ORACLE

OPTIMIZED SOLUTIONS

An Oracle Technical White Paper
March 2012

How To Improve PeopleSoft HCM Performance and Reliability with an Oracle Optimized Solution

Architecture and Best Practices for PeopleSoft on SPARC SuperCluster

Introduction

The modular Oracle PeopleSoft Human Capital Management (HCM) suite offers powerful and comprehensive Human Resource (HR) solutions, however, implementing sophisticated HR applications typically involves extensive research and planning. Information Technology (IT) staff must carefully design a robust, highly available architecture, choose hardware and software, test the design, and deploy and tune the components to perform optimally. These efforts can consume valuable resources and entail a lengthy time to ROI. Enterprises need to shorten planning and deployment time and reap application benefits quickly.

Oracle addresses these concerns with the new Oracle Optimized Solution for Oracle's PeopleSoft Human Capital Management, providing a new way of implementing datacenter infrastructure. Integrated in a pre-racked and pre-configured Oracle SPARC SuperCluster system, the solution provides a complete and optimized infrastructure built with best-of-breed servers, storage, and networking. Already tested, the system is delivered ready to run the Oracle PeopleSoft HCM modules in a reliable, scalable, and efficient way, providing IT managers with known and quantified results.

Oracle tested Oracle PeopleSoft Enterprise HCM running on Oracle's SPARC SuperCluster to demonstrate solution performance, reliability, and scalability. This paper details the testing effort and discusses best practices for successfully implementing the Oracle Optimized Solution for PeopleSoft HCM in a greatly reduced timeframe.

The Oracle Optimized Solution for PeopleSoft Human Capital Management on SPARC SuperCluster

Designed with fast service deployment in mind, the Oracle Optimized Solution for PeopleSoft Human Capital Management (HCM) eliminates the risk associated with implementing HR solutions. Best-of-breed components — Oracle's PeopleSoft HCM, Oracle WebLogic server, Oracle Real Application Clusters (Oracle RAC), Oracle Database 11*g* Enterprise Edition, Oracle Solaris operating system, Oracle Solaris Cluster, and Oracle's SPARC SuperCluster — come together in the solution to deliver the most integrated and advanced tools for human capital management.

Based on Oracle SPARC SuperCluster — a complete enterprise infrastructure — the solution reduces the time needed to get from concept to service deployment. Oracle's proven, tested, application-to-disk infrastructure runs in the pre-assembled and pre-tested SPARC SuperCluster, eliminating the need for additional hardware configuration. Because the entire environment is engineered and optimized to work together, IT organizations can get services up and running faster. Predictable, high-performance applications take advantage of a highly reliable, available, and serviceable platform to eliminate the potential pitfalls and time-consuming troubleshooting associated with typical enterprise implementations.

The Integrated Enterprise Application Stack

The Oracle Optimized Solution for PeopleSoft HCM is based on a PeopleSoft module that provides rich self-service capabilities and extensive business process automation. The solution also includes testing with the PeopleSoft Payroll for North America module, which offers integral legal and practical processes for the enterprise while streamlining payroll processes and minimizing the burden on IT managers and payroll staff. The PeopleSoft modules utilize Oracle Database 11*g* running on the Oracle Solaris operating system. The integrated components of the software stack include:

- **Oracle Solaris.** The Oracle Solaris operating system versions 10 and 11 include innovative, built-in functionality, such as near wire-speed networking throughput and high availability features that deliver industry-leading performance. Virtualization features help to optimize resource utilization, and advanced security features provide the isolation and control required by governments, financial institutions, and HR environments. In this solution, Oracle Solaris 11 runs in the database virtual machine (VM) in order to take full advantage of the outstanding performance of Exadata storage cells. Oracle Solaris 10 runs in the application VM to maintain compatibility for applications that were not certified for Oracle Solaris 11 when these tests were performed.

- **Oracle Real Application Clusters (Oracle RAC), with Oracle Database 11*g* Enterprise Edition.** Oracle RAC enables a single database to run across a cluster of servers, providing high availability, performance, and scalability without application modification. Oracle Database 11*g* is designed for rapidly evolving datacenter environments and provides efficient, reliable and secure data management for mission-critical transactional applications, query-intensive data warehouses, and mixed workloads. With new self-managing capabilities, the database also eliminates time-consuming, error-prone administrative

tasks, so database administrators can focus on strategic business objectives instead of performance and availability fire drills.

- **Oracle PeopleSoft Enterprise PeopleTools.** PeopleSoft Enterprise PeopleTools provides a comprehensive development toolset that supports the creation of PeopleSoft Enterprise HCM applications. Using built-in toolset efficiencies, developers can build and customize PeopleSoft applications quickly and easily. PeopleSoft Enterprise PeopleTools includes robust middleware options and the PeopleSoft Integration Broker to support a Service-Oriented Architecture (SOA), a development standard.

- **Oracle PeopleSoft Enterprise Human Resources.** PeopleSoft Enterprise Human Resources delivers comprehensive HR capabilities, from workforce management to compensation and talent management. At its foundation, PeopleSoft HR offers an enterprise-wide human resources database to support myriad business processes, maintain historical records and job related data, report progress on critical HR functions, and facilitate better decision making.

- **Oracle PeopleSoft Payroll for North America.** Oracle's PeopleSoft Enterprise Payroll for North America ensures that payroll is accurate, on time, and processed efficiently. It offers full database integration with other PeopleSoft products such as Time and Labor, Compensation, Recruiting, Project Costing, and General Ledger to facilitate accurate and timely payroll expense management and reporting.

- **Oracle WebLogic Server.** Oracle WebLogic Server Enterprise Edition is the application server of choice for demanding customer environments the world over. Comprehensive management capabilities enable administration of sophisticated systems via a well-designed graphical console. Proven clustering technology, cross-domain management, and comprehensive diagnostic tools are standard with Oracle WebLogic Server Enterprise Edition, and are well integrated with the Oracle portfolio for superior interoperability and support across the technology stack.

- **Oracle Tuxedo.** Oracle Tuxedo is Oracle Fusion Middleware's strategic transaction processing product and is a leading platform for distributed transaction processing. It provides mainframe-class scale and performance on open, distributed systems, and is the premier platform for re-hosting mainframe applications on mainstream hardware. Oracle Tuxedo provides cost-effective reliability and extreme scalability up to hundreds of thousands of transactions per second and preserves investments by extending the life of existing IT assets through support for standard architectures such as SOA.

- **Oracle Solaris Cluster.** Oracle Solaris Cluster offers solutions for mission-critical applications in physical and virtual environments, and for local to global datacenters. Leveraging Oracle Solaris Containers to consolidate multi-tier applications and databases, Oracle Solaris Cluster raises the bar of high availability in virtualized environments. These virtualization features can help organizations reduce software licensing fees, system management overhead, and hardware costs for mission-critical application deployments while ensuring security isolation, resource management, and fault isolation. Oracle Solaris Cluster enables Oracle's SPARC SuperCluster to further enhance high levels of availability with a simple, high-performance solution to meet enterprise service-level agreements.

Oracle SPARC SuperCluster T4-4

Oracle SPARC SuperCluster is a complete, pre-engineered, and pre-tested high-performance enterprise infrastructure solution that is easy to deploy right out of the box. Based on Oracle SPARC T4-4 servers and Oracle Exadata Storage Servers interconnected over an InfiniBand fabric, SPARC SuperCluster T4-4 offers both the Exadata highly optimized database machine and the Exalogic highly optimized system for running applications and middleware. Designed to run a wide range of enterprise applications including database, middleware, and Oracle and custom applications, the solution is ready to go on day one and eliminates much of the integration effort, cost, and deployment time typically associated with clustered solutions.

SPARC SuperCluster T4-4 Components

- **Oracle's SPARC T4-4 Servers.** In SPARC SuperCluster, the Oracle Optimized Solution for PeopleSoft HCM is configured with Oracle's SPARC T4-4 servers. Built on the new SPARC T4 processor, each server in the SPARC SuperCluster is configured with four 3.0 GHz SPARC T4 processors accompanied by 1 TB of RAM, two Sun PCIe QDR InfiniBand Host Channel Adapters, and two Sun multithreaded 10 Gb Ethernet Networking cards. Oracle's SPARC T4-4 servers set performance records with PeopleSoft HCM, outperforming the previous PeopleSoft HRMS Self-Service 8.9 world record with 15,000 concurrent users performing typical HR transactions such as viewing a paycheck, and promoting or hiring an employee. Oracle's SPARC T4-4 servers running PeopleSoft HCM also outperformed the previous PeopleSoft Enterprise Payroll (N.A.) 9.1 world record, processing a 500,000-employee payroll in 30.84 minutes.

