

ORACLE®

OPTIMIZED SOLUTIONS

An Oracle Technical White Paper
October 2011

Oracle Optimized Solution for JD Edwards EnterpriseOne

Delivering Performance and Efficiency Using Oracle's SPARC
T4-1 Server

ORACLE®

Introduction	1
A Brief Overview.....	2
Solution Architecture Overview.....	3
A New Approach to Optimizing Storage.....	4
Solution Software Components	4
Licensing of the Software Stack.....	5
Solution Hardware Components	6
Storage Architecture	7
Workload Description	8
Test Environment	9
SPARC T4-1 Server Configuration	9
Test Results and Analysis	13
Test Run 1: Best Achievable Response Time During DIL Testing	13
Test Runs 2 and 3: Extended Response Time Testing with Increased Short UBE Load	14
Test Run 4: Maximizing Short UBE Throughput with 0.5 Second Response Times.....	14
Test Run 5: Maximum Long UBE Throughput with Acceptable Response Times.....	14
Test Runs 6 and 7: Best Compromise Performance for Long and Short UBEs with Acceptable Response Times	14
Final Analysis	15
Real-World Sizing Based on the DIL Test Workload.....	15
Room for Consolidation	15
Sizing Guidelines.....	16
Conclusion	19
References.....	19
Appendices	20
A. Configuring the Sun Storage 6 Gb SAS PCIe RAID HBA Card	20
B. DIL Benchmark Tools Configuration.....	22
C. Tuning and Best Practices	25
D. Workload Details.....	36

Introduction

Whether organizations are deploying a new application service or modernizing existing business critical services, selecting and implementing a solution architecture that will perform optimally and meet service level agreements is a challenging process. An optimized solution must balance technical requirements against business needs such as budget constraints, performance and reliability requirements, and future scalability needs.

This paper describes Oracle's extensive testing of a hardware and software combination optimized to efficiently host the JD Edwards EnterpriseOne suite of Enterprise Resource Planning (ERP) tools from Oracle. The paper demonstrates how Oracle's servers enhanced with carefully tuned conventional and advanced flash-based storage technology provide an excellent price/performance combination for small and medium enterprise deployments.

Server consolidation through virtualization is key to this solution, allowing a very dense and yet high-performance solution at low cost and low data center impact. The result is significant reductions in both initial acquisition and ongoing operating costs to the enterprise.

The JD Edwards EnterpriseOne stack consists of a core set of tools supplemented by more than seventy application modules addressing industry-specific enterprise resource planning needs. In the testing carried out as part of this Oracle Optimized Solution, commonly used modules are included in a standardized "Day in the life" (DIL) workload that is used to obtain performance metrics for particular system configurations to characterize their behavior.

As part of this paper, a sizing guide and tuning best practices section are provided as well as information regarding next steps should customers need to scale within and beyond the capabilities of the chosen platform.

A Brief Overview

Today's enterprises are faced with the need to track assets and ensure legal compliance in ever more complex environments. JD Edwards EnterpriseOne is a cost-effective and practical way of ensuring proper tracking and compliance. The suite of software consists of three core functional blocks, on top of which application modules can be loaded. These core functions are listed in Table 1.

TABLE 1. CORE FUNCTIONAL BLOCKS IN JD EDWARDS ENTERPRISEONE SOFTWARE SUITE

SERVER/FUNCTION	DESCRIPTION
A database server	A conventional Oracle Database 11g Release 2 Standard Edition server for transaction storage and asset tracking
A Web server	A conventional Oracle WebLogic Server 10.3.2 Standard Edition Web server providing the user interface to the optional modules
The JD Edwards EnterpriseOne application server and optional modules	The JD Edwards EnterpriseOne 8.98/9.0 application server core installation and optional business functionality modules

The JD Edwards EnterpriseOne core installation requires only Standard Edition versions of Oracle Database and Oracle WebLogic server software. Upgrades to Enterprise Edition versions are possible to take advantage of more advanced capabilities. However, these are not necessary for the configurations discussed in this document.

Taking advantage of the Oracle Solaris virtualization features, each of these 'servers' is hosted inside an Oracle Solaris Container hosted on a single SPARC T4-1 server from Oracle. This approach provides several advantages to the user:

- Separation of the applications' working environments simplifies many aspects of security and management.
- Portability and backup options are improved through integration with ZFS.
- Oracle Solaris Containers enable fine-grained control over access to system resources for tuning purposes.
- Systems consolidation and interconnect elimination significantly reduces data center footprint and simplifies management.

Compared to competitive configurations consisting of multiple servers and storage devices, the consolidated approach offered by the Oracle Optimized Solution for JD Edwards EnterpriseOne provides greatly simplified deployment and management.

Solution Architecture Overview

The solution architecture takes advantage of Oracle Solaris Containers and optimizes the use of onboard storage to take best advantage of the resources and bandwidth available within the platform. Figure 1 shows how the storage devices are mapped to specific containers within the SPARC T4-1 server.

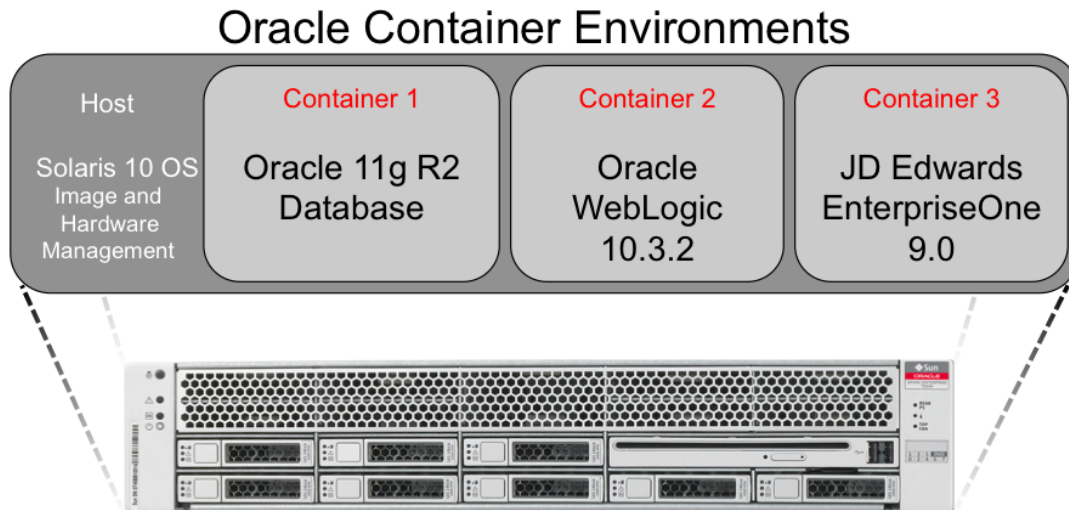


Figure 1. Technical architecture for the Oracle Optimized Solution for JD Edwards EnterpriseOne.

The SPARC T4-1 server performs four primary roles, as identified by the four functional areas listed in Figure 1. The base platform provides the underlying capabilities upon which three Oracle Solaris Containers are created to carry out the operations required of the full JD Edwards EnterpriseOne stack.

In the base platform, the disk configuration is managed and configured in such a way as to optimize performance of the storage subsystems for each of the Oracle Solaris Containers that provide the application functionality. In this base platform, the base operating system is installed, which consists of Oracle Solaris 10 10/09, and all relevant patches and any required additional platform-wide facilities are configured and managed. In addition, the required security profile is applied to lock down the base installation, automatically adding security to the virtualized Oracle Solaris Containers installations.

A New Approach to Optimizing Storage

The key area requiring optimization in this solution is the disk I/O performance. This can be seen in analysis of the performance results later in this document. The classic solution to this problem has been to utilize external high-performance storage devices, such as SAN arrays, or to over-provision storage and under-utilize the capacity (often called ‘short stroking’ disk drives) to achieve a high spindle count and wastefully stripe data across the higher-speed portion of these disk to boost aggregate I/O operations per second (IOPS). These classical approaches to increasing storage speed are either very expensive (especially in the case of SAN storage) or wasteful in the sense that both capacity and power consumption are excessive when large numbers of sparsely populated ‘short-stroked’ disks/arrays are deployed.

The traditional solutions described above have long been used to address the problem, but the increasing cost pressures of administering these large disk farms (for example, floor space, increasing energy consumption, and heat generation) make it increasingly difficult to justify the cost/performance trade-off.

Oracle has been able to achieve optimal performance from the onboard disk storage of the SPARC T4-1 server through careful selection of an appropriate disk controller Host Bus Adapter (HBA) that provides performance acceleration via an onboard battery-backed write cache. This offers a cost-effective alternative to achieving sufficiently high performance in this application, with the added benefit of radically reducing the size, power, and management overhead often associated with external storage arrays, and it is an excellent way to accelerate databases and disk I/O intensive applications.

Solution Software Components

The primary software components that comprise this Oracle Optimized Solution are described in the subsections that follow.

Oracle Solaris 10 10/09 Operating System

Oracle Solaris provides industry-leading performance and numerous advanced features, including:

- **Virtualization**—Oracle Solaris Containers enable optimized resource utilization to deliver predictable service levels.
- **ZFS file system**—ZFS radically simplifies the management of complex storage subsystems and provides seamless data integrity and portability functions at no cost to the user. Performance optimizations leverage flash-based storage to give extreme performance with minimal administrative overhead.
- **Networking**—The open, virtualizable, and programmable Oracle Solaris networking stack delivers high throughput and the option to utilize integrated hardware encryption at zero cost.

