ORACLE
OPTIMIZED SOLUTIONS

An Oracle White Paper
July 2013

Increasing the Performance and Efficiency of Siebel CRM — A Technical White Paper

ORACLE

Executive Overview

Founded on a service-oriented architecture, Oracle's Siebel Customer Relationship Management (Siebel CRM) software allows businesses to build scalable standards-based applications that can help to attract new business, increase customer loyalty, and improve profitability. As companies deliver more comprehensive and rich customer experiences through CRM tools, demand can scale rapidly, forcing data centers to expand system resources quickly to meet increasing workloads. Data center resources can be scaled horizontally (with more servers added at each tier), vertically (by adding more powerful servers), or both. As servers are added at the Siebel Web Clients, Web Server, Gateway/Application Server, and Database Server tiers, a frequent result is server sprawl. Over time, this can result in negative consequences — greater complexity, poor utilization, increased maintenance fees, and skyrocketing power and cooling costs.

Consolidating tiers is one approach that can help to contain server sprawl and reduce costs. Recognizing the need to grow efficiently while scaling Siebel CRM capabilities, Oracle created the Oracle Optimized Solution for Siebel CRM, which consolidates the Web, Gateway, Application, and Database tiers on a single SPARC T-Series server from Oracle, limiting the number of physical machines needed to effectively deploy applications and improving the bottom line. As shown in testing exercises using a well-known Siebel CRM workload and virtualization technologies built into SPARC T-Series servers, the solution scales easily to accommodate user load.

Because Siebel CRM applications support business profit centers, they often operate under stringent availability requirements and necessitate demanding service levels. For this reason, Oracle engineers designed this consolidated solution with high availability (HA) capability. Software tiers were again consolidated using built-in virtualization technologies — with a clustered server configuration at the virtual level that provides HA.

Introduction

To safely and securely consolidate Siebel CRM application tiers, SPARC T-Series servers offer a choice of built-in, no-cost virtualization technologies:

- Oracle Solaris Containers. Containers are an integrated virtualization mechanism that can isolate application services within a single Oracle Solaris instance. Faults in one container have no impact on applications or service instances running in other Containers.

- Oracle VM Server for SPARC. Native to Oracle's SPARC T-Series processors, this hypervisor technology allows multiple tiers to be consolidated within isolated domains, without imposing additional cost. Each domain runs an independent copy of Oracle Solaris, and there are no licensing fees for additional OS copies.

Using one or both of these virtualization technologies, Siebel CRM services in each tier can run in isolation, without impacting service execution in other tiers. System resources can be allocated and reassigned to each tier as needed. Compared to other competitive and proprietary virtualization technologies, using Oracle Solaris Containers, Oracle VM Server for SPARC, or both can provide significant cost savings when consolidating a Siebel CRM infrastructure. In addition, Oracle guarantees binary compatibility for applications running under Oracle Solaris, whether the OS runs natively as the host OS or as a guest OS in a virtualized environment.

Oracle engineers configured different Siebel CRM tiers in virtualized environments on SPARC T-Series servers. In the test, engineers consolidated tiers on a single server, configuring each Siebel CRM tier in a separate container or domain. To facilitate HA capability, engineers implemented Oracle Solaris Cluster (which supports both Containers and domains) on two SPARC T-Series servers to simulate mission-critical Siebel CRM application workloads in a consolidated yet resilient virtualized environment.

The test workload was extracted from the well-established Siebel Platform Sizing and Performance Program (PSPP) benchmark, which simulates real-world environments using some of the most popular Siebel CRM modules. Engineers looked at system resource utilization, response time, and throughput metrics as they scaled the number of users under typical application workloads. This paper shows the test results and clearly documents best practices, which can help system architects more effectively size and optimize the Siebel CRM application on SPARC T-Series servers.

The test results demonstrate how no-cost virtualization technologies in SPARC T-Series servers — combined with Oracle Solaris Cluster software — can optimize scalability while reducing data center complexity, lowering operating costs, and delivering high availability for business-critical CRM services.

Key Solution Technologies

The tested solution was based on Oracle's massively scalable SPARC T-Series servers, the Oracle Solaris 11 operating system, and Oracle's storage technologies, as shown in Figure 1**Error! Reference source not found.**. Built-in, no-cost virtualization technologies — Oracle Solaris Containers or Oracle VM Server for SPARC — reside at the heart of the solution architecture and enable a flexible infrastructure for consolidation. Oracle Solaris Cluster (and often third-party management tools) is typically added to enhance business continuity and simplify resource allocation tasks for virtualized environments.

Oracle engineers constructed a full solution based on a pair of SPARC T5-2 servers, which each features SPARC T5 multithreading processors. With such advanced thread density; a single SPARC T5-2 server is a powerhouse for consolidating a Siebel CRM infrastructure. To demonstrate this point, Oracle engineers ran a series of scalability tests using both container and domain virtualization technologies. As the test results show, the consolidated solution on a single SPARC T-Series server exhibited good scalability, providing reasonable response times and high throughput rates.

In SPARC T-Series servers, chip-multithreading (CMT) technology in SPARC T-Series processors enables effective scalability. CMT technology applies the available transistor budget to achieve up to eight cores within a single processor. Each core can switch between threads on a clock cycle, helping to keep the processor pipeline active while lowering power consumption and heat dissipation. Because of the advanced thread density, the SPARC T-Series server scales well to provide headroom to support growth while minimizing power use.

Oracle engineers used a clustered configuration of two SPARC T-Series servers. Each SPARC T-Series server houses two SPARC T5 processors for a maximum of 128 threads per server. In an economical clustered configuration like that used in the HA testing model, two servers provide enough multithreading capability to minimize degradation during failover operations. The clustered configuration also demonstrated good scalability, reasonable response times, and high levels of throughput, at the same time enabling highly available Siebel CRM application services.

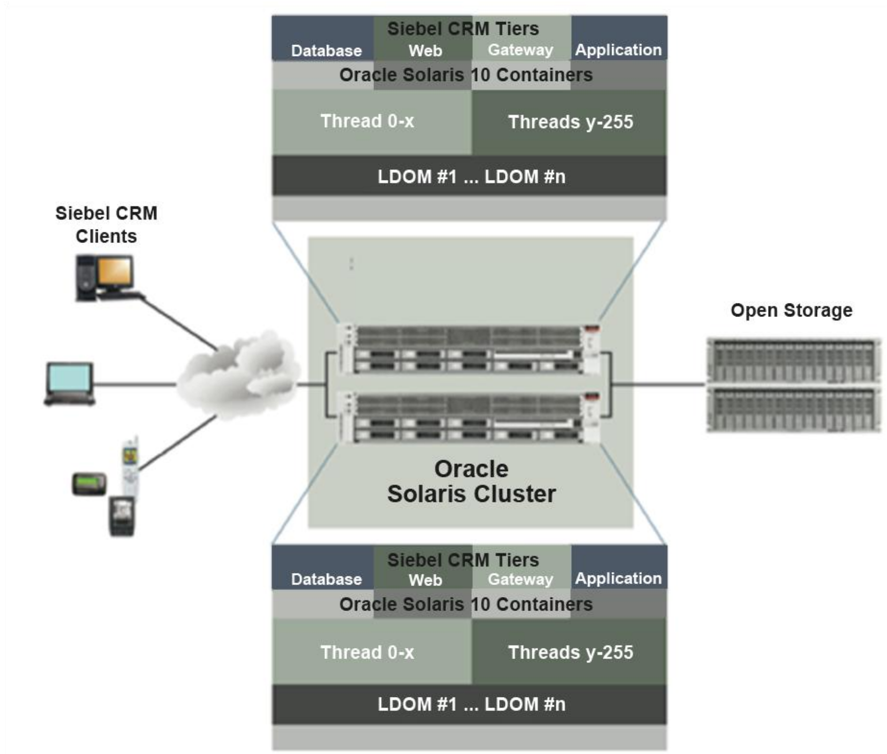

Figure 1. Oracle engineers implemented Oracle Solaris Cluster on two SPARC T-Series servers in a consolidated, clustered HA configuration.

An Overview of Oracle's Siebel CRM Application Architecture

The Siebel CRM application suite includes the following tiers (see Figure 2):

- *Web Clients*. Web clients provide user interface functionality and can encompass a variety of types (Siebel Web Client, Siebel Wireless Client, Siebel Mobile Web Client, Siebel Handheld Client, and so on). Oracle Application Testing Suite simulated the load generated by the different sized end-user populations.
- *Web Server*. This tier processes requests from Web clients and interfaces to the Gateway/Application Server tier. In the scalability testing that was performed, Oracle engineers installed the Siebel Web Server Extension and configured the Oracle iPlanet Web Server at this tier.
- *Gateway/Application Server*. This tier provides services on behalf of Siebel Web Clients. It consists of two sublayers: the Siebel Enterprise Server and the Siebel Gateway Server.
- *Database Server*. While the Siebel file system stores data and physical files used by Siebel CRM Web and dedicated clients and Siebel Enterprise Server, the Siebel Database Server stores Siebel CRM database tables, indexes, and seed data.

- In a multiple-server deployment, the Siebel Enterprise Server includes a logical grouping of Siebel servers. (However, in a small configuration, the Siebel Enterprise Server might contain a single Siebel server.) The Siebel Gateway coordinates the Siebel Enterprise Server and its set of Siebel servers. It also provides a persistent backing store of Siebel Enterprise Server configuration information.

Each Siebel server is a flexible and scalable application server that supports a variety of services such as data integration, workflow, data replication, and synchronization services for mobile clients. The Siebel server also includes logic and infrastructure for running different Siebel CRM modules, as well as providing connectivity to the Database Server. The Siebel server consists of several multithreaded processes that are commonly known as Siebel Object Managers.

