An Oracle White Paper
September 2013

Oracle ZFS Storage Appliance and Oracle IT: Use Cases and Benefits

ORACLE

Executive Overview

Oracle's broad IT organizations handle an incredible volume of data across a vast spectrum of storage workloads. These workloads include everything from storage of internal corporate e-mail and collaborative workspaces to mission-critical databases and HPC environments supporting microprocessor development. Oracle's internal-facing IT business units must meet the IT needs of well over 100,000 employees, many of whom are engaged in software and hardware development activities. At the same time, Oracle IT also manages Oracle's customer-facing cloud businesses as well as various internal-facing mission-critical databases. From a storage workload perspective, Oracle IT is one of the most varied, demanding, and crucial environments imaginable. The Oracle ZFS Storage Appliance serves Oracle IT as an extremely versatile, high-performance storage platform—the ideal tool to meet the storage needs of this highly diverse and demanding IT environment. Today, Oracle has more than 225 PBs of storage on the Oracle ZFS Storage Appliance.

Introduction

Over the past few years, Oracle IT has been working toward standardization and consolidation of a wide variety of NAS workloads on Oracle ZFS Storage Appliance, including mission-critical Oracle Database storage. In August 2012, Oracle IT surpassed 225 PBs of storage on the Oracle ZFS Storage Appliance platform. Thousands of database instances and millions of users both inside and outside of Oracle run on Oracle ZFS Storage Appliance within Oracle IT every day. As part of Oracle IT's ongoing initiative to transition from legacy storage systems (primarily from NetApp and EMC) and standardize on the Oracle ZFS Storage Appliance, significant performance and efficiency benefits have been realized. For instance, replacing EMC Symmetrix systems in a compute farm that supports application development produced a 12x performance increase. In another example, transaction times were slashed by 23 percent to 66 percent in a mission-critical OLTP database used for gatekeeping and managing downloads from a patch repository. But these performance improvements are just a part of the story. Management efficiency benefits as a result of the Oracle ZFS Storage Appliance's advanced management interface, unprecedented analytics tools, and powerful scripting capabilities have enabled massive data expansion with nearly a 2:1 improvement in headcount per TB metrics.

From a product development and product improvement perspective, Oracle IT is one of the most demanding and comprehensive proving grounds. It has enabled Oracle's storage business to aggressively and continuously refine and improve the Oracle ZFS Storage Appliance. Storage customers who select the Oracle ZFS Storage Appliance over competing options do so with the confidence that the product has been rigorously designed for and proven in a plethora of the most demanding environments within Oracle IT.

Scope and Intent

The purpose of this paper is to highlight some interesting examples of the use of the Oracle ZFS Storage Appliance within Oracle IT and to document interesting reference architectures for those use cases. While it is hoped that this paper will serve as a good reference for Oracle ZFS Storage Appliance customers in terms of suitability of the product for the various workloads contained herein, the particulars of the architectures of each of the implementations are intended to be looked at as examples only. Each organization has its own sets of specific requirements and constraints that need to be considered when architecting storage solutions. The Oracle ZFS Storage Appliance is a highly versatile and configurable family of storage appliances, so most customers will find it sufficiently flexible to meet their particular needs in a wide variety of scenarios.

Use Case Example: Mission-Critical OLTP Oracle Database and Patch Repository

An extremely compelling example of a mission-critical OLTP workload within Oracle IT is the Oracle patches and updates and Oracle Software Cloud Delivery system database. This system handles a variety of important tasks, including patch deployment to more than three million registered users. Each week, this database handles more than six million requests and 400,000 patch downloads.

The primary storage for both the database and the patch file repository associated with this system is an Oracle ZFS Storage Appliance cluster system with 10 GbE connectivity, 512 GB DRAM, 12 SSDs for read cache, and eight 6-core processors. The eight disk shelves contain one hundred seventy-six 7,200 RPM SAS HDDs for storage and 16 SSDs for write cache. The entire database resides on two storage pools (one per head), each with four separate file systems for the following:

1. Log files with a 128 K record size and log cache device usage set to "metadata only" mode
2. Database data files with an 8 K record size
3. OCR/voting files (for Oracle Real Application Clusters [Oracle RAC]) with 128 K record size
4. Patch repository files