JD Edwards EnterpriseOne 8.98/9.0

JD Edwards EnterpriseOne is an integrated application suite of comprehensive enterprise resource planning software that combines business value, standards-based technology, and deep industry experience into a business solution with a low total cost of ownership. A collection of prebuilt Business Intelligence (BI) applications available as modules for JD Edwards EnterpriseOne provide organizations the ability to implement and integrate more quickly, with less risk, and at a fraction of the cost required for building traditional BI solutions. JD Edwards EnterpriseOne provides organizations with the ability to transform information into business actions and enjoy a strategic advantage over less-nimble competitors.

Oracle Database 11g Release 2 Standard Edition

Designed for data center environments that are rapidly evolving and changing to keep up with the demands of the business, Oracle Database 11g provides efficient, reliable, and secure data management for mission-critical transactional applications, query-intensive data warehouses, and mixed workloads. With new self-managing capabilities, Oracle Database 11g also eliminates time-consuming, error-prone administrative tasks, so database administrators can focus on strategic business objectives instead of on performance and availability fire drills.

Oracle WebLogic Server 10.3.2 Standard Edition

Oracle WebLogic Server Standard Edition is an application server designed for building and running enterprise applications and services. It fully implements the latest Java Platform, Enterprise Edition (Java EE) standards and offers choice in development frameworks and tooling. Comprehensive and accessible management capabilities enable administration of sophisticated systems via a well-designed graphical console and/or automation. All users benefit from Oracle WebLogic Server's reliability and performance, which has been tested over years of enterprise-grade production use in demanding customer environments the world over.

Licensing of the Software Stack

Software licensing has a material impact on the total cost to deploy integrated systems. In addition to the end-user applications themselves, the applications often carry prerequisites, such as database management systems and Web middleware. This is true of JD Edwards EnterpriseOne, which at a minimum requires a database management system and a Java (Web) server.

In the context of the Oracle Optimized Solution for JD Edwards EnterpriseOne, the database is Oracle Database and the Java server is Oracle WebLogic Server. To provide “complete solution licensing” that mimics the “Optimized Solution engineering,” Oracle offers the licensable product called Oracle Technology Foundation for JD Edwards EnterpriseOne. This product provides limited-use licensing of Oracle Database Standard Edition, Oracle WebLogic Server Standard Edition, and several additional components of Oracle Fusion Middleware for use with JD Edwards EnterpriseOne deployments.

Because Oracle Technology Foundation for JD Edwards EnterpriseOne is licensed per application user, not per CPU, these components can be installed on an unlimited number of processors, providing some protection of software licensing costs due to CPU sprawl. An additional benefit is the ability to implement multiple, identical, hardware deployments for no additional license fees. This allows customers to provide development, QA, and testing instances of their deployments in an extremely cost-effective manner while further reducing development times and enhancing speed-to-deployment.

In cases where customers find value in the extended feature set offered by the Enterprise Edition of either the Oracle Database or Oracle WebLogic Server products, or if the broader implementation calls for use of these products by applications beyond JD Edwards EnterpriseOne, customers have the option of acquiring normal, full-use licenses instead.

Solution Hardware Components

The subsections below describe the server and storage components that were selected for this Oracle Optimized Solution.

SPARC T4-1 Server

Oracle's SPARC T4 processor-based servers are the platform of choice for a full range of data center applications and workloads. The SPARC T4-1 server was selected for the Oracle Optimized Solution for JD Edwards EnterpriseOne because of the key characteristics described in Table 2.

TABLE 2. KEY CHARACTERISTICS OF THE SPARC T4-1 SERVER THAT MAKE IT THE RIGHT CHOICE

SPARC T4-1 SERVER SYSTEM CHARACTERISTIC	KEY BENEFITS FOR THE SOLUTION
SPARC T4 processor	A single processor with 8 cores and 64 threads based on the latest SPARC core with enhanced single-threaded performance allows very effective provisioning of resources and enables the SPARC T4-1 server to run many virtual servers, helping drive up utilization, lower IT costs, and keep server sprawl to a minimum.
High memory capacity and bandwidth	Deployed with typically 64 GB of memory (but capable of up to 256 GB), the SPARC T4-1 server provides the capability of holding large volumes of in-memory data, optimizing Web and database response times, and providing a large workspace to the ZFS storage subsystem, enhancing its performance.
High-capacity disk and I/O subsystem	An eight-disk-capacity, local-storage subsystem within the chassis allows up to 4.8 TB of raw storage capacity that can be managed through onboard or enhanced hardware RAID controllers for ultimate data protection.

Sun Storage 6 Gb SAS PCIe RAID HBA

The Sun Storage 6 Gb SAS PCIe RAID HBA from Oracle provides very high data reliability, availability, and performance to businesses that are facing storage challenges driven by unprecedented data growth. With data transfer rates of up to 6 Gb/sec per port, the RAID adapter brings users improved performance and capabilities compared to onboard controllers. Designed for Oracle's latest high-speed PCI Express-based servers, Sun Storage 6 Gb SAS PCIe RAID HBAs couple advanced RAID data protection with 8-port, 6 Gb/sec SAS connectivity and 512 MB of Battery Backed Write Cache (BBWC) to provide secure, scalable, performance-optimized, enterprise-class data storage and retrieval solutions.

Storage Architecture

The underlying hardware RAID disk configuration is optimized to enhance both speed of access and reliability of data on disk. A 6 Gb SAS PCIe RAID HBA is used. This card includes a 512 MB write cache and battery back-up features that combine to improve data integrity and significantly accelerate write transactions by reducing latency and allowing very rapid, battery-backed, write-completion acknowledgement through NVRAM write cache. This HBA is used to provide three logical disk volumes to the platform optimized for write performance and data protection.

The logical volumes provided by the hardware RAID controller are managed by Oracle Solaris under the control of ZFS, providing greatly enhanced features for rapid import, export, and check pointing/rollback of any file system, which provides rapid paths to recovery in the event of failure and implicit data integrity checking and correction.

By appropriate positioning of data on these logical volumes, considerable performance gains have been achieved compared to standard disk configurations. Performance gains are such that the use of the in-chassis SAS2 storage devices provides sufficient performance without the need for the added expense and complication of external storage arrays.

TABLE 3. STORAGE CONFIGURATION

VOLUME	CONFIGURATION	DESCRIPTION
VOLUME 1	2 x 300 GB SAS2 disks RAID 1 mirrored	Exposed as a single logical drive, provides the Global Zone root OS boot disk for the platform and root images for the containers for the other virtual servers
VOLUME 2	3 x 300 GB SAS2 disks RAID 5 mirrored	Exposed to the operating system as a single logical drive; used by ZFS as a storage pool for Oracle Database
VOLUME 3	3 x 300 GB SAS2 disks RAID 5 mirrored	Exposed as a single logical drive; used by ZFS as a storage pool for Oracle Database Redo Log files Server data

Workload Description

In order for performance and scalability to be adequately characterized, a standard workload was used to stress the system under test and achieve the optimum utilization of the architecture under peak loads. The results provide a suitable basis for performance comparisons with alternative competitive architectures and allow customers to draw accurate comparisons based on their specific workloads during ‘real-life’ operations.

For this purpose, the JD Edwards EnterpriseOne product team provides a “Day-in-the-Life” (DIL) workload that uses a cross-section of JD Edwards EnterpriseOne features in a typical mid-scale manufacturing enterprise environment to simulate a typical day’s workflow. The workload can be configured to allow the maximum capability of the platform under test to be ascertained, while at the same time ensuring that various system performance metrics are captured. A key metric is the response time to user queries issued via the Web-based GUI as the number of simulated interactive users is increased during the testing. The workload can be throttled to maintain maximum performance with acceptable response times to allow tuning for the highest possible user count with an acceptable response time.

Interactive transactions, such as those in the DIL workload, typically generate only moderate load on a system such as the SPARC T4-1 server. To stress the system to determine its true capabilities, batch-type workloads must also be used. Batch processes are non-interactive and are often used for functions such as General Ledger processing and Sales Invoice Print Runs for which user interaction is not required. Such batch tasks can be broadly separated into short-running (less than a few tens of seconds to complete) and long-running (many minutes to hours to complete) tasks.

The JD Edwards EnterpriseOne terminology for these batch workloads is “Universal Batch Engine processes” (UBEs). UBEs are a common occurrence in the day-to-day operation of an enterprise, encompassing anything from short batch jobs, such as lightweight database queries to prepare and print a PDF document, to much longer running jobs, such as Inventory or Payroll reconciliation, which may run for hours and heavily exercise the database and, consequently, the disk storage system.

With the goal of demonstrating a real system workload, scripts have been developed to ensure that a significant level of load is generated on the system by running significant levels of both short and long batch tasks during testing of the interactive response times to the queries of the target 1200 interactive users of the system. In order to more realistically gauge overall system performance, a mixture of long- and short-running UBE workloads are simultaneously launched throughout the testing of the DIL workload by some custom scripts running on the JD Edwards EnterpriseOne server itself. During testing, these may launch thousands of short UBEs/hr and a configurable number of long-running UBEs throughout the period of the benchmark run.