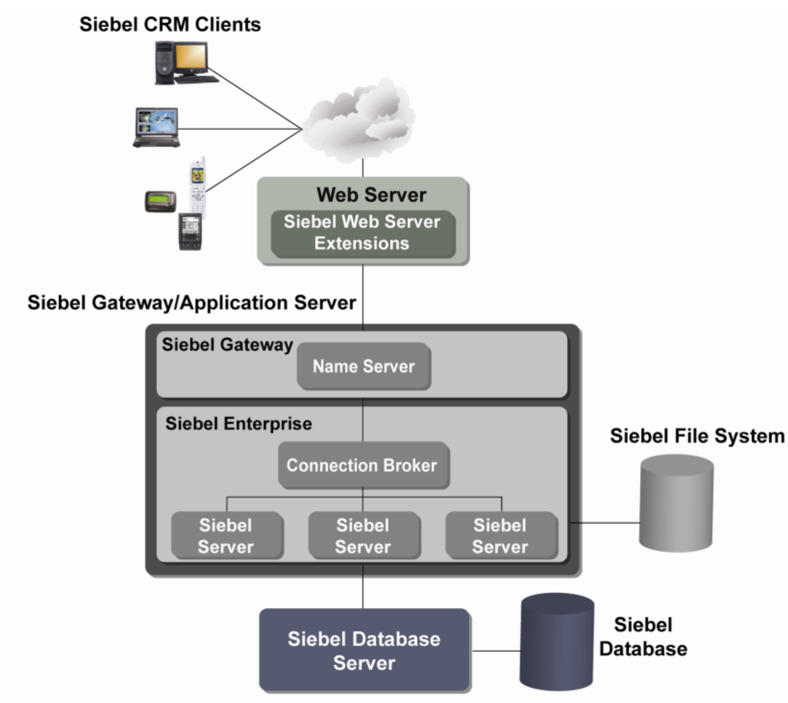

Figure 2. This high-level overview of the Siebel CRM application architecture shows the tiered software architecture.

To provide high availability to all four tiers of Oracle Siebel CRM 8, Oracle Solaris Cluster software is deployed to support mission-critical application availability (see "Configuring for HA Using Oracle Solaris Cluster Software", page 11). During testing, engineers analyzed performance and scalability with Siebel CRM workloads in an HA configuration, using clustered Containers to support each software tier.

Workload Description

CRM systems often require customization — typically, more frequently than other business applications. Common changes include adding or removing certain application modules, modifying the function of existing modules, or integrating the CRM application with other business applications and processes. While application performance varies according to the particulars of any deployment, testing a configuration's scalability with a well-defined workload helps to provide a useful starting point for defining appropriate configurations and sizing.

For the purposes of scalability testing, engineers used a workload extracted from the well-known Siebel Platform Sizing and Performance Program (PSPP) workload. This workload is based on scenarios derived from large Siebel customers and replicates real-world, concurrent, thin-client requirements of typical end users. The PSPP 8.1.1.4 workload is based on user populations who repeatedly perform functions that can be performed with Siebel Financial Services Call Center. The Siebel Financial Services Call Center software provides a comprehensive solution for sales and service, helping customer service and telesales representatives to provide world-class customer support, improve customer loyalty, and increase revenues through cross-selling and up-selling opportunities.

In the testing, the PSPP workload simulates the following task mix for the functions listed above:

Siebel Financial Services Call Center Business Transactions was used to execute three complex business transactions simultaneously for 30,000 concurrent users. The ratios of these three call center scenarios were 30%, 40%, and 30% each. Between each user operation and the next one, the "think time" (a synthetic delay simulating the typical pause between a user's actions) averaged approximately 10, 13, and 35 seconds, respectively.

The following is a high-level description of the use cases tested using the Incoming Call Creates Opportunity, Quote, and Order (CC1) test script:

- Create a new contact, create a new opportunity for that contact, add two products to the opportunity, navigate to opportunities in the Quotes View.
- Click the **AutoQuote** button to generate a quote. Enter the quote name and price list.
- Drill down on the quote name to go to the Quote - Line Items View and specify a discount.
- Click the **Reprice All** button.
- Update the opportunity.
- Navigate to the Quotes - Orders View. Click the **AutoOrder** button to automatically generate an order.
- Navigate back to the opportunity.

The following is a high-level description of the use cases tested using the Incoming Call Creates Service Request (CC2) test script:

- Create a new service request.
- Associate the contact and account for that service request.
- Click the **Verify** button to bring up the pick applet.
- Select **Entitlement**.
- Query and select the policy.
- Select the product and add the product to the service request.

- Save the service request.
- Go to the Service Request Activity Plan.
- Select the activity plan and save the service request.

The following is a high-level description of the use cases tested using the Incoming Call Updates Service Request (CC3) test script:

- Click the **Service Screen** tab.
- Go to My Service Request.
- Drill down on the service request and go to SR Activity.
- Navigate to SR – Related SR.
- Select **Add Solution** and save the service request.
- Update the service request, set it to pending, and save it.
- Navigate back to Service Request Activity.

Business Transaction Types

Based on the Siebel PSPP benchmark workload described above, Oracle Application Testing Suite generated loads to simulate different user populations while simultaneously executing complex business transactions. Between each user operation, “think time” averaged approximately 15 seconds. The following paragraphs characterize core business transaction types used in the testing.

Web Services — Find and Then Submit a New Service Request and Update the Service Request

This transaction simulates a Web service that interfaces to a hypothetical legacy application to find or create a service request. The Web service acts as a delivery mechanism for integrating heterogeneous applications through internet protocols. A Web service can be specified using Web Services Description Language (WSDL) and is then transported via Simple Object Access Protocol (SOAP), a transport protocol based on XML. Since the PSPP benchmark suite has no UI presentation layer, the load generator simulates a Java Platform, Enterprise Edition (Java EE) Web application to send a Web service request to a Siebel server (`EAIObjMgr_enu`) to invoke Siebel CRM business services.

The Siebel Web Services framework generates WSDL files to describe the Web services hosted by the Siebel CRM application. Also, this framework can call external Web services by importing a WSDL document as an external Web service (using the WSDL import wizard in Siebel Tools). Each Web service exposes multiple methods, such as Query Service Request, Create Service Request, and Update Service Request.

Web service authentication is done through a session token. The ServerDetermine session type is used and a session token is maintained to avoid a Login process for each request. To use the ServerDetermine session type, a login Web service call (`SessionAccessPing`) retrieves the session token before calling other Web services. At the end of the transaction, a logout call (`SessionAccessPing`) makes the session token unavailable.

Test Environment

The test environment was designed to determine scalability and availability using a clustered configuration of two SPARC T5-2 servers. These test environments are not representative of typical production deployments but are simplified proof-of-concept configurations designed for test and development.

HA Test Environment

Figure 3 shows the HA test environment.

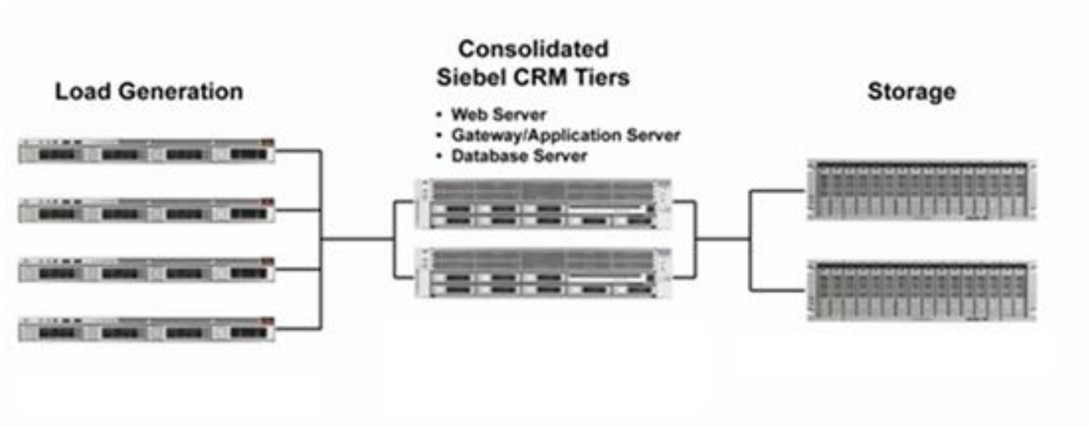

Figure 3. The HA test environment implemented Siebel CRM tiers on two clustered SPARC T-Series servers.

The test used the following hardware and software components:

- Hardware
  - Two SPARC T5-2 servers, each with two SPARC T5-2 processors and 256 GB of RAM
  - A Sun ZFS Storage Appliance 7320
  - Four servers from Oracle for load generation
- Software
  - Oracle Solaris 11 (SPARC)
  - Oracle Database 11*g* Release 2 for the database server
  - Siebel CRM Release 8.1.1.4 Industry Applications
  - Oracle iPlanet Web Server 7.0u9
  - Oracle Solaris Cluster 3.2u3

Testing High Availability — Implementing HA in a Consolidated Environment

Highly available clusters provide nearly continuous access to data and applications by keeping systems running through failures that would normally bring down a single server. In mission-critical clustered systems, no single failure — whether it is a hardware, software, or network failure — can cause a cluster to fail. Recognizing the need to keep business-critical Siebel CRM applications up and running (and to support disaster planning scenarios), Oracle conducted testing using a clustered HA configuration for Siebel CRM 8.1.1.4 workloads. Oracle's clustering products — in particular, Oracle Solaris Cluster software — enable highly available solutions that can meet stringent business continuity requirements for Siebel CRM deployments.

