution provides 584 TB of raw storage compared to 423 TB with the EMC solution.

TCO Comparison Detail

Figure 3 provides an analysis of the cash flows (all US$) over a five-year period for the Oracle and EMC configurations described in Table 7.[8] The analysis includes the initial acquisition costs from Table 7 as "Year 0." The costs shown for the subsequent years are annual costs for hardware service, software licenses, power, cooling, and data center floor space. In "Year 5" of the analysis, the cost model assumes the purchase and deployment of additional hardware to accommodate growth. The analysis of undiscounted cash flows shows that the Oracle solution TCO is nearly 2.8x lower than the five-year TCO for EMC. From another perspective, the savings over five years is 64 percent with Oracle.

[8] TCO analysis was done using industry-accepted formulas. Power consumption and floor space requirements sourced from EMC publically available product specification sheets.

Cash Flows for EMC Solution ($K)

Cash Flow	Year 0	Year 1	Year 2	Year 3	Year 4	Year 5	Total
Hardware Acquisition	2,394	0	0	96	3	73	2,566
Hardware Service	0	103	152	164	227	240	887
Software - Hardware License	752	136	136	280	159	302	1,766
Power and Cooling	0	18	19	21	23	25	107
Floor Space	0	12	13	14	15	16	70
Integration and Deployment	117	0	0	7	0	8	133
Total (undiscounted)	**3,262**	**270**	**321**	**583**	**427**	**665**	**5,529**

Cash Flows for Oracle Optimized Solution ($K)

Cash Flow	Year 0	Year 1	Year 2	Year 3	Year 4	Year 5	Total
Hardware Acquisition	694	0	0	89	11	47	841
Hardware Service	0	80	87	102	108	119	496
Software - Hardware License	159	35	35	73	42	42	385
Power and Cooling	0	10	11	13	14	15	62
Floor Space	0	19	20	22	24	27	112
Integration and Deployment	19	0	0	4	0	1	24
Tape Media	40	0	5	5	5	5	60
Total (undiscounted)	**912**	**144**	**158**	**308**	**204**	**256**	**1,982**
Annual Savings (undiscounted)	**$2,350,000**	**$126,000**	**$163,000**	**$275,000**	**$223,000**	**$409,000**	**$3,547,000**

5-Year TCO for EMC	**$5,529,000**	
5-Year TCO for Oracle Optimized Solution	**$1,982,000**	
TCO Savings with Oracle	**64%**	**(2.8x lower than EMC)**

Figure 3. Detailed comparison of TCO for Oracle Optimized Solution for Tiered Storage Infrastructure versus EMC solution.

Making Data Archives Affordable Without Giving Up Performance or Capacity
February 2013

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

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200

oracle.com

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

Hardware and Software, Engineered to Work Together