Demonstrating the batch UBE workload allows Oracle to ascertain the achievable level of interactive users while simultaneously maintaining a more realistic level of batch workload, giving a more realistic metric for overall system capability.

Further details about the DIL and UBE workloads are provided in Appendix D.

Test Environment

The test environment was implemented as described previously in the “Solution Architecture Overview” section of this document along with the following two additional servers that were required for the benchmark configuration:

- **Deployment server**—A small configuration x86 server running the JD Edwards EnterpriseOne Deployment Server for controlling software updates to the JD Edwards EnterpriseOne Server. This server does not take an active part in the testing of this configuration, but is required for the initial deployment and application code maintenance of the JD Edwards EnterpriseOne environment.
- **Load generation server**—An x86 server acting as a load generation device, load testing tool, and performance evaluation server load.

In-depth descriptions of solution component configurations are provided in the subsections that follow.

SPARC T4-1 Server Configuration

The system under test represented a mid-range SPARC T4-1 server configuration with an 8-disk internal storage subsystem. A larger storage capacity configuration may utilize 600 GB disk drives in place of the tested 300 GB drives.

Memory Configuration

The SPARC T4-1 server has 16 memory DIMM slots. For maximum performance, it is recommended that all slots be filled with identical density memory DIMMS. In this configuration, 4 GB DIMMs were chosen providing a total of 64 GB of RAM. To further increase RAM, it is recommended that 8 GB or 16 GB DIMMs be substituted for a total capacity of up to 256 GB of memory. If a small configuration system is chosen with only 32 GB of memory, it is recommended that only eight 4 GB DIMMs be installed allowing additional DIMMs to be installed at a later time, if needed.

Storage Configuration

The SPARC T4-1 server incorporates an 8-disk storage chassis. A doubling of available storage capacity can be achieved by adopting the larger 600 GB capacity disk drives available for this platform; however, for the DIL workload, the 300 GB disk drives provide more than enough capacity for these tests.

The SPARC T4-1 server is provided with two onboard disk controllers; however, in this configuration, these are not used and instead they are replaced by a single higher-performance advanced RAID HBA with Battery Back Write Cache. This provides a key benefit in that implementing a large write cache allows the OS to receive acknowledgement of a confirmed write-to-disk operation very rapidly. The card accepts writes to cache before the data is committed to disk. The data is written to disk at some point in the near future as time allows. Very rapid confirmation of writes significantly reduces write latency and, therefore, increases application response times and throughput considerably. In the event of a power failure before the acknowledged data is written to disk, the battery-backed cache will be maintained until power returns, and the data will be immediately committed to disk following the restoration of power to the server.

The use of the RAID HBA also allows more advanced RAID data protection configurations, in addition to providing performance benefits (particularly on data reads) due to data striping across multiple devices.

Storage Logical Volume Configurations

The RAID HBA is initialized and partitioned into three independent logical volumes to allow different RAID configurations for different uses and to separate the different storage areas to allow for more granular backup and restore processes.

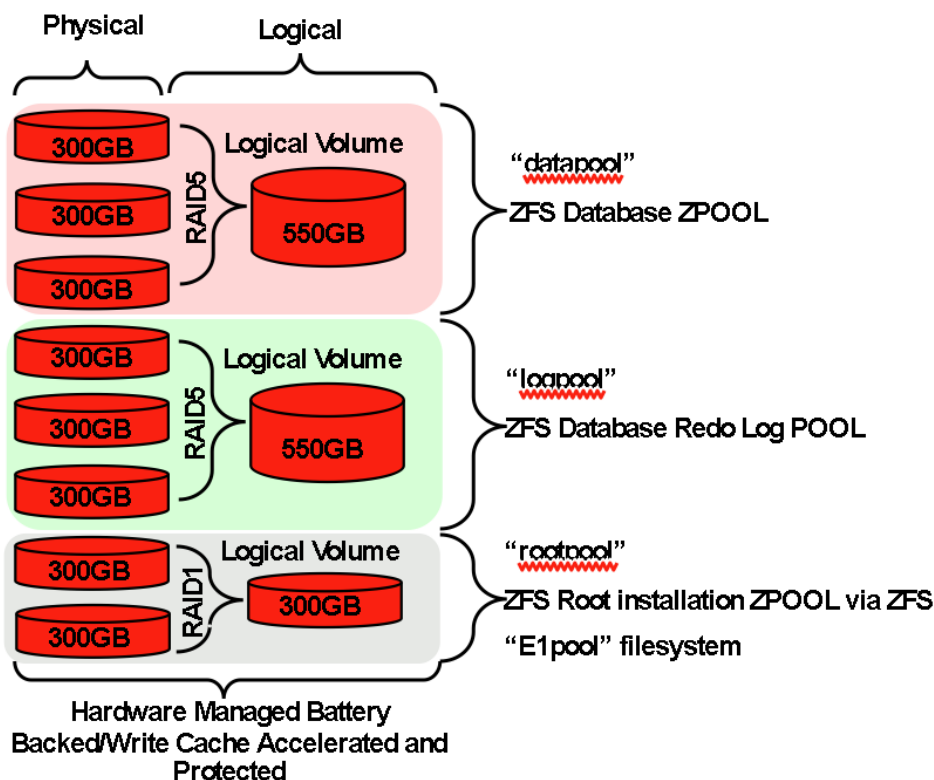


Figure 2. Logical and physical volume configurations for the storage infrastructure.

TABLE 4. ZFS STORAGE POOL CONFIGURATIONS

ZFS POOL NAME	HARDWARE CONFIGURATION	ZFS CONFIGURATION	DESCRIPTION
datapool	<ul style="list-style-type: none"> 3 x 300 GB SAS2 drives managed by the Sun Storage 6 Gb SAS PCIe RAID HBA. Configured as a single 550 GB RAID 5 protected logical volume. Write performance is significantly enhanced through the NVRAM write cache provided by the HBA. 	The logical volume is used to directly host the datapool ZFS pool. No additional data protection is required.	Holds all Oracle Database files and working data. Database performance benefits significantly from the rapid write acknowledgements provided by the NVRAM write cache.
logpool	<ul style="list-style-type: none"> 3 x 300 GB SAS2 drives managed by the Sun Storage 6 Gb SAS PCIe RAID HBA Configured as a single 550 GB RAID 5 protected logical volume. Write performance is significantly enhanced through the NVRAM write cache provided by the HBA. 	The logical volume is used to directly host the logpool ZFS pool. No additional data protection is required.	Holds all Oracle Database redo log files. Database performance benefits significantly from the separation of the log files onto a separate file system and also from the rapid write acknowledgements provided by the NVRAM write cache.
rootpool	<ul style="list-style-type: none"> 2 x 300 GB SAS2 drives managed by the Sun Storage 6 Gb SAS PCIe RAID HBA Configured as a single 300 GB RAID 1 protected logical volume. Write performance is significantly enhanced through the NVRAM write cache provided by the HBA. 	<p>The logical volume is used to directly host the rootpool ZFS pool. No additional data protection is required.</p> <p>The E1pool is also housed in a file system on this volume.</p>	<p>Holds the OS installation and virtual machine container disk images and OS-specific configuration files. Benefits significantly from the rapid write acknowledgements provided by the NVRAM write cache.</p> <p>The E1pool is used to store JD Edwards EnterpriseOne server data. This data is characterized as 'read-mostly.'</p>

The physical layout of the storage of the SPARC T4-1 server is illustrated in Figure 3. Note that the RAID HBA controller is not seen in this image. It is installed in the rear of the server and does not connect externally.




Figure 3. Physical layout of storage for an 8-disk SPARC T4-1 server.

As can be seen in the diagram, three groups of disks have been configured. Physical location is not critical since all disks are 'equally close' to the disk controller and no disk slots are 'favored' in terms of performance. For convention, planning, and management purposes, these groupings have been adopted. For planning purposes, it is recommended that disks in the same group be kept physically close to each other, though this is not required for performance reasons.

Test Results and Analysis

Multiple test cases were run to test several dimensions of server performance with varying load configurations of JD Edwards EnterpriseOne. Seven tests were completed and raw data was collected, as shown in Table 5.