Configuring for HA Using Oracle Solaris Cluster Software

A cluster is two or more servers (or nodes) that work together as a single, continuously available system to provide applications, system resources, and data to users. Each cluster node is a fully functional standalone system. However, in a clustered environment, an interconnect bridges the nodes, which work together as a single entity to provide increased availability and performance. The interconnect carries important cluster information (data as well as a heartbeat) that allows cluster nodes to monitor the health of other cluster nodes. High availability using clustered systems is achieved through a combination of both hardware and software.

Oracle Solaris Cluster software enables business continuity and global disaster recovery solutions to meet evolving data center needs. In a nutshell, the clustering software

- Makes use of proven availability and virtualization features in Oracle Solaris 11 and in SPARC processor-based systems, including those in SPARC servers
- Supports an industry-leading portfolio of commercial applications, including Oracle Database, Siebel CRM, and Web server technologies
- Is certified with a broad range of storage arrays and SPARC and x64/x86 platforms

The most recent release of Oracle Solaris Cluster software implements high availability for consolidated environments that use container or domain virtualization technologies, such as the Siebel CRM proof-of-concept solution described in this paper. Oracle Solaris Cluster software supports Oracle Solaris Containers for fault isolation, security isolation, and resource management. Oracle Solaris Cluster can also help to protect virtualized environments that use Oracle VM Server for SPARC domains, lowering risk for servers that provide multiple application services.

When consolidating Siebel CRM tiers in this way, Oracle Solaris Cluster provides high availability agents to monitor components running in different virtualized environments (see Table 1). Available Oracle Solaris Cluster agents include software to support services such as Oracle Database, Siebel services, NFS, DNS, the Oracle iPlanet Web Server, the Apache Web Server, and so forth. Oracle Solaris Cluster software provides configuration files and management methods to start, stop, and monitor these application services.

TABLE 1. ORACLE SOLARIS CLUSTER AGENTS

SOLUTION COMPONENT	PROTECTED BY
Web Server	Oracle Solaris Cluster HA for Oracle iPlanet Web Server
Siebel Gateway	Oracle Solaris Cluster HA for Siebel (resource type: SUNW.sblgtwy)
Siebel Server	Oracle Solaris Cluster HA for Siebel (resource type: SUNW.sblsrvr)
Oracle Database	Oracle Solaris Cluster HA for Oracle Database

Figure 4 depicts the HA proof-of-concept configuration used as the basis of the testing. The HA configuration uses Oracle Solaris Cluster's Zone Cluster feature to consolidate the entire solution stack on two physical machines by deploying the Web server, Gateway, Application, and Database tiers in four separate "virtual clusters."

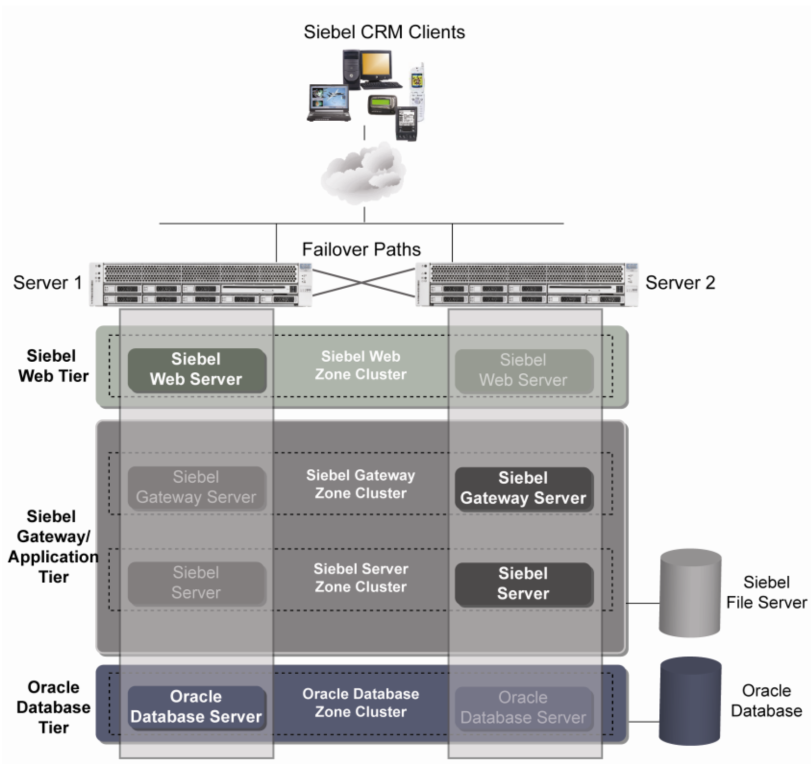

Figure 4. Oracle Solaris Cluster can help to deliver highly available Siebel CRM services.

Designed as a failover environment, the Web server and database are deployed on one machine, and the Gateway and Siebel servers are deployed on the other. This distributes the workload across the two machines. If one machine fails, all services are hosted on the surviving machine. When the failed machine is restored, Oracle Solaris Cluster can automatically restore application distribution across the two machines, or an operator can do it manually.

This HA configuration is intended to retain operational capability during a failure, with as little user impact as possible. As a result, optimization of the servers is biased for maximum concurrent user performance with sufficient computing power kept in reserve to elegantly facilitate transition to failover mode.

Using the GUI management tool shown in Figure 5, each virtual cluster is assigned appropriate system resources, and each environment operates independently of the others. Appendix C includes configuration information for the zone clusters. Note that the proof-of-concept configuration, while useful for the purposes of this testing, is not necessarily typical of a production Siebel CRM environment.

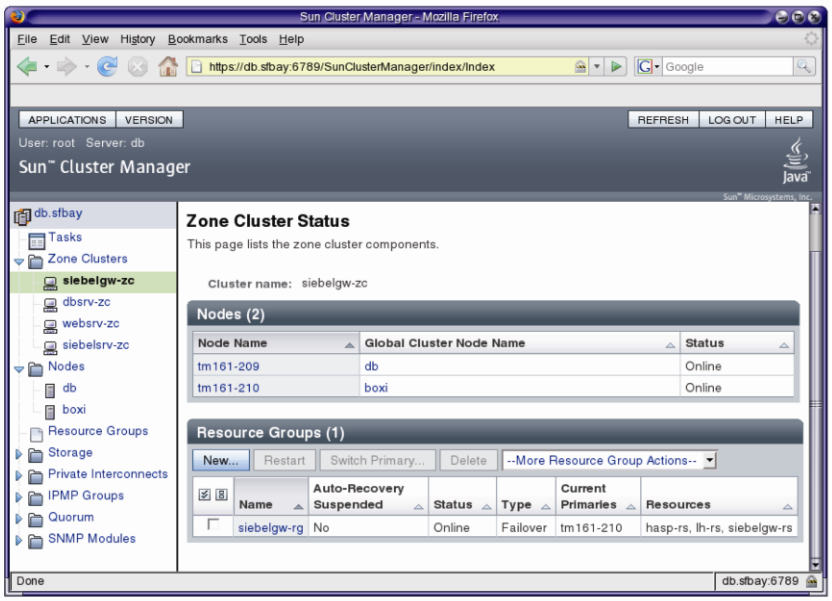

Figure 5. Oracle's Sun Cluster Manager is used to configure and monitor clustered resources for each zone cluster.

In conjunction with highly reliable solution components (such as SPARC servers, Oracle's Sun Storage products, and Oracle Solaris), Oracle Solaris Cluster helps to construct HA solutions that can deliver reliable and resilient Siebel CRM application services. Figure 6 illustrates a large-scale deployment environment — Gateway and database services are clustered and redundant Web and Siebel servers are deployed to achieve high levels of availability.

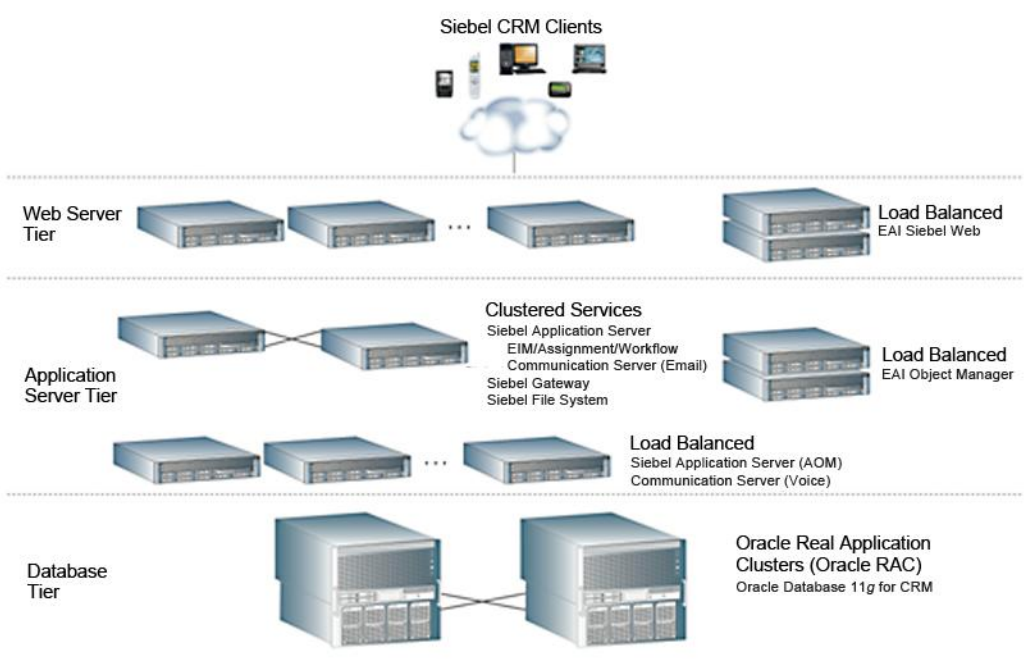

Figure 6. A typical large-scale deployment of clustered servers creates a reliable environment for Siebel CRM services.