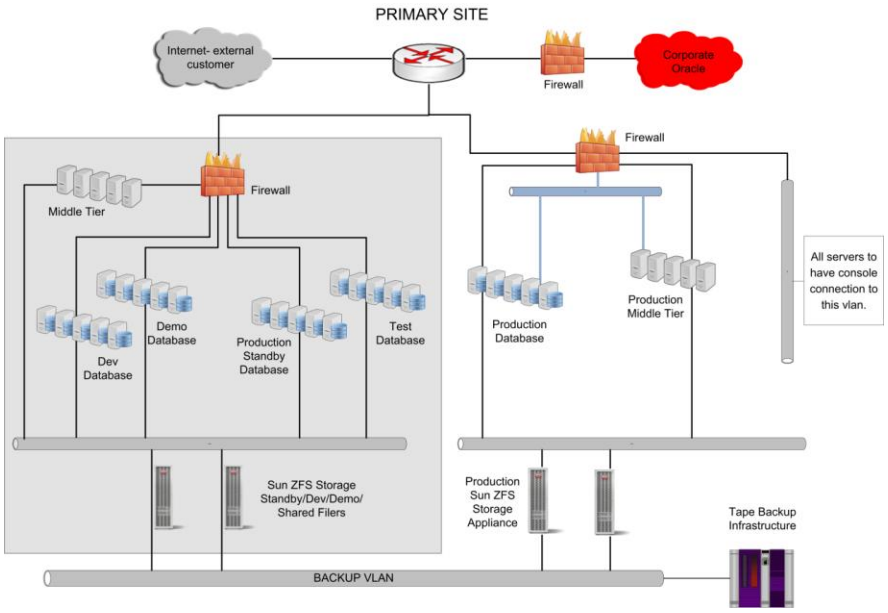

Figure 2. Oracle patches and updates and Oracle Software Cloud Delivery system primary site.

The NFS protocol is used for communication with the Oracle Linux-based Oracle Database with Oracle RAC server environment. LZJB compression is used for all file systems mirroring for data protection in order to obtain optimal performance. LZJB reduces the footprint of this data significantly. LZJB compression requires less CPU overhead than other compression options, making it perfect for this performance-sensitive environment. In fact, in many cases, the use of LZJB actually enhances performance because it reduces throughput requirements through the back-end SAS interfaces.

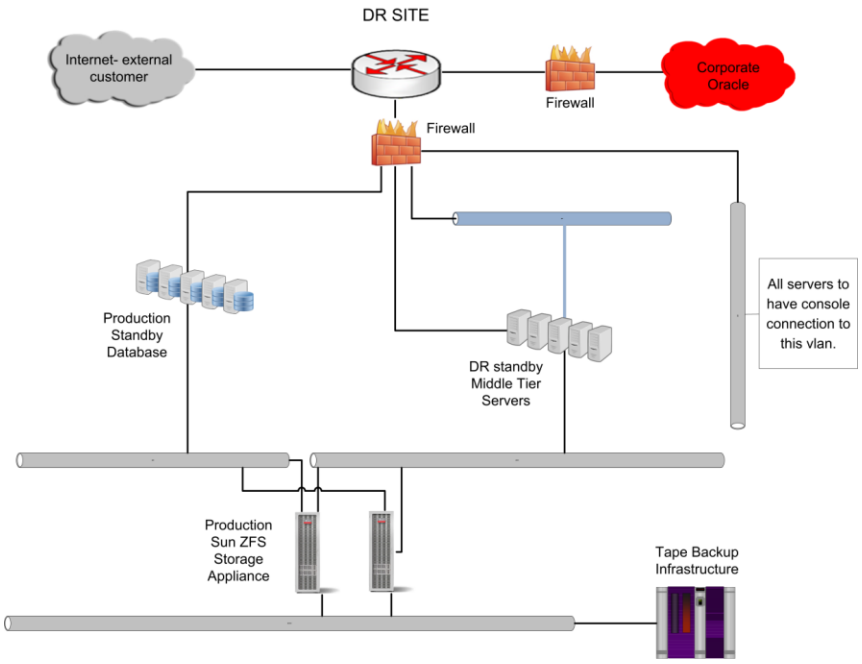

Figure 3. Oracle patches and updates and Oracle Software Cloud Delivery system DR site.