- **Oracle Exadata Storage Servers.** With its combination of smart Oracle Exadata Storage Server software, complete and intelligent Oracle Database software, and the latest industry-standard hardware components, the Oracle Exadata portion of the SPARC SuperCluster machine delivers extreme database performance in a highly available, highly secure environment.

- **Oracle's Sun ZFS Storage 7320 Appliance.** Providing 40 TB of disk capacity for shared file systems, the Sun ZFS Storage 7320 Appliance uses Flash-enabled Hybrid Storage Pools to dramatically decrease application response times. Easy-to-use DTrace Analytics optimize performance with minimal intervention, and powerful storage controllers run multiple data services, increasing efficiency and deployment flexibility. Oracle Solaris ZFS and self-healing technologies provide superior data integrity, while cluster failovers and Flash-based write caches increase data availability.

- **Oracle's Sun Datacenter InfiniBand Switch 36.** Oracle's SPARC SuperCluster uses an InfiniBand fabric for rapid exchange of data among the cluster components. The high-speed low-latency InfiniBand fabric utilizes a pair of redundant Sun QDR InfiniBand Switches to interconnect the Oracle SPARC T4-4 servers, Exadata Storage Cells, and Sun ZFS Storage 7320 Appliances. The Sun Datacenter InfiniBand Switch 36 is designed for application clusters comprised of Oracle's Sun rackmount servers and Oracle's Sun storage and delivers the extreme scale, application isolation, and elasticity needed to consolidate and virtualize core enterprise business applications such as PeopleSoft HCM.

Solution Architecture

All hardware and network components are pre-configured to provide redundancy and increase application availability (See Figure 1).

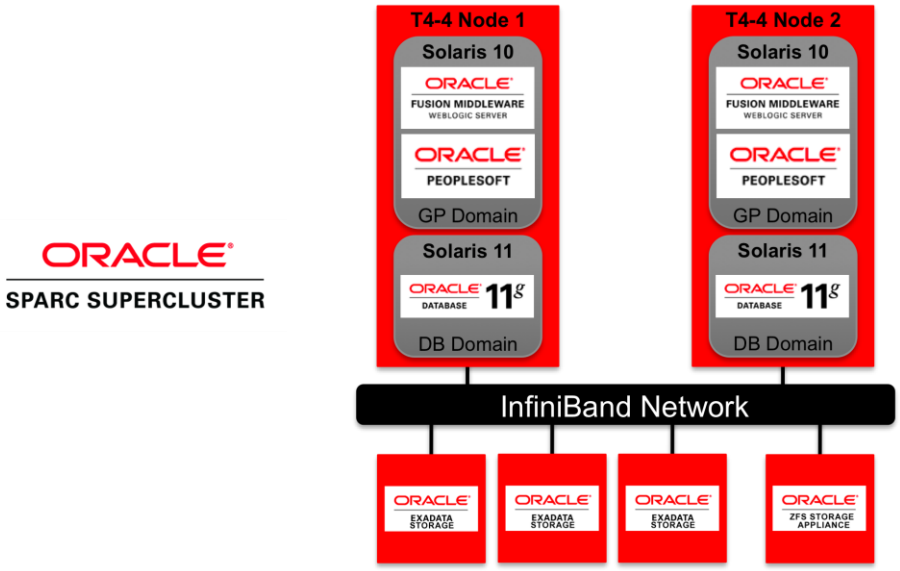

Figure 1. Architecture of the Oracle Optimized Solution for PeopleSoft HCM on SPARC SuperCluster.

Oracle VM Server for SPARC (previously called Sun Logical Domains or LDoms) provides highly efficient, enterprise-class virtualization capabilities for Oracle's SPARC T-Series servers. In this solution, two Oracle SPARC T4-4 servers are each partitioned into two virtual servers using Oracle VM Server for SPARC. The first virtual server, the DB Domain (aka database VM), runs Oracle Solaris 11 and hosts Oracle RAC with the Oracle Exadata software (Figure 2). The second virtual server, the GP Domain (aka applications VM), hosts the Oracle WebLogic and PeopleSoft applications on Oracle Solaris 10. In order to accommodate various types of workloads, the hardware resources are partitioned equally between the two VMs, each VM being allocated half the hardware resources. Depending on the applications and workloads being run, the CPU and memory resources can be migrated between the two VMs.

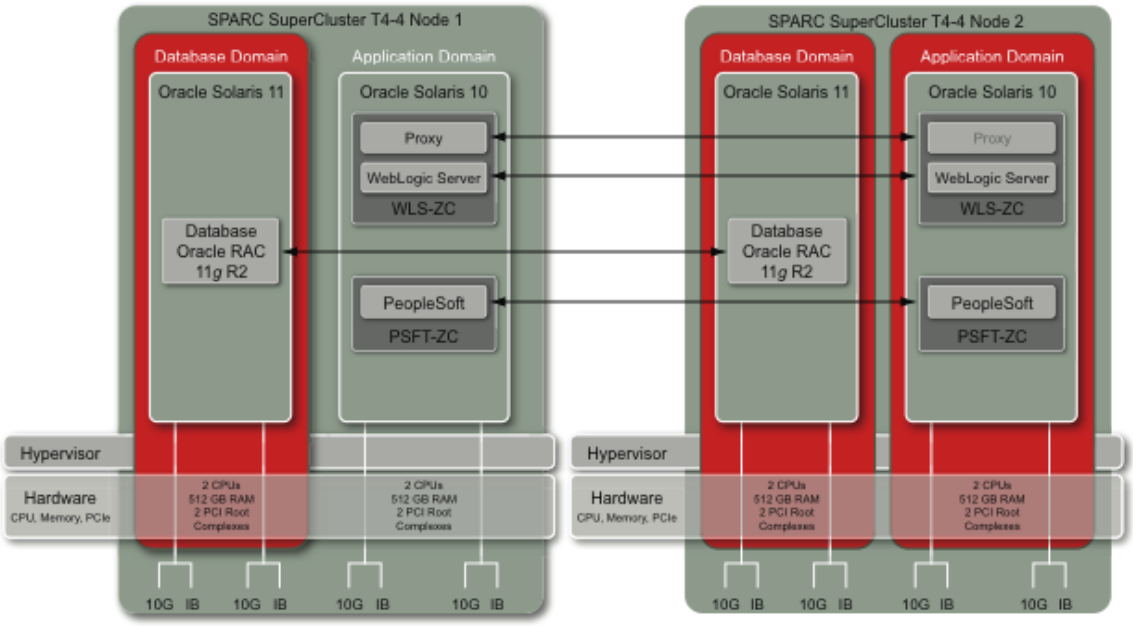

Figure 2. The Oracle SPARC T4-4 servers are first partitioned using Oracle VM Server for SPARC.

A two node Oracle RAC database was configured to be active-active in database VMs running Oracle Solaris 11. The applications VM is subdivided into two Oracle Solaris Containers to isolate the Web tier (Oracle WebLogic Server) from the PeopleSoft application tier. Requiring very little overhead, Oracle Solaris Containers leave all the compute power for application and Web tiers. Containers provide great granularity in allocating resources such as CPU and memory, and for capping I/O.

To ensure high availability, the zone clusters are configured so that each one can be failed over individually across the two Oracle SPARC T4-4 servers (Figure 3). This partitioning scheme enables the workloads to be distributed across servers providing better control over resource utilization. The SPARC SuperCluster T4-4 half-rack configuration is based on two SPARC T4-4 servers, each running instances of Oracle PeopleSoft Application Server and Oracle WebLogic Server in an active-active clustered configuration. Oracle Solaris Cluster provides high availability with failover between nodes. The Reverse Proxy Server (RPS) instance running in the WebLogic Server receives HTTP requests from Web clients. The Reverse Proxy Server tries to balance the load among multiple web server instances running on both server nodes in a WebLogic Server Cluster. In the event of a failure, the Oracle Solaris Cluster agent transparently switches over the Reverse Proxy Server to the standby node.