Testing Scenarios

Engineers executed three test scenarios once each with 30,000, 40,000 and 60,000 active users using an HA configuration and clustered Containers defined on the two SPARC T-Series servers. Table 2 shows the Siebel CRM server configurations for the three user population scenarios.

TABLE 2. CONFIGURATION OF SERVICES FOR HA TESTING

NUMBER OF CONCURRENT USERS	NUMBER OF WEB SERVERS	NUMBER OF SIEBEL SERVERS	TOTAL NUMBER OF SIEBEL OBJECT MANAGERS	NUMBER OF ORACLE DATABASE INSTANCES
30,000	2	2	15	1
40,000	2	2	24	1
60,000	2	2	36	1

Performance and Scalability Results with Oracle Solaris Cluster

Oracle Solaris 11 on each server was configured with four clustered Containers, in addition to the global zone. Each clustered zone isolated a different Siebel CRM tier — Web, Gateway, Application, or Database. Table 3 shows how system resources were dedicated to each tier. This design represents a reasonable and likely deployment scenario.

TABLE 3. RESOURCES ALLOCATED TO EACH TIER AND CONTAINER IN TESTING

TIER AND CONTAINER	VIRTUAL CPUS (VCPUS)[2]	MEMORY
Web tier	30 vCPUs	6 GB
Application tier	75 vCPUs	48 GB
Gateway tier	3 vCPUs	3 GB
Database tier	36 vCPUs	36 GB

[2] Since the SPARC T5-2 server has two SPARC T-Series processors with 16 cores, there is a maximum of 512 vCPUs possible per system, for a total of 1024 vCPUs in this configuration.

In this round of testing, data was also collected from UNIX system performance tools, Load Runner (the workload generator software), and Oracle Automatic Workload Repository. The following pages contain metrics for testing of the HA configuration, including

- CPU utilization (as a percentage)
- Memory utilization (in GB)
- Business transaction throughput (in number of transactions per hour)
- Average transaction response time (in seconds)
- Transaction throughput (in number of transactions per hour) and response time (in seconds)
- Power consumption (in watts)

Business Transaction Throughput (Clustered Configuration)

Table 4 lists the throughput rates. As the user population increases from 8,000 to 14,000 to 26,000 users, throughput increases almost linearly.

TABLE 4. TRANSACTION THROUGHPUT (TRANSACTIONS/HOUR)

BUSINESS TRANSACTION TYPE	30,000 USERS	40,000 USERS	60,000 USERS
Siebel Financial Services Call Center	170,000	180,000	200,000

Average Transaction Response Time (Clustered Configuration)

Table 5 lists the average response time in seconds for each transaction type. For the purposes of the testing exercise, response times are measured at the Web server instead of at the end user. (This is because response times at the end user depend on a number of other variables, such as network latency, the bandwidth between the Web server and the browser, and the time for content rendering by the browser.)

TABLE 5. AVERAGE TRANSACTION RESPONSE TIME (SECONDS)

SIEBEL CRM TIER	30,000 USERS	40,000 USERS	60,000 USERS
Siebel Financial Services Call Center	0.20	0.22	0.30

Transaction Throughput and Response Time (Clustered Configuration)

Performance and scalability are inextricably linked. For this reason, it is important to examine throughput and response time metrics together when analyzing application performance and configuration scalability. As application load increases, response time must remain within acceptable bounds. As a rule of thumb, as the number of concurrent users increases, if there is a linear increase in throughput, the increase in response times should also be within an acceptable limit.

Power Consumption (Clustered Configuration)

During the testing of the HA configuration, power consumption was not explicitly measured. Estimated power consumption for a SPARC T5-2 server supporting 30,000 concurrent Siebel CRM users is around 778 watts, which is approximately 8.9 users per watt.

Failover Testing with Oracle Solaris Cluster

In addition to performance and scalability testing, Oracle engineers conducted failover testing. Using the same test configuration shown in Figure 4, in which one server node hosts primary instances of the Web and Database services while a second node hosts primary instances of the Gateway and Seibel servers, Oracle engineers conducted four separate failover tests.

The failover tests executed under a workload simulating 1000 concurrent users and consisted of these four scenarios:

- Failover of the primary Gateway server on node 2. After the simulated workload reached 1000 active users, engineers killed all processes associated with the Gateway server on node 1. As a result, Oracle Solaris Cluster restarted the Gateway resource group on node 2. Once the Gateway server came online, workload generation resumed. Throughput and response time were measured to examine whether these metrics were consistent both before and after the failover.
- Reboot of the primary Web server on node 1. With 1000 simulated concurrent users, engineers rebooted the zone cluster on node 1 supporting the Web server. Oracle Solaris Cluster then failed over the Web server resource group to the second node. Once the Web server came online, the workload simulator resumed load generation and engineers measured throughput and response time to determine consistency before and after the fault.
- Reboot of the Database server instance on node 1. After the simulated workload reached 1000 active users, engineers rebooted the zone cluster on node 1 with the Database server. Oracle Solaris Cluster failed over the Database server resource group to the second node. Once the Database server came online, workload generation resumed. Throughput and response time were measured to determine consistency before and after the failover.
- Complete power loss of node 2. In this scenario, after the simulated workload reached 1000 users, engineers powered off node 2 via the server's built-in service processor. In response, Oracle Solaris Cluster restarted the Gateway and Siebel Server resource groups on node 1. Again, throughput and response time were measured for consistency before and after the node failure.

In all four scenarios, throughput and response times were consistent before and after failover. Table 7 shows metrics for the 1000-user workload, including baseline values measured prior to testing.

TABLE 7. TRANSACTION THROUGHPUT AND RESPONSE TIME IN FAILOVER SCENARIOS

FAILOVER TEST SCENARIO	# USERS	THROUGHPUT (TPH)	RESPONSE TIME (IN SECONDS)	DETECTION (D) AND RECOVERY (R) TIMES (IN MINUTES AND SECONDS)
Baseline (All tiers, nodes 1 and 2)	400 Financial	3791 8999	0.21 0.11	N/A
Failover of primary Gateway server on node 2	400 Financial	3793 8980	0.21 0.11	Gateway: D = 1s, R = 1mn17s Siebel: R = 26s Total stack: D+R = 1mn44s

Failover of primary Web server on node 1	400 Financial	3777 9052	0.21 0.11	Web: D = 14s, R = 1mn57s Total: D+R = 2mn11s
Failover of primary Database server on node 1	400 Financial	3793 8971	0.21 0.12	Database: D = 17s, R = 1mn1s Total: D+R = 1mn18s
Failover of node 2 (power-off)	400 Financial	3784 8930	0.21 0.12	D = 16s Gateway: R = 23s Siebel: R = 1mn24s Total stack: D+R = 2mn3s

Best Practices and Recommendations

Prior to testing the solution, engineers made several optimizations to the Siebel CRM configuration. Summarized below, these settings and modifications can help customers optimize performance and scalability when consolidating Siebel CRM Web, Gateway/Application, and Database tiers on a server. Sizing recommendations are included at the end of this section and can be tailored to site-specific requirements. Oracle consultants are experienced in designing optimal solutions for Siebel CRM applications and knowledgeable about best practices. By engaging these consultants in application and system architectural design, customers can achieve optimal configurations to help meet business and site requirements.

Server/Operating System Optimizations

Best practices for optimizing the server and operating system include the following:

- Make sure the server firmware is up to date. Check the System Firmware Release site (http://www.oracle.com/technetwork/systems/patches/firmware/release-history-jsp-138416.html) for the latest firmware release.
- Install the latest release of Oracle Solaris 11. Customers running Siebel CRM applications on Oracle Solaris 11 5/08 should apply kernel patch 137137-09 from http://www.oracle.com/technetwork/systems/patches/solaris/index.html. Later releases incorporate an equivalent workaround for this critical Siebel-specific bug, so no additional patching is required. Eventually Oracle will fix this bug in their code base, but in the meantime, the Oracle Solaris 11 10/08 release (or the patch for the earlier Oracle Solaris version) addresses this issue for Siebel applications (and other 32-bit applications that include memory allocators that return unaligned mutexes). For more information, see Sun RFE 6729759 ("Need to accommodate non-8-byte-aligned mutexes") or Oracle's Siebel support document #735451.1.
- Optimize Oracle Solaris 11 settings in `/etc/system`. Enable 256 MB memory page sizes on all nodes. By default, the latest update of the Solaris 11 OS uses a default maximum of 4 MB memory pages even when 256 MB pages are a better application fit. To set a 256 MB page size, change the setting in `/etc/system` as follows:

```
set max uheap lpsize=0x10000000
```

- To avoid running into the standard input/output (`stdio`) limitation of 256 file descriptors, add the following lines to `start_server` in the Siebel CRM Gateway/Application tier:

```
ulimit -n 2048
LD_PRELOAD_32=/usr/lib/extendedFILE.so.1
export LD_PRELOAD_32
```

- The default file descriptor limit in a shell is 256 and the maximum limit is 65,536. However, 2,048 is a reasonable limit from the application's perspective.
- Improve scalability with a MT-hot memory allocation library: `libumem` or `libmtmalloc`. To improve the scalability of the multithreaded workloads, preload an MT-hot, object-caching memory allocation library such as `libumem(3lib)` or `mtmalloc(3malloc)`. To preload the `libumem` library, set the `LD_PRELOAD_32` environment variable in the shell (`bash/ksh`) as shown below.

```
Export LD_PRELOAD_32=/usr/lib/libumem.so.1:$LD_PRELOAD_32
```

  Web and Application servers in the Siebel CRM enterprise stack are 32 bit. However, Oracle Database 10*g* or 11*g* on Oracle Solaris 11 for SPARC processor-based servers is 64 bit. Hence, the path to the `libumem` library in the `PRELOAD` statement differs slightly in the Database tier, as shown below.

```
Export LD_PRELOAD_64=/usr/lib/sparcv9/libumem.so.1:$LD_PRELOAD_64
```

  Be aware that the trade-off is an increase in memory footprint — there can be a resulting 5 percent to 20 percent increase in the memory footprint with an MT-hot memory allocation library preloaded. In previous Siebel CRM 8 code testing, there was around a 5 percent improvement in CPU utilization with a 9 percent increase in the memory footprint with a load of 400 users.

