ORACLE
OPTIMIZED SOLUTIONS

Oracle Optimized Solution for JD Edwards EnterpriseOne

Delivering High Availability, Performance, Security and Efficiency Using Oracle's SPARC T5-2 Servers

ORACLE WHITE PAPER | JUNE 2015

ORACLE

Table of Contents

Introduction

Whether organizations are deploying a new resource planning service or modernizing existing business-critical resource planning services, selecting and implementing a solution architecture that will perform optimally and meet service-level agreements (SLAs) is a challenging process. Increasingly businesses are also critically reliant on the operation of their Enterprise resource Planning (ERP) platform for business operations, not only from a continuity standpoint but also to allow them to adapt their businesses to the ever-changing demands of their marketplaces. The Oracle Optimized Solution for JD Edwards EnterpriseOne aims to balance business requirements (cost and functionality) and technical requirements (high availability, security and performance). With the optimized solution, organizations can achieve a low-risk and low-cost implementation that surpasses their requirements and provides a clear and smooth path forward into the future.

This paper describes the implementation of Oracle Optimized Solution for JD Edwards EnterpriseOne, including the configuration process and the results of performance and high availability (HA) testing of the full architecture described in the available Implementation Guide. The results demonstrate that Oracle's SPARC T5-2 servers—enhanced with carefully configured HA hardware and software—can provide an excellent availability, price, and performance combination for small and medium enterprise deployments. The Oracle Optimized Solution approach unites these systems into an integrated and highly secure software and hardware infrastructure combination for JD Edwards EnterpriseOne applications.

Server consolidation and duplication through virtualization is key to this solution, reducing the cost of hardware and achieving high density, reliability, security and performance in a fully documented solution at low cost and with minimal data center impact. The result is a significant savings in the initial acquisition costs, an ongoing reduction in operating costs to the enterprise, lower deployment risk, quicker time to operation, and increased service availability.

Oracle's JD Edwards EnterpriseOne software consists of a core set of tools supplemented by more than 80 application modules addressing industry-specific ERP needs. In the testing carried out as part of this Oracle Optimized Solution, commonly used modules are exercised by a standardized JD Edwards EnterpriseOne "day in the life" (DIL) workload that is used to obtain performance metrics for particular system configurations to characterize their behavior under a range of load types and use cases.

In addition to this technical paper, a full Implementation Guide has been prepared describing all aspects of the installation and configuration recommendations and best practices. Security, sizing and tuning best practices are discussed as well as information regarding next steps in case organizations need to scale within or beyond the capabilities of the demonstrated platforms.

A Brief Overview of JD Edwards EnterpriseOne

Today, enterprises are faced with the need to operate within ever more complex and dynamic environments, often under legally imposed guidelines for business operations, asset management, and tracking obligations. In addition to these external pressures, increasingly, businesses are often critically reliant on the continuity of their day-to-day operations, which are coordinated through their ERP platform. Outages and unreliability have a significant and material impact on continuity and the ability to operate, and in the worst case, outages and unreliability can cause the viability of a business to suffer.

JD Edwards EnterpriseOne is a cost-effective and practical way of ensuring proper management, tracking, and compliance, which in turn enables streamlining of business models and operations. Increasingly, these requirements form significant, if not vital, parts of the day-to-day operations of a company. As such, they become “business critical” and whenever possible, they must be available 24/7 with the highest availability and integrity possible.

Tackling these requirements of a critical ERP platform deployment is a complex exercise and is often not something that a business is best equipped to solve on its own. Oracle has provided the Oracle Optimized Solution for JD Edwards EnterpriseOne in order to guide businesses through their deployments, providing them with considerable benefits in terms of cost reduction, risk reduction, and accelerated time to deployment.

The JD Edwards EnterpriseOne software stack consists of three core functional blocks, on top of which business-specific application modules can be loaded. Implementing HA and data integrity on top of these layers adds complication; if not performed correctly, it can, in fact, reduce overall system availability.

The three core functional blocks of JD Edwards EnterpriseOne are listed in Table 1. In a classic deployment, each function would be housed on a separate server, and in an HA deployment, each server would be duplicated through some form of clustering with additional duplication for load balancing an option.

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

Functional Block	Description
Database server	An instance of Oracle Database for data storage and tracking of assets and operations
Web server	An Oracle WebLogic Server enabling the web-based presentation of the user interface for both the core and optional JD Edwards EnterpriseOne modules
JD Edwards EnterpriseOne application server	The JD Edwards EnterpriseOne application server core installation and optional business logic functionality modules

The Oracle Database and Oracle WebLogic Server portions of the application stack are available in both standard and enterprise license levels, which provide basic and enhanced functionality, respectively. The choice of license has a very significant impact on overall system cost and can be very significantly reduced by using the licenses bundled in the Oracle Technology Foundation for JD Edwards EnterpriseOne pack. This pack includes the standard edition versions of Oracle Database and Oracle WebLogic Server software tied specifically to a JD Edwards EnterpriseOne deployment, and is attractively priced based only on the number of users required in the deployment.

One of the key goals of Oracle Optimized Solution for JD Edwards EnterpriseOne was to control cost when deploying an HA architecture by ensuring the careful use of the capabilities included with the Standard Edition licenses and some additional low-cost Oracle Solaris Cluster licenses to achieve a very substantial cost savings on the licensing of the software required. For customers requiring additional functionality, upgrades to the Enterprise Edition version of the required database and webserver are possible, however moving to the enterprise licenses might cause a considerable increase in deployment costs. It is the demonstrated goal of this solution to minimize costs where appropriate, with minimal sacrifice to functionality, therefore, this document discusses the most cost effective means of taking full advantage of the included Standard Edition licenses for both performance and HA.

When considering the security of the deployment, care has been taken to ensure that where decisions can be made that might affect security aspects of the deployment. With the optimized solution, best practices are called to the attention of the customer and recommendations made. The goal of this approach is to properly address possible security issues at every layer of the solution..

This solution architecture as demonstrated, achieves excellent performance, capable of serving upwards of 2000 Interactive users, with interactive response time of less than 0.1 seconds, and very high Universal Batch Engine(UBE) throughput. This is achieved along with maintaining excellent High Availability characteristics and very low overall system utilization levels.

Taking advantage of Oracle's SPARC server and Oracle Solaris virtualization features, which are available at no additional cost, allows "free" implementation of each of the required "servers" inside Oracle VM Server for SPARC or Oracle Solaris Zones virtual machines (VMs). Redundant pairs of virtual servers are hosted on two separate SPARC T5-2 servers. Requirements for shared storage are satisfied by using an Oracle ZFS Storage ZS3-2 appliance connected to the SPARC T5-2 servers via two 24-port, 10 GbE Oracle Switch ES1-24 switches. This approach provides the following advantages:

» Separating the applications' working environments into discrete VMs simplifies many aspects of security and virtual server management.
» Portability and backup options are improved through the use of the ZFS file system and tools for VM disk image management.
» The proper use of Oracle VM Server for SPARC resource capping allows full compliance with Oracle's hard partitioning requirements for the minimization of licensing costs.
» Virtual server consolidation, interconnect consolidation, and use of VLANs reduces the data center footprint and simplifies management of both security and traffic flows.
» Oracle virtualization technologies are fully qualified to host clustering environments for all of the required applications used throughout JD Edwards EnterpriseOne.

Compared to competitive configurations that do not provide HA at all levels of Oracle's JD Edwards application stack, the consolidated approach offered by Oracle Optimized Solution for JD Edwards EnterpriseOne provides unique and compelling benefits to customers.

Overview of the Solution Architecture

The solution architecture takes advantage of Oracle VM Server for SPARC domains and Oracle Solaris Zones as well as shared storage, and it optimizes the size and placement of VMs to take best advantage of the available low cost licenses to achieve the best performance and availability within the configuration. Figure 1 shows how the servers, VMs, and storage devices are configured across the SPARC T5-2 servers and the Oracle ZFS Storage ZS3-2 appliance (shared storage array).

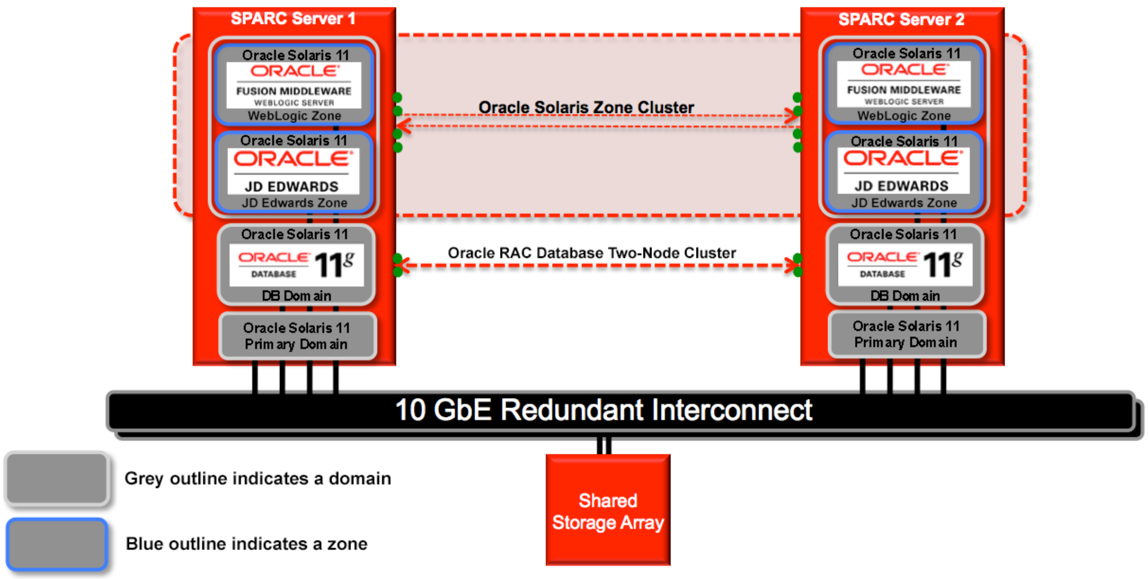

Figure 1: Technical architecture for Oracle Optimized Solution for JD Edwards EnterpriseOne

The SPARC T5-2 servers are partitioned into domains that fulfill two roles. Each SPARC T5-2 server contains the following:

- An Oracle VM Server for SPARC domain dedicated to running one node of a two-node Oracle Real Application Clusters (Oracle RAC) database cluster
- An Oracle VM Server for SPARC domain acting as an Oracle Solaris global zone hosting multiple clustered application zones for the various required JD Edwards EnterpriseOne and Oracle WebLogic Server tasks including
    - One or more zones hosting Oracle WebLogic servers for load balancing and availability
    - One or more zones hosting JD Edwards EnterpriseOne servers

The domain hosting the Oracle RAC database is clustered with its equivalent on the alternate SPARC T5-2 server, and the domain hosting the zones acts as a zone host in which multiple Oracle Solaris Zones are clustered with equivalent Oracle Solaris Zones on the other SPARC T5-2 server. Clustering ensures that the hosted zones are provided with a suitable HA framework to operate inside, which enables HA at all tiers of the JD Edwards EnterpriseOne deployment.

When no failures are present and no HA failover is underway, both SPARC T5-2 servers are able to render service and provide the highest possible overall throughput. In the event of any single failure, at most only half of the service capacity is lost and service can be continued with minimal interruption on the remaining server or zones. During system failures, the level of interruption to users or operations depends on the nature of the failure. In the most common failure scenarios, user impact is minimal and recovery operations are automatic and imperceptible. More severe types of failure might result in short interruptions in service in the order of seconds or minutes and, at worst, result in users having to log back into their accounts to continue working, or possibly re-start previously dispatched UBE tasks.