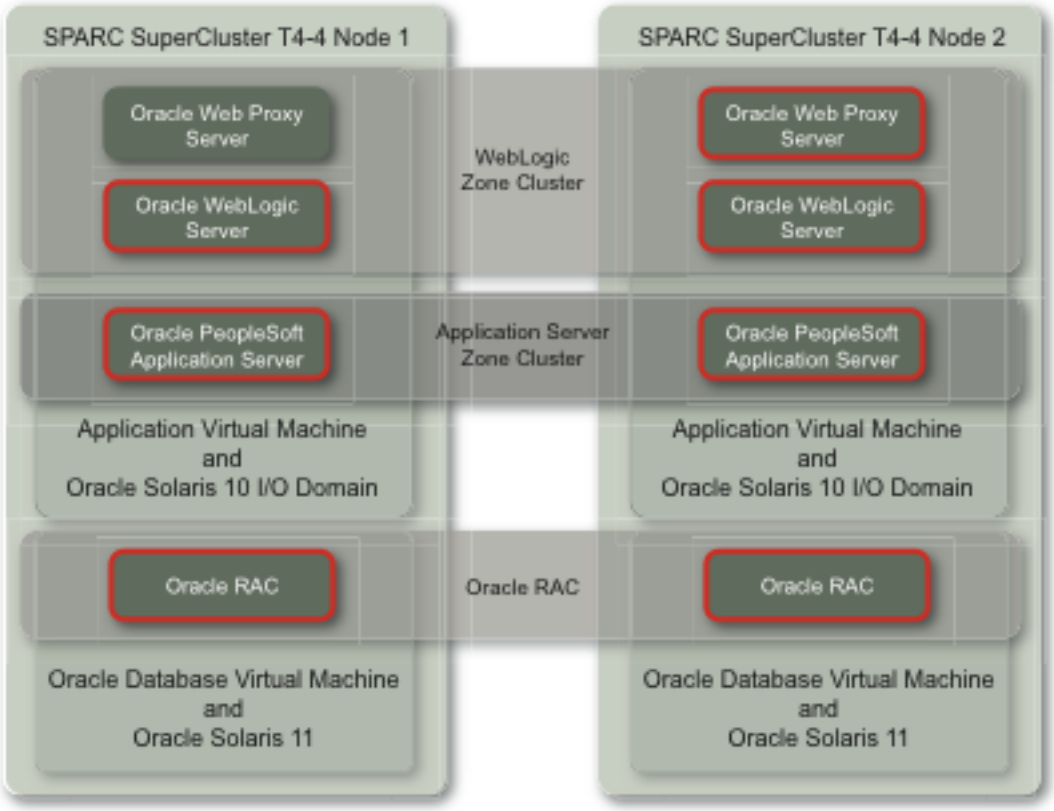

Figure 3. Oracle Solaris Containers and Oracle Solaris Cluster are used to ensure that applications are highly available.

Each container is connected to the InfiniBand and 10 Gigabit Ethernet networks for full communication with the cluster interconnect and storage, and for external client access (Figure 4 below). Physical connections have a fixed hostname and address on each of the InfiniBand and 10 Gigabit Ethernet networks, and each zone cluster has a number of logical hostnames on the 10 Gigabit Ethernet network for failover situations.

The Web tier receives the client traffic over the 10 Gigabit Ethernet network and connects to the application tier and the Sun ZFS Storage 7320 appliance over InfiniBand network partition 8503.

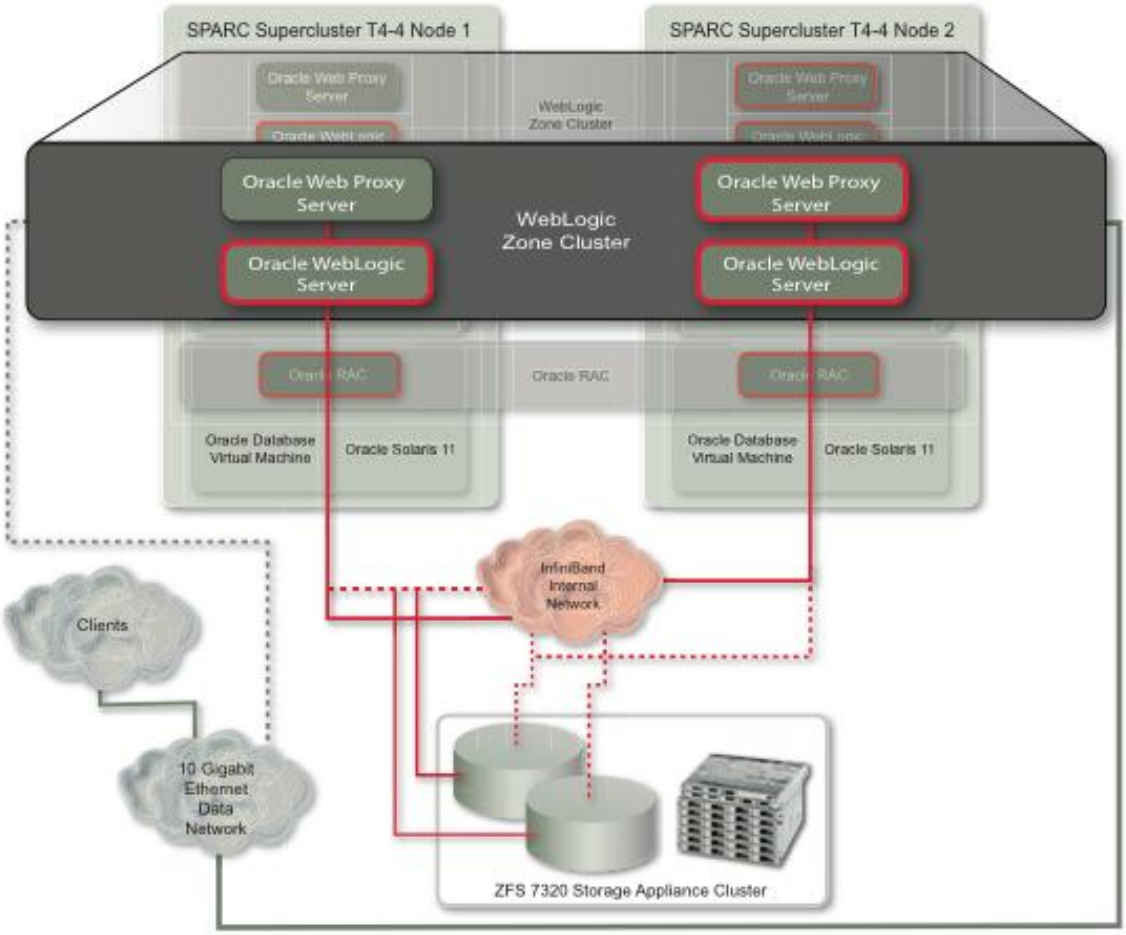

Figure 4. The Web containers are connected to both the storage InfiniBand network and the 10 Gigabit Ethernet network.

The application tier connects to the database tier and the Sun ZFS Storage 7320 appliance over InfiniBand network partition 8503. It is also possible to connect to the application tier over the 10 Gigabit Ethernet network (not shown in Figure 5 below) to provide direct access for software management without requiring access to the global zone of the Oracle Solaris 10 domains.