- Tune the TCP/IP network stack by modifying these settings:

```
ndd -set /dev/tcp tcp_time_wait_interval 60000
ndd -set /dev/tcp tcp_conn_req_max_q 1024
ndd -set /dev/tcp tcp_conn_req_max_q0 4096
ndd -set /dev/tcp tcp_ip_abort_interval 60000
ndd -set /dev/tcp tcp_keepalive_interval 900000
ndd -set /dev/tcp tcp_rexmit_interval_initial 3000
ndd -set /dev/tcp tcp_rexmit_interval_max 10000
ndd -set /dev/tcp tcp_rexmit_interval_min 3000
ndd -set /dev/tcp tcp_smallest_anon_port 1024
ndd -set /dev/tcp tcp_slow_start_initial 2
ndd -set /dev/tcp tcp_xmit_hiwat 799744
ndd -set /dev/tcp tcp_recv_hiwat 799744
ndd -set /dev/tcp tcp_max_buf 8388608
ndd -set /dev/tcp tcp_cwnd_max 4194304
ndd -set /dev/tcp tcp_fin_wait_2_flush_interval 67500
ndd -set /dev/udp udp_xmit_hiwat 799744
ndd -set /dev/udp udp_recv_hiwat 799744
ndd -set /dev/udp udp_max_buf 8388608
```

I/O Best Practices

The Siebel CRM 8 PSPP workload is moderately sensitive to disk I/O. For example, when all 30,000 concurrent users are online, the database writes about 7.5 MB worth of data per second (out of 7.5 MB, approximately 3 MB is written to the redo logs), and it reads about 18.5 kB per second. The Oracle Database server writes data randomly into the data files (because the tables are scattered), whereas writes to the redo logs are largely sequential. For the purpose of testing, the database resided on a UFS file system.

Best practices relating to I/O include the following:

- Store the data files separately from the redo log files. If the data files and redo log files are stored on the same disk drive and the disk drive fails, the redo files cannot be used in the database recovery procedure.
- Size the online redo logs to control the frequency of log switches. In the tested configuration, two online redo logs were configured each with 10 GB of disk space.
- Eliminate double buffering by forcing the file system to use direct I/O. Oracle Database caches data in its own cache within the shared global area (SGA) known as the database block buffer cache. Database reads and writes are cached in block buffer cache so that subsequent accesses for the same blocks do not need to reread data from the operating system. In addition, UFS file systems in Oracle Solaris default to reading data though the global file system cache for improved I/O. This is why, by default, each read is potentially cached twice — one copy in the operating system's file system cache and the other copy in Oracle Database's block buffer cache. In addition to double caching, extra CPU overhead exists for the code that manages the operating system file system cache. The solution is to eliminate double caching by forcing the file system to bypass the OS file system cache when reading from and writing to the disk. To implement direct I/O and eliminate double caching, mount the UFS file systems (that hold the data files and the redo logs) with the `forcedirectio` option:

```
mount –o forcedirectio /dev/dsk/<partition> <mountpoint>
```

Web Tier Best Practices

Best practices for the Web tier include the following:

- Upgrade to the latest service pack of the Oracle iPlanet Web Server (formerly Sun Java Web Server).
- Run the Web server in multiprocess mode by setting the MaxProcs directive in `magnus.conf` to a value greater than 1. In multiprocess mode, the Web server can handle requests using multiple processes with multiple threads in each process. With a value greater than 1 for MaxProcs, the Web server relies on the operating system to distribute connections among multiple Web server processes. However, many modern operating systems (including Oracle Solaris) do not distribute connections evenly, particularly when there are a small number of concurrent connections. For this reason, tune the parameter for the maximum number of simultaneous requests by setting the RqThrottle parameter in `magnus.conf` to an appropriate value. A value of 1024 was used in the test.

Siebel Application Tier Best Practices

Best practices for the Siebel Application tier include the following:

- Comment out the following lines in `$SIEBEL_HOME/siebsrvr/bin/siebmtshw`.

```
# This will set 4M page size for Heap and 64 KB for stack
# MPSSHEAP=4M
# MPSSSTACK=64K
# MPSSERRFILE=/tmp/mpsserr
# LD_PRELOAD=/usr/lib/mpss.so.1
# export MPSSHEAP MPSSSTACK MPSSERRFILE LD PRELOAD
```

  All SPARC T-Series systems support a 256 MB page size. However Siebel CRM's `siebmtshw` script restricts the page size to 4 MB and 64 kB for stack unless indicated lines are commented out in the script.

- Experiment with a smaller number of Siebel CRM Object Managers. Configure the Object Managers in such a way that each Object Manager handles at least 200 active users. Siebel CRM's standard recommendation of 100 or fewer users per Object Manager is suitable for conventional systems but not ideal for CMT systems such as the SPARC T-Series server. Oracle's SPARC T-Series systems are ideal for running multithreaded processes with numerous lightweight processors (LWPs) per process. With fewer Siebel Object Managers, there is also usually a significant improvement in the overall memory footprint.

Oracle Database Tier Best Practices

Best practices for the Oracle Database tier include setting the following initialization parameters:

- Set the Oracle initialization parameter, `DB_FILE_MULTIBLOCK_READ_COUNT`, to an appropriate value, such as 8. The `DB_FILE_MULTIBLOCK_READ_COUNT` parameter specifies the maximum number of blocks read in one I/O operation during a sequential scan. In the testing, `DB_BLOCK_SIZE` was set to 8 kB. Since average reads are around 18.5 kB per second, setting `DB_FILE_MULTIBLOCK_READ_COUNT` to a higher value does not necessarily help to improve I/O performance.

- Explicitly set the database initialization parameter *`enableNUMA`*`optimization` to FALSE for SPARC T-Series servers. On these multisocket servers, the parameter *`enableNUMA`*`optimization` is set to TRUE by default. During the 14,000-user test, intermittent shadow process crashes occurred with the default. There were no additional gains with the default NUMA optimizations.

Best Practices for High Availability Configurations

Oracle Solaris Cluster HA for Siebel provides fault monitoring and automatic failover for the Siebel Gateway and Siebel Server. However, in a Siebel CRM cluster deployment, any physical node running the Oracle Solaris Cluster agent for Siebel cannot also run the Resonate agent. (Resonate and Oracle Solaris Cluster can coexist in the same Siebel enterprise, but not on the same physical server. For more information, see the *Oracle Solaris Cluster Data Service for Siebel Guide for Solaris OS* at http://docs.oracle.com/cd/E19680-01/html/821-1539/index.html.)

Load balancing is a technique to spread the workload between two or more instances of the same application to increase throughput and availability. The Web tier can be load balanced for high availability in an N+1 architecture, for example, by having multiple Containers or domains housing the Web server with Siebel Web Server Extensions along with a hardware load balancer.

Additionally, Oracle Solaris Cluster can load balance the Web server. An Oracle Solaris Cluster feature called Shared Address Resource for Scalable Services allows multiple instances of the same application (such as the Web server) on each node to listen and process requests sent to the same IP address and port number. However, when the Cluster HA agent for the Web server is used together with the Cluster HA agent for Siebel Server, Oracle Solaris Cluster can provide failover service only to the Web server.

To provide disaster recovery over unlimited distances, Oracle Solaris Cluster Geographic Edition provides a multisite, multicluster disaster recovery solution to manage application availability across geographically remote clusters. In the event that a primary cluster fails, Oracle Solaris Cluster Geographic Edition enables administrators to initialize business services with replicated data on a secondary cluster, as depicted in Figure 9.

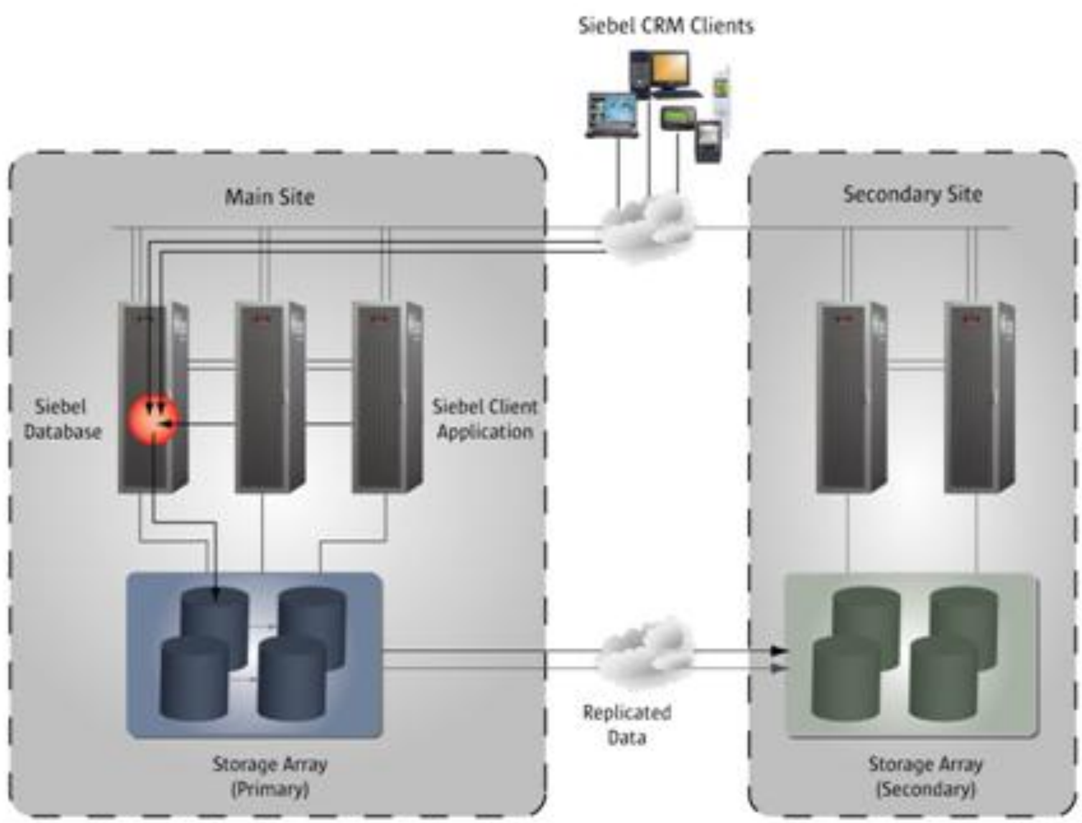

Figure 9. Oracle Solaris Cluster Geographic Edition enables disaster recovery solutions over long distances for Siebel CRM services.