The use of Oracle VM Server for SPARC version 3.1 (or later) and Oracle Solaris 11.1 (or later) zones in the configuration of the servers is central to the architecture, allowing segregated hosting of the qualified versions of JD Edwards EnterpriseOne, Oracle RAC, and Oracle WebLogic Server.

Enhancing Security with an Oracle Optimized Solution Approach

Security has become paramount for ERP systems as they expand to integrate and unify more and more business functions and core processes. Business critical and confidential data too is aggregated as a part of this process, including intellectual property, financial data, and personal information. The concentration of critical data and key business processes means that security must be a major concern when designing and deploying any ERP environment. Protecting the internal infrastructure from intrusions and ensuring the integrity of digital information has become a priority and a core business function for most organizations.

Oracle Optimized Solutions are designed to address key security challenges, including:

» **Complexity**. Complexity breeds insecurity, and the security of an ERP implementation as a whole will only be as strong as its weakest component. It can be very difficult to understand how to securely implement the myriad products included in an heterogeneous ERP system. Oracle Optimized Solutions offer guidelines and recommendations to simplify ERP implementations using consolidation and virtualization technologies.

» **Implementation flaws**. Secure software is important, but not sufficient for comprehensive systemic security. Most security vulnerabilities stem from flawed implementation or architecture, improper configuration, improper access control, lack of patch management, unencrypted communications, or inadequate policies and processes. Oracle Optimized Solutions provide proven and tested architecture recommendations that follow existing security best practices and recommendations, dramatically reducing risk.

» **Performance and cost**. A number of security settings and requirements can have a significant negative impact on the ERP system’s performance and cost. For example, on-the-fly encryption—required for effective security—can significantly impact performance for production systems, or can require expensive computational add-ons. Oracle Optimized Solutions offer high performance security by using Cryptographic Instruction Accelerators directly implemented into SPARC processor cores. The result is wire speed security capabilities without the performance penalties and cost barriers typically associated with real-time secure computing.

Because Oracle can innovate, integrate, and test at all levels of the technology stack, it is in a unique position to provide a comprehensive security approach at the solution level.

» Operating system security includes features in Oracle Solaris and Oracle Solaris Cluster such as Oracle Solaris Secure by Default, and tier isolation utilizing Oracle VM Server for SPARC and Oracle Solaris Zones.

» Network security is provided with network isolation via partitioning (VLANs), isolation of storage data traffic and application data traffic, a physically separated systems management network, and dedicated cluster interconnects.

» Storage security incorporates use of Network File System (NFS) exceptions to restrict access, target groups and initiator groups to secure block devices (iSCSI LUNs), and isolated data an administrative traffic using separate network interfaces.

» General security recommendations are provided that treat security as a priority during the architecture definition and design phases of the ERP system, rather than an afterthought during or after implementation.

HA Configurations of the Database and Zones

Extending the first generation of this solution to incorporate an HA configuration required substantial changes to the architecture in order to introduce redundant hardware at all levels and to include clustering software and clustering configurations of interconnect and storage.

If budgets are not a significant concern, HA architectures typically deploy the most-capable Enterprise Edition versions of all of the required software and hardware available, at considerable increase in cost. In order to control costs and maximize return on investment (ROI), Oracle has put considerable effort into optimizing the use of the available JD Edwards and Oracle Solaris Clustering license packages providing access to low/no additional cost licenses. As a result, the architecture has been optimized to take best advantage of the Standard Edition Database and WebLogic licenses included with the Oracle Technology Foundation JD Edwards EnterpriseOne pack, with only the addition of low-cost Oracle Solaris Cluster licenses needed to extend HA functionality to zone clusters.

Oracle RAC Database

Oracle Database 11*g* Release 2 Standard Edition provides database cluster functionality for clusters equipped with up to four CPU sockets, this allows two, dual socket, T5-2 servers to be clustered in the solution using these licenses. A further requirement is the use of Oracle Clusterware software and Oracle Automatic Storage Manager to be compliant with the Standard Edition database license rules.

Oracle WebLogic Server

Oracle WebLogic Server 11*g* (10.3.6) is used to provide the solution's web and transaction processing engine, even Enterprise Edition WebLogic Server requires additional WebLogic Cluster licensing for native clustering functionality and this is at considerable additional expense. The Oracle optimized Solution for JD Edwards therefore implements high availability, using the Standard Edition of Oracle WebLogic Server, and clusters this by using Oracle Solaris Cluster 4.2.

Oracle Solaris Cluster

Oracle Solaris Cluster 4.2 is used to cluster the Oracle WebLogic Server and JD Edwards EnterpriseOne VMs. Oracle Solaris Cluster requires the purchase of licenses based on processor core counts. These are low-cost licenses, and by taking advantage of Oracle's hard partitioning license rules, compliant VM resource capping can be implemented to control cost and ensure that only the required numbers of cores are licensed according to performance or cost requirements.

Oracle has provided, at no additional charge, Oracle Solaris Cluster agents for both Oracle WebLogic Server and JD Edwards EnterpriseOne to improve the ability to cluster these applications and to provide considerable additional HA benefits, such as implementing automatic failover and recovery and awareness of service dependencies. This ensures that if, for example, an instance of JD Edwards EnterpriseOne fails, it would be brought back online or failed over to a standby instance only after the Oracle Database and Oracle WebLogic Server instances required to service it were verified and also available.

Shared Storage and Connectivity

Shared storage is essential for cluster implementations whether they are Oracle RAC database clusters or Oracle Solaris Clusters, and inter-cluster member communications are also vital to ensure proper cluster awareness and failure detection.

Shared storage is implemented by using the Oracle ZFS Storage ZS3-2 appliance in a 'no single point of failure' (NSPF) configuration, requiring two clustered controllers and two storage shelves.

10 Gigabit Ethernet is used for both standard network traffic and storage traffic and provides all of the connectivity within the architecture, storage protocols include both the NFS and iSCSI protocols. Data security for traffic in transit is assured by traffic separation using VLANs separated by dedicated switch ports and by optional encryption if desired. To ensure high availability, two 24-port, 10 GbE Oracle Switch ES1-24 switches are used, and all network communications are implemented using Oracle Solaris IP multipathing (IPMP) to provide redundant paths over both switches to all devices. Careful device choices are made to ensure the maximum separation of paths (both physically over dedicated VLAN ports and logically according to traffic type) to network devices to further enhance availability.

Oracle Virtualization Technologies

Both Oracle VM Server for SPARC (previously called Logical Domains or LDoms) and Oracle Solaris Zones virtualization technologies are used in the solution and complement each other with their capabilities. These technologies are included at no additional charge with all Oracle SPARC servers and the Oracle Solaris operating system.

The SPARC T5-2 servers are each partitioned into active domains using Oracle VM Server for SPARC, administered by a third control domain. One active domain operates as an Oracle RAC database node, and the other active domain operates as a global zone host for multiple zone clusters of the JD Edwards EnterpriseOne and Oracle WebLogic Server applications. The two active domains are administered and serviced by the control domain.

Secure Solaris installation standards are provided for each of the required domains and Zones to ensure that only required software is installed and potential security issues are properly addressed.

Each domain is provided with a pair of mirrored boot disks on hosting SPARC T5-2 server. Each SPARC T5-2 server also implements eight virtual switch services that each provide connectivity to one of the specifically identified and isolated networks that are required for this architecture. In most cases, two virtual switches are created for each network, each switch representing an independent path of connectivity between devices. In each domain, virtual network devices are connected to each virtual switch and are locally combined in IPMP device pairs to provide redundancy at the domain level.

In Figure 2, each of the identified separate networks is color coded according to VLAN separation, and all Oracle Switch ES1-24 ports and cables that correspond to a specific VLAN are colored alike. Virtual network devices in each domain are connected by dotted lines to the virtual switch service that connects them to the appropriate physical network.

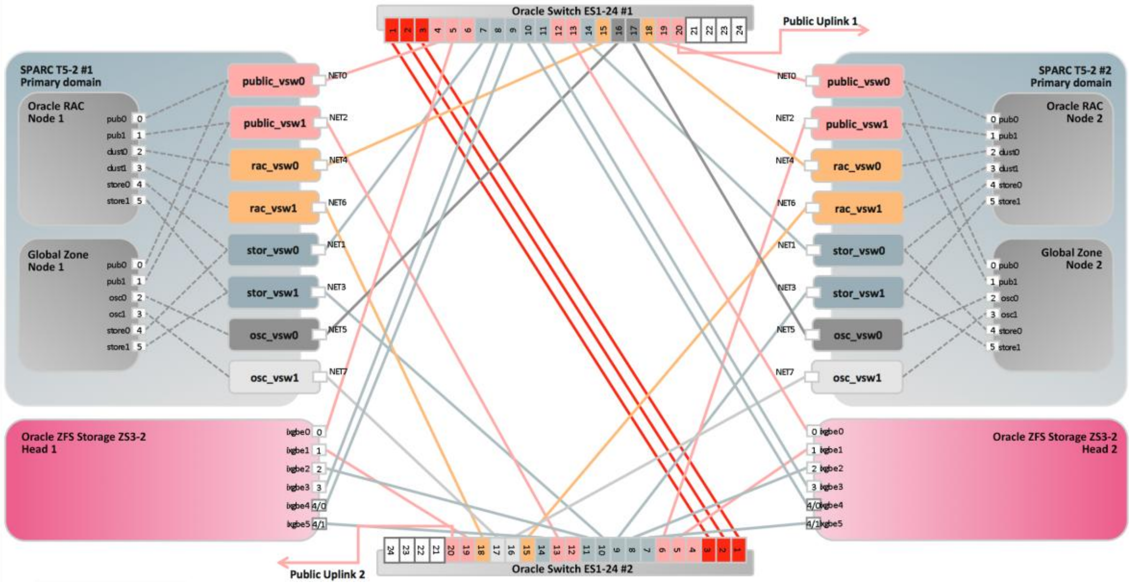

Figure 2: Physical-to-virtual network device and service mapping

Virtual storage devices are also configured according to redundancy and connectivity requirements in the domains. Local boot devices are mirrored, and remote or shared devices from the Oracle ZFS Storage ZS3-2 appliances are specifically defined and shared to individual clients according to their use case, and with the intention of limiting access to storage to only thise domains or zones that require access for their operation. Local devices appear as simple local SCSI devices inside each domain, and remote devices can be provided as either NFS mounted filesystems or as iSCSI LUNs as appropriate connected to a private network that is not visible externally from the architecture. Visibility of remote devices or mounts is limited to only the required clients through specific configuration of the sharing behavior of the Oracle ZFS Storage ZS3-2 appliance and port configuration of the network switches. In addition, performance characteristics are also tuned for each mount or LUN in terms of optimal block size, mount options, compression, throughput, latency optimization, and so on.

It has been determined, by previous testing, that the key areas governing database performance are the disk I/O performance, the capacity of the storage solution, and—to a lesser extent—the bandwidth of the cluster network interconnect. In order to maximize the performance of the disk storage provided in the solution, each Oracle VM Server for SPARC domain hosting an Oracle RAC database node is connected to the shared storage by dedicated redundant, high-speed 10 GbE network devices. Care is taken to ensure that all single points of failure (SPoFs) are removed from the storage network connections internally by using separate NICs on separate PCIe busses. In addition, all SPoFs are removed externally by using redundant paths over two separate 10 GbE switches.