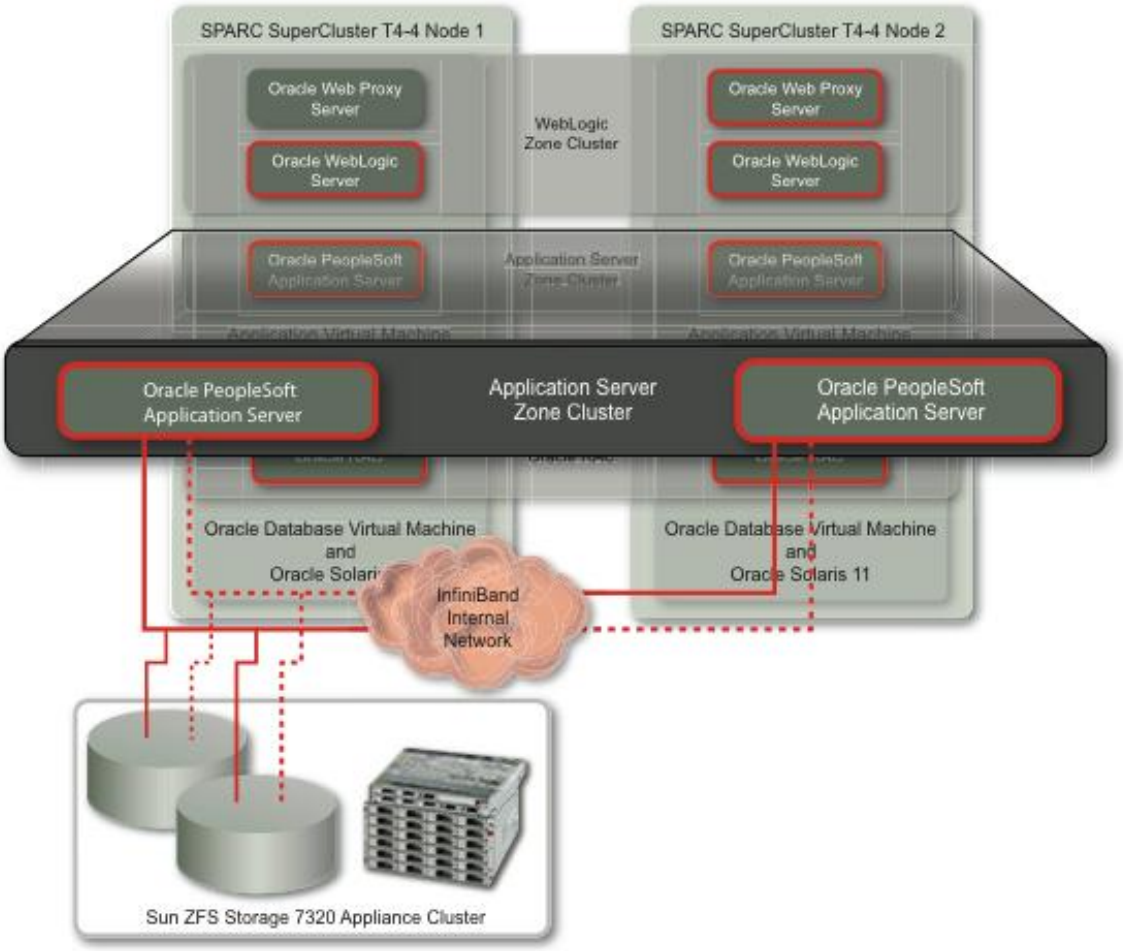

Figure 5. The applications containers are connected to the storage InfiniBand network.

The database tier connects to the Exadata Storage Servers over the InfiniBand network partition FFFF and to the Sun ZFS Storage 7320 appliance over InfiniBand network partition 8503 as illustrated in Figure 6.

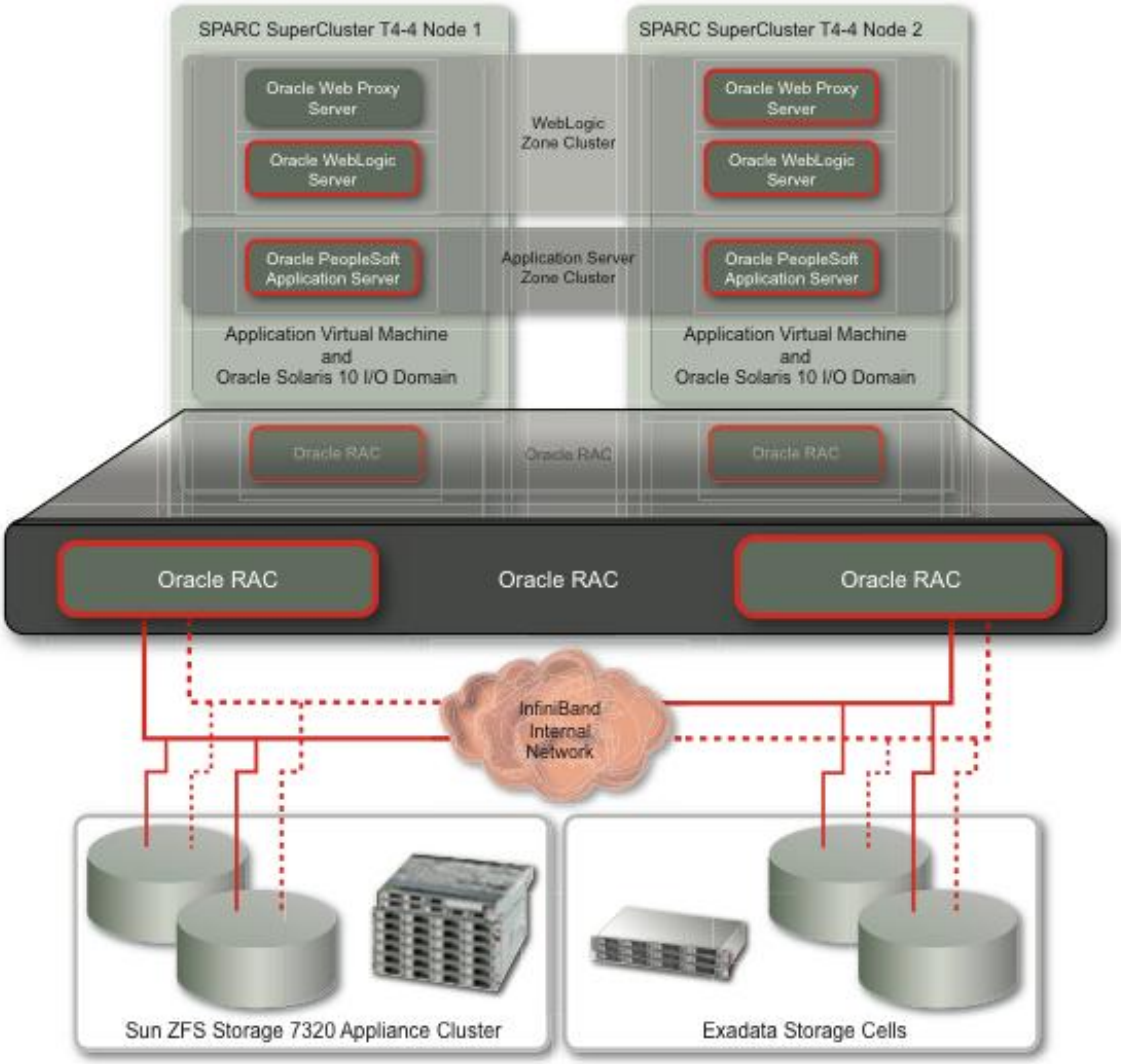

Figure 6. The Oracle RAC virtual machines are connected to the storage and Exadata InfiniBand networks.

Each server in the Oracle SPARC SuperCluster T4-4 is configured with two InfiniBand Host Channel Adapters and two 10 Gigabit Ethernet Network Interface Cards. InfiniBand cables are wired from the servers, storage cells, and the Sun ZFS Storage 7320 Appliance to the two leaf switches for redundant connections to storage and the cluster interconnects. The leaf switches are connected to each other and to the spine switch. All tiers in the architecture communicate using the InfiniBand network while incoming client connections to the Web servers utilize 10 Gigabit Ethernet interfaces.

The Oracle Solaris Cluster software requires a quorum device. To create the quorum device, an iSCSI volume is exported from the Sun ZFS Storage 7320 Appliance and is accessed over IP over IB (IPoIB) using the InfiniBand links.

The servers and storage use the IPoIB protocol to communicate across the InfiniBand network internally. The Oracle RAC database accesses the data stored on the Oracle storage cells using the Reliable

Datagram Sockets (RDS) protocol on a private Exadata InfiniBand network (FFFF). Applications access the shared volumes on the Sun ZFS Storage 7320 Appliance using the NFS protocol transported using IPoIB.

The solution architecture is based on redundant network components and links to ensure application availability.