TABLE 5. DATA COLLECTED FROM SEVEN TESTS							
DESCRIPTION	NO UBE'S	LOW SHORT UBE'S	MEDIUM SHORT UBE'S	HIGHEST SHORT UBES	8 LONG UBE'S ONLY	4 LONG UBE'S PLUS MEDIUM SHORT UBE'S	2 LONG UBE'S PLUS HIGH SHORT UBE'S
TEST RUN	1	2	3	4	5	6	7
TEST DURATION (H:MM:SS)	1:11:56	1:11:46	1:11:45	1:11:55	1:11:59	1:11:49	1:11:50
JD EDWARDS ENTERPRISEONE SERVER % CPU UTILIZATION	2	5	13	26	7	15	17
JD EDWARDS ENTERPRISEONE SERVER % MEMORY UTILIZATION	20	20	21	23	19	19	19
DATABASE SERVER % CPU UTILIZATION	0	1	2	5	2	4	5
DATABASE SERVER % MEMORY UTILIZATION	30	31	31	31	31	32	32
ORACLE WEBLOGIC SERVER % CPU UTILIZATION	2	3	5	6	3	4	4
ORACLE WEBLOGIC SERVER % MEMORY UTILIZATION	22	21	22	21	21	21	21
DATABASE SERVER DISK I/O UTILIZATION % AVG/MAX	5/78	6/23	9/86	20/100	100/100	100/100	81/100
DATABASE SERVER REDO LOG DISK I/O UTILIZATION % AVG/MAX	1/2	2/17	2/4	4/34	14/83	17/86	16/81
SHORT UBE'S COMPLETED	0	2730	7349	8522	0	1741	2024
LONG UBE'S RUNNING/COMPLETED	0	0	0	0	8	4	2
AVERAGE INTERACTIVE RESPONSE TIME (SECS)	0.270	0.312	0.391	0.519	0.503	0.583	0.588

Test Run 1: Best Achievable Response Time During DIL Testing

The purpose of this test was to determine the best possible response time that can be achieved with the SPARC T4-1 server under 1200-user load levels. This test is carried out with no UBE load running and demonstrates a 0.27 response time average for all interactive queries. Note that no system resources were exhausted during this test.

Test Runs 2 and 3: Extended Response Time Testing with Increased Short UBE Load

These test cases extend the testing of the response time to interactive queries during the addition of considerable short-running UBE job loads. This test case demonstrates that the server can achieve linear growth in short-running UBEs with corresponding small increases in overall interactive response times since the I/O subsystem of the server is not saturated during testing.

Testing demonstrated the ability to scale from 2730 short UBEs with 0.312 second response times up to 7349 short UBEs with a 0.391 second response time.

Test Run 4: Maximizing Short UBE Throughput with 0.5 Second Response Times

This test case finds the maximum capability of the system to cope with short UBE loads with acceptable interactive response times. During testing, the maximum short UBE throughput achieved was 8522 jobs completed during the test while maintaining a 0.519 second average interactive response time. This result demonstrates a maximum achievable result because the disk I/O subsystem is reaching peaks of 100% utilization and the response becomes non-linear near this data point.

Test Run 5: Maximum Long UBE Throughput with Acceptable Response Times

The purpose of this test was to test the platform's capability to process a maximum number of long UBE loads during testing while maintaining an acceptable 0.503 second interactive response time. Eight long-running UBE jobs were maintained during the test period. Some jobs were completed and restarted during the test, and others ran throughout the duration of the testing. During this time, no short UBE loads were tested.

It can be seen that the long UBE loads heavily taxed the database and, consequently, the disk subsystem of the server, maintaining peak and average disk utilization of 100% throughout the test duration. This demonstrates that the maximum UBE load that can realistically be maintained is no higher than eight jobs concurrently.

Test Runs 6 and 7: Best Compromise Performance for Long and Short UBEs with Acceptable Response Times

In real-world environments, it is expected that users will generate a mix of both long and short UBE loads during operation and the system will need to retain acceptable interactive user response times. Tests 6 and 7 demonstrate overall system capabilities when four or two long UBEs are maintained along with a maximal level of short UBEs while maintaining acceptable system response times close to 0.5 seconds.

With four long UBEs and a throughput of 1741 short UBEs, response time was maintained at 0.583 seconds and the system was 100% utilizing disk I/O capabilities throughout the test run.

During test run 7, two long UBEs were maintained with a throughput of 2024 short UBEs and a response time of 0.588 seconds. The system has very little spare disk I/O capacity during this test, averaging 81% and peaking at 100%.

Test 7 demonstrates the maximum DIL throughput of the system under the expected conditions of mixed long and short UBE loads.

Final Analysis

The SPARC T4-1 server is extremely capable of handling the workload at these levels, the single limiting factor being disk I/O capacity on the integrated disk sub-system. To achieve higher results requires additional hardware in the form of additional higher-performance disk systems or dedicated off-load of the database server and database storage to a standalone platform.

At no time during testing were the CPU or memory of the server more than 27% (CPU) or 74% (memory) utilized, which indicates that if the disk I/O capacity limitations were improved, as discussed earlier, overall significant performance increases are easily achievable with the SPARC T4-1 server.

Real-World Sizing Based on the DIL Test Workload

To account for the fact that the DIL workload cannot accurately predict real-world customer use cases and performance requirements, conservative hardware sizing guidelines are outlined in this section.

The hardware sizing guides target substantially lower real-world user counts to avoid any potential for performance issues once the solution is deployed and to allow reasonable accommodation for growth and peak demands that cannot be accurately predicted.

Real-world sizing, therefore, covers 300 user configurations based on testing carried out at the 1200-users level under the DIL workload. Such conservative actual deployment sizing is common in JD Edwards EnterpriseOne deployments with sizing guides based on testing using the DIL kit.

Room for Consolidation

As indicated in the previous discussion regarding real-world sizing, considerable de-rating has been built into the proposed real-world deployment sizes based on the performance seen under DIL testing workloads.

One of the primary benefits that may be realized as a result of this conservative sizing is that substantial headroom may be available for further unrelated workload consolidation or for integration of connected applications that can augment the capabilities of JD Edwards EnterpriseOne.

Such additional consolidation, if undertaken, should consider the full range of JD Edwards EnterpriseOne workload conditions that are possible in the actual deployment to ensure that acceptable performance under worst-case conditions can be maintained before additional consolidation is performed. If such worst-case situations can be satisfied, then considerable further value-add and ROI benefits are readily available to customers choosing to deploy and consolidate on this platform.

Sizing Guidelines

For sizing guidelines, hardware configurations are provided for the workloads previously described as a starting point for average (medium) and large deployments. For more detailed sizing information, contacting a local Oracle representative is encouraged and recommended.

The configurations outlined below refer to the systems and configurations available at the time of the test. It is recommended that the most recent options and system configurations available at the time of implementation be deployed in the customer environment¹.

Average Configuration—Up to 300 Real-World Users

Table 6 provides the recommended hardware components for the medium configuration.

TABLE 6. RECOMMENDED HARDWARE COMPONENTS FOR THE MEDIUM CONFIGURATION

SERVER VARIANT	COMPONENT	CONFIGURATION
SPARC T4-1 Server	Processor: SPARC T4, 2.85 GHz, 8-core, 8 threads/core	Factory pre-installed.
	Memory: 64 GB RAM, utilizing 16 x 4 GB DIMMs	As indicated in service manual.
	Disks: 8 x 300 GB SAS2 2.5-inch disk drives	Install disks in all available slots.
	HBA: Sun Storage 6Gb SAS PCIe RAID HBA, internal, 8-port with 512 MB memory	As indicated in service manual and installation guide. PCI slot location and cable routing options must be followed. All eight ports must be cabled for maximum performance.

¹ LEGAL DISCLAIMER - Expected performance characteristics are based on laboratory test implementations and are provided as is without warranty of any kind. The entire risk of the use of information provided herein remains with the recipient and in no event shall Oracle be liable for any direct, consequential, incidental, special, punitive, or other damages including, without limitation, damages for loss of business profits, business interruption, or loss of business information.

For the purposes of the DIL testing at the 1200-user level, an 8-disk chassis fully populated with 300 GB disks was found to provide sufficient capacity. However, for actual customer deployments, larger-capacity disk storage may be required. It should be noted that is not possible to re-allocate disks to vary storage in the various logical volumes provided by the RAID HBA since all eight disks are required in the indicated configurations to provide the minimum required number of disks present in each of the RAID volume types. (RAID 1 requires a minimum of two disks, and each RAID 5 volume requires a minimum of three disks.)

The server configuration suggests 64 GB RAM be installed as sixteen 4 GB DIMMs. This configuration fully populates the DIMM slots in the server with memory and provides maximum performance at this memory density. If room for memory expansion is required, it is suggested that the server initially be populated with only eight 8 GB DIMMs, half populating the system, leaving easy incremental space for an additional eight DIMMS when required. This will impose a small performance penalty due to the non-fully populated memory system not achieving its full bandwidth, but this is not likely to be an issue since memory is not a limiting factor in system performance.

For storage management, the Sun Storage 6GB SAS RAID PCIe RAID HBA is selected and should be installed according to standard service or installation documentation.

Larger Configurations—Above 300 Real-World Users

Larger configurations of the system were not directly part of the testing for the Oracle Optimized Solution for JD Edwards EnterpriseOne. However, the SPARC T4-2 server has been tested in higher-capability configurations and has demonstrated capacities as high as 10,000 interactive users.

Further testing and performance enhancements described here will be carried out in a subsequent version of this solution that mirrors this configuration using SPARC T4-1 servers to provide a stepping stone configuration that achieves substantially higher performance than a standalone SPARC T4-1 server, providing a range of architectures to choose from depending on the requirements of the particular deployment. Though not tested as part of this solution, it is discussed here for informational purposes.

Table 7 shows the recommended hardware components for the large configuration.