Sizing Guidelines

Under the Siebel CRM 8 PSPP testing workload, engineers set virtual CPU (vCPU) and memory allocations for Oracle Solaris Containers, as shown in Table 8.

TABLE 8. ACTUAL RESOURCE ALLOCATIONS FOR OVER 30,000 USERS ON SPARC T-SERIES SERVER

TIER	VCPUS	MEMORY	ACTUAL USAGE IN TESTED CONFIGURATION
Web tier	25 vCPUs	32 GB	CPU: 78.21% Memory: 5.5 GB
Application tier	220 vCPUs	128 GB	CPU: 76.29% Memory: 80 GB
Database tier	45 vCPUs	64 GB	CPU: 71.33% Memory: 28 GB

TABLE 9. RECOMMENDED RESOURCE ALLOCATIONS FOR 40,000 USERS

TIER	VCPUS	MEMORY	ACTUAL USAGE IN TESTED CONFIGURATION
Web tier	10 vCPUs	32 GB	With 22 vCPUs, 10 GB RAM: CPU: 13.67% Memory: 1.1 GB
Application tier	55 vCPUs	256 GB	With 98 vCPUs, 50 GB RAM: CPU: 10.75% Memory: 19 GB
Database tier	15 vCPUs	72 GB	With 19 vCPUs, 25 GB RAM: CPU: 14.22% Memory: 12 GB

TABLE 10. RECOMMENDED RESOURCE ALLOCATION FOR 60,000 USERS

TIER	VCPUS	MEMORY	ACTUAL USAGE IN TESTED CONFIGURATION
Web tier	15 vCPUs	64 GB	With 22 vCPUs, 8 GB RAM: CPU: 33% Memory: 2 GB
Application tier	98 vCPUs	256 GB	With 98 vCPUs, 42 GB RAM: CPU: 28% Memory: 36 GB

Database tier	19 vCPUs	128 GB	With 19 vCPUs, 32 GB RAM: CPU: 30% Memory: 15 GB

Given resource allocations in Table 9 and Table 10, a SPARC T5-2 server could potentially be configured as summarized in Table 11.

TABLE 11. POSSIBLE CONFIGURATIONS FOR SPARC T-SERIES SERVER

NUMBER OF USERS	DESCRIPTION	TOTAL VPCUS	PHYSICAL CPUS	TOTAL MEMORY
30,000	Entry T5	128	2	128 GB
40,000	Medium T5	256	4	256 GB
60,000	Large T5	512	8	512 GB

Of course, actual resource configurations depend specifically on site requirements. In small-to-medium deployments, one strategy is to deploy a server with a greater number of physical resources than what is minimally required for Siebel CRM applications and to use excess resources and additional virtualized environments to support other (non-Siebel) application workloads. This enables tremendous flexibility as growth occurs. Another alternative is to deploy the Siebel CRM solution on a smaller server, such as the SPARC T4-2 server. Using a smaller server lowers the cost of deploying an HA configuration by implementing a second server, as in the HA test configuration.

Baseline Configurations

Expected performance characteristics are based on proof-of-concept test implementations and are provided as is without warranty of any kind. The entire risk of using information provided herein remains with the reader 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.

Based on the testing described in this paper, the remainder of this section outlines recommended hardware configurations as a starting point for a range of deployment sizes deployments. For optimal sizing information, contact your local Oracle representative.

Entry Level HA Configuration — Up to 8,000 Users

For a highly available configuration supporting up to 8,000 concurrent users, the following hardware components should be considered:

- Servers — Two SPARC T4-1 servers, each with 1 CPU and 128 GB of RAM.

Small HA Configuration — Up to 14,000 Users

For a medium-sized HA configuration supporting up to 14,000 users, these hardware components are recommended for deployment:

- Servers — Two SPARC T4-2 servers, each with 2 CPUs and 128 GB of RAM.

Medium HA Configuration (T5) — Up to 30,000 Users

For a highly available configuration supporting up to 26,000 concurrent users, consider the following hardware components:

- Servers — Two SPARC T5-2 servers, each with 2 CPUs and 512 GB of RAM. Since the SPARC T4-Series servers support up to 1 TB RAM, this configuration enables memory expansion in support of additional applications or to enhance available processing resources.

Large HA Configuration (T5) — Up to 40,000 Users

For a highly available configuration supporting up to 26,000 concurrent users, consider the following hardware components:

- Servers — Two SPARC T5-4 servers, each with 4 CPUs and 512 GB of RAM. Since the SPARC T5-4 servers support up to 2 TB RAM, this configuration enables memory expansion in support of additional applications or to enhance available processing resources.

Extra Large HA Configuration (T5) — Up to 60,000 Users

For a highly available configuration supporting up to 26,000 concurrent users, consider the following hardware components:

- Servers — Two SPARC T5-8 servers, each with 8 CPUs and 512 GB of RAM. Since the SPARC T5-8 servers support up to 4 TB RAM, this configuration enables memory expansion in support of additional applications or to enhance available processing resources.

Conclusion

Virtualization allows Siebel CRM applications to be consolidated securely and effectively on a single server, offering many benefits over the use of multiple physical machines — better resource utilization, smaller data center footprint, and lower power consumption. The advanced thread density of a single SPARC T-Series server allows throughput to scale almost linearly for small, medium, and large user populations, while at the same time achieving reasonable response times.

Testing confirmed scalability of Siebel CRM workloads when HA technology is deployed in conjunction with virtualization technologies built into SPARC servers. By implementing Oracle Solaris Cluster HA products on two servers, Oracle engineers observed good scalability using virtualized Siebel CRM tiers for up to 60,000 users. Thus, a clustered configuration of economical SPARC T-

Series servers offers a scalable and resilient platform for deploying mission-critical Siebel CRM services.

By taking advantage of the advanced thread density and scalability of Oracle's SPARC servers, customers can build fail-sale virtualized environments that enable remote failover, allowing IT managers to meet SLAs and satisfy stringent disaster recovery requirements for Siebel CRM applications.

In configuring a server for a Siebel CRM deployment, Oracle consultants can help to define an effective architectural model, determine optimal sizing, decide what virtualization technologies to use, and recommend initial resource allocations. For more information on engaging experienced Oracle experts to design an agile Siebel CRM environment for your business, see the Oracle Advanced Customer Support Services Website at www.oracle.com/us/support/systems/advanced-customer-services/index.html.

References

TABLE 13. REFERENCES

WEBSITES	
Oracle's SPARC Servers	www.oracle.com/us/products/servers-storage/
Oracle's Siebel CRM software	http://www.oracle.com/us/products/applications/siebel/overview/index.html
PAPERS AND DOCUMENTATION	
"Using Sun Systems to Build a Virtual and Dynamic Infrastructure"	http://www.oracle.com/partners/en/build-virtual-dynmc-infrastructr-164103.pdf
"Oracle VM Server for SPARC: Enabling A Flexible, Efficient IT Infrastructure"	www.oracle.com/us/oraclevm-sparc-wp-073442.pdf
"Best Practices For Network Availability With Oracle VM Server for SPARC"	http://www.oracle.com/technetwork/articles/systems-hardware-architecture/vmsrvrsparc-availability-163930.pdf
Sun Cluster Data Service for Siebel Guide for Solaris OS	http://download.oracle.com/docs/cd/E19787-01/819-2986/index.html

Appendix A — Configuration of Containers

Each Siebel CRM server ran on a non-global zone as follows:

- `siebelweb` for the Web server
- `siebelapp` for the Gateway/Application servers
- `siebeldb` for the Database server

Virtual CPUs (vCPUs) and memory were allocated to the `siebelweb` and `siebelapp` Containers. Only memory was allocated to the `siebeldb` zone, leaving the `siebeldb` zone to use necessary vCPUs from the global zone. Since all database processes ran in the `siebeldb` non-global zone, there was a negligible consumption of CPU resources in the global zone during the test. The configuration of each zone is shown using the `zonecfg` command.