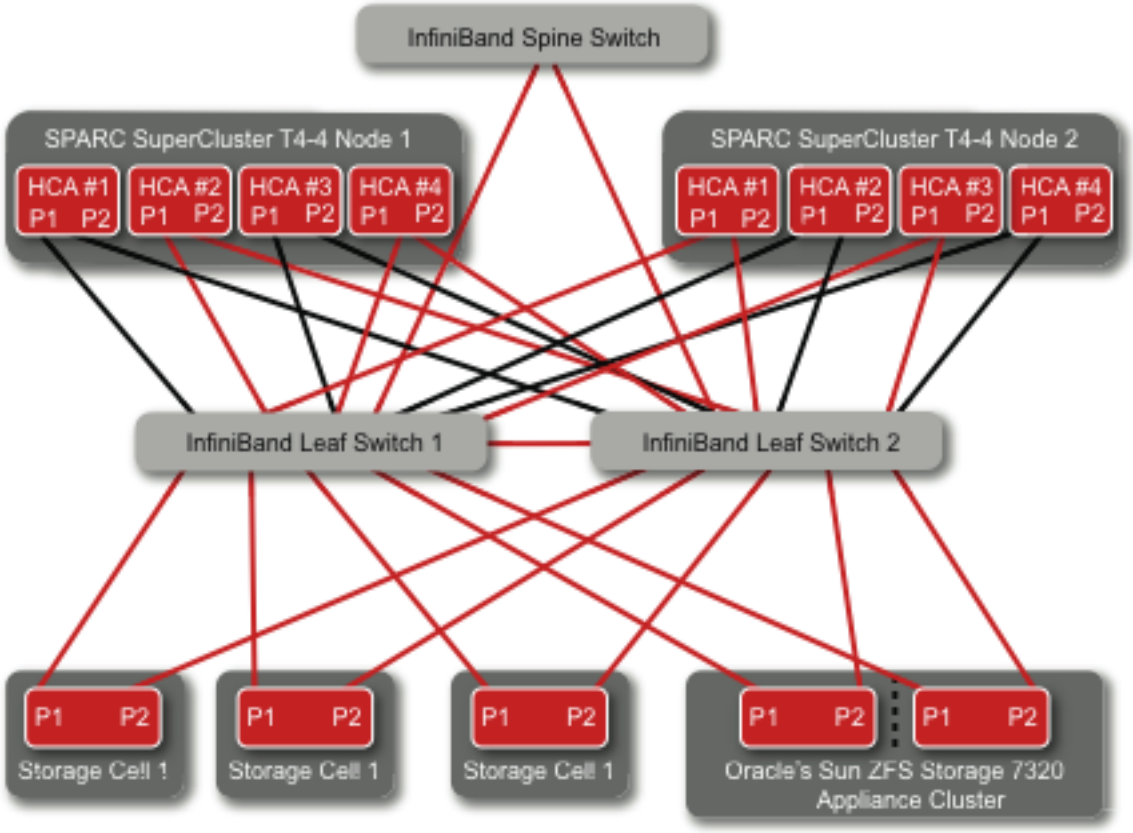

Figure 7. The solution architecture ensures high application availability by using pre-configured, redundant components and a redundant InfiniBand network.

To isolate the InfiniBand network traffic, the solution utilizes the following four partitions (Table 1).

TABLE 1. PARTITIONS USED TO ISOLATE INFINIBAND NETWORK TRAFFIC

PARTITION	TRAFFIC
FFFF (Default)	• Between Oracle RAC and Exadata storage cells
8501 (SC33)	• Oracle Solaris Cluster private traffic
8502 (SC33)	• Oracle Solaris Cluster private traffic
8503 (ZFS/APP)	• Between Oracle RAC and the Sun ZFS Storage 7320 Appliance • Between the application VM and the Sun ZFS Storage 7320 Appliance • Between the application VM and Oracle RAC

Analyzing Solution Characteristics

To understand the behavior of the architecture under peak load conditions, determine optimum utilization, and verify the scalability of the solution, Oracle engineers recently tested the Oracle Optimized Solution for PeopleSoft Human Capital Management running on Oracle's SPARC SuperCluster T4-4. Engineers tested the solution with a PeopleSoft Payroll for North America workload. The payroll workload requires very intensive compute and throughput (I/O) capabilities in order to complete the heavy batch jobs being executed.

Test Configuration and Workloads

Table 2 lists the hardware and software configuration used during the testing efforts.

TABLE 2. ORACLE OPTIMIZED SOLUTION FOR PEOPLESOFT HCM TEST HARDWARE AND SOFTWARE CONFIGURATION	
SERVER HARDWARE CONFIGURATION	
Oracle SPARC SuperCluster	• Half Rack SPARC SuperCluster T4-4
LOAD GENERATORS (HARDWARE)	
	• 4 x86 clients running Windows Server 2003 SP2
PEOPLESOFT SERVER INFRASTRUCTURE (PRE-RELEASE)	
	• PeopleTools 8.51.11 • PeopleSoft HRMS 9.10.303 • WebLogic 10.3.3 • Tuxedo 10.3.0.0 (RP031 patch) • Oracle Solaris Cluster 3.3u1 • Oracle Solaris 10 u10 (Applications VM) • Oracle Solaris 11 (DB VM) • Oracle Database 11.2.0.3
LOAD CONTROLLER	
	• Oracle Application Testing Suite (OATS) 9.21.0030

The solution is delivered pre-installed and pre-configured. During testing, the storage, operating system, and database were configured and tuned as shown in Table 3.

TABLE 3. PEOPLESOFT ENTERPRISE HR TEST CONFIGURATION	
STORAGE POOLS	
App-Pool	• Triple mirrored pool, no single point of failure • 5 disk drives from each of 2 disk shelves (10 total) • 1 logzillas from the 2 shelves closest to the head, mirrored and striped (2 total) • 2 striped readzillas on the primary active head for this pool • No readzillas when failed over
APP-POOL PROJECTS AND SHARES	
App Project	• All shares with default record size, read-flash enabled, logbias = latency • 1 NFS share for PeopleSoft web-tier (WebLogic + PeopleTools + HCM application) - approximately 25 GB • 1 NFS share for PeopleSoft app-tier (Tuxedo + PeopleTools + HCM application) – ~30G
Cluster-Quorum Project	• iSCSI LUN for quorum device - 1GB, 8k volume block size, read-flash enabled, logbias = latency

Table 4 lists additional settings used during the testing process.

TABLE 4. ADDITIONAL SETTINGS		
ORACLE DATABASE		
PARAMETERS	LOCATION	NOTES
optimizer_dynamic_sampling=0 _gby_hash_aggregation_enabled=false _unnest_subquery=false _ignore_desc_in_index=true	pfile or spfile	Oracle recommends setting these database initialization parameters for all PeopleSoft application modules in the environment's designated database initialization file

PeopleSoft Payroll for North America Workload

The North American Payroll batch workload measures five Payroll application business process runtimes for one database model representing a large organization. The five Payroll processes used are:

- **Paysheet Creation.** Generates data worksheets consisting of standard payroll information for each employee for the given pay cycle. The Paysheet process can be run separately from the other two tasks, usually before the end of the pay period.
- **Payroll Calculation.** Looks at Paysheets and calculates checks for employees. Payroll Calculation can be run any number of times throughout the pay period. The first run performs most of the processing, while each successive run updates only the calculated totals of changed items. This iterative design minimizes the time and processing resources required to calculate payroll. In this workload Payroll Calculation was run only once, simulating the end of a pay period.
- **Payroll Confirmation.** Takes the information generated by Payroll Calculation and updates employee balances with the calculated amounts. The system assigns check numbers at this time and creates direct deposit records. Payroll Confirmation can only be run once, and therefore must be run at the end of the pay period.
- **Print Advice Forms.** Takes the information generated by Payroll Calculation and Confirmation and produces an Advice for each employee to report earnings, taxes, deductions, net pay, and bank accounts where Net Pay was sent.
- **Create Direct Deposit File.** Takes the information generated by Payroll Calculation and Confirmation and produces an electronic transmittal file used to transfer payroll funds directly into an employee's bank account.

PeopleSoft Human Resources Self Service Online Workload

The Human Resources (HR) Self Service workload consists of measured client response times for 2,000, 4,000 and 8,000 concurrent users. The 13 PeopleSoft Enterprise HR business processes tested in this workload are listed in Table 5:

TABLE 5. PEOPLESOFT ENTERPRISE HR SELF-SERVICE WORKLOADS	
EMPLOYEE SELF-SERVICE (60% OF TOTAL LOAD)	
TRANSACTION	**PERCENTAGE OF USERS**
Update Home Address: Update address in Personal Data section	3
Update Home Phone: Update phone number in Personal Data section	3
View Benefits Summary: View overall benefits enrollment data	10
Benefits Change Life: View benefits and alter the beneficiaries' allocations in the Basic Life Plan	2
View Paycheck: View current paycheck information	78
Update Direct Deposit Info: Add a direct deposit directive	2
Add Job Profile	2
MANAGER SELF-SERVICE (20% OF TOTAL LOAD)	
TRANSACTION	**PERCENTAGE OF USERS**
View Employee Info: View job and personal information	50
Initiate Termination: Initiate a termination by recording an effective date and reason for termination	20
Initiate Promotion: Initiate a promotion by entering a new job title and salary	10
Initiate Employee Salary Change: Process a salary change for a single employee	20
HR ADMINISTRATION (20% OF TOTAL LOAD)	
TRANSACTION	**PERCENTAGE OF USERS**
Add a person and their biographical details	20
Enter the specified job data and work location, followed by the payroll and compensation details	40
Add job details to an existing employee	40

The online services portion of the testing makes use of database tables in a very different way than the payroll process. Users log in to the services and update information randomly. In the Payroll process, multiple application tables are updated sequentially or randomly by multiple Payroll job streams. As a result, it is difficult to predict which tables are best placed in the Flash storage devices.

Solution High Availability Test Results

To validate high availability of the overall solution, numerous rigorous failure tests were performed under load on the App and Web architecture tiers, as well as on the whole Solaris 10 domain (Applications VM) and on a complete SPARC T4-4 server node. The following tables list the measured results. The numbers provided in this document are for indication of activities duration and based on the cluster configured using default settings.