Oracle Active Data Guard software is used in conjunction with the Oracle ZFS Storage Appliance Replication feature to provide disaster recovery capability at a remote site. Oracle Active Data Guard is used to replicate the live production OLTP database while Remote Replication is used to replicate the patch repository and Oracle Recovery Manager (Oracle RMAN) backup files to the remote site. Snapshot and cloning data services are employed to facilitate 100 separate test/dev environments (on separate appliances of Oracle ZFS Storage Appliance family) where any proposed changes to the system can be tried and refined before production implementation.

Prior to implementing the Oracle ZFS Storage Appliance, this database resided on EMC Clariion storage. Once replaced with the Oracle ZFS Storage Appliance, the time required for all SQL transactions dropped significantly. Overall reduction in elapsed per execution time dropped from 23 percent to 66 percent, depending the specific SQL transaction.

In addition to the performance improvements, significant efficiency improvements have been realized, as well. The database administrator noted an approximate 50 percent reduction in storage space that the database occupies due to the LJZB compression while realizing these performance improvements.

Migration from the EMC system to the Oracle ZFS Storage Appliance was accomplished using Oracle Active Data Guard. The team was able to accomplish the entire transition with less than one hour of downtime.

This high-performance, mission-critical database is a great example of applicability of the Oracle ZFS Storage Appliance to tier 1 enterprise environments as well as related disaster recovery and test/dev environments. The high levels of performance, superb management tools with advanced analytics, and

advanced data services make the Oracle ZFS Storage Appliance an ideal choice for this type of environment.

Use Case Example: My Oracle Support Test/Dev Environment

Many Oracle customers are familiar with the My Oracle Support system, which is Oracle's primary service and support portal for downloading updates and patches, and obtaining services as well as technical documentation and best practices for Oracle products. This system is being continuously expanded and improved as a part of general service improvement initiatives. In order to facilitate these ongoing development activities, a test/dev environment is employed that uses Oracle VM with 35 virtual servers to handle the application, middleware, and database stacks. Storage for this Oracle VM-based test/dev environment is standardized on Oracle ZFS Storage Appliance. Snapshots and clones are employed to rapidly deploy workable "copies" of the production systems that occupy very little storage capacity—including the virtual machine (VM) files themselves along with all of the application files. In this large-scale environment, the ability of the ZFS file system to support a virtually unlimited number of snapshots and clones is a compelling feature. This system allows a full test/dev environment to be up and running in about three hours. (Under the prior system that leveraged full copies, deploying a test/dev environment could take several days, as a point of comparison.) The general setup for the storage is one pool per head, each with one project and several file systems. There are 35 file systems total, each corresponding to one of the 35 VMs, all of which can be cloned multiple times to support various test/dev instances. These 35 base file systems contain all of the Oracle VM virtual server files, along with all of the OS, database, and application files that are necessary to support a test/dev instance upon cloning. All pools are either triple-mirrored or double-mirrored, depending on the level of criticality. LZJB compression is employed with 128 K record size for all file systems. All systems are HA clusters with NSPF and use 7,200 K RPM drives. Ten GbE connectivity and NFS protocol are employed.

Use Case Example: Logic and Circuit Simulation for Microprocessor Development

Oracle IT manages a dedicated facility that handles the HPC requirements for Oracle's SPARC processor design. This facility employs the Oracle ZFS Storage Appliance to handle the storage needs of the microprocessor development teams. They use a number of Oracle ZFS Storage Appliance systems of varying configurations, but all have 7.2 K RPM HDDs and employ NFS protocol via 10 GbE connectivity. Most systems also have mirrored write cache SSDs except for a few that experience read workloads almost exclusively. For optimal performance, mirroring is used for data protection, and LZJB compression to reduce footprint and increase back-end SAS effective throughput. Versus NetApp, the Oracle ZFS Storage Appliance provides ample CPU capacity and large memory with efficient in-memory read caching, which allow for the use of the cost-effective, less energy-hungry 7.2 K RPM HDDs.

Use Case Example: Oracle Managed Cloud Services

Oracle Managed Cloud Services (formerly Oracle On Demand) is a business within Oracle IT that handles operation, administration, and management of customer-facing IT resources. It is responsible for administration and management of a complete Oracle stack architecture in cloud deployments. Oracle Managed Cloud Services uses Oracle ZFS Storage Appliance extensively for mission-critical storage for this environment. Cloud application offerings that reside on Oracle ZFS Storage Appliance include Oracle E-Business Suite, Oracle Express, Oracle Email Center, Oracle iLearning, Oracle's Agile product lifecycle management applications, Oracle Beehive, Oracle's PeopleSoft product portfolio, Oracle's Siebel products, Oracle Hyperion enterprise performance management products, Oracle's JD Edwards EnterpriseOne, and many others, including some third-party software offerings.