In order to provide private and secure networks for cluster interconnects and public networks for data communications to other elements of the solution eight physical network devices (NICs) are required, four on-board devices and four additional NICs via two dual NIC PCIe cards. These NICs are accessed through virtual switch services provided by the primary domain on each server and connected only to the relevant domains. Where appropriate, network connections to the servers use Oracle Solaris IPMP for maximum availability.

The resulting logical configuration of the services provided to the Oracle RAC database and global zone host domains (called “RAC_Node” and “GZ_Node,” respectively) is shown in Figure 3. The same domain configuration is used on both SPARC T5-2 servers. Virtual switch services (labeled *xxx*_vsw0 and *xxx*_vsw1) are created and attached to the physical external ports of the server and are used in domains to create redundant IPMP connections, and internal storage is provided by the “jde_vds” virtual disk service, which provides both boot and mirror virtual disks to each domain where the ZFS filesystem is used to mirroring, use is made of both onboard local disk controllers for added performance and redundancy. Each domain is also provided with a virtual DVD drive for straightforward initial OS installation.

The primary (control) domain is also provided with mirrored root disks and can utilize the native NET0–NET7 devices to communicate to the various VLANs, if needed, though for security its access can be restricted.

The SPARC T5-2 server physical NIC ports NET0–NET7 are used by the eight virtual switch services (VSWs) to allow the Oracle RAC and global zone domains to access the relevant networks via local virtual network devices (VNICs) in each domain (pub0, pub1, clust0, clust1, etc). The NET0–NET7 ports are externally connected through redundant 10 GbE switches to which implement port based VLAN separation of networks for access to the shared Oracle ZFS Storage ZS3-2 storage, cluster interconnects to the other SPARC T5-2 server, or out to the public network. VSW and VNIC names are chosen to reflect their purpose indicating ‘pub’lic networks, ‘clust’er networks, ‘stor’age networks, or ‘osc’ Oracle Solaris Cluster networks.

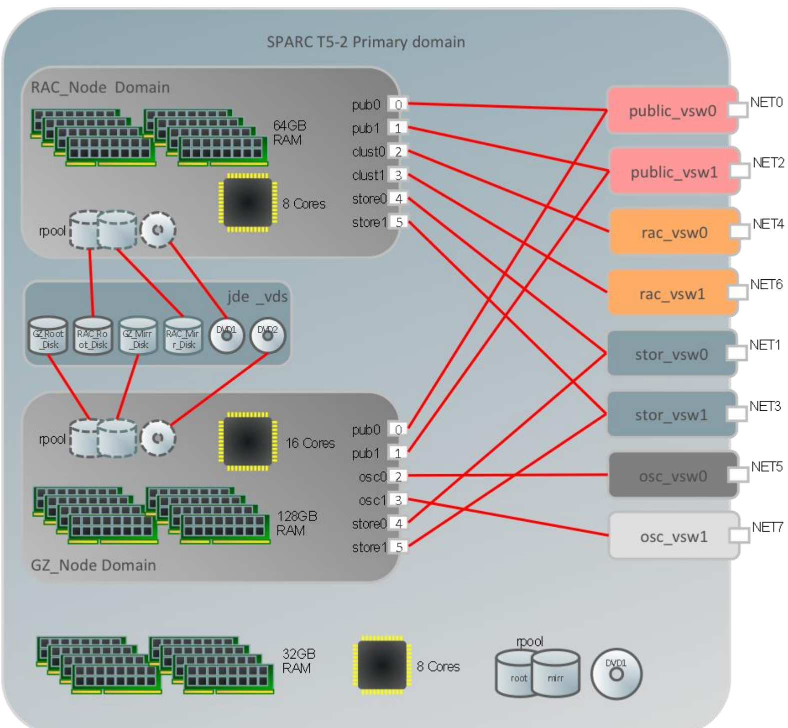

Figure 3: Domain configuration and logical-to-physical connectivity

Figure 3 also shows the partitioning of the CPU and Memory resources of the SPARC T5-2 server into the primary, RAC_Node, and GZ_Node domains. The partitioning of the SPARC T5-2 resources results in allocations of dedicated CPU cores and memory to each domain, some core facilities such as virtual console ports used to communicate with the active domains are not shown.

Detailed Domain Configuration

The allocation of CPU cores and memory shown in the Figure 3 demonstrates the LARGE configuration of a SPARC T5-2 server recommended for this solution (see the sizing guide portion of this document for additional configuration options). Smaller configurations can be used and will enable savings in license costs as well as potentially create space for further domains to be created for additional consolidation work.

Resources for the primary domain are allocated as shown in Table 2.

TABLE 2. PRIMARY DOMAIN RESOURCE ALLOCATION AND DERIVATION

Resource Type	Amount	Dedicated or Shared	Service Provided	Device Used	Minimum/ Maximum Quantity	Notes
CPU Core	8	Dedicated	—	—	2 (min) 8 (max)	Dedicated CPU
Memory	16 GB	Dedicated	—	—	8 GB(min) 32 GB(max)	Dedicated memory
Local Disks	2	Dedicated		SYS/SASHBA0 (HDD0+HDD1)	1 disk (no HA) 2 disks (mirror)	Local boot disks for the primary domain
Local Network	8	Shared	NET0	SYS/MB	—	Local network devices via on motherboard NICs or via 2 x dual NIC PCIe cards in slots 2 and 6
			NET1	SYS/MB	—	
			NET2	SYS/MB	—	
			NET3	SYS/MB	—	
			NET4	PCIE6	—	
			NET5	PCIE6	—	
			NET6	PCIE2	—	
			NET7	PCIE2	—	
Virtual Switch Server	8	Shared	public_vsw0	NET0	—	Virtual switch services provided through local network ports connected to an Oracle Switch ES1-24 switch for access to port-based VLANs
			public_vsw1	NET2	—	
			rac_vsw0	NET4	—	
			rac_vsw1	NET6	—	
			osc_vsw0	NET5	—	
			osc_vsw1	NET7	—	
			stor_vsw0	NET1	—	
			stor_vsw1	NET3	—	
Virtual Disk Server	1	Shared	RAC_Root_Disk	SYS/SASHBA0/HDD2	1	Virtual boot disk
			RAC_Mirr_Disk	SYS/SASHBA1/HDD4	1	Virtual mirror disk
			dvd1	sol-11_1-text-sparc.iso	1	Virtual CD-ROM
			GZ_Root_Disk	SYS/SASHBA0/HDD3	1	Virtual boot disk
			GZ_Mirr_Disk	SYS/SASHBA1/HDD5	1	Virtual mirror disk
			dvd2	sol-11_1-text-sparc.iso	1	Virtual CD-ROM

The choice of physical devices used to provide the virtual services was defined to ensure that when the guest domains use derived virtual devices, maximum availability is maintained by maximum separation of underlying physical devices.

For example, when a guest domain wishes to connect to a public network, it should do so by using the two NICs attached to the two “public_vsw” services. These services are, in turn, provided over independent ports driven by separate controllers on separate PCIe busses. The guest domain would implement IPMP over the two virtual NICs and be assured of maximum availability and path independence.

Similarly, guest domain virtual disks are also chosen to be provided by separate local disk controller HBAs on the SPARC T5-2 server’s motherboard to ensures that if the guest mirrors the two local disks it is allocated, it is assured of maximum path separation for each disk and maximum uptime from the resulting mirrored volume. This also provides some performance improvements by spreading disk mirroring loads over two separate HBA's.

Resources for the RAC_Node domain are allocated as shown in Table 3.

TABLE 3. RAC_NODE DOMAIN RESOURCE ALLOCATION AND DERIVATION

Resource Type	Amount	Dedicated or Shared	Service Provided	Device Used	Min/Max Quantity	Notes
CPU Core	8	Dedicated	-	—	4 (min) — (max)	Dedicated CPU; no limit for maximum number of cores
Memory	64 GB	Dedicated	-	—	8 GB (min) 128 GB (max)	Dedicated memory
Local Disks	2	Dedicated	rpool /u01	RAC_Root_Disk and RAC_Mirr_Disk —	2 disks (mirror)	Virtual mirrored boot disks for the RAC_Node domain NFS mount of shared zone data storage
Local Network	6	Shared	NET0 NET1 NET2 NET3 NET4 NET5	public_vsw0 public_vsw1 rac_vsw0 rac_vsw1 stor_vsw0 stor_vsw1	— — — — — —	Virtual network devices via virtual switches connected to relevant VLANs

Resources for the GZ_Node domain are allocated as shown in Table 4.

TABLE 4. GZ_NODE DOMAIN RESOURCE ALLOCATION AND DERIVATION

Resource Type	Amount	Dedicated or Shared	Service Provided	Device Used	Minimum/ Maximum Quantity	Notes
CPU Core	16	Dedicated	—	—	4 (min) — (max)	Dedicated CPU; no limit for maximum number of cores
Memory	128 GB	Dedicated	-	—	64 GB (min) — (max)	Dedicated memory Server capacity limits ultimate maximum memory allocations
Local Disks	2	Dedicated	Rpool /u01	GZ_Root_Disk and GZ_Mirr_Disk /iscsi/c0t6*xxxx..xxxx* —	2 disks (mirror) 1	Virtual mirrored boot disks for the GZ_Node domain iSCSI cluster quorum disk provided over IPMP (NET4 and NET5) NFS mount of shared zone data storage
Local Network	6	Shared	NET0 NET1 NET2 NET3 NET4 NET5	public_vsw0 public_vsw1 osc_vsw0 osc_vsw1 stor_vsw0 stor_vsw1	— — — — — —	Virtual network devices via virtual switches connected to relevant VLANs Shared storage access is through IPMP over NET4 and NET 5

The Solution's Software Components

The primary software components that comprise Oracle Optimized Solution for JD Edwards EnterpriseOne are described in the subsections that follow.

Oracle Solaris 11 Operating System

Oracle Solaris provides industry-leading performance and numerous advanced features, all at no additional cost. Features include the following:

» Operating system and platform virtualization—Oracle VM Server for SPARC and Oracle Solaris Zones enable optimized resource utilization to deliver excellent performance and availability levels.

» ZFS file system—ZFS radically simplifies the management of complex storage subsystems and provides seamless data integrity and portability functions at no cost. 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.

Each VM used in the solution runs Oracle Solaris 11 (at the time of writing, the latest version was Oracle Solaris 11.1 SRU 19).

JD Edwards EnterpriseOne 9.1.4

JD Edwards EnterpriseOne is an integrated application suite of comprehensive ERP software that combines business value, standards-based technology, and deep industry experience into a business solution with a low TCO. A collection of prebuilt business intelligence (BI) applications available as modules for JD Edwards EnterpriseOne provide organizations with 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 11*g* Release 2 Standard Edition

Designed for data center environments that are rapidly evolving to keep up with the demands of businesses, Oracle Database 11*g* 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 11*g* 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 10.3.6 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 a choice for development frameworks and tooling. Comprehensive and accessible management capabilities enable the administration of sophisticated systems via a well-designed graphical console, automation, or both. 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.

Securing the Software Components

At every stage of the software installation process (as described in the Implementation Guide document) security considerations and best practices are described and recommended. These procedures detail software installation requirements, supporting tools and packages and their installation, as well as configuration and removal post-installation if necessary. In addition, references are also provided for established security best practice documents that discuss specific use cases and less common scenarios in more detail. References are also available for third party partners that are experienced in specific security issues relating to the JD Edwsards software stack and in particular customized modules and software.