TABLE 7. RECOMMENDED HARDWARE COMPONENTS FOR THE LARGE CONFIGURATION

SERVER VARIANT	COMPONENT	CONFIGURATION
SPARC T4-1 Server	Processor: 1 x SPARC T4, 2.85 GHz, 8-core, 8 threads/core	Factory pre-installed.
	Memory: 128 GB RAM, utilizing 16 x 8 GB DIMMs	As indicated in service manual.
	Disks: 2 x 300 GB SAS2 2.5-inch disk drives	Install disks in slots 0 and 1 and mirror using the onboard SAS HBA.
SPARC T4-1 Server	Processor: 1 x SPARC T4, 2.85 GHz , 8-core, 8 threads/core	Factory pre-installed.
	Memory: 64 GB, utilizing 16 x 4 GB DIMMs	As indicated in service manual.
	Disks: 2 x 146 GB SAS2 2.5-inch disk drives	Install in slots 0 and 1 and mirror for Root disk protection.
	HBA: Emulex Dual-Port Fibre Channel FC_AL HBA	As indicated in service manual.
Oracle's Sun Storage 2540-M2 Array	Dual RAID Controllers, 12 x 300 GB 3.5-inch FC Disks	As indicated in service manual.

The large configuration is provided to demonstrate that the SPARC T4-1 server remains central to the solution and provides the capability to scale the JD Edwards EnterpriseOne capacity up to an anticipated 4,000 to 5,000 users by offloading the database portion of the workload to an additional dedicated SPARC T4-1 server with a dedicated Fibre Channel storage array. The choice of putting the database on an additional server is indicated by the maximum I/O and memory utilization data captured in the medium configuration when running with UBE workloads during testing.

If a customer environment were to be challenged by the ultimate performance capabilities of a medium configuration, the upgrade path is incremental and does not require a complete replacement of already acquired assets, improving ROI for this solution.

Conclusion

Testing shows that a single SPARC T4-1 server from Oracle can provide extremely competitive performance for small and medium JD Edwards EnterpriseOne deployments. During testing, performance optimization was achieved through careful arrangement of the I/O subsystem, which involved separating and optimizing the placement and use of the various data resources on the server's local storage system. By optimizing the use of the local storage, it is possible to provide high performance while consuming a very small footprint. This enables a very cost-effective solution, maximizing customer ROI as well as minimizing TCO.

The solution leverages conventional storage in the most beneficial modes of operation and avoids the need to use more expensive flash-based or external storage devices, such as Fibre Channel arrays.

References

TABLE 8. REFERENCES

DESCRIPTION	URL
JD Edwards EnterpriseOne Product page	http://www.oracle.com/us/products/applications/jd-edwards-enterpriseone/index.html
SPARC T4-1 server page	http://www.oracle.com/us/products/servers-storage/servers/sparc-enterprise/t-series/sparc-t4-1-ds-487858.pdf
Oracle's flash technology page	http://www.oracle.com/us/products/servers-storage/storage/flash-storage/index.html
Oracle JD Edwards Partner Portal	http://www.go9withjde.com/
Oracle JD Edwards EnterpriseOne Documentation	http://www.oracle.com/technetwork/documentation/jdedent-098169.html

Appendices

A. Configuring the Sun Storage 6 Gb SAS PCIe RAID HBA Card

The initial configuration of the Sun Storage 6 Gb SAS PCIe RAID HBA card requires that the server be initially configured at the OBP prompt, before an operating system is booted (or installed).

From the OBP ok prompt...

```
{0} ok show-devs
```

check the output for a line that looks like

```
/pci@400/pci@2/pci@0/pci@c/Megaraid,sas@0. This is the root path to the controller.
```

Select that device; note that the initial space at the front of the quoted text is important.

```
{0} ok " /pci@400/pci@2/pci@0/pci@c/Megaraid,sas@0" select-dev
```

Now get a report on the status of the card; there is a large amount of output to collect. Ensure you have a large history or scroll-back buffer on the terminal you are using. You will refer back to this information later.

```
{0} ok cli -AdpAllInfo
```

Check to see that all the physically installed drives are visible to the controller and reporting correctly. All drives should report as 6 Gbps SAS2 devices and 300 GB drives will report with a real-world capacity of 279 GB.

```
{0} ok cli -PdList
```

Check to see if any logical volumes have already been configured on the server. Check the manual for the required procedures needed to destroy any pre-existing logical volumes.

```
{0} ok cli -LdPdInfo
```

You can now create the required set of logical volumes. There will be three volumes created.

1. RAID1 mirrored Root disk volume
2. RAID5 mirrored Data volume for the database
3. RAID5 mirrored Log volume for the Database Server redo log files

To create a logical volume, the syntax is similar to this example:

```
{0} ok cli -CfgLdAdd -r1[252:0,252:1]
```

`r1` means we are creating a RAID1 volume. `252` is the enclosure number, which can be determined from the output of the earlier `PdList` command, and `0` and `1` are the disk slots in the enclosure that are to be used to create a mirrored pair of drives. This command RAID1-mirrors drives `0` and `1` into a logical drive that will then be visible to the operating system that will be loaded later.

To create the required three logical volumes as described earlier, on an 8-disk server using all 8 disks, the commands required would be:

Make a RAID1 volume out of drives `0` and `1`:

```
{0} ok cli -CfgLdAdd -r1[252:0,252:1]
```

Make a RAID5 volume out of drives `2,3,4`:

```
{0} ok cli -CfgLdAdd -r5[252:2,252:3,252:4]
```

Make a RAID5 volume out of drives `5,6,7`:

```
{0} ok cli -CfgLdAdd -r5[252:5,252:6,252:7]
```

Note that it is important to create the logical volumes in this order since when Oracle Solaris is installed, it will, by default, install on the first disk it sees (unless configured otherwise), and that should be the RAID1-mirrored root volume.

List the details and status of the logical drives by running this command:

```
{0} ok cli -LdPdInfo
```

Next create an alias to the bootable drive that was made. This is not strictly necessary but makes future administration easier by reducing the amount of typing and possibilities for mistakes that accompany long path names.

First look to see what drives OBP currently sees. This may take a few seconds on systems with lots of drives and especially if Fibre Channel drives have been connected.

```
{0} probe-scsi-all
```

Drives that are under the control of the RAID card will be listed under a path similar to this:

```
/pci@400/pci@2/pci@0/pci@c/Megaraid,sas@0
Target 0
  Unit 0  Disk      LSIDellMR9261-8i2.0.    583983104 Blocks, 298 GB
Target 1
  Unit 0  Disk      LSIDellMR9261-8i2.0.    1167966208 Blocks, 597 GB
Target 2
  Unit 0  Disk      LSIDellMR9261-8i2.0.    1167966208 Blocks, 597 GB
```

In this case, only three drives are visible since there are now only three logical volumes visible from the RAID card.

The OBP disk device paths that correspond to each of the individual logical volumes are therefore easy to derive and in this case would be as follows:

```
/pci@400/pci@2/pci@0/pci@c/Megaraid,sas@0/disk  
/pci@400/pci@2/pci@0/pci@c/Megaraid,sas@1/disk  
/pci@400/pci@2/pci@0/pci@c/Megaraid,sas@2/disk
```

Target 0 will be the boot device. Create a device alias of `disk` that can be more easily used as the default boot device:

```
{0} ok nvalias disk /pci@400/pci@2/pci@0/pci@c/Megaraid,sas@0/disk
```

Next set the server to automatically boot from the newly created device:

```
{0} ok setenv boot-device disk
```

Note that when/if setting up an Oracle Solaris JumpStart client for this machine, the logical volumes appear as (in this case) `c2t0d0` (and `c2t2d0`, `c2t3d0`).

Once completed, Oracle Solaris can be installed with and it should see only three disks, as have been configured above. It is strongly encouraged that ZFS be used as the file system on all three disks, including the root disk, because future management of the server will be greatly simplified through the extensive feature set offered by ZFS.

B. DIL Benchmark Tools Configuration

The JD Edwards EnterpriseOne DIL workload is flexible and configurable. Optimizing the use of the workload is the key to driving the solution to obtain the best possible results that most closely provide a parallel to real-world workloads. In order to optimize the workload, specific configurations of load generator and test scripts were prepared. The result of this tuning exercise is documented in the next paragraphs.

The interactive portions of the workload are generated during testing through the load generation server. To allow the 1200-user load level to be achieved, changes were made to the default configuration supplied as part of the DIL package to increase the number of users who are able to execute each task.

It was important to test all aspects of the SPARC T4-1 server configuration, so considerable effort was put into generating appropriate UBE batch workloads during the testing of the interactive response of the system. Several variations of batch workloads were prepared and tested.

Short-Running UBE Workloads

The short-running UBE workloads, as discussed previously, were also tested. These workloads are created and executed to simulate the actions of users initiating non-interactive tasks on the system, such as initiating the preparation and printing of a PDF document such as an invoice. These typically are initiated and run for 15–30 seconds to run to completion. These jobs are used in these tests to generate consistent ‘background’ workloads on the database and JD Edwards EnterpriseOne server to ensure that normal daily processing tasks do not have an adverse affect on users’ interactive response times.

In order to simulate ongoing creation and execution of these tasks, a script was created in the `E1Server` zone that resubmits batches of short UBE tasks every 15–20 seconds to ensure a consistent level of short UBE load.