Web Server

```
# zonecfg -z siebelweb
zonecfg:siebelweb> info
zonename: siebelweb
zonepath: /zones2/webserver
brand: native
autoboot: false
bootargs:
pool:
limitpriv:
scheduling-class:
ip-type: shared
inherit-pkg-dir:
        dir: /lib
inherit-pkg-dir:
        dir: /platform
inherit-pkg-dir:
        dir: /sbin
inherit-pkg-dir:
        dir: /usr
net:
        address: 18.1.1.4.236
        physical: nxge2
        defrouter not specified
dedicated-cpu:
        ncpus: 22
capped-memory:
        physical: 8G
```

Application Server

```
# zonecfg -z siebelapp
zonecfg:siebelapp> info
zonename: siebelapp
zonepath: /zones3/appserv
brand: native
autoboot: false
bootargs:
pool:
limitpriv:
scheduling-class:
ip-type: shared
inherit-pkg-dir:
    dir: /lib
inherit-pkg-dir:
    dir: /platform
inherit-pkg-dir:
    dir: /sbin
inherit-pkg-dir:
    dir: /usr
net:
    address: 18.1.1.4.29
    physical: nxge1
    defrouter not specified
dedicated-cpu:
    ncpus: 196
capped-memory:
    physical: 88G
```

Database Server

```
# zonecfg -z siebeldb
zonecfg:siebeldb> info
zonename: siebeldb
zonepath: /zones/dbserver
brand: native
autoboot: false
bootargs:
pool:
limitpriv:
scheduling-class:
ip-type: shared
inherit-pkg-dir:
    dir: /lib
inherit-pkg-dir:
    dir: /platform
inherit-pkg-dir:
    dir: /sbin
inherit-pkg-dir:
    dir: /usr
net:
    address: 18.1.1.4.237
    physical: nxge3
    defrouter not specified
device
    match: /dev/dsk/c6t0d0s6
device
    match: /dev/dsk/c8t0d0s6
capped-memory:
    physical: 32G
```

Appendix B — Configuration of Oracle VM Server for SPARC

The `ldm list` command shows the three domains used for testing.

```
# ldm list
NAME              STATE    FLAGS   CONS    VCPU  MEMORY
primary           active   -n-cv   SP      38    32G
siebelapp         active   -n---   15001   196   89600M
siebelweb         active   -n---   15000   22    8G
```

Details on the three domain configurations are given below.

Primary Domain

```
Domain Name: primary

VARIABLES
    boot-device=/pci@400/pci@0/pci@1/scsi@0/disk@0,0:a disk net

IO
    DEVICE            PSEUDONYM        OPTIONS
    pci@400           pci
    pci@500           pci
    pci@600           pci
    pci@700           pci

VCC
    NAME              PORT-RANGE
    primary-vcc0      15000-15010

VSW
    NAME              MAC                NET-DEV   DEVICE     MODE
    primary-vsw0      00:14:4f:fb:64:21 nxge3     switch@0
    primary-vsw1      00:14:4f:fb:49:d2 nxge2     switch@1

VDS
    NAME              VOLUME         OPTIONS      DEVICE
    primary-vds0      vol1                        /dev/dsk/c3t40d1s2
    primary-vds1      vol2                        /dev/dsk/c2t40d1s2

VCONS
    NAME              SERVICE                     PORT
    SP
```

Based on measurements from the test, if the Database server is run in a Guest Domain instead of the Primary Domain, then some resources should be reassigned to its Guest Domain, but leaving at least 1 vCPU and 0.5 GB of RAM assigned to the Primary Domain.

Siebel Application Server Domain

```
Domain Name:  siebelapp

VARIABLES
    auto-boot?=false
    boot-device=/virtual-devices@100/channel-devices@200/disk@0

NETWORK
```

```
     NAME              SERVICE                     DEVICE     MAC
     vnet2             primary-vsw1@primary        network@0
00:14:4f:f8:8f:13

 DISK
     NAME              VOLUME                      TOUT DEVICE
SERVER
     vdisk2            vol2@primary-vds1                disk@0
primary

 VCONS
     NAME              SERVICE                     PORT
     siebelapp         primary-vcc0@primary        15001
```

Siebel Web Server Domain

```
 Domain Name:  siebelweb

 VARIABLES
     auto-boot?=false
     boot-device=/virtual-devices@100/channel-devices@200/disk@0
     nvramrc=devalias vnet0 /virtual-devices@100/channel-
devices@200/network@0

     use-nvramrc?=true

 NETWORK
     NAME              SERVICE                     DEVICE     MAC
     vnet1             primary-vsw0@primary        network@0
00:14:4f:fb:01:50

 DISK
     NAME              VOLUME                      TOUT DEVICE
SERVER
     vdisk1            vol1@primary-vds0                disk@0
primary

 VCONS
     NAME              SERVICE                     PORT
     siebelweb         primary-vcc0@primary        15000
```

Appendix C — Configuration of Zone Clusters

In testing, engineers configured zone clusters for each Siebel CRM server instance, as shown in Figure 4. The zone clusters were:

- `websrv-zc` for the Web server
- `siebelgw-zc` for the Gateway server
- `siebelsrv-zc` for the Application server
- `dbsrv-zc` for the Database server

Below, the `clzc` command shows status information for the zone clusters and the `clrg` command shows status information for cluster resource groups. In subsequent pages, the `clzc` command displays configuration details for each zone cluster.

```
 # clzc status

=== Zone Clusters ===

--- Zone Cluster Status ---

Name            Node Name   Zone HostName   Status   Zone Status
----            ---------   -------------   ------   -----------
siebelsrv-zc    db          tm161-207       Online   Running
                boxi        tm161-208       Online   Running

siebelgw-zc     db          tm161-209       Online   Running
                boxi        tm161-210       Online   Running

websrv-zc       db          tm161-211       Online   Running
                boxi        tm161-212       Online   Running

dbsrv-zc        db          tm161-205       Online   Running
                boxi        tm161-206       Online   Running


# clrg status -Z all

=== Cluster Resource Groups ===

Group Name                  Node Name   Suspended   Status
----------                  ---------   ---------   ------
siebelsrv-zc:siebelsrv-rg   tm161-207   No          Offline
                            tm161-208   No          Online

siebelgw-zc:siebelgw-rg     tm161-209   No          Offline
                            tm161-210   No          Online

websrv-zc:websrv-rg         tm161-211   No          Online
                            tm161-212   No          Offline

dbsrv-zc:dbsrv-rg           tm161-205   No          Online
                            tm161-206   No          Offline
```

Web Server

```
 # clzc show -v websrv-zc

=== Zone Clusters ===

Zone Cluster Name:                                   websrv-zc
  zonename:                                          websrv-zc
  zonepath:                                          /zone/websrv-zc
  autoboot:                                          TRUE
  brand:                                             cluster
  bootargs:                                          <NULL>
  pool:                                              <NULL>
  limitpriv:                                         <NULL>
  scheduling-class:                                  <NULL>
  ip-type:                                           shared
  enable_priv_net:                                   TRUE

  --- Solaris Resources for websrv-zc ---

  Resource Name:                                     net
    address:                                         tm161-216
    physical:                                        auto

  Resource Name:                                     fs
    dir:                                             /siebel/web
    special:                                         /dev/global/dsk/d8s6
    raw:                                             /dev/global/rdsk/d8s6
    type:                                            ufs
    options:                                         []

  Resource Name:                                     sysid
    name_service:
DNS{domain_name=sfbay.sun.com name_server=129.145.155.220}
    nfs4_domain:                                     dynamic
    security_policy:                                 NONE
    system_locale:                                   C
    terminal:                                        xterms
    timezone:                                        US/Pacific

  Resource Name:                                     capped-memory
    physical:                                        3G
    swap:                                            4G

  Resource Name:                                     capped-memory
    swap:                                            4G

  Resource Name:                                     inherit-pkg-dir
    dir (0):                                         /lib
    dir (1):                                         /platform
    dir (2):                                         /sbin
    dir (3):                                         /usr

  Resource Name:                                     inherit-pkg-dir
    dir (1):                                         /platform
    dir (2):                                         /sbin
    dir (3):                                         /usr

  Resource Name:                                     inherit-pkg-dir
    dir (2):                                         /sbin
    dir (3):                                         /usr

  Resource Name:                                     inherit-pkg-dir
    dir (3):                                         /usr
```

```
  Resource Name:                                  dedicated-cpu
    ncpus:                                        16
    importance:                                   20

  Resource Name:                                  dedicated-cpu
    importance:                                   20

  Resource Name:                                  rctl
    name:                                         zone.max-swap
    priv:                                         privileged
    limit:                                        4294967296
    action:                                       deny

  --- Zone Cluster Nodes for websrv-zc ---

  Node Name:                                      db
    physical-host:                                db
    hostname:                                     tm161-211

    --- Solaris Resources for db ---

    Resource Name:                                net
      address:                                    10.6.161.211
      physical:                                   nxge0
      defrouter:                                  <NULL>

  Node Name:                                      boxi
    physical-host:                                boxi
    hostname:                                     tm161-212

    --- Solaris Resources for boxi ---

    Resource Name:                                net
      address:                                    10.6.161.212
      physical:                                   nxge0
      defrouter:                                  <NULL>
```