This Oracle Managed Cloud Services infrastructure consists of more than 17,000 VMs running on more than 4,000 physical servers. Storage for the environment is NFS, and more than 5,300 customer instances representing more than 2,400 customers now utilize Oracle ZFS Storage Appliance for their storage infrastructures. The typical customer environment involves a complete software stack, from the OS level (Oracle Linux or Oracle Solaris) with Oracle Database (including Oracle RAC, in some cases), Oracle Fusion Middleware, and Oracle Fusion Applications running on top. Each customer gets his or her own VM or physical server, but storage for multiple customers may be consolidated on Oracle ZFS Storage Appliance via NFS protocol and 10 GbE infrastructure. To ensure complete customer isolation and security, Oracle VM machine files and OS files for any given customer are compartmentalized, and storage access is controlled at the file system level. This allows multitenancy on the Oracle ZFS Storage Appliance while preventing data access across customer accounts.

Storage for databases, middleware, and individual applications is segregated from one another and from VM files and OS storage by project to optimize performance and speed deployment. Separate file systems are used for different data types to segregate. For example, the database log file is segregated from database data files and from the application files themselves. Project and file system settings on the Oracle ZFS Storage Appliance are highly customizable, allowing great flexibility to optimally tune for a variety of environments. A project dedicated to the Oracle E-Business Suite might include the following shares:

- Oracle Database data files
- Oracle Database index files
- Oracle Database temporary tablespace files
- Oracle Database backup files
- Oracle Database redo log files
- Oracle Fusion Middleware product files
- Application product code tree
- Application log and temp files

Storage of database data files for the ERP system will use an 8 K record size setting at the file system level whereas the database log file settings under the same project will use the project's default 128 K record size in a separate file system. In that same project, there will be other file systems for the application files that run the ERP system application itself, which inherit the project's default 128 K record size.

All Oracle ZFS Storage Appliance systems involved in this environment are HA cluster configurations and are set up with NSPF. All use mirrored write flash accelerator devices and striped read flash cache devices. The extensive processor power and DRAM scalability of the Oracle ZFS Storage Appliance allows the attachment of a large number of disk spindles, and typically 7,200 RPM SAS HDDs are used for back-end storage. In certain database instances with extremely high-performance requirements under random I/O, 15,000 RPM SAS drives are employed. Mirroring is used for data protection in order to achieve optimal performance, and triple mirroring is used in some cases for higher redundancy data protection.

Backup and archive for the production environment is accomplished using NDMP from the Oracle ZFS Storage Appliance to an Oracle Secure Backup server and ultimately on to tape. Many customers also elect to have a remote DR site. For disaster recovery of Oracle Database, Oracle Active Data Guard is used to replicate the live database to the remote storage via the database server. DR for patch repositories, scripts, and application files is accomplished using Remote Replication as a storage data service in the Oracle ZFS Storage Appliance.

Prior to Oracle ZFS Storage Appliance adoption, the Oracle Managed Cloud Services business relied primarily on NetApp filers. Throughout the ongoing migration process, many benefits of the Oracle ZFS Storage Appliance have been realized. By leveraging the superior performance and management efficiency benefits along with the industry's best storage analytics environment (DTrace Analytics, which is a feature of Oracle ZFS Storage Appliance), overall performance has improved while operating expenses associated with management time have decreased. In fact, one recent study within Oracle IT found nearly a 2:1 improvement in headcount/GB efficiency after a major storage upgrade replacing NetApp systems with Oracle ZFS Storage Appliance. The storage analytics offer unprecedented visibility to help resolve bottlenecks to optimize system performance quickly, particularly in this extensive VM environment, as compared to what competing options would offer. For example, if a client were to experience slow performance, DTrace Analytics would allow the storage admin to determine quickly if the bottleneck is related to storage or if it is an issue related to client tuning. Unlike with competing analytics packages, it is possible with DTrace Analytics to drill down easily from the physical interface level through the entire storage stack all the way to the file or disk level.

Migration from NetApp to Oracle ZFS Storage Appliance has been ongoing, and is now nearing completion. Moving an environment of this scale and complexity is obviously a significant undertaking that takes time. Fortunately, NFS migration to Oracle ZFS Storage Appliance has been made as straightforward and unobtrusive as possible by the use of the Shadow Migration feature.