Long-Running UBE Workloads

The long-running UBE jobs are started once at the beginning of a test cycle and allowed to run throughout the testing. Some batch jobs complete within the approximately hour-long processing of the DIL workload, so they are restarted, if necessary, to maintain the target number of long UBE jobs running during particular tests. To provide further information on the behavior of the system under long-running UBE batch workloads, all the DIL long-running UBE tasks were started and allowed to run to completion to determine runtimes with no other load (interactive or short UBE) on the system. The details of the operations carried out by these tasks are also provided to allow users to determine applicability of these results to their own scripts and UBE tasks.

TABLE 9. UBE WORKLOAD COMPONENTS

UBE NAME	T4-1 PLATFORM RUNTIME WITHOUT INTERACTIVE LOAD (MINUTES)	T3-1 PLATFORM RUNTIME WITHOUT INTERACTIVE LOAD (MINUTES)	DESCRIPTION
R42565	36.2	91.71	Sales Order Invoicing Batch Job
			The R42565 batch process was configured to initiate document invoice orders located in the F4211 table (Sales Order Detail File).
			There were 150 documents processed in the testing and it included 47,436 records in the F42119 table (Sales Order History File).

R42800	20.8	0.52	Sales Order Invoicing
<p>The R42800 batch process was configured to initiate selected document numbers located in the F4211 table (Sales Order Detail File) with a next status code of 620 and within a single business unit.</p> <p>There were 12,900 records processed with this data selection criterion.</p>			
R43500	62.36	155.23	Purchase Order Print
<p>The R43500 batch process was configured to initiate document invoice numbers located in the F4311 table (Purchase Order Detail File) with a next status code of 280 and within a single business unit.</p> <p>There were 255,481 records processed with this data selection criterion.</p>			
R4891	25.44	45.81	Freight Update
<p>The R4891 batch process was configured to initiate shipment numbers located in the F4215 table (Shipment Header) with a next status of 70 and within a single business unit.</p> <p>There were 800 records processed with this data selection criterion.</p>			
R09801	29.7	78.16	General Ledger Post
<p>The R09801 batch process was configured to initiate batch numbers located in the F0011 table (Batch Control Records) with a batch status of "A" and batch type of "G."</p> <p>There were 990,099 records processed with this data selection criterion.</p>			
R31410	9.94	17.13	Work Order Processing
<p>The R31410 batch process was configured to initiate document invoice numbers in the F4801 (Work Order Master File).</p> <p>There were 28,751 records processed with this data selection criterion.</p>			
R31802A	29.07	12.06	Manufacturing Account Journal
<p>The R31802A batch process was configured to initiate document invoice numbers in the F4801 (Work Order Master File) with a status code of "95."</p> <p>There were 1,501 records processed with this data selection criterion.</p>			
R42520	16.61	49.11	Print pick Slips Batch
<p>The R42520 batch process was configured to initiate shipment numbers in tables F4201 (Sales Order Header File), F4211 (Sales Order Detail File), and F49T211 (Sales Order Detail File - Secondary Tag Table) and a next status code of "520."</p> <p>There were 37,154 records processed with this data selection criterion.</p>			

Note: The status codes are defined by the activity rules of the JD Edwards EnterpriseOne environment. Default status codes were used whenever possible.

Note: Some database layout changes were implemented between testing on SPARC T3-1 and SPARC T4-1 servers. This is expected to be the reason for the increased processing times for jobs R42800 and R31802A. Even so, the average performance gain seen when using the SPARC T4-1 server is approximately 2x compared to the previous-generation server.

C. Tuning and Best Practices

Setting Up the Three Application Zones

Zones are an effective and efficient means to provide virtual OS image isolation between application stacks, and are a standard feature of the Oracle Solaris OS. For the Oracle Optimized Solution for JD Edwards EnterpriseOne, three zones are configured based on the global zone (the initial base OS image installed on the SPARC T4-1 server). Instructions for the creation and configuration of each zone are provided in this section of the document.

Global Zone Configuration

Prior to configuring the OS on the global zone, the low-level disk configuration must be completed in the OBP environment. The procedures for this are described in detail in Appendix A. Once the disk hardware configuration is completed, a conventional installation of Oracle Solaris can be carried out on the RAID 1-mirrored logical volume and the remaining portions of the tuning and best practices carried out.

The global zone represents the original base OS installation on the platform, and it must be configured appropriately to host the three application zones that will be used: the Oracle Database zone, the JD Edwards EnterpriseOne zone, and the Oracle WebLogic Web server zone.

The ZFS file system was used throughout this configuration because it provides considerable ease of configuration and portability of zone images once configured.

It is also recommended that the following setting is placed in the `/etc/system` file in the global zone, as a security feature.

```
set noexec_user_stack=1
```

To simplify network communications between the zones, host names (on private subnets) should be populated into the global zone `/etc/hosts`. Aliases to the name of the zone alongside the actual host name adopted by the zone will make management easier.

```
10.6.222.241  jds-zone1      E1Server
10.6.222.84   jds-zone2      WebServer
10.6.222.79   jds02          DBServer
```

Mount additional storage devices for use in the zones as data storage.

From the output of the `format` command, the three logical volumes can be determined, as required, to allow proper configuration of the required ZFS pools and subsequent ZFS file systems that will be used by the zones.

```
-bash-3.2# format
Searching for disks...done

AVAILABLE DISK SELECTIONS:
  0. c1t0d0 <LSI-MR9261-8i-2.50 cyl 65533 alt 2 hd 287 sec 31>
     /pci@400/pci@2/pci@0/pci@c/LSI,mrsas@0/sd@0,0
  1. c1t1d0 <LSI-MR9261-8i-2.12-556.93GB>
     /pci@400/pci@2/pci@0/pci@c/LSI,mrsas@0/sd@1,0
  2. c1t2d0 <LSI-MR9261-8i-2.12-556.93GB>
     /pci@400/pci@2/pci@0/pci@c/LSI,mrsas@0/sd@2,0
Specify disk (enter its number): ^C
-bash-3.00#
```

In this listing, the three devices containing the LSI-MR9261 naming are the RAID configured Logical Volumes provided by the SAS RAID card.

The SPARC T4-1 test system was installed with Oracle Solaris 10 8/11 (pre-release build) and configured with 64 GB of swap .

The OS should be installed on `c1t0d0`, which is a RAID1-mirrored pair of disks. The output of the `zpool status rootpool` command can be used to confirm this. The rootpool should be hosted on the `c1t0d0s0` device.

```
-bash-3.00# zpool status rootpool
pool: rootpool
state: ONLINE
scrub: none requested
config:

   NAME        STATE      READ WRITE CKSUM
   rootpool    ONLINE    0     0     0
     c1t0d0s0  ONLINE    0     0     0

errors: No known data errors
-bash-3.00#
```

Now create the additional zpools to hold the data for the various zones that will be created.

```
zpool create datapool c1t1d0
zpool create logpool c1t2d0
```

The Global Zone has been installed on ZFS using `rootpool`.

Create a file system on the `rootpool` to house the JD Edwards EnterpriseOne server data.

```
zfs create rootpool/Elserverdata
chmod 700 /rootpool/Elserverdata
```

Create file systems to host the zone root images that will be used for each zone.

```
zfs create rootpool/DBserverroot
chmod 700 /rootpool/DBserverroot
zfs create rootpool/Elserverroot
chmod 700 /rootpool/Elserverroot
zfs create rootpool/Webserverroot
chmod 700 /rootpool/Webserverroot
```

Each zone now has a defined `serverroot` directory and also a data storage pool constructed on top of appropriate types of storage with suitable performance characteristics based on the type of data and access patterns that are most likely.

In Oracle Solaris zones, it is possible to assign CPU resources to each zone, but this was not found to be necessary in this configuration since during testing, CPU resources were not a limiting factor to performance.

To take advantage of the performance enhancements available in JDK1.6, the system-level link `/usr/java` should be checked to ensure that it points to `jdk1.6.0` (by default, `/usr/java` points to `jdk1.5.0` for compatibility with some older applications).

Creating and Configuring the JD Edwards EnterpriseOne Server Zone

Create a standard (default, sparse) zone based on the global zone. In this case, two network devices are used, one for external access (on the 10.6.162.x subnet) and another for internal private network access between zones (on the 10.6.222.x subnet).

```
zonecfg -z ElServer
zonecfg:ElServer> create
zonecfg:ElServer> set zonepath=/rootpool/Elserverroot/ElServer
zonecfg:ElServer> add net
zonecfg:ElServer:net> set physical=igb0
zonecfg:ElServer:net> set address=10.6.162.23/22
zonecfg:ElServer:net> set defrouter=10.6.163.254
zonecfg:ElServer:net> end
zonecfg:ElServer> add net
zonecfg:ElServer:net> set physical=igb1
zonecfg:ElServer:net> set address=10.6.222.241/24
zonecfg:ElServer:net> set defrouter=10.6.222.254
zonecfg:ElServer:net> end
zonecfg:ElServer> add fs
zonecfg:ElServer:fs> set dir=/db
zonecfg:ElServer:fs> set special=rootpool/Elserverdata
```

```
zonecfg:ElServer:fs> set type=zfs
zonecfg:ElServer:fs> end
zonecfg:ElServer> verify
zonecfg:ElServer> commit
zonecfg:ElServer> info
zonecfg:ElServer> exit
```