Gateway Server

```
 # clzc show -v siebelgw-zc

=== Zone Clusters ===

Zone Cluster Name:                                siebelgw-zc
  zonename:                                       siebelgw-zc
  zonepath:                                       /zone/siebelgw-zc
  autoboot:                                       TRUE
  brand:                                          cluster
  bootargs:                                       <NULL>
  pool:                                           <NULL>
  limitpriv:                                      <NULL>
  scheduling-class:                               <NULL>
  ip-type:                                        shared
  enable_priv_net:                                TRUE

  --- Solaris Resources for siebelgw-zc ---

  Resource Name:                                  net
    address:                                      tm161-215
    physical:                                     auto

  Resource Name:                                  fs
    dir:                                          /siebel/gateway
```

```
    special:                                            /dev/global/dsk/d10s6
    raw:                                                /dev/global/rdsk/d10s6
    type:                                               ufs
    options:                                            []

  Resource Name:                                        sysid
    name_service:
DNS{domain_name=sfbay.sun.com name_server=129.145.155.220}
    nfs4_domain:                                        dynamic
    security_policy:                                    NONE
    system_locale:                                      C
    terminal:                                           xterms
    timezone:                                           US/Pacific

  Resource Name:                                        capped-memory
    physical:                                           1G
    swap:                                               1G

  Resource Name:                                        capped-memory
    swap:                                               1G

  Resource Name:                                        inherit-pkg-dir
    dir (0):                                            /lib
    dir (1):                                            /platform
    dir (2):                                            /sbin
    dir (3):                                            /usr

  Resource Name:                                        inherit-pkg-dir
    dir (1):                                            /platform
    dir (2):                                            /sbin
    dir (3):                                            /usr


  Resource Name:                                        inherit-pkg-dir
    dir (2):                                            /sbin
    dir (3):                                            /usr

  Resource Name:                                        inherit-pkg-dir
    dir (3):                                            /usr

  Resource Name:                                        dedicated-cpu
    ncpus:                                              2
    importance:                                         20

  Resource Name:                                        dedicated-cpu
    importance:                                         20


Resource Name:                                          rctl
    name:                                               zone.max-swap
    priv:                                               privileged
    limit:                                              1073741824
    action:                                             deny

  --- Zone Cluster Nodes for siebelgw-zc ---

  Node Name:                                            db
    physical-host:                                      db
    hostname:                                           tm161-209

    --- Solaris Resources for db ---

    Resource Name:                                      net
      address:                                          10.6.161.209
      physical:                                         nxge0
```

```
      defrouter:                                        <NULL>

  Node Name:                                            boxi
    physical-host:                                      boxi
    hostname:                                           tm161-210

    --- Solaris Resources for boxi ---

    Resource Name:                                      net
      address:                                          10.6.161.210
      physical:                                         nxge0
      defrouter:                                        <NULL>
```

Application Server

```
 # clzc show -v siebelsrv-zc

=== Zone Clusters ===

Zone Cluster Name:                                      siebelsrv-zc
  zonename:                                             siebelsrv-zc
  zonepath:                                             /zone/siebelsrv-zc
  autoboot:                                             TRUE
  brand:                                                cluster
  bootargs:                                             <NULL>
  pool:                                                 <NULL>
  limitpriv:                                            <NULL>
  scheduling-class:                                     <NULL>
  ip-type:                                              shared
  enable_priv_net:                                      TRUE

  --- Solaris Resources for siebelsrv-zc ---

  Resource Name:                                        net
    address:                                            tm161-214
    physical:                                           auto

  Resource Name:                                        fs
    dir:                                                /siebel/server
    special:                                            /dev/global/dsk/d12s6
    raw:                                                /dev/global/rdsk/d12s6
    type:                                               ufs
    options:                                            []

  Resource Name:                                        sysid
    name_service:
DNS{domain_name=sfbay.sun.com name_server=129.145.155.220}
    nfs4_domain:                                        dynamic
    security_policy:                                    NONE
    system_locale:                                      C
    terminal:                                           xterms
    timezone:                                           US/Pacific

  Resource Name:                                        capped-memory
    physical:                                           34G
    swap:                                               43G

  Resource Name:                                        capped-memory
    swap:                                               43G

  Resource Name:                                        inherit-pkg-dir
    dir (0):                                            /lib
    dir (1):                                            /platform
```

```
    dir (2):                                          /sbin
    dir (3):                                          /usr

  Resource Name:                                      inherit-pkg-dir
    dir (1):                                          /platform
    dir (2):                                          /sbin
    dir (3):                                          /usr

  Resource Name:                                      inherit-pkg-dir
    dir (2):                                          /sbin
    dir (3):                                          /usr
  Resource Name:                                      inherit-pkg-dir
    dir (3):                                          /usr

  Resource Name:                                      dedicated-cpu
    ncpus:                                            70
    importance:                                       20

  Resource Name:                                      dedicated-cpu
    importance:                                       20

  Resource Name:                                      rctl
    name:                                             zone.max-swap
    priv:                                             privileged
    limit:                                            46170898432
    action:                                           deny

  --- Zone Cluster Nodes for siebelsrv-zc ---

  Node Name:                                          db
    physical-host:                                    db
    hostname:                                         tm161-207

    --- Solaris Resources for db ---

    Resource Name:                                    net
      address:                                        10.6.161.207
      physical:                                       nxge0
      defrouter:                                      <NULL>

  Node Name:                                          boxi
    physical-host:                                    boxi
    hostname:                                         tm161-208

    --- Solaris Resources for boxi ---

    Resource Name:                                    net
      address:                                        10.6.161.208
      physical:                                       nxge0
      defrouter:                                      <NULL>
```

Database Server

```
 # clzc show -v dbsrv-zc

=== Zone Clusters ===

Zone Cluster Name:                                    dbsrv-zc
  zonename:                                           dbsrv-zc
  zonepath:                                           /zone/dbsrv-zc
  autoboot:                                           TRUE
```

```
  brand:                                                cluster
  bootargs:                                             <NULL>
  pool:                                                 <NULL>
  limitpriv:                                            <NULL>
  scheduling-class:                                     <NULL>
  ip-type:                                              shared
  enable_priv_net:                                      TRUE
  --- Solaris Resources for dbsrv-zc ---

  Resource Name:                                        net
    address:                                            tm161-213
    physical:                                           auto

  Resource Name:                                        fs
    dir:                                                /oradata/redo
    special:                                            /dev/global/dsk/d9s6
    raw:                                                /dev/global/rdsk/d9s6
    type:                                               ufs
    options:                                            []
  Resource Name:                                        fs
    dir:                                                /oradata/control
    special:                                            /dev/global/dsk/d13s6
    raw:                                                /dev/global/rdsk/d13s6
    type:                                               ufs
    options:                                            []


  Resource Name:                                        fs
    dir:                                                /oradata/data
    special:                                            /dev/global/dsk/d7s6
    raw:                                                /dev/global/rdsk/d7s6
    type:                                               ufs
    options:                                            []

  Resource Name:                                        sysid
    name_service:
DNS{domain_name=sfbay.sun.com name_server=129.145.155.220}
    nfs4_domain:                                        dynamic
    security_policy:                                    NONE
    system_locale:                                      C
    terminal:                                           xterms
    timezone:                                           US/Pacific

  Resource Name:                                        capped-memory
    physical:                                           24G
    swap:                                               40G
    locked:                                             24G

  Resource Name:                                        capped-memory
    swap:                                               40G
    locked:                                             24G

  Resource Name:                                        capped-memory
    locked:                                             24G

  Resource Name:                                        inherit-pkg-dir
    dir (0):                                            /lib
    dir (1):                                            /platform
    dir (2):                                            /sbin
    dir (3):                                            /usr

  Resource Name:                                        inherit-pkg-dir
    dir (1):                                            /platform
    dir (2):                                            /sbin
    dir (3):                                            /usr
```

```
  Resource Name:                                    inherit-pkg-dir
    dir (2):                                        /sbin
    dir (3):                                        /usr

  Resource Name:                                    inherit-pkg-dir
    dir (3):                                        /usr

  Resource Name:                                    dedicated-cpu
    ncpus:                                          32
    importance:                                     20
  Resource Name:                                    dedicated-cpu
    importance:                                     20

  Resource Name:                                    rctl
    name:                                           zone.max-locked-memory
    priv:                                           privileged
    limit:                                          25769803776
    action:                                         deny

  Resource Name:                                    rctl
    name:                                           zone.max-swap
    priv:                                           privileged
    limit:                                          42949672960
    action:                                         deny

  --- Zone Cluster Nodes for dbsrv-zc ---

  Node Name:                                        db
    physical-host:                                  db
    hostname:                                       tm161-205

    --- Solaris Resources for db ---

    Resource Name:                                  net
      address:                                      10.6.161.205
      physical:                                     nxge0
      defrouter:                                    <NULL>

  Node Name:                                        boxi
    physical-host:                                  boxi
    hostname:                                       tm161-206

    --- Solaris Resources for boxi ---

    Resource Name:                                  net
      address:                                      10.6.161.206
      physical:                                     nxge0
      defrouter:                                    <NULL>
```

Below, the `clrs` command reports resource status for `dbsrv-zc` zone cluster.

```
# clrs status -Z dbsrv-zc

=== Cluster Resources ===

Resource Name       Node Name      State        Status Message
-------------       ---------      -----        --------------
hasp-rs             tm161-205      Online       Online
                    tm161-206      Offline      Offline

lh-rs               tm161-205      Online       Online - LogicalHostname
```

```
online.
                    tm161-206      Offline      Offline

db-rs               tm161-205      Online       Online
                    tm161-206      Offline      Offline

lsr-rs              tm161-205      Online       Online
                    tm161-206      Offline      Offline
```

About the Author

Albert Prucha has over 20 years of professional computing experience ranging from coding to data center design. Much of his experience derives from leading professional services practices in virtualization, security, fractional computing, and telepresence. Albert makes an effort to train and certify in competing technologies and products in order to more fairly evaluate their qualities. He is most familiar working with academic, state government, manufacturing, and public utility clients where Information Technology seeks every possible optimization. Not content with constraining his technical pursuits to computers, Albert also enjoys working with microbrewing, hydroponics, and Stirling engines.

Acknowledgements

The author would like to recognize the following individuals for their contributions to this article:

- Gia-Khanh Nguyen, Oracle Solaris Cluster Engineering
- Michael D. Hernandez, Oracle Data Center Client Solutions
- Giri Mandalika, ISV engineering
- Uday Shetty, ISV engineering
- Jenny Chen, ISV engineering

ORACLE®

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

Increasing the Performance and Efficiency of Siebel CRM — A Technical White Paper
July 2013
Author: Albert 'Chad' Prucha

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

Hardware and Software, Engineered to Work Together