Use Case Example: Oracle University

Another interesting workload is the Oracle University infrastructure. Oracle University provides both internal and customer training and education services. It uses databases and content streaming to provide an interactive, online training experience. This is a critical system in that it is externally available and customer facing—it must remain online at all times in order to support training requirements of Oracle customers and employees around the world. This system used to be run with NetApp filers, but now runs with an Oracle ZFS Storage Appliance, along with a second similar system at a DR site where Oracle ZFS Storage Appliance software's Remote Replication feature is used to move the data. The Oracle ZFS Storage Appliance family consists of HA clusters with NSPF and one storage pool per head. They are set up with mirroring as data protection. File systems typically use LZJB compression with the default 128 K record size. The Oracle Database data file shares are the exception, and use an 8 K record size. Since transitioning from NetApp, performance has increased significantly, while the compression enables a smaller storage footprint.

Shadow Migration was used to perform the migration seamlessly from NetApp with minimal downtime. Shadow Migration is a feature of the Oracle ZFS Storage Appliance that allows it to serve as primary storage while simultaneously moving data from a third-party NFS-capable NAS system, such as NetApp.

This workload is interesting in that the storage must provide a high-performance, high-availability interactive streaming content experience in the form of 41,000 classes to more than 350,000 students across the globe annually. Any customer interested in any type of streaming media workload consolidated with database storage can be assured that the Oracle ZFS Storage Appliance has proven itself in one of the most rigorous environments of this type.

Use Case Example: Development of Service Solutions and Interop Testing

An internal-facing group within IT tests interoperability and develops solutions for Oracle Fusion Applications running on a variety of third-party servers and operating systems. Each environment runs an OS on either a physical server or VM that has three Oracle Databases—one that is transactional and two for identity management. More than 50 Oracle WebLogic Server instances are used to develop application changes or patched for apps such as Oracle Business Intelligence Suite, Oracle Access Manager, and Oracle Fusion Applications. The storage for this environment is provided primarily by Oracle ZFS Storage Appliance via a 10 GbE environment. A combination of NFS shares and iSCSI LUNS are used for the storage of the database workloads, VM files, and OS images, and code changes that are being developed in a clone-based test/dev environment. This group's charter is to resolve current customer issues as needed and deploy patches or application changes rapidly to support the service organization, making storage performance and availability critical.

Use Case Example: IT Automation Tool

Oracle has an IT automation tool that is used by both internal and external parties to assist in infrastructure management. This tool gives information about the devices on which applications are running and provides monitoring tools along with a change approval process complete with user rules and roles. The IT automation tool is linked with the access provisioning system to gain access to a single source for user rules and roles. Storage for this environment is consolidated on Oracle ZFS Storage Appliance.

The compute environment consists of hundreds of VMs. Storage for this environment runs on multiple appliances from the Oracle ZFS Storage Appliance family using 10 GbE infrastructure and NFS to mount the VMs. NFS is also used for the file shares for the applications that run on the VMs. VM clients of the storage run primarily Oracle Linux instances within the Oracle VM Hypervisor. Some VM clients also support specific legacy systems running VMWare or Oracle VM VirtualBox and Microsoft Windows Server 2008 instances. The Oracle ZFS Storage Appliance family appliances are HA clusters with NSPF and one storage pool per head. They are set up with mirroring and data protection. Most projects use LZJB compression with the default 128 K record size. The Oracle Database data file shares are the exception, and use 8 K record size. The Remote Replication data service is used to replicate to a DR site.

Many advantages were realized by migrating to the Oracle ZFS Storage Appliance from NetApp. First, superior performance is obtained compared with competing products. The performance advantage is due to the variable block sizes, mirroring as a data protection option, large flash-based cache, and high-performance controller hardware, The superior management tools, including the advanced browser user interface and DTrace Analytics package, simplify and speed management, maintenance, and deployment activities, all of which reduce effective operating costs over time. The analytics package allows unprecedented visibility into performance bottlenecks, including into the storage network stack, which is particularly helpful in optimizing and troubleshooting large-scale VM environments. Additionally, a comprehensive scripting interface is leveraged to develop standardized workflows to automate repetitive tasks. In some cases, these scripts even involve user inputs of basic parameters but still can run entirely from the appliance itself in an automated fashion. This advanced scripting functionality reduces complexity and cost by eliminating the need found in many non-Oracle storage environments to have centralized administration hosts that are dedicated to executing scripts for users. The Oracle ZFS Storage Appliance's Shadow Migration feature makes migrating from other NFS capable systems straightforward and seamless, minimizing downtime requirements for transitions.