Table 6 shows the results of failing one of the zone cluster nodes in the Web tier. The fault was injected to the node running the Web Admin Server but not the Reverse Proxy Server. The faulted zone is automatically rebooted back by the Solaris OS. Oracle Solaris Cluster then automatically restarts the Web tier services in the zone.

TABLE 6. WEB TIER FAILURE TEST OBSERVATIONS

PROCESS STAGE	TIME (IN SECONDS)
Failure detection	13
Failover completed	11
End to end fault to services recovery	13 + 11 = 24
Recovery to full redundancy	68

- Failure detection = Time for the Oracle Solaris Cluster framework to detect the zone failure.
- Failover completed = Time for the Web tier processing to recover from the failed zone cluster node and run with just the remaining node.
- End to end fault to services recovery = Total time for detection and failover
- Recovery to full redundancy = Time for the failed zone cluster node to restart and for both nodes to resume normal active/active processing operations, after failover completion.

Table 7 shows the results of failing one of the zone cluster nodes in the Application tier. The faulted zone is automatically rebooted back by the Solaris OS. Oracle Solaris Cluster then automatically restarts the App tier services in the zone.

TABLE 7. APPLICATION TIER FAILURE TEST OBSERVATIONS

PROCESS STAGE	TIME (IN SECONDS)
Failure detection	6
Failover completed	4
End to end fault to services recovery	6 + 4 = 10
Recovery to full redundancy	81

- Failure detection = Time for the Oracle Solaris Cluster framework to detect the zone failure.
- Failover completed = Time for the App tier processing to recover from the failed zone cluster node and run with just the remaining node.
- End to end fault to services recovery = Total time for detection and failover
- Recovery to full redundancy = Time for the failed zone cluster node to restart and for both nodes to resume normal active/active processing operations, after failover completion.

Table 8 shows the results of failing one of the Solaris 10 domains containing two zone cluster nodes (Web and App). The fault was injected to the domain on the SPARC T4-4 server node running the Reverse Proxy Server (RPS).

TABLE 8. SOLARIS 10 DOMAIN FAILURE TEST OBSERVATIONS	
PROCESS STAGE	**TIME (IN SECONDS)**
Failure detection	8
Failover completed (App tier only)	8
Failover completed (App and Web tiers, including RPS)	79
End to end fault to services recovery	8 + 79 = 87

- Failure detection = Time for the Oracle Solaris Cluster framework to detect the failure of the App domain and its two zone cluster nodes (Web and App).
- Failover completed (App tier only) = Time for the App tier processing to recover from the failed zone cluster node and run with just the remaining node.
- Failover completed (App and Web tiers, including RPS) = Time for App and Web tiers processing to recover from the failed zone cluster node and run with just the remaining node. This includes the failover of the Reverse Proxy Server to the other node.
- End to end fault to services recovery = Total time for detection and failover

Table 9 shows the results of failing one of the SPARC T4-4 servers. The power failure was injected on the server that is not running the Reverse Proxy Server (RPS).

TABLE 9. COMPLETE NODE FAILURE TEST OBSERVATIONS	
PROCESS STAGE	**TIME (IN SECONDS)**
Application Domain (Solaris 10)	
Failure detection	9
Failover completed (App tier only)	12
Failover completed (App and Web tiers)	13
End to end fault to services recovery	9 + 13 = 22
Database Domain (Solaris 11)	
Failure detection	28
Failover completed	10
End to end fault to services recovery	28 + 10 = 38

The results were measured for both the Application domain and the Database domain.

In the Application domain:

- Failure detection = Time for the Oracle Solaris Cluster framework to detect the failure of the App domain and its two zone cluster nodes (Web and App).
- Failover completed (App tier only) = Time for the App tier processing to recover from the failed zone cluster node and run with just the remaining node.

- Failover completed (App and Web tiers) = Time for App and Web tiers processing to recover from the failed zone cluster node and run with just the remaining node.
- End to end fault to services recovery = Total time for detection and failover

In the Database domain:

- Failure detection = Time for Oracle Clusterware to detect the failure of the Database domain.
- Failover completed = Time for the Database tier processing to recover from the failed RAC node and run with just the remaining node.
- End to end fault to services recovery = Total time for detection and failover

PeopleSoft Payroll for North America Workload Test Results

The test measures the performance and scalability of PeopleSoft Payroll for North America running on Oracle's SPARC SuperCluster using three Exadata Storage servers tuned for the workload at hand. For this workload, tests simulate an environment of 500,480 employees, with a total 750,720 payments using 96 job streams. The Process Scheduler is running in the application tier (container), and the payroll database is running in archive log mode using the flashback option. The results are shown in Table 10.

TABLE 10. SUPERCLUSTER: ORACLE PEOPLESOFT NORTH AMERICA PAYROLL 9.1 (UNICODE) TEST RESULTS

		PERCENT CPU UTILIZATION			MEMORY FOOTPRINT IN GB		
BUSINESS PROCESS NAME	CHECKS PROCESSED/ HOUR[1]	RAC DB NODE1	RAC DB NODE2	PROCESS SCHEDULER	RAC DB NODE1	RAC DB NODE2	PROCESS SCHEDULER
Paysheet Creation	14,817,000	15	13	18	99	99	19
Payroll Calculation	2,334,000	12	16	32	110	112	37
Payroll Confirmation	2,322,000	17	12	18	101	100	28
NET PAYMENTS PER HOUR 1,080,000							
Print Advice Forms	2,596,000	9	9	24	100	100	16
Create Direct Deposit File	36,035,000	39	32	1	101	102	17

Note: The performance results illustrated in this document were obtained on a highly available Oracle SPARC SuperCluster T4-4 system as it could be deployed at a customer site, and not on a special-purpose, performance-tuned benchmark setup. There are many critical differences between real-world and one-time benchmark configurations. Designed to reduce latency and overhead, increase throughput

[1] Per hour throughputs are rounded to the nearest thousand

and bandwidth, and simplify configurations overall, most benchmarks generally do not implement high availability features. For example, in benchmarks, most data typically is stored on striped local disks instead of using highly available shared storage. For these reasons, the performance results disclosed in this document are designed to provide general guidance for real world deployments, not to be compared with published benchmark results from Oracle or other vendors.

PeopleSoft Human Resources Self Service Online Test Results

The duration of each test run documented below was approximately 50 minutes at steady-state.

Table 11 lists the resource allocation resource utilization, and average transaction response times for the PeopleSoft HR Self-Service module with 2,000 concurrent users.

TABLE 11. PEOPLESOFT ON SPARC SUPERCLUSTER USE-CASE HRMS 9.0 SELF-SERVICE ONLINE TESTING RESULTS WITH 2,000 CONCURRENT USERS

RESOURCE ALLOCATIONS				
VIRTUAL MACHINE (LDOM)	**VIRTUAL CPUS PER NODE**		**MEMORY PER NODE**	
Database	128		256 GB	
Web and Application	128		256 GB	
RESOURCE UTILIZATION				
TIER	**NODE 1 CPU PERCENTAGE**	**NODE 2 CPU PERCENTAGE**	**NODE 1 MEMORY (GB)**	**NODE 2 MEMORY (GB)**
Web	7	9	22	26
Application	6	6	24	24
Database	2	3	78	78
AVERAGE TRANSACTION RESPONSE TIMES				
TRANSACTION TYPE				
Average Search	Under 0.5 second			
Average Save	Under 0.5 second			

Table 12 lists the resource allocation resource utilization, and average transaction response times for the PeopleSoft HR Self-Service module with 4,000 concurrent users.

TABLE 12. PEOPLESOFT ON SPARC SUPERCLUSTER USE-CASE HRMS 9.0 SELF-SERVICE ONLINE TESTING RESULTS WITH 4,000 CONCURRENT USERS

RESOURCE ALLOCATIONS				
VIRTUAL MACHINE (LDOM)	**VIRTUAL CPUS PER NODE**		**MEMORY PER NODE**	
Database	128		256 GB	
Web and Application	128		256 GB	
RESOURCE UTILIZATION				
TIER	**NODE 1 CPU PERCENTAGE**	**NODE 2 CPU PERCENTAGE**	**NODE 1 MEMORY (GB)**	**NODE 2 MEMORY (GB)**
Web	12	7	26	22
Application	13	10	29	29
Database	4	4	78	78
AVERAGE TRANSACTION RESPONSE TIMES				
TRANSACTION TYPE				
Average Search	Under 0.5 second			
Average Save	Under 0.5 second			

Table 13 lists the resource allocation resource utilization, and average transaction response times for the PeopleSoft HR Self-Service module with 8,000 concurrent users.