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 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 Oracle Optimized Solution engineering, Oracle offers a 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.

Oracle Technology Foundation for JD Edwards EnterpriseOne is licensed per application user, not per CPU, and covers multiple deployments so can be installed on an unlimited number of processors, providing some protection against increased 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 as well as production 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 the use of these products by applications other than JD Edwards EnterpriseOne, customers have the option of purchasing Enterprise Edition licenses where required.

The Solution’s Hardware Components

The following subsections describe the server and storage components that were selected for Oracle Optimized Solution for JD Edwards EnterpriseOne.

SPARC T5-2 Servers

Oracle’s SPARC T5 processor-based servers are the platform of choice for a full range of data center applications and workloads. The SPARC T5-2 server was selected for Oracle Optimized Solution for JD Edwards EnterpriseOne because of the key characteristics described in Table 5.

TABLE 5. KEY CHARACTERISTICS OF THE SPARC T5-2 SERVER

SPARC T5-2 Server System Characteristic	Key benefits for the solution
SPARC T5 processor	The SPARC T5-2 server is available in two SPARC T5 processor configurations. Each SPARC T5 processor has 16 cores and supporting 8 threads each, giving a total of 128 threads per CPU and up to 256 threads per system. SPARC T5 processors are based on the latest SPARC core with enhanced single-threaded performance, at a clock speed of 3.6 GHz. They integrate virtualization technology, allowing very efficient, high-performance system partitioning and virtualization. Systems equipped with the SPARC T5 processor are able 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 the solution-recommended 256 GB of memory (but capable of up to 1 TB of memory), the SPARC T5-2 server provides the capability of hosting large, very capable VMs; 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	Deployed in this solution are six SAS2 disks of 300 GB or 600 GB capacity each, managed by the two onboard SAS disk controllers. Disks are configured in pairs in each Oracle VM Server for SPARC VM using ZFS RAID mirroring. Performance and data protection are maximized in each VM, in local, mirrored 600 GB of protected space.

Onboard SAS Controller Utilization on SPARC T5-2 Servers

Each SPARC T5-2 server has two onboard SAS2 disk controllers, with controller SASHBA0 owning disk slots 0–3 and controller SASHBA1 owning disk slots 4 and 5. In order to maximize availability of the boot disks for each domain, mirroring of all local storage is recommended. The choice of which disk controller and disks are supplied to each VM can improve both availability and performance by ensuring maximum separation of both workload and device paths across independent controllers.

In this solution, the SPARC T5-2 platform is split into three domains: a primary domain that manages the other two domains, an Oracle RAC database node domain, and a global zone host domain. The primary domain does not have significant local disk performance requirements, and its HA requirements are secondary to the HA requirements of the domains it manages (a failure that affects the primary domain will also affect the domains it manages, but those domains are active members of clusters and can rely on their cluster peers to take over for them if they fail). As a consequence, the root disks in the primary domain are mirrored but not spread over both of the available disk controllers.

The Oracle RAC and global zone domains are provided with two virtual disks by the primary domain 'jde_vds' virtual disk server service. This service is configured to provide essentially a pass-through connection of physical disks to the guest domains. The choice of which disks are provided to each domain enables load spreading and availability dual pathing. In the guest domains, these virtual disks appear as normal, local SCSI disks and are mirrored using ZFS for the domains' root ZFS pool (rpool). In addition, each domain is provided with a virtual CD-ROM, this is a read-only connection to an Oracle Solaris installation DVD ISO file on the primary domain, which appears to the guests as a local CD-ROM device. The primary domain has a real physical DVD-ROM device that can be used in the conventional manner to install the OS on the primary domain.

Figure 4 shows the local storage virtualization model, which both SPARC T5-2 servers follow.

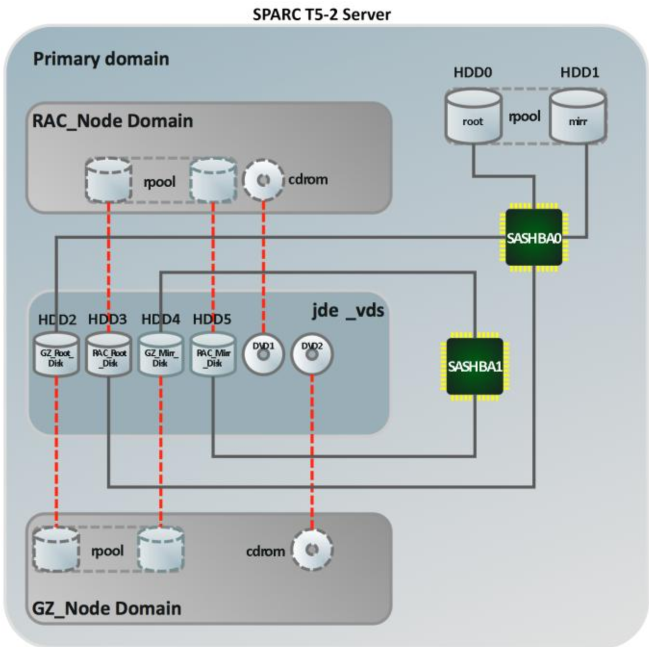

Figure 4: Local physical and virtual storage configuration

Remote/Shared Storage Architecture

The Oracle RAC database and Oracle Solaris Cluster rely on shared storage between their cluster member nodes in order to operate. In this architecture, all shared storage is provided by an Oracle ZFS Storage ZS3-2 appliance, which is configured in a redundant fashion with two controller heads and two storage trays.

Figure 5 is a simplified diagram of the storage connectivity and services.

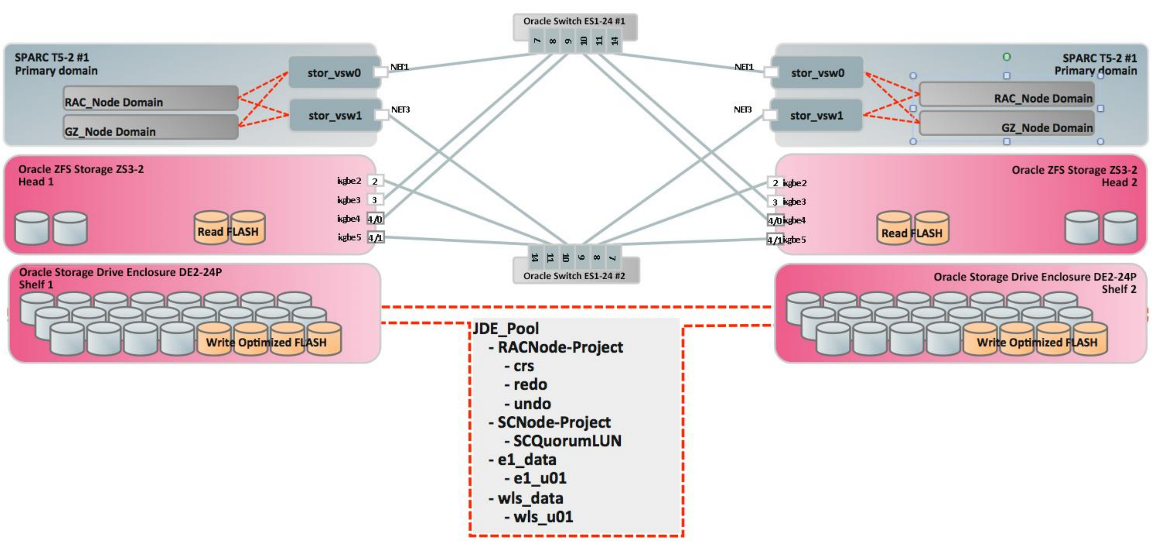

Figure 5: Shared storage configuration

The central area in Figure 5 illustrates the Oracle ZFS Storage ZS3-2 appliance's storage pool, storage project, and individual share/LUN mappings that have been created for this solution. These items are discussed in more detail later in this document.

Oracle ZFS Storage Appliances allocate storage based on a hierarchy of storage "pools" (consisting of collections of raw storage devices—such as hard disk drives [HDDs]—and read- or write-optimized solid state disks [SDDs]), "projects" that describe the general properties of shares or LUNs for classes of use, and individual "shares" that are the actual shared file system or LUN and have either properties inherited from the project or individually varied properties for specific requirements.

The optimal configuration of storage pools varies as the number of available raw devices is increased. In the minimal configuration that is required for having no SPoF in this HA configuration (two storage shelves and two controller heads), only a single pool is recommended for maximum performance. In this solution, the pool is named 'JDE_Pool'.

In general, the specific properties and "sharedness" of each file system, or LUN, are carefully chosen based on testing and best practices to optimize performance and security for the specific use case of each share. Small changes in, for example, NFS mount options used in clients can have significant effects on overall system performance or functionality, so it is very important to establish correct sharing and mounting properties to ensure the optimal behavior of each share.

The projects and shares established for this solution are outlined in the Table 6 through Table 9. Only a single storage pool is used and comprises all of the available storage devices from both storage shelves.

TABLE 6. RACNODE_PROJECT SHARE CONFIGURATION

Share	Description	Configuration
tables	The database data files (Oracle Automatic Storage Management disk files)	8 kb record size, write-throughput optimized, no compression, cache data and metadata
redo	Database redo log files	128 kb record size, write-latency optimized, no compression; do not cache
index	Database index files	8 kb record size, write-throughput optimized, no compression, cache data and metadata
undo	Database undo logs files	128 kb record size, write-throughput optimized, no compression; do not cache
temp	Database temporary files	128 kb record size, write-throughput optimized, no compression; do not cache
archive	Database archive files	128 kb record size, write-throughput optimized, GZIP compression; do not cache

TABLE 7. SCNODE_PROJECT LUN CONFIGURATION

LUN	Description	Configuration
SCQuorum_LUN	Oracle Solaris Cluster quorum voting LUN	1 GB LUN accessible only by the Oracle Solaris Cluster member nodes

TABLE 8. JD EDWARDS ENTERPRISEONE FILE SYSTEM SHARE CONFIGURATION

Share	Description	Configuration
e1_data	JD Edwards EnterpriseOne application data	128 kb record size, write-throughput optimized, no compression; do not cache

TABLE 9. ORACLE WEBLOGIC SERVER FILE SYSTEM SHARE CONFIGURATION

Share/lun	Description	Configuration
WLS_data	Oracle WebLogic server application data	128 kb record size, write-throughput optimized, no compression; do not cache

Mounting Options for Shared NFS File Systems

In testing, it was determined that specific NFS mounting options and technologies (in the case of the database files) are required to provide functionality, and they offer significant performance advantages when used correctly. Improper or default options can result significant performance degradation.

In the case of database tables under Oracle Automatic Storage Management control it is strongly recommended to convert the default OS NFS client in to the Oracle-provided dNFS client (a lower-overhead, multithreaded client specially developed for Oracle Database use).

Specific client mounting options are required by Oracle Databases and can be tested using conventional NFS clients before upgrading to the dNFS client.

TABLE 10. DATABASE CLIENT NFS/DNFS MOUNT OPTIONS

Share	Mount Options
Data files (Oracle Automatic Storage Management disks)	rw,suid,bg,hard,nointr,rsize=1048576,wsize=1048576,proto=tcp,noac,forcedirectio,vers=3
Binaries/files	rw,suid,bg,hard,nointr,rsize=1048576,wsize=1048576,proto=tcp,noac, vers=3

For other NFS client mounting options relating to the e1_data or WLS_data shares, conventional Oracle Solaris NFS clients are sufficient and should be configured with similar mount options.