Use Case Example: Compute Farm for Software Development

Oracle IT supports internal software development activities by providing a large-scale, high-performance test/dev environment. This environment, which is managed entirely by Oracle Enterprise Manager, uses 15,000 hosts for compute to support development needs, while the Oracle ZFS Storage Appliance provides the storage for this environment. This model enables parallel processing of builds and tests and offloads local development and QA systems. It creates a test baseline as part of the label creation process. (A label, as used here, is defined as a string, set of files, or a snapshot of a

development process at any given point in time.) It includes both static and dynamic farm systems and has a capacity of 350,000 compute hours per day and can complete 180,000 farm jobs per day. The farm is kept busy by global development on a 24/7 basis—even on weekends the queue is rarely less than 20,000 jobs.

Within the dynamic farm, one of the main goals is to be able to shift between VM images for 1,500 VMs quickly. The storage for the VM images (mostly Oracle Linux and some MS Windows), and OS files resides on an Oracle ZFS Storage Appliance.

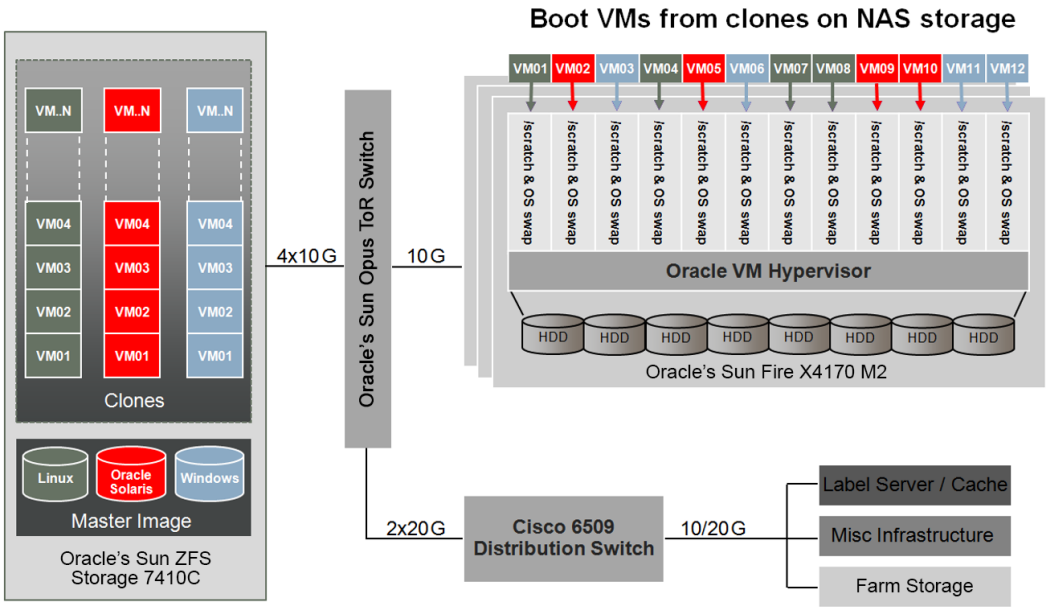

Figure 4. Compute farm architecture.

Clones are used to replicate base images from an image repository project to the live VM project. Scripting is used extensively to automate tasks for efficiency and consistency. In order to facilitate pre-release compute starvation and post-release compute availability, there are separate input Oracle ZFS Storage Appliance appliances from which application code is drawn and output Oracle ZFS Storage Appliance appliances outside of the dynamic farm where development products are placed. This is all accomplished with NFS storage on an Ethernet-based infrastructure.