TABLE 13. PEOPLESOFT ON SPARC SUPERCLUSTER USE-CASE HRMS 9.0 SELF-SERVICE ONLINE TESTING RESULTS WITH 8,000 CONCURRENT USERS

RESOURCE ALLOCATIONS				
VIRTUAL MACHINE (LDOM)	**VIRTUAL CPUS PER NODE**		**MEMORY PER NODE**	
Database	128		256 GB	
Web and Application	128		256 GB	
RESOURCE UTILIZATION				
TIER	**NODE 1 CPU PERCENTAGE**	**NODE 2 CPU PERCENTAGE**	**NODE 1 MEMORY (GB)**	**NODE 2 MEMORY (GB)**
Web	24	31	26	23
Application	42	33	52	51
Database	10	9	78	78
AVERAGE TRANSACTION RESPONSE TIMES				
TRANSACTION TYPE				
Average Search	Under 0.5 second			
Average Save	Under 1 second			

Results Analysis

The Oracle Optimized Solution for PeopleSoft Human Capital Management based on Oracle SPARC SuperCluster is ideal for datacenters looking for rapid, easy PeopleSoft HR application deployments with solid performance and availability characteristics. IT staff can be confident that all components are fully tested and work together on a fully integrated and optimized solution platform. The datacenter environment is further enhanced by the ability to use other management tools from Oracle, including Oracle Enterprise Manager Ops Center and Oracle Enterprise Manager. The result is a swiftly deployed, complete, managed environment requiring a minimum of implementation time and effort.

Best Practices for Optimizing Workload Performance

Workload optimization is vital to achieving the best possible deployment results. Testing revealed several best practices that organizations can use to achieve optimal deployment availability and performance. The following sections outline the particular settings addressed and apply not only to Payroll but also to the majority of batch and online PeopleSoft workloads.

The values provided in this section should be adapted and tuned according to the environment and the configuration.

Best Practices for Configuring PeopleSoft Payroll for North America

One recommended practice for configuring workload processes is to create concurrent batch jobs. The following steps set up batch jobs that execute concurrently:

- Increase Max Concurrent of Cobol SQL processes to desired number of concurrent Payroll streams
- Create multiple `PSAESRV` processes on the process scheduler server.
- Set the value of `Max Instances` to the desired number in the `[PSAESRV]` section of the `psprcs.cfg` configuration file.
- Log in to PeopleSoft Internet Architecture (PIA) and set the same value for the `Max API Aware` field in the Process Scheduler -> Servers -> PSUNX -> Server Definition Web interface. The `Max API Aware` value should be larger than, or equal to, the total of `Max Concurrent` values set for all of the process types including `Application Engine` and Cobol SQL.
- Split the workload into multiple batch jobs.

Temporary Files

While executing business transactions, PeopleSoft Application Server creates and destroys temporary files dynamically and transparently in the /var/tmp directory on Oracle Solaris. This behavior may keep the underlying I/O device busy under load. On a very busy system, the I/O device could become a bottleneck and significantly slow transaction execution. In such situations, replacing the I/O device with a high performing device or pointing the /var/tmp to the temporary filesystem, /tmp, can improve the performance of the application server.

For example:

```
# mv /var/tmp /var/tmp_orig ; ln -s /tmp /var/tmp
```

Optimize Oracle Database

There are a few recommended database practices when running PeopleSoft Payroll for North America workloads.

Data Partitioning

Partitioning is a data volume management technique that addresses the key problem of supporting very large tables and indexes by allowing tables and indexes to be broken down into smaller and more manageable pieces called partitions. It is a good idea to partition hot tables that contain millions of records. Once partitions are defined, SQL statements can access and manipulate the data in smaller partitions rather than the entire table or index. It is advisable to choose a partition key that will allow each concurrent job stream to operate on an individual partition.

Optimize Statistics

By default, Oracle Database 10*g* and 11*g* database use the Oracle Cost-Based Optimizer. When using the Oracle Cost-Based Optimizer, table and index statistics play a vital role in query performance. Maintaining these statistics is critical to optimal database and query performance. From time to time, gather statistics for objects where the statistics become stale over time because of changing data volumes or changes in column value. New statistics should be gathered after a schema object's data or structure are modified in ways that make the previous statistics inaccurate. In general we generate statistics without histogram.

For example, after loading a significant number of rows into a table, collect new statistics on the number of rows. After updating data in a table, new statistics on the average row length are needed. PeopleSoft recommends gathering the statistics for the whole schema and for each individual table with the following SQL queries:

```
DBMS_STATS.GATHER_SCHEMA_STATS (ownname => [table_owner], ESTIMATE_PERCENT =>
DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt => 'FOR ALL INDEXED COLUMNS SIZE 1');

DBMS_STATS.GATHER_TABLE_STATS (ownname => [table_owner], tabname => table_name],
ESTIMATE_PERCENT => DBMS_STATS.AUTO_SAMPLE_SIZE, method_opt=> 'FOR ALL INDEXED COLUMNS
SIZE 1', CASCADE => TRUE);
```

Database Indexes

Indexes are vital to database operation. Several best practices can help to optimize database indexes:

- **Exercise caution when creating new indexes.** Performance issues may or may not be caused by a missing index. Be careful about adding new indexes to resolve a performance issue. If a new index appears to resolve a critical performance problem, monitor the overall performance of the database, not just the targeted query. It is critical to ensure that the new index does not cause unintended side effects.

- **Keep indexes from affecting DML operation performance.** If Oracle Data Manipulation Language (DML) statements that modify data—such as `INSERT`, `UPDATE`, or `DELETE`—are being executed many times on a table, make sure that the addition of a new index on the same table does not negatively affect the performance of the DML operations. This is generally not a problem if the SQL queries being executed are retrieving—but not adding or modifying—the existing data. In all other cases, there is some index maintenance overhead. For example, if 10 indexes are created on a table named `TEST`, adding a new row of data to the `TEST` table may require the database management system to update all 10 indexes.

- **Instrument indexes on heavily operated tables.** From time to time, the indexes on heavily operated tables should be instrumented and the index usage monitored for periods of bursty database activity. To instrument an index to monitor usage, run the following SQL command:

```
ALTER INDEX <INDEX_NAME> MONITORING USAGE;
```

To stop monitoring the index usage once database activity declines, run:

```
ALTER INDEX <INDEX_NAME> NOMONITORING USAGE;
```

- **Validate instrumentation.** Once instrumented, query the V$OBJECT_USAGE view in the database intermittently to check whether the instrumented index is being used in executing any SQL queries.

```
SELECT USED, START_MONITORING
    FROM V$OBJECT_USAGE
    WHERE INDEX_NAME LIKE '%<INDEX_NAME>%'
    /
```

- **Function-based indexes.** Function-based indexes with descending columns are being recreated as regular indexes with ascending columns.

- **Object Privileges.** Ensure the database user PSADMIN has the privilege to select rows from V$MYSTAT and V$SESSION views. Missing the SELECT privilege on those views could lead to a number of cursor invalidations and cause the database to eventually exhibit poor performance under load on the system.

```
GRANT SELECT ON V_$MYSTAT to PSADMIN;
GRANT SELECT ON V_$SESSION to PSADMIN;
```

Database Initialization Parameters

For best performance, Oracle recommends setting certain database initialization parameters to match the specific PeopleTools and Oracle Database versions running in the environment. These parameters should be set in a PFILE or a Server Parameter File (SPFILE), depending on which is the designated Oracle Database initialization parameter file. For more information, see Oracle Support document ID 747587.1, *PeopleSoft Enterprise PeopleTools Certifications*.

Optimizing the PeopleSoft Human Resource Management System

Some parameters affect different components of the Oracle PeopleSoft HR environment and should be modified to obtain the best possible results.