TABLE 11. ORACLE WEBLOGIC SERVER AND JD EDWARDS ENTERPRISEONE FILE SYSTEM MOUNT OPTIONS

Share	Mount Options
Data files (Oracle Automatic Storage Management disks)	rw,bg,hard,nointr,rsize=131072,wsize=131072,proto=tcp,vers=3,noatime

By adhering to these mount option recommendations, significant performance improvements can be achieved.

Clustering in the Solution

Clustering technology is central to the ability to provide HA. Oracle RAC clustering is used for the database tier, and Oracle Solaris Cluster is used in the application tier to provide easy-to-configure and capable application reliability.

Oracle RAC Database Availability

Conventional Oracle RAC database clustering is used in the database tier with essentially no other tunings beyond standard best practices and the specific mount point configurations outlined earlier in this document.

Oracle Solaris Zone Clusters

Oracle Solaris Cluster simplifies the configuration of complex clustered application instances by leveraging the Oracle Solaris Zones virtualization feature to provide isolated environments for the application instance to operate within, and then it provides a framework by which those zones can be replicated and clustered to another host.

These zone clusters provide simplicity in configuration and operation and allow best practices to be easily applied for the clustering requirements of diverse application stacks.

In addition to providing the zone clustering technology, Oracle application-specific cluster agents are provided at no additional cost. These enhance the applications' "cluster awareness" and cluster capability. With the release of Oracle Solaris Cluster 4.2, an agent is now available for JD Edwards EnterpriseOne. An agent was already available for Oracle WebLogic Server before, and both agents are used to obtain the best functionality for their respective applications.

The key advantages provided by the use of zone clusters include the following:

» Application auto-failover to a peer cluster member in the event of failure

» Automatic resource dependency checking and assurance (for example, JD Edwards EnterpriseOne is not brought up until and unless both the database and Oracle WebLogic Server virtual servers are also available)

» Resource group definitions for both storage and networking resources

» Close integration with Oracle ZFS Storage Appliance functionality through custom cluster-related "work flows"

A minimum of two zone clusters are required: one encompassing the JD Edwards EnterpriseOne application and the other encompassing the Oracle WebLogic Server application, as demonstrated in the implementation guide documentation. However, for additional performance or load balancing, additional zone clusters can be created, as needed, to accommodate individual customer requirements. For example, a dedicated JD Edwards EnterpriseOne cluster can be created to cope with batch workloads (UBEs), or additional JD Edwards modules such as One View Reporting (OVR), and so on.

Other zone clusters can also be easily accommodated for unrelated application stacks to take advantage of unused capacity on the pair of SPARC T5-2 servers, potentially further increasing the ROI available from consolidation on these platforms.

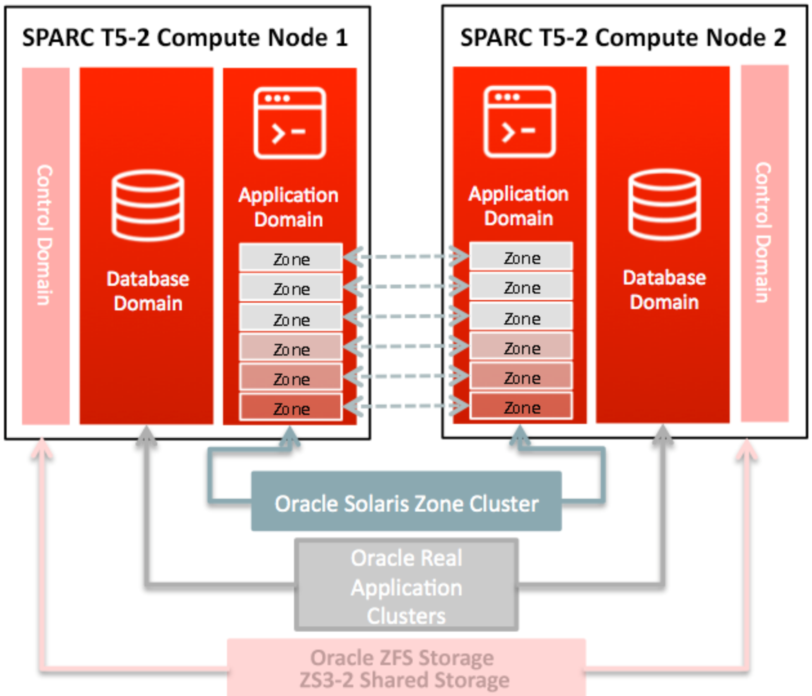

Figure 3: Illustration of zone clusters and Oracle RAC database cluster logical layout

Network Interconnect

In order to implement HA, all devices that need to communicate with another device must be able to do so through redundant paths; this includes both servers (and the VMs hosted on them) and the storage heads of the Oracle ZFS Storage ZS3-2 appliance. In order to provide redundant paths between components, interconnected redundant switches are required and these must be configured appropriately to allow each switch to take over communications if its partner fails. Each device attached to the network implements suitable network redundancy using technologies such as Oracle Solaris IPMP, and ensures that each network interface (NIC) is connected to paths across alternate switches. Where possible, each NIC is not reliant on other infrastructure components, such as PCI busses, PCI cards, or onboard devices.

In addition to providing redundancy, the network switches also ensure proper traffic separation and security for the various data flows that are required for the solution. These data flows consist of the following:

» Public network traffic
» Storage network traffic
» All private cluster interconnects

In order to separate traffic flows, VLANs are implemented on the switches and configured on selected ports (using switch-based port VLAN IDs (PVIDs) and port-based VLAN configurations), essentially allowing only traffic from elected VLANs to pass in or out of the switch on dedicated ports. The switches are further interconnected by an inter-switch “trunk” that allows all traffic from each VLAN on a switch to be visible to the equivalent VLAN on the other switch.

Figure 7 shows the full network diagram with all servers, storage devices, and switches and the color-coded VLANs used to separate traffic.

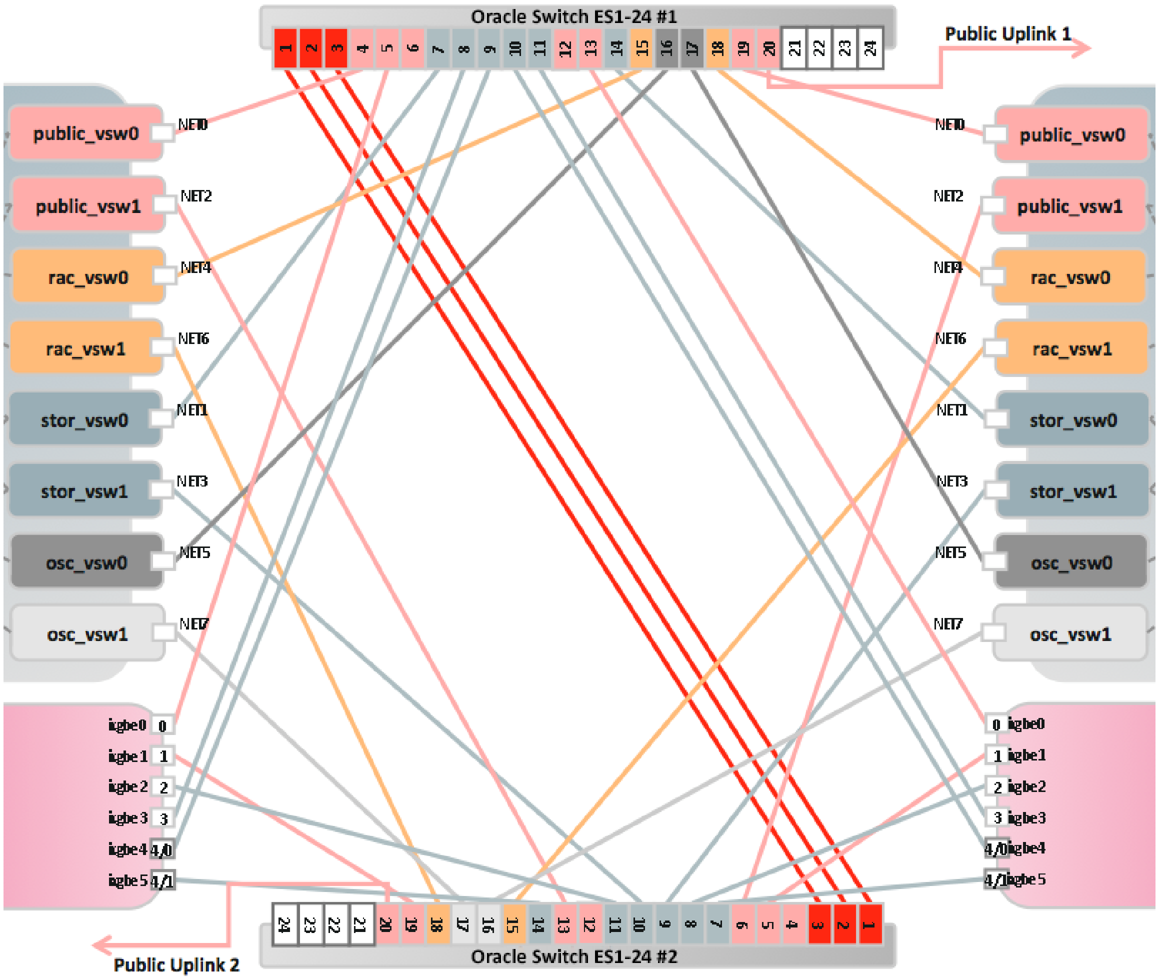

Figure 7: Network diagram showing device ports and color-coded VLAN connectivity

Figure 8 shows a simplified network diagram that represents the connections required for a single VLAN. Each server on the network uses two NICs that are combined into an IPMP redundant device. Each NIC is connected to separate switches, so that if a switch fails, the other device in the IPMP pair can continue to provide connectivity through the remaining switch. Two switches are required, and the switches are interconnected by redundant inter-switch links. Each switch is also provided with an uplink to a public network (only if required; cluster interconnects and storage are regarded as private and are not uplinked from the switches for security). It can be seen from the figure that if any single port, cable, or switch fails, there is always a communication path available between the two servers.

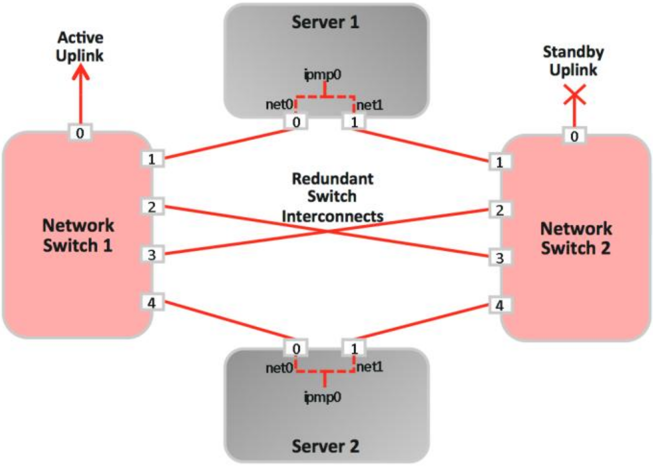

Figure 8: Simplified network diagram showing redundant connections

Each VLAN shown in Figure 7 can be considered similar to the simplified diagram in Figure 8, with the VLAN configurations being effectively superimposed on the Oracle Switch ES1-24 switches and using dedicated ports for VLAN separation. Exceptions to this description are required for the Oracle Solaris Cluster interconnect, which does not implement IPMP but instead manages redundancy itself using two NICs for its inter-cluster private interconnect. In the case of Oracle Solaris Cluster, NIC0 on each server is connected to switch 1, and NIC1 on each server is connected to Switch 2. These interconnects require separate VLANs on each switch to separate the traffic in interconnect 1 from interconnect 2. Oracle Solaris Cluster then manages redundancy itself.