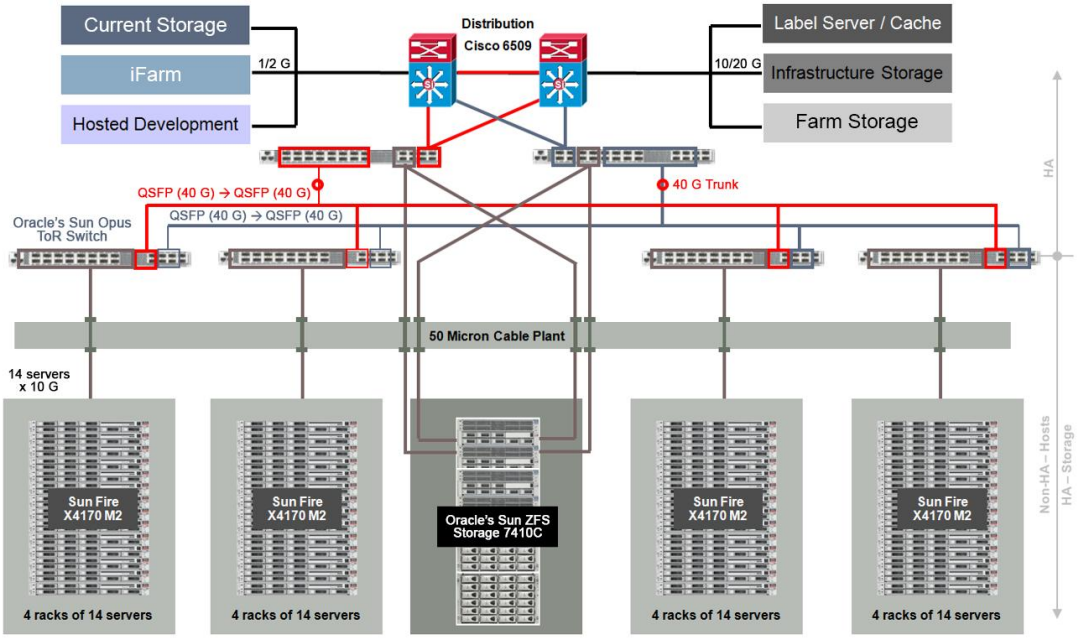

Figure 5. Dynamic farm architecture.

To date, about 350 NetApp filers and 70 EMC arrays have been replaced in this environment with appliances from the Oracle ZFS Storage Appliance family. With the replacements, Oracle IT found that approximately four NetApp filers could be replaced by one appliance from the Oracle ZFS Storage Appliance family, given the performance and compression advantages of Oracle ZFS Storage Appliance. Oracle IT realized approximately an 8:1 disk consolidation ratio, and performance increased to about 191,000 IOPS going to the Oracle ZFS Storage Appliance.

Conclusion

Oracle IT is nearing completion of a successful transition to the Oracle ZFS Storage Appliance from a variety of storage vendors. This transition process has been aided by Oracle tools such as Oracle Active Data Guard for Oracle Database and by the Shadow Migration feature of the Oracle ZFS Storage Appliance. Oracle IT has realized enormous benefits since upgrading its storage infrastructure to the Oracle ZFS Storage Appliance. Performance has improved markedly, sometimes many times over the incumbent systems, as a result of the Oracle ZFS Storage Appliance's powerful hardware and intelligent Hybrid Storage Pool data management technology. Hybrid Storage Pool is a feature of Oracle ZFS Storage Appliance. Efficiency has improved significantly as well, both in terms of storage capacity efficiency as a result of LZJB compression and in terms of management efficiency as a result of the advanced management interface, powerful analytics, sophisticated scripting language, and deep Oracle integration. In fact, headcount per TB under management has improved by nearly a 2:1 factor

in some areas. These performance and efficiency advantages are apparent across a wide variety of storage workloads, including mission-critical OLTP databases, massive software test/dev environments, and VM environments for Oracle's customer-facing cloud businesses.

While the Oracle ZFS Storage Appliance serves Oracle IT beautifully, Oracle IT also serves the Oracle ZFS Storage Appliance by providing one of the most rigorous and varied product proving grounds imaginable. Customers considering a move to the Oracle ZFS Storage Appliance need only look to Oracle's own IT operations to substantiate its selection for the most demanding storage workloads.

Oracle ZFS Storage Appliance and Oracle IT: Use Cases and Benefits
September 2013
Author: Bryce Cracco (Storage Product Mgmt)
Contributing Authors: Matthew Barnson, Gary Burch, Priyobrato Chatterjee, Robert Lashley, Ryan MacDonald, Chris W Martin, Torrey McMahon, Ryan Mikuni, Erik Neilsen, Andres Noriega, Jaideep Paidimari, Gowthaman Ranganathan, Burra Sarma

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