If Oracle Solaris DTrace will later be used to investigate the behavior of the zone, this additional setting should be issued.

```
zonecfg -z ElServer
zonecfg:ElServer> set limitpriv=default,dtrace_proc,dtrace_user
zonecfg:ElServer> verify
zonecfg:ElServer> commit
zonecfg:ElServer> exit
```

Now the zone is ready to be booted and further configurations can be carried out inside the zone OS image.

```
zoneadm -z ElServer install
zoneadm list -cv
zoneadm -z ElServer boot
```

Log in to the server and carry out the following operations to create the required user profiles.

```
zlogin -C ElServer
usermod -s /usr/bin/bash root
```

If you root should be allowed to directly log in to the zone, make the following changes:

Edit the file `/etc/ssh/sshd_config`, and change the line:

```
PermitRootLogin no
```

To read:

```
PermitRootLogin yes
```

And then restart the ssh daemon with the command:

```
svcadm restart ssh
```

Create the users required for the JD Edwards EnterpriseOne installation, and ensure that the group and user IDs are the same as those used in the Oracle Database zone.

```
groupadd -g 98194049 jde900
groupadd -g 98194050 oinstall
groupadd -g 98194051 dba
useradd -u 237422 -g jde900 -d /JDE -m -s /bin/ksh jde900
passwd jde900 jde900
useradd -u 237445 -g oinstall -G dba -d /export/oracle -m -s \ /usr/bin/bash
oracle
passwd oracle oracle
```

Note these example settings configure insecure passwords; more secure passwords should be provided.

Log in to the `oracle` user account and add the following paths to the `.profile` file found in the `oracle` user's home directory.

```
ORACLE_BASE=/db/app/oracle; export ORACLE_BASE
ORACLE_HOME=$ORACLE_BASE/product/11.1.0/db_1; export ORACLE_HOME
ORACLE_SID=orcl; export ORACLE_SID
JDE_HOME=/db/jdedwards/JDE_HOME/bin; export JDE_HOME
PATH=$PATH:/opt/SUNWspro/bin:$ORACLE_HOME/bin:/db/usr/local/bin:/usr/ucb:.;
export PATH
# for Verity software
VDK612_HOME=/db/jdedwards/e900/verity612/solaris; export VDK612_HOME
VDK612_PATH=$VDK612_HOME/_ssol26/bin; export VDK612_PATH
PATH=$VDK612_PATH:$PATH
LD_LIBRARY_PATH=$VDK612_PATH:$LD_LIBRARY_PATH; export LD_LIBRARY_PATH
```

As `root` in the `E1Server` domain, change the following project settings to increase the maximum number of shared memory segments. The default is 100.

```
jds-zone1> projmod -s -K 'project.max-shm-ids=(privileged,1024,deny)' \ default
```

Increase the maximum size of a shared memory segment. Default = $\frac{1}{4}$ of physical memory.

```
jds-zone1> projmod -s -K 'project.max-shm-memory=(privileged,12gb,deny)' \
default
```

Increase the maximum number of semaphores allowed per semaphore set. Default = 512.

```
jds-zone1> projmod -s -K 'process.max-sem-nsems=(privileged,4096,deny)' \
default
```

Increase the maximum number of message queue identifiers. Default = 128.

```
jds-zone1> projmod -s -K 'project.max-msg-ids=(privileged,2048,deny)' default
```

Increase the maximum number of semaphores identifiers. Default = 128.

```
jds-zone1> projmod -s -K 'project.max-sem-ids=(privileged,15320,deny)' \ default
```

Increase the number of bytes for a message on the message queue. Default = 64k.

```
jds-zone1> projmod -s -K 'process.max-msg-qbytes=(privileged,16834,deny)' \
default
```

Increase the maximum number of message on a message queue. Default = 8192.

```
jds-zone1> projmod -s -K 'process.max-msg-messages=(privileged,16384,deny)' \
default
```

For security, add the following parameter to the `/etc/system` file.

```
set noexec_user_stack=1
```

And finally reboot the zone to make these changes occur.

```
reboot
```

Once the domain is rebooted, the JD Edwards EnterpriseOne software can be installed.

Creating and Configuring the WebLogic Web Server Zone

Create the sparse zone based on the global zone OS. Do this in the global zone.

```
zonecfg -z WebServer
zonecfg:WebServer> create -b
zonecfg:WebServer> set zonepath=/rootpool/Webserverroot/WebServer
zonecfg:WebServer> add net
zonecfg:WebServer:net> set physical=igb0
zonecfg:WebServer:net> set address=10.6.162.24/22
zonecfg:WebServer:net> set defrouter=10.6.163.254
zonecfg:WebServer:net> end
zonecfg:WebServer> add net
zonecfg:WebServer:net> set physical=igb1
zonecfg:WebServer:net> set address=10.6.222.84/24
zonecfg:WebServer:net> set defrouter=10.6.222.254
zonecfg:WebServer:net> end
zonecfg:WebServer> verify
zonecfg:WebServer> commit
zonecfg:WebServer> info
zonecfg:WebServer> exit
```


If DTrace is required to investigate the performance of this zone, this additional setting should also be issued.

```
zonecfg -z WebServer
zonecfg:WebServer> set limitpriv=default,dtrace_proc,dtrace_user
zonecfg:WebServer> verify
zonecfg:WebServer> commit
zonecfg:WebServer> exit
```

The zone can now be installed and booted.

```
zoneadm -z WebServer install
zoneadm -z WebServer boot
```

Log in to the server to create and modify its OS-specific settings.

```
zlogin WebServer
usermod -s /usr/bin/bash root
```

If root should be allowed to directly log in to the zone, make the following changes:

Edit the file `/etc/ssh/sshd_config`, and change the line:

```
PermitRootLogin no
```

To read:

```
PermitRootLogin yes
```

And then restart the `ssh` daemon with the command:

```
svcadm restart ssh
```

Add the required users and groups to the OS image:

```
groupadd -g 98194050 oinstall
groupadd -g 98194051 dba
useradd -u 237445 -g oinstall -G dba -d /export/home/oracle -m \
-s /usr/bin/bash oracle
passwd oracle oracle
```

Log in to the `oracle` user's home directory using `su - oracle` and add the following paths to the `.profile` file in the home directory.

```
export PATH=/usr/bin:/usr/sbin:/usr/ucb:/etc:/usr/openwin/bin:.
export WL_HOME=/export/home/weblogic/wlserver_10.3
export MSLOGS=/export/home/weblogic/user_projects/domains/wl1_base_domain
export NODEMGR_HOME=${WL_HOME}/common/nodemanager
export DOMAIN_HOME=/export/home/weblogic/user_projects\
/domains/wl1_base_domain
```

For security, add the following parameter to the `/etc/system` file.

```
set noexec_user_stack=1
```

Next reboot the domain to let these settings take effect.

```
reboot
```

And finally install the Oracle WebLogic server software according to the paths and user names that have been established above.

Once the Oracle WebLogic software is installed, some further tuning should be carried out.

Five JVMs were used per instance on the Oracle WebLogic server during the 1200-user test. Each JVM used a minimum and maximum heap size of 4,096 MB. `NewSize` and `MaxNewSize` were set to 1344m. These values were modified through the Oracle Enterprise Manager Console, which can be accessed via the following URL:

```
http://<WebLogicServerHostname_or_IP>:7001/console/login
```

Further modifications are made to the server configuration through the Server Manager GUI. The configuration changes will be reflected in the files `jas.ini` and `jdbj.ini`.

JAS.INI changes:

```
PathCodes=('PD900')
# MAXUser default is 150
MAXUser=1000

UserSession=36000000

[PerfMON]
transactionTimeout=60000

# Time out value for requests sent to the Enterprise Server
# default is 90000
enterpriseServerTimeout=180000

# maxPoolSize= default value is 100
maxPoolSize=35
```

JDBJ.INI changes:

```
# maxConnection default is 150
maxConnection=1000
connectionTimeout=1800000

resultSetTimeout=60000
```

Creating and Configuring the Oracle Database Server Zone

Create a standard (default, sparse) zone based on the global. In this case, two network devices are used, one for external access (on the 10.6.162.x subnet) and another for internal private network access between zones (on the 10.6.222.x subnet).

```

zonecfg -z DBServer
zonecfg:DBServer> create
zonecfg:DBServer> set zonepath=/rootpool/DBserverroot/DBServer
zonecfg:DBServer> add net
zonecfg:DBServer:net> set physical=igb0
zonecfg:DBServer:net> set address=10.6.162.22/22
zonecfg:DBServer:net> set defrouter=10.6.163.254
zonecfg:DBServer:net> end
zonecfg:DBServer> add net
zonecfg:DBServer:net> set physical=igb1
zonecfg:DBServer:net> set address=10.6.222.79/24
zonecfg:DBServer:net> set defrouter=10.6.222.254
zonecfg:DBServer:net> end
zonecfg:DBServer> add fs
zonecfg:DBServer:fs> set dir=/db
zonecfg:DBServer:fs> set special=datapool/DBserverdata
zonecfg:DBServer:fs> set type=zfs
zonecfg:DBServer:fs> end
zonecfg:DBServer> verify
zonecfg:DBServer> commit
zonecfg:DBServer> info
zonecfg:DBServer> exit

```

If DTrace will later be used to investigate the behavior of the zone, this additional setting should be issued.

```

zonecfg -z DBServer
zonecfg:DBServer> set limitpriv=default,dtrace_proc,dtrace_user
zonecfg:DBServer> verify
zonecfg:DBServer> commit
zonecfg:DBServer> exit

```

Now the zone is ready to be booted and further configurations can be carried out inside the OS image.

```

zoneadm -z DBServer install
zoneadm list -cv
zoneadm -z DBServer boot

```

Log in to the server and carry out the following operations to create the required user profiles.

```

zlogin -C DBServer
usermod -s /usr/bin/bash root

```

If root should be allowed to directly log in to the zone, this can be enabled with the following changes:

Edit the file `/etc/ssh/sshd_config`, and change the line:

```
PermitRootLogin no
```

To read:

```
PermitRootLogin yes
```

And then restart the ssh daemon with the command:

```
svcadm restart ssh
```

Create the users required for the JD Edwards EnterpriseOne installation. Ensure that the group and user IDs are the same as those used in the other zones.

```
groupadd -g 98194050 oinstall
groupadd -g 98194051 dba
useradd -u 237445 -g oinstall -G dba -d /export/home/oracle -m -s \
/usr/bin/bash oracle
passwd oracle oracle
```

Note: these settings configure default (insecure) passwords purely as examples; more secure passwords should be provided.