Note that the switch interconnects for each VLAN could potentially represent a network loop, and connecting such a configuration of switches to a network will cause significant network problems while automatic protocols attempt to break traffic loops. This problem can be avoided by careful configuration of spanning tree protocols, port weights and costs on the uplink and trunk ports on the switches, which ensures that the trunk always forwards traffic between switches while only one or other of the uplink ports is ever enabled in order to break loops (during normal operation, the other uplink port is automatically “culled” from the network by the spanning tree).

Spanning tree protocols are potentially an issue on customer networks because two “standards” exist: the IEEE MSTP standard, which is a true standard (and is the default protocol used by Oracle Switch ES1-24), and the Cisco PVRST protocol, which is configured by default on Cisco switches (and often is left configured in deployed networks). These two network loop resolution protocols are incompatible; only one protocol should be present on a network. *It is vitally important that network administrators use only one protocol in their networks.* Suitable switch configurations are provided for either protocol in the full Deployment Guide documentation for the solution.

Workload Description

In order for performance and scalability to be adequately characterized on this proposed architecture, the JD Edwards EnterpriseOne day-in-the-life (DIL) workload was used to stress the system and determine performance, utilization, and availability of the architecture under various load conditions. The test results provide a basis for performance comparisons with alternative competitive architectures and allow customers to draw accurate comparisons based on their specific workloads during their real-world operations.

For this purpose, the JD Edwards EnterpriseOne product team provides a 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.

Key metrics obtained from the testing include the response time for user queries issued via the virtual users' web-based GUI to JD Edwards EnterpriseOne as the number of simulated interactive users is increased to its peak during 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.

Typical load scenarios are then used to test HA metrics during normal load failure events to determine the performance and availability impact of various types of failures throughout the architecture and also to determine the time and (if necessary) the steps required to recover back to full HA capacity again.

Workload Details

Interactive user transactions, such as those in the DIL workload, typically generate only moderate load on a system such as the SPARC T5-2 server. To stress the system to determine its true capabilities, batch-type workloads must also be used. Batch processes are non-interactive and vary from brief operations that are completed in seconds to high-load operations that might require many hours to be completed. Such batch tasks are broadly classified in these tests 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 general ledger, inventory, or payroll reconciliation, which might run for hours and heavily exercise the database and, consequently, the disk storage system.

With the goal of demonstrating a real system workload, test scripts were developed to ensure that a significant level of load is generated on the system by running various combinations of both short and long batch tasks during testing of the interactive response times for the target 2,000 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 these scripts running on the JD Edwards EnterpriseOne server itself. During testing, these might launch thousands of short UBEs per hour 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 workloads, 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 "Overview of the Solution Architecture" section of this document along with the following additional servers, which were required for the benchmark configuration:

- Deployment server—A small-configuration x86 server for controlling software installation and updates to the JD Edwards EnterpriseOne server. This server does not take an active part in the testing of this configuration or in the day-to-day operation of JD Edwards EnterpriseOne, but it is required for the initial deployment and for the maintenance of application code for the JD Edwards EnterpriseOne environment. This server, in the proposed solution, is actually a VM running Microsoft Windows Server 2008 in an Oracle VM VirtualBox on a low-cost x86 server.

» Load generation server—An x86 server acting as a load generation device, load testing tool, and performance evaluation server.

Test Results and Analysis

Seven test cases were executed to test several dimensions of system performance with varying load configurations of JD Edwards EnterpriseOne. Test data was collected, as shown in Table 12.

TABLE 12. DATA COLLECTED FROM LOAD TESTS

Description	NUMBER OF INTERACTIVE USERS AND UBES/HR WORKLOAD							
	500	750	1000	1500	2000	1000 + 5307	1000 + 8456	1000 + 10277
Test Run	1	2	3	4	5	6	7	8
Interactive Response Time (Seconds)	0.056	0.055	0.054	0.074	0.098	0.124	0.134	0.152
JD Edwards EnterpriseOne Server % CPU utilization	0	0	1	1	1	2	2	3
JD Edwards EnterpriseOne Server % Memory Utilization	3	5	9	13	15	9	10	10
Database Server % CPU utilization	0	1	0	0	1	1	1	1
Database Server % Memory utilization	23	25	28	31	34	44	44	44
Oracle WebLogic Server % CPU Utilization	0	1	1	1	5	4	1	7
Oracle WebLogic Server % Memory utilization	6	6	10	14	20	10	10	10

These results are shown charted in the figures 9 thru 14 and are described and analyzed in the detail.

Interactive User workload test results with no UBE load

Testing did not go beyond 2000 users as it is clear that system load levels are minimal at even these load levels and provide plenty of headroom for additional consolidation or corner case customer workloads. In the case of the UBE batch workload testing, an interactive load of 1000 users was applied and the amount of UBE workloads launched was ramped from approximately 5000 UBE's per hour to over 10000 UBE's per hour and again load levels remained low and indicate substantial headroom for consolidation or high customer UBE workloads.

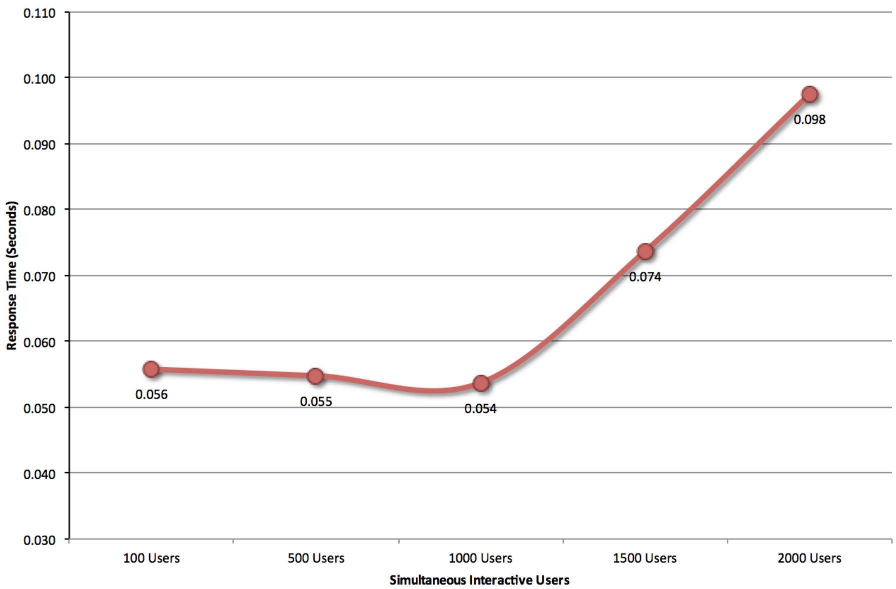

Figure 9 : Pure Interactive User average Response Times in seconds at various user counts, no UBE load

Figure 9 shows Interactive response times with no UBE workloads for between 500 and 2000 users. It can be seen that the minimum response time is constant at loads below 1500 users and indicates the minimum response time of the overall architecture, loads above 1000 users begin to load the system and demonstrate a linear increase in response time up to 0.098 seconds for 2000 users.

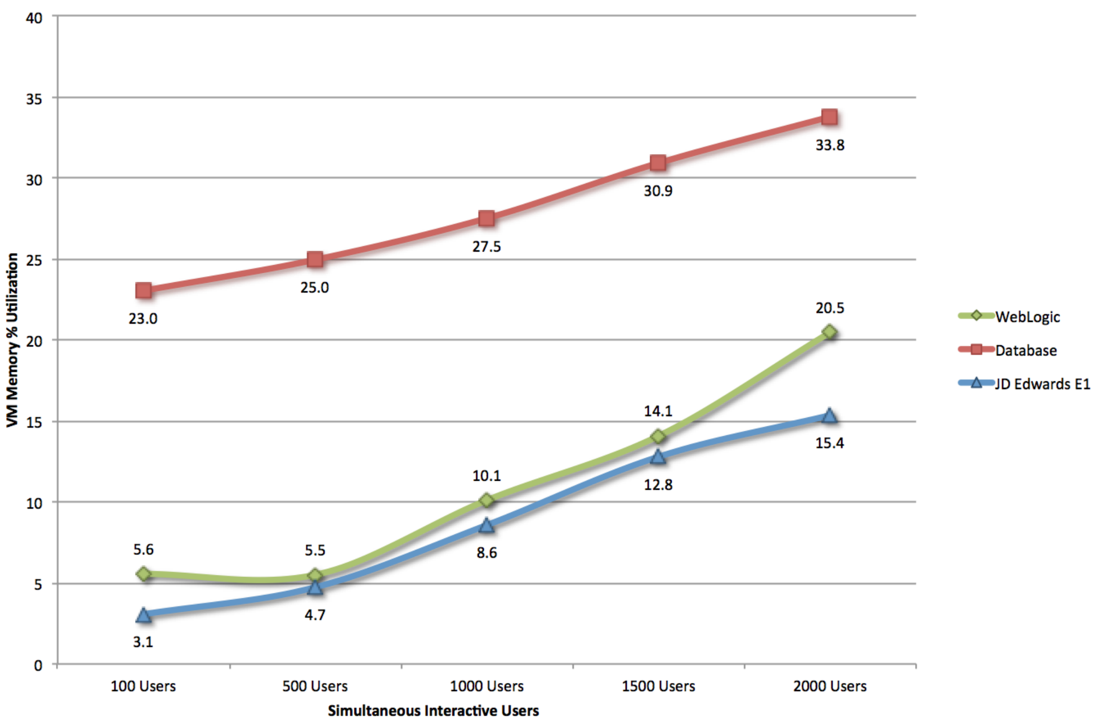

Figure 10 : Interactive Workload Server Memory Utilization, no UBE load

In Figure 10, the memory utilization of the Database, JDE E1 and WebLogic server virtual machines is shown, no UBE workload is present in this test. It can be seen that the memory utilization rats are very low rising to a maximum of 33.8% for the Database domain at the 2000 user test level. These numbers indicate significant headroom is available on the platform even in the tested 256GB configuration, especially bearing in mind that the T5-2 platform can be upgraded to accommodate up to 1TB or RAM.