Best Practices for Configuring PeopleSoft HR Application and Web Servers

For optimal performance, the following settings in Table 14 should be modified in the `psappsrv.cfg` file:

TABLE 14. PEOPLESOFT HR APPLICATION AND WEB SERVER PARAMETERS IN `PSAPPSRV.CFG`	
KEY JOLT LISTENER SETTINGS	
Min Handlers=XX	The minimum number of handlers should be at least the expected concurrency divided by **Max Clients per Handler**. For example, if 100 users is the expected concurrency, **`Min Handlers`** should be set to a value superior to 100/10 = 10.
Max Handlers=XY	A few additional handlers should be configured to handle peak loads.
Max Clients per Handler=XZ	The default value for **`Max Clients per Handler`** is 40, for our tests we used a value of 10. In a production environment with large numbers of concurrent users this value should be higher than 10.
KEY `PSAPPSRV` UBBGEN SETTINGS	
Min Instances=YX	This value should be monitored and adjusted depending on the workload. The general value for minimum instances used in this test is one PSAPPSERV process per 100 concurrent users. For example, if the expected concurrency is 500 users, set **Min Instances** = **Max Instances** = 5.
Max Instances=YY	**Max Instances** should always be set to a value superior to the value of **Min Instances**.
OTHER KEY `PSAPPSRV` SETTINGS	
Recycle Count=ZX	To minimize disruptions, **`PSAPPSERV recycle count`** was set to zero (0) during our testing. However, in a production environment, this value should always be higher than zero and adjusted depending on the frequency of restarts of the domain and on the amount of memory.

Oracle also recommends monitoring the `PSAPPSRV` queue size. Using the `psadmin` command line utility, monitor the `PSAPPSRV` queue size and increase or decrease the number of `PSAPPSRV` instances as needed.

PeopleSoft Application Server Cache

In order to improve responsiveness, the PeopleSoft Application Server tries to minimize database access by caching metadata and objects in memory and when necessary on the file system. On a busy system, the I/O device(s) that holds the actual cache files could become a bottleneck, causing significant delays in server responses to clients. In such a situation, consider spreading the application server cache files on a high performing I/O device. During this test the cache files were stored in shared storage on the ZFS Storage appliance.

Oracle WebLogic Server Settings

There are a number of recommended best practices for configuring the Oracle WebLogic Server to obtain optimal performance.

- **Distribute the user workload for the Web Server domain.** Configure one Web server domain to handle 400 to 500 concurrent users at most. For example, if the expected concurrency is 1200 users, configure at least three identical web domains.

- **Enable Web server load balancing.** If multiple application servers are to be used for scalability or reliability purposes, edit `<web_domain>/applications/peoplesoft/PORTAL/WEB-INF/psftdocs/ps/configuration.properties` to make the Web server balance the load. The following parameter sets the Web server to load balance among `server1`, `server2`, through `servern` in a round-robin fashion:

```
psserver=<server1>:<jslport>,<server2>:<jslport>,..,<servern>:<jslport>
```

- **Modify file descriptor settings in the `.profile` file.** For both the Oracle WebLogic application and Web servers, the file descriptor settings should be modified in the `.profile` file to set the shell limit for the number of open files. In addition, modify the value of the parameter `MAX_FILE_DESCRIPTORS` to 4096 in the `PS_HOME/webserv/<domain_name>/bin/setEnv.sh` script to reflect the file descriptor limit configured in `.profile`.

```
/* For the application servers */
ulimit -SH -n 1024

/* For the Web servers */
 ulimit -SH -n 4096
```

- **Set the shell limit for open files in `$WL_HOME/common/bin/commEnv.sh`.** If the solution environment uses the Oracle WebLogic Web server, the same Web server shell limit for open files should be set in `$WL_HOME/common/bin/commEnv.sh`. The `ulimit -n` statement is found in the `resetFd(){}` sub-function in `commEnv.sh`.

```
ulimit -SH -n 4096
```

Java Virtual Machine Settings

There are also recommended best practices to achieve optimal performance for the Java Virtual Machine.

- **Set shell variables in the `setEnv.sh` file.** Edit the `JAVA_OPTIONS_SOLARIS` variable to tune the Java Virtual Machine (JVM) options that are appropriate for the deployment.

For example:

```
JAVA_OPTIONS_SOLARIS="-server -Xms512m -Xmx1024m -XX:PermSize=256m -
XX:MaxPermSize=256m -
Dtoplink.xml.platform=oracle.toplink.platform.xml.jaxp.JAXPPlatform -
Dcom.sun.xml.namespace.QName.useCompatibleSerialVersionUID=1.0 -XX:+DisableExplicitGC
-Xnoclassgc -Xrs -Xss256k -Xverify:none -XX:+UseParallelGC -XX:ParallelGCThreads=4 -
XX:LargePageSizeInBytes=256m -Dweblogic.threadpool.MinPoolSize=50 -
Dweblogic.threadpool.MaxPoolSize=75 -Dweblogic.SocketReaders=10 -
Dweblogic.MuxerClass=weblogic.socket.NIOSocketMuxer -
Dweblogic.management.discover=false -
Dweblogic.diagnostics.debug.DebugLogger.DISABLED=true -Dweblogic
.GatheredWritesEnabled=true -Dweblogic.ScatteredReadsEnabled=true -
Dweblogic.configuration.schemaValidationEnabled=false -
Djava.net.preferIPv4Stack=true"
```

- **Monitor JVM garbage collection (GC) activity.** Add the following JVM options to the above parameters:

```
-verbose:gc –XX:+PrintGCDetails –Xloggc:/tmp/gc.$$.log
```

Inter-Process Communication (IPC) Settings

The Oracle Solaris project facility lets administrators control resources used by processes and simplify certain parameter settings. For this reason, projects should be used for setting certain parameters such as shared memory size and the maximum number of message queues.

For example, sometimes the default value for the maximum number of message queues (128) cannot handle an application server instance configured for a large number of users. When all message queues are filled, the following error message occurs as the application server processes are booting.

```
Booting server processes ...
exec PSSAMSRV -A -- -C psappsrv.cfg -D CS90SPV -S PSSAMSRV :
       Failed.
113954.ben15!PSSAMSRV.29746.1.0: LIBTUX_CAT:681: ERROR: Failure to create message
queue
113954.ben15!PSSAMSRV.29746.1.0: LIBTUX_CAT:248: ERROR: System init function
failed, Uunixerr = :
msgget: No space left on device
113954.ben15!tmboot.29708.1.-2: CMDTUX_CAT:825: ERROR: Process PSSAMSRV at ben15
failed with /T
tperrno (TPEOS - operating system error)
```

- **Set the project resource control to a value larger than the default of 128.** In the above situation, Oracle recommends setting the project resource control `project.max-msg-ids` to a value larger than the default of 128. As a best practice, Oracle recommends setting the value to double that amount. To increase the `max-msg-ids` to a value of 256, perform the following steps:
- Get the project ID.

```
% id -p
uid=222227(psft) gid=2294(dba) projid=3(default)
```

- Increase the maximum value for the message queue identifiers to 256 using the `prctl` utility.

```
# prctl -n project.max-msg-ids -r -v 256 -i project 3
```

- Verify the new maximum value for the message queue identifiers.

```
# prctl -n project.max-msg-ids -i project 3
project: 3: default
NAME    PRIVILEGE      VALUE   FLAG    ACTION  RECIPIENT
project.max-msg-ids
        privileged     256     -       deny    -
        system         16.8M   max     deny    -
```

- The following command makes this setting persistent across system reboots.

```
# projmod -p 3 -c "PeopleSoft  IPC tuning" -K "project.max-msg-
ids=(priv,256,deny)" default
```

- **Increase the size of `project.maxshm-memory`.** By default, the maximum shared memory segment size is set to 25% of installed physical memory. If a large System Global Area (SGA) is needed for the database, the `project.maxshm-memory` resource control should be increased so that the database can start successfully. If `project.max-shm-memory` is set too low, Oracle Database outputs the following error during startup:

```
ORA-27102: out of memory
```

The following example illustrates how to increase the maximum allowed shared memory to 32 GB under project `default`.

```
# projmod -p 3 -c "Fix to ORA-27102 out of mem" -K "project.max-shm-
memory=(priv,32G,deny)" default
```

For More Information

For more information on the Oracle Optimized Solution for PeopleSoft Human Capital Management, see the references listed in Table 15.

TABLE 15. REFERENCES FOR MORE INFORMATION

Oracle Optimized Solutions	http://oracle.com/optimizedsolutions
Oracle SPARC SuperCluster	http://www.oracle.com/supercluster
Oracle SPARC T-series Servers	http://www.oracle.com/goto/tseries
Oracle Solaris	http://www.oracle.com/solaris
Oracle's Sun ZFS Storage Appliance	http://www.oracle.com/us/products/servers-storage/storage/unified-storage/
PeopleSoft Best Practices Center	http://www.oracle.com/technetwork/middleware/fmw4apps/peoplesoft/index.html

Oracle Optimized Solution for Oracle's PeopleSoft Human Capital Management March 2012

Author: Pierre Reynes, Jacques Bessoudo, Uday Shetty, Giri Mandalika, Tom Duell, Gia-Khanh Nguyen

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

SOFTWARE. HARDWARE. COMPLETE.