Log in to the `oracle` user account and add the following paths to the `oracle` user's `.profile` file, which is found in the `oracle` user's home directory.

```
ORACLE_BASE=/db/app/oracle; export ORACLE_BASE
ORACLE_HOME=$ORACLE_BASE/product/11.1.0/db_1; export ORACLE_HOME
ORACLE_SID=orcl; export ORACLE_SID
JDE_HOME=/db/jdedwards/JDE_HOME/bin; export JDE_HOME
PATH=$PATH:/opt/SUNWsprow/bin:$ORACLE_HOME/bin:/db/usr/local/bin:/usr/ucb:.export
PATH
# for Verity software
VDK612_HOME=/db/jdedwards/e900/verity612/solaris; export VDK612_HOME
VDK612_PATH=$VDK612_HOME/_ssol26/bin; export VDK612_PATH
PATH=$VDK612_PATH:$PATH
LD_LIBRARY_PATH=$VDK612_PATH:$LD_LIBRARY_PATH; export LD_LIBRARY_PATH
```

In the DBServer domain, change the following project settings:

```
jds-zone1> projmod -s -K 'project.max-shm-ids=(privileged,1024,deny)' default
```

Increase the maximum size of a shared memory segment. Default = ¼ of physical memory.

```
jds-zone1> projmod -s -K 'project.max-shm-memory=(privileged,12gb,deny)' \
default
```

Increase the maximum number of semaphores allowed per semaphore set. Default = 512.

```
jds-zone1> projmod -s -K 'process.max-sem-nsems=(privileged,4096,deny)' \
default
```

Increase the maximum number of message queue identifiers. Default = 128.

```
jds-zone1> projmod -s -K 'project.max-msg-ids=(privileged,2048,deny)' default
```

Increase the maximum number of semaphore identifiers. Default = 128.

```
jds-zone1> projmod -s -K 'project.max-sem-ids=(privileged,15320,deny)' \
default
```

Increase the number of bytes for a message on the message queue. Default = 64k.

```
jds-zone1> projmod -s -K 'process.max-msg-qbytes=(privileged,16834,deny)' \
default
```

Increase the maximum number of message on a message queue. Default = 8192.

```
jds-zone1> projmod -s -K 'process.max-msg-messages=(privileged,16384,deny)' \
default
```

For security, add the following parameter to the `/etc/system` file.

```
set noexec_user_stack=1
```

And finally reboot the zone to make these changes occur.

```
reboot
```

Once the domain is rebooted, the Oracle Database software can be installed.

Note : The following changes are stated, but must be performed as a user that is appropriate to the customer's Oracle Database security requirements.

Please make the following changes to the Oracle Database configuration through `sqlplus` as the `sysdba` user:

```
alter system set processes = 3000 scope-spfile;
alter system set open_cursors=2000 scope-spfile;
alter system set session_cached_cursors=1000 scope-spfile;
alter system set pga_aggregate_target=0 scope-spfile;
alter system set sga_target=6G scope-spfile;
alter system set sga_max_size=7G scope-spfile;
alter system set memory_target=11G scope-spfile;
alter system set memory_max_target=11G scope-spfile;
alter system set recyclebin=OFF scope-spfile;
```

Default sizes for the standard database files were increased in the following ways:

- The `system` and `sysaux` tablespaces were resized to 1500 M and set to `autoextend` on next 10 M.
- The `tempfile` tablespace was resized to 7000 M.
- The number of redo logs was increased to 6 of 250 M sizes each.

D. Workload Details

There are two primary workloads provided by the DIL Workload used for the benchmark tests:

- Interactive Transactions—These are user transactions that reflect interactive Web-based operations.
- UBE batch jobs—Short- and long-running UBE batch job types are run during the test to increase the realism of the workload.

The details of each of these workloads are described in the following section.

DIL Workload

Oracle's DIL workload consists of a suite of scripts that simulate user interactions through a load generation system. The simulated transactions exercise common operations of JD Edwards EnterpriseOne applications, including business processes such as various interactions related to payroll, sales orders, purchase orders, work orders, and manufacturing processes such as inventory tracking. These are sometimes generically referred to by industry acronyms such as SCM, CRM, HCM, SRM, and FMS.

The DIL workload uses the JD Edwards EnterpriseOne interactive transaction types described in Table 10.

TABLE 10. JD EDWARDS ENTERPRISEONE INTERACTIVE TRANSACTIONS

	TRANSACTION TYPE	DESCRIPTION	VIRTUAL USERS PER TRANSACTION TYPE AT 1200-USER LEVEL
1	P03B102	Apply Receipts	110 users
2	P04111	Supplier Ledger Inquiry	110 users
3	P051191	Daily Time Entry	24 users
4	P17500	Case Management Add	60 users
5	P31114	Work Order Completion	36 users
6	P3411	MRP Messages (WO Orders)	24 users
7	P3411	MRP Messages (OP Orders)	24 users
8	P3411	MRP Messages (OT Orders)	24 users
9	P4113	Inventory Transfer	60 users
10	P42101	Sales Order Entry – 10 Line Items	340 users
11	P42101	Sales Order Update	60 users
12	P4310	Purchase Order Entry – 25 Line Items	220 users
13	P4312	Purchase Order Receipts	24 users
14	P4314	Voucher Match	24 users
15	P4915	Ship Confirmation – Approval only	36 users
16	P4915	Ship Confirmation – Confirm/Ship only	12 users
17	P4915	Ship Confirmation – Confirm and Change Entry	12 users

Note that the ‘Users’ column indicates the number of users who may simultaneously be carrying out that particular kind of transaction. The DIL test is provided with an expectation of running 1000 users in total. This must be increased to 1200 users to fully test this platform. To reach this maximum level, the number of users running each transaction was increased by 20% in all cases compared to default DIL configurations.

The short-running UBEs generate consistent load on the JD Edwards EnterpriseOne server and are continuously re-submitted upon completion. Long-running UBEs tend to generate substantial load on the database server since they typically involve more complex SQL queries. Once started, these long-running UBEs continue to run throughout the test runs. Tests runs typically last one hour and some long-running UBEs run for more than three hours before completion, representing further continuous load on the database server.

Such a combination of interactive, short-UBE and long-UBE workloads provide a good workload, stressing all aspects of the solution's configuration. Tables 11 and 12 list the transaction types for the long-running and short-running UBE jobs, respectively.

TABLE 11. LONG-RUNNING UBE JOBS

TRANSACTION TYPE	DESCRIPTION
R09801	General Ledger Post
R31802A	Manufacturing Account Journal
R31410	Work Order Processing
R3483	MRP Processing
R42520	Print Pick Slips Batch
R42800	Sales Order Update
R43500	Purchase Order Print

TABLE 12. SHORT-RUNNING UBE JOBS

TRANSACTION TYPE	DESCRIPTION
R0004P	UDC Records Types Print
R0006P	Business Unit Report
R00067	Business Unit Translation Report
R0008P	Date Patterns Report
R0010P	Company Constants Report
R0012P1	AAI Report
R0014	Payment terms Report
R0018P	Tax Detail Report
R00425	Organization Structure Report
R014021	One Line per Address Report
R01402W	Who's Who Report
R01403	Full Address Report with Codes
R03B155	A/R Summary Analysis
R03B31	Activity Log Report
R03B461	Collection Report
R04413	Open A/P Summary Report
R04602	Supplier Analysis Report
R09017	Account Translation Report
R09205P	Account Master Report
R09410	Trial Balance Report
R098201	Annual Close Report
R41411	Select Items Cost Count
R41543	Item Ledger/Account Integrity
R42072	Price Category Print
R46990	Warehouse Location Integrity
R48096	Cost Plus Markup Listing



Oracle Optimized Solution for JD Edwards
EnterpriseOne
October 2011, Version 1.2
Authors: Peter Wilson, Laura Verdibello

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

Copyright © 2011, Oracle and/or its affiliates. All rights reserved. This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark licensed through X/Open Company, Ltd. 0611

Hardware and Software, Engineered to Work Together