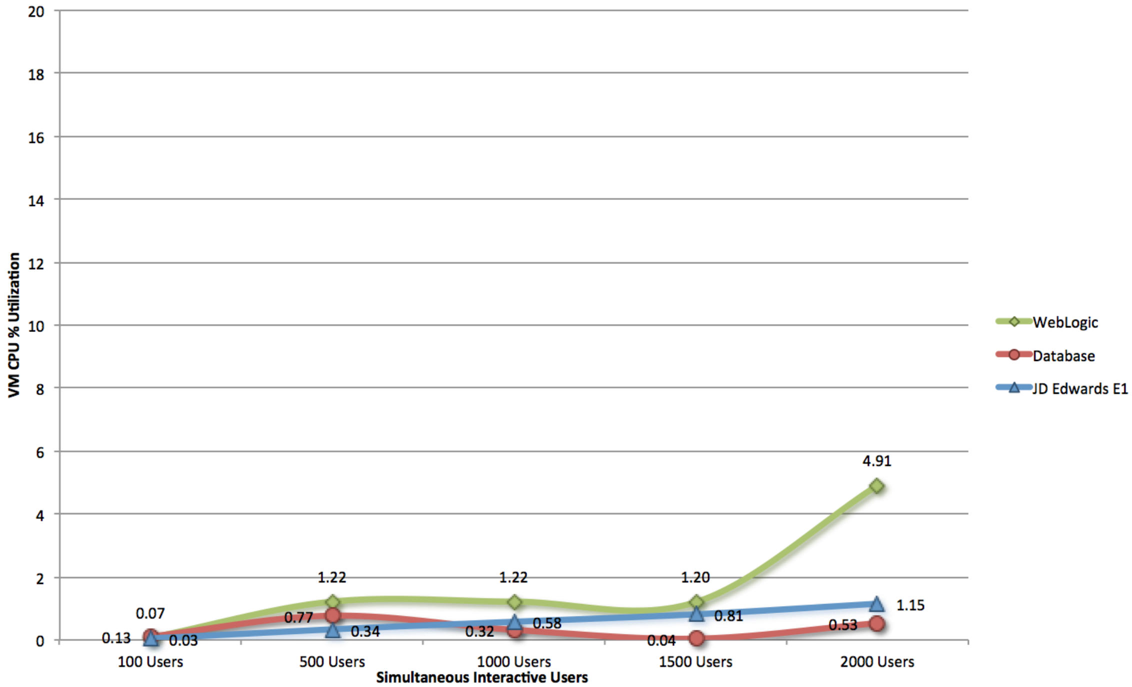

Figure 11 : Interactive User Server CPU Utilzation during testing, no UBE load

Figure 11 shows the utilization of CPU resources in each of the Database, JDE E1 and WebLogic virtual machines during Interactive User load testing. These utilization numbers are very low with a maximum load of no more than 4.91% in the WebLogic Zone indicating significant headroom for further consolidation or customer specific workloads. Utilization numbers this low are essentially showing quiescent servers, and indicate that the number of CPU's dedicated to the domain hosting the JDE E1 and WebLogic non-global zones could be reduced to minimize Oracle Solaris Cluster licensing costs in the solution. However a customer may consider using the existing Clustered Zone infrastructure to host other consolidation zones in in a clustered fashion, essentially leveraging the already configured cluster to benefit their additional consolidated workloads.

Interactive User with UBE's workload test results

UBE workloads are tested while maintaining a 'typical' Interactive user load of 1000 users, the goal being to maintain the Interactive user load with typical response times below 0.5 seconds, while ramping up UBE workload and monitoring the effect on response times, CPU utilization and Memory Utilization. The UBE workload consists of a mixture of both short and long running UBE tasks that are kicked off by scripts and maintain a strady mix of Long running UBE's (that run throughout the test interval), and a constantly refreshed queue of short running UBE's that are started and tracked to completion until the goal throughput has been achieved over the 1 hour steady state of the test interval.

The following charts demonstrate UBE loads from 5307 UBE's per hour (88 UBEs/min) up to 10277 UBEs per hour (171 UBEs/min).

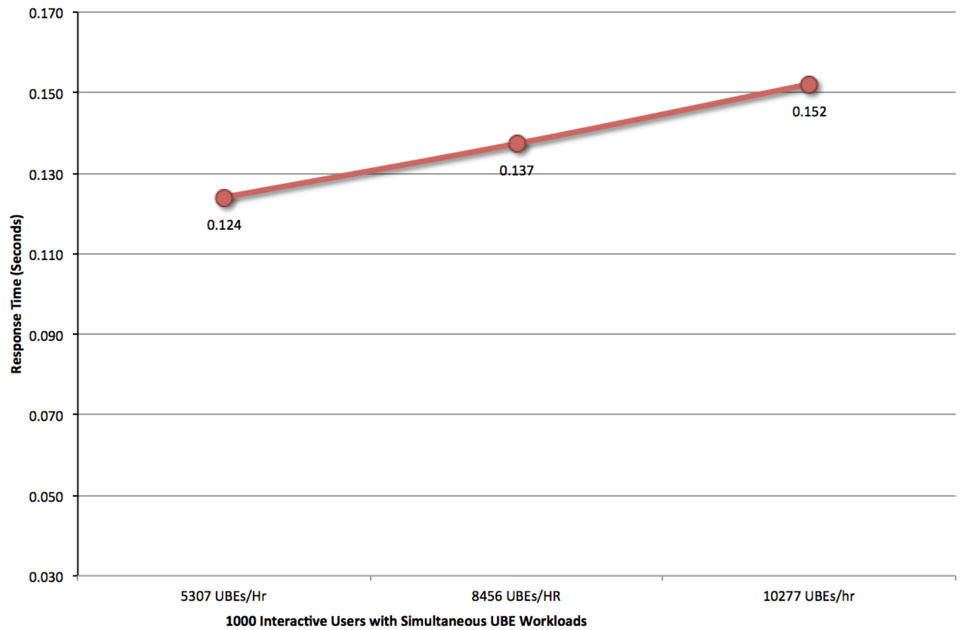

Figure 12 : Interactive Response time for 1000 users during UBE load testing

Figure 12 shows slightly increased average response times during significant UBE load testing, between 0.124 seconds and 0.152 seconds as UBE load increases from 5207 to 10277 UBEs per hour. This demonstrates significant throughput of UBEs and at the same time demonstrates considerable headroom is still available for either additional consolidation or increased UBE workload on the platform.

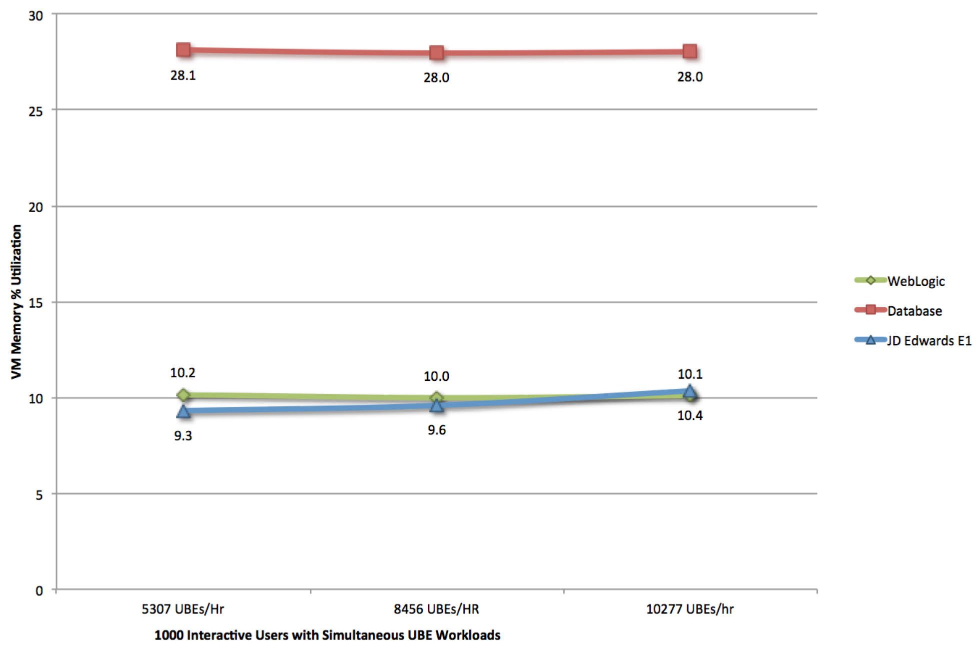

Figure 13 : Server Memory Utilization during UBE workload testing

Figure 13 shows the memory utilization for the DB, JDE E1 and WebLogic virtual machines during the UBE load testing. A maximum memory utilization of 28.1% was seen in the DB domain demonstrating considerable headroom for additional load or consolidation.

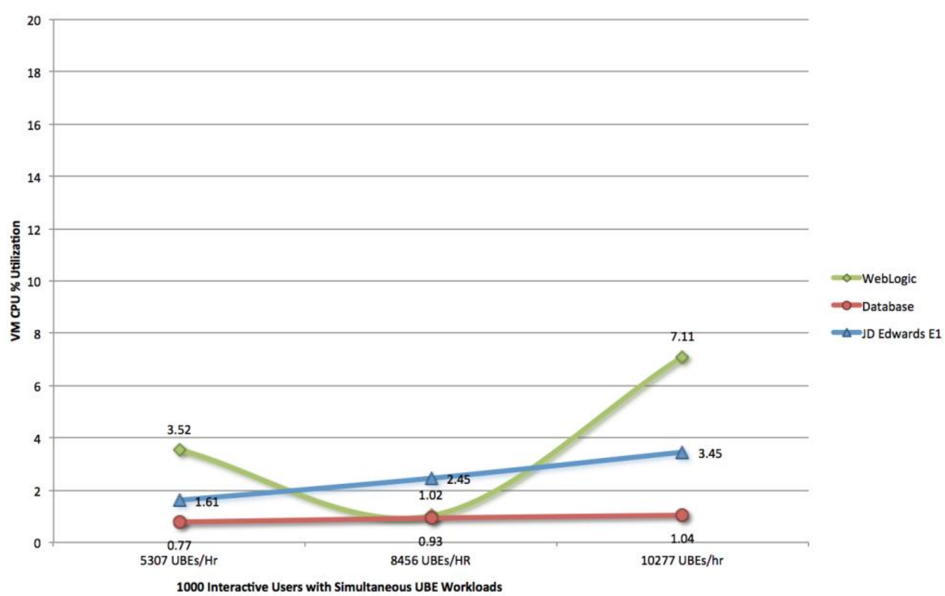

Figure 14 : Server CPU utilization during UBE load testing

Figure 14 shows CPU utilization for the DB, JDE E1 and WebLogic virtual machines during UBE load testing. CPU loads remained low during testing with a maximum of 7.11% used on the WebLogic Zone, this load varied throughput testing and is thought to be a result of load balancing issues with the test configuration, even so loads are very low and indicate considerable headroom for additional consolidation or workload.

Failure Scenarios

Failures can take many forms and protecting the operation of the application during such failures can be a challenge. The first step in protection consists of duplication of hardware hardening and/or duplication. By choosing proven duplicated SPARC platforms and best practice duplicated switches, together with resilient ZFSSA storage servers hardware reliability is assured.

Once hardware resiliency and duplication is assured, software must be configured to take advantage of the hardware to transfer services from failed hardware to alternate devices. Clustering software technology is deployed here to assure availability of the Database and also the JDE E1 and WebLogic tiers of the solution.

In most common failure scenarios no impact is expected to the operation of the JD Edwards services with failures being masked via duplication and appropriate multi-pathing.

In the case of less common and more severe failures some downtime is unavoidable in HA architectures as the faults are recognized and failover and recovery actions are automatically instituted by the clustering technology. In these situations customer can expect s brief period of lack of response from the service followed at worst by a need to log back into the service and possible re-submit UBE batch jobs that will be flagged as incomplete or failed in their normal workflow. In these worst case scenarios it is expected that this brief period of downtime will be resolved within 3-5 minutes.

» Minor failures have no noticeable impact on performance and just require timely correction to reduce the possibility of additional failures causing more serious problems (such as failed cables, power supplies, network devices or disks)

» Major failures have a noticeable impact on the users' experience and might require manual corrective actions to recover fully. Example include storage head or shelf failures, virtual or physical server failures, and network switch failures.

Final Analysis

The SPARC T5-2 cluster is extremely capable of handing the workloads described in this paper, and at the current testing load levels, no upper limit has been discovered in the capability of the platform. To achieve higher throughput would require further development of the testing environment, to allow it to scale to provide higher levels of load to the system under test.

At no time during testing were the CPU or memory of the servers more than 8% percent (CPU) or 35% percent (memory) utilized. This indicates that even in the tested configuration with only 256GB of installed memory in each T5-2 significant headroom is available in the platform for additional workload growth through added JDE modules or through consolidation of other workloads onto the platform, perhaps even taking advantage of the already configured HA Zone Cluster environment to provide these additional workloads with clustered redundancy.

Real-World Sizing Based on the DIL Test Workload

To account for the fact that the DIL workload cannot accurately predict real-world customer customizations or individual performance requirements, and the fact that during DIL testing very low system loads were seen it is reasonable to size the real world loads on the system based on the DIL testing results.

Real world sizing therefore accommodates up to 2,000 interactive user configurations based on testing carried out at the 2000 users level under the DIL workload. Such sizing leaves significant headroom on the platform to cope with actual deployment loads with real customer JD Edwards EnterpriseOne deployments incorporating additional JDE modules and customizations that are typical.

Room for Consolidation

The primary benefit that can be realized as a result of the available substantial headroom is that capacity is available for customers to increase utilization of the platform by adding other workloads. Such workloads do not need to be related to JD Edwards.

Such additional consolidation, if undertaken by users, should consider the full range of JD Edwards EnterpriseOne workload conditions that are possible in the customer's actual deployment to ensure that acceptable performance under worst-case conditions can be maintained before additional consolidation is performed. If such worst-case load conditions 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

The tested configuration of servers and storage has some considerable flexibility for modification based on a customer's specific performance, cost, or footprint requirements. When diverging from the tested configuration, the following guidelines and caveats are proposed. For more detailed sizing information, contacting a local Oracle representative is recommended.

The configurations outlined below refer to the systems and configurations available at the time of the publication of this document[1].

NOTE: When sizing a JD Edwards EnterpriseOne deployment based on DIL workload testing, it is important to understand that DIL interactive users are approximations of real-world users and they stress only the core modules of JD Edwards EnterpriseOne. In real-world scenarios, each individual deployment of JD Edwards EnterpriseOne will have unique workload characteristics, employ additional JD Edwards EnterpriseOne modules, or have non-standard use cases that might significantly alter the workload characteristics seen in individual deployments. This is especially true if users have heavily customized their JD Edwards EnterpriseOne environments. These sizing guides are provided on that basis and are intended as indicators only.

Large Configuration—Up to 2000 Real-World Users

TABLE 14. RECOMMENDED HARDWARE COMPONENTS FOR THE LARGE CONFIGURATION

Server Variant	Component	Configuration
SPARC T5-2 Server	Processor: SPARC T5, dual 3.6 GHz, 16-core, 8 threads/core	Factory pre-installed.
	Memory: 256 GB RAM, utilizing 16 x 16 GB DIMMs	As indicated in service manual.

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 the 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.

	Disks: 6 x 600 GB SAS2 2.5-inch disk drives	Install disks in all available slots.
	2 x PCIe card: Sun Dual Port 10 GbE PCIe 2.0 ExpressModule, Base-T card from Oracle	As indicated in service manual and installation guide. PCIe slot locations are recommended in the Implementation Guide for optimal card performance and availability
Oracle ZFS Storage ZS3-2 Appliance	2 x controller: 2 x Intel 2.1 GHz 8-core, 256 GB RAM, cluster-capable, 2 x SAS HDD, 2 x SSD 1.6 TB read flash SSD	Factory pre-installed.
	2 x storage shelves: Oracle Storage Drive Enclosure DE2-24P storage shelves, 20 x 900 GB HDDs, 4 x write flash SSD	As indicated in service manual.
10 GbE Oracle Switch ES1-24	2 x 10 GbE Oracle Switch ES1-24	As indicated in Solution Implementation Guide.

The large configuration easily accommodates upwards of 2000 real-world users (as tested at 2000 DIL interactive users) and UBE batch throughput of over 10000 UBE's per hour.

Medium Configuration—Up to 1500 Real-World Users

TABLE 15. RECOMMENDED HARDWARE COMPONENTS FOR THE MEDIUM CONFIGURATION

Server Variant	Component	Configuration
SPARC T5-2 Server	Processor: SPARC T5, dual 3.6 GHz, 16-core, 8 threads/core	Factory pre-installed.
	Memory: 256 GB RAM, utilizing 16 x 16 GB DIMMs	As indicated in service manual.
	Disks: 6 x 300 GB SAS2 2.5-inch disk drives	Install disks in all available slots.
	2 x PCIe card: Sun Dual Port 10 GbE PCIe 2.0 ExpressModule, Base-T card from Oracle	As indicated in service manual and installation guide. PCI slot locations are recommended in the Implementation Guide for optimal performance and availability
Oracle ZFS Storage ZS3-2 Appliance	2 x controller: 2 x Intel 2.1 GHz 8-core, 256 GB RAM, cluster-capable, 2 x SAS HDD, 1 x SSD 1.6 TB read flash SSD	Factory pre-installed.

	2 x storage shelves: Oracle Storage Drive Enclosure DE2-24P storage shelves, 20 x 300 GB HDDs, 2 x write flash SSD	As indicated in service manual.
10 GbE Oracle Switch ES1-24	2 x 10 GbE Oracle Switch ES1-24	As indicated in Solution Implementation Guide.

The medium configuration accommodates upwards of 1500 real-world users (1500 DIL interactive users) and UBE batch workloads of up to 7500 UBE's per hour.

Small: Configuration—Up to 1000 Real-World Users

TABLE 16. RECOMMENDED HARDWARE COMPONENTS FOR THE SMALL CONFIGURATION

Server Variant	Component	Configuration
SPARC T5-2 Server	Processor: SPARC T5, single 3.6 GHz, 16-core, 8 threads/core	Factory pre-installed.
	Memory: 128 GB RAM, utilizing 8 x 16 GB DIMMs	As indicated in service manual.
	Disks: 6 x 300 GB SAS2 2.5-inch disk drives	Install disks in all available slots.
	2 x PCIe card: Sun Dual Port 10 GbE PCIe 2.0 ExpressModule, Base-T card from Oracle	As indicated in service manual and installation guide. PCI slot locations are recommended in the Implementation Guide for optimal performance and availability
Oracle ZFS Storage ZS3-2 Appliance	2 x controller: 2 x Intel 2.1 GHz 8-core, 256 GB RAM, cluster-capable, 2 x SAS HDD	Factory pre-installed.
	2 x storage shelves: Oracle Storage Drive Enclosure DE2-24P storage shelves, 20 x 300 GB HDDs	As indicated in service manual.
10 GbE Oracle Switch ES1-24	2 x 10 GbE Oracle Switch ES1-24	As indicated in Solution Implementation Guide.

The small configuration accommodates upwards of 1000 real-world users (1000 DIL interactive users) and UBE batch workloads up to 5000 UBE's per hour.

Sizing Configuration Notes

The Large configuration was tested and achieved the performance discussed in this document.

The Medium configuration reduces disk quantities in the Oracle ZFS Storage ZS3-2 appliance for both write- and read-optimized flash devices.

The small configuration utilizes the smallest SPARC T5-2 server configuration available, which has only a single SPARC T5 processor installed and, as a consequence, provides only 128 GB of RAM in the chosen small server configuration. Additionally, the Oracle ZFS Storage ZS3-2 appliance's configuration completely removes both read- and write-acceleration SSDs from the configuration. If HA can be compromised in the customer's environment, it would be possible to remove one storage controller and one storage shelf from the configuration to substantially reduce cost, but at the expense of removing HA from the storage components of the solution.

Conclusion

Testing shows that the combination of a single SPARC T5-2 server from Oracle and an Oracle ZFS Storage ZS3-2 appliance can provide extremely good performance with very high levels of high availability for JD Edwards EnterpriseOne deployments up to at least 2,000 real-world users while maintaining capacity for additional load or consolidation on the platform. During testing, performance and availability optimization was demonstrated through careful arrangement of the I/O subsystems, networking, and clustering technologies involving the use and documentation of many best practices. This enabled a cost-effective, reliable, low-risk solution, maximizing customer ROI and minimizing TCO.

References

TABLE 17. REFERENCES

Description	URL
JD Edwards EnterpriseOne product page	http://www.oracle.com/us/products/applications/jd-edwards-enterpriseone/index.html
SPARC T5-2 server product page	http://www.oracle.com/us/products/servers-storage/servers/sparc/oracle-sparc/t5-2/overview/index.html
Oracle ZFS Storage ZS3-2 appliance product page	https://www.oracle.com/storage/nas/zs3-2/index.html
Oracle Switch ES1-24 product page	http://www.oracle.com/us/products/networking/ethernet/switch-es1-24/overview/index.html
Oracle's JD Edwards Partner Portal	http://www.go9withjde.com/
JD Edwards EnterpriseOne documentation	http://www.oracle.com/technetwork/documentation/jdedent-098169.html

Appendix A: Additional Workload Details

There are two primary types of 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—These are short- and long-running UBE batch jobs that are run during the test to increase the realism of the workload.

The details for each of these types of 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, as well as manufacturing processes such as inventory tracking. These are sometimes generically referred to by industry terms such as supply chain management (SCM), customer relationship management (CRM), human capital management (HCM), supplier relationship management (SRM), and financial management solutions (FMS).

The DIL workload uses the JD Edwards EnterpriseOne interactive transaction types described in Table 18.

TABLE 18. JD EDWARDS ENTERPRISEONE INTERACTIVE TRANSACTIONS

	Transaction Type	Description
1	P03B102	Apply Receipts
2	P0411I	Supplier Ledger Inquiry
3	P051191	Daily Time Entry
4	P17500	Case Management Add
5	P31114	Work Order Completion
6	P3411	MRP Messages (WO Orders)
7	P3411	MRP Messages (OP Orders)
8	P3411	MRP Messages (OT Orders)
9	P4113	Inventory Transfer
10	P42101	Sales Order Entry – 10 Line Items
11	P42101	Sales Order Update
12	P4310	Purchase Order Entry – 25 Line Items
13	P4312	Purchase Order Receipts
14	P4314	Voucher Match
15	P4915	Ship Confirmation – Approval Only
16	P4915	Ship Confirmation – Confirm/Ship Only

17	P4915	Ship Confirmation – Confirm and Change Entry

The short-running UBEs generate consistent load on the JD Edwards EnterpriseOne server and are continuously resubmitted upon completion. Long-running UBEs tend to generate substantial load on the database server, because they typically involve more complex SQL queries. Once started, these long-running UBEs continue to run throughout the test runs. Test runs typically last one hour, some long-running UBEs may not complete during this 1 hour test.

Such a combination of interactive, short-running UBE and long-running UBE workloads provide a good workload, stressing all aspects of the solution's configuration. Table 19 and Table 20 list the transaction types for the long-running and short-running UBE jobs, respectively.

TABLE 19. LONG-RUNNING UBE JOBS

Transaction Type	Description
R42565	Print Invoices
R31410	Work Order Processing
R3483	MRP Processing
R43500	Purchase Order Print

TABLE 20. 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
R01402W	Who's Who Report
R03B155	A/R Summary Analysis
R03B31	Activity Log Report
R41411	Select Items Cost Count
R42072	Price Category Print

Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

Oracle Optimized Solution for JD Edwards EnterpriseOne, Delivering High Availability,
Performance, and Efficiency Using Oracle's SPARC T5-2 Servers
June 2015
Author: Peter Wilson
Contributing Authors: Laura Verdibello