ORACLE
OPTIMIZED SOLUTIONS

An Oracle Technical White Paper
October 2011

Oracle Optimized Solution for Oracle WebCenter Content on Oracle's SPARC SuperCluster

ORACLE

ORACLE

Introduction

Oracle WebCenter Content offers a truly unified enterprise content management system (ECM) that offers tremendous scale and a full array of rich functionality, including document and imaging management, content publishing, digital asset management, and records and retention management. However, deploying sophisticated content management solutions can require care in designing a highly-available architecture, choosing hardware and software, testing the design, and tuning the components to perform in an optimal fashion. Unfortunately, these efforts can consume valuable resources, and delay the time to reach successful deployment and realize returns on investments.

The Oracle Optimized Solution for Oracle WebCenter Content on Oracle's SPARC SuperCluster is designed to address the need for unified enterprise content management infrastructure. Integrated in a pre-racked and pre-configured SPARC SuperCluster system, the solution provides complete and optimized infrastructure built with best-of-breed servers, storage, networking, and software. Already assembled and tested, the system is delivered ready to run Oracle WebCenter Content in a reliable, scalable, and efficient manner, providing organizations with the confidence they need to deploy effective content management solutions.

The SPARC SuperCluster platform is designed to provide complete enterprise infrastructure, offering the best application performance, running all enterprise applications, and running all database and middleware – all the time. The SPARC SuperCluster offers extreme performance and highly available operation to applications through the incorporation of all hardware components needed to run both the application binaries as well as any database, middleware, and storage the application requires. This paper specifically focuses on the hardware and software component architectures used to test and evaluate Oracle WebCenter Content on the SPARC SuperCluster.

Oracle Optimized Solution for Oracle WebCenter Content on Oracle's SPARC SuperCluster

Designed to allow the rapid deployment of proven unified enterprise content management infrastructure, the Oracle Optimized solution for Oracle WebCenter Content on SPARC SuperCluster helps eliminate the cost and risk associated with deploying disparate solutions. To this end, best-of-breed components were combined to create a high volume content management platform upon which organizations could run a number of high performance content-centric use cases employing Oracle WebCenter Content. Based on the SPARC SuperCluster, the solution reduces the time needed to get from concept to service deployment. Oracle's proven and tested application-to-disk infrastructure runs in this pre-assembled and pre-tested SuperCluster, eliminating additional time needed for tuning the database or related infrastructure that can lead to significant deployment delays. The platform allows for multiple business units in a company to talk to each other, sharing information and leveraging large amounts of data between different collaborative needs.

As described in this document, the Oracle Optimized Solution for Oracle WebCenter Content on SPARC SuperCluster provides an ideal deployment environment, offering efficient ingest and reuse of

content while incorporating the highest levels of security at all stages. Use cases tested in this document include:

- *High Volume Content Ingestion (HVCI).* Businesses need to handle ever-increasing amounts of data or aggregate multiple disparate systems into one centralized platform for better management and analytics, and the ability to ingest massive amounts of data is essential.
- *Full Text Search (FTS).* Once the data has been ingested into the content management system, organizations need to be able to search through those data for actionable business intelligence.
- *Secure Sockets Layer (SSL) Encryption.* Regulatory and compliance mandates often dictate running a secure business workflow, though technology such as encryption is often disabled due to performance concerns preempting secure operation.

SPARC SuperCluster

SPARC SuperCluster is a complete and pre-engineered infrastructure solution for running a wide range of enterprise applications, including database, middleware, and custom applications. The SPARC SuperCluster solution integrates best-of-breed Oracle components, including Oracle's SPARC T4-4 servers, Oracle Exadata database server and storage cells, Oracle's Sun ZFS Storage appliances, QDR InfiniBand switching, and Oracle Solaris 11.

Hardware components that play a central role in the SuperCluster include:

- *Oracle SPARC T4-4 Servers.* Each server in the SuperCluster is configured with four 3.0 GHz SPARC T4 processors accompanied by 1 TB of DDR3 RAM (memory), four Sun PCIe QDR InfiniBand Host Channel Adapters, and four Sun multithreaded 10 Gb Ethernet networking cards, plus internal disks for local operating system installation. Each SPARC T4 processor features eight cores with eight hardware threads per core, with integrated hardware cryptographic acceleration built into each processor core.
- *Oracle Exadata Software and Storage Cells*. The Oracle Database 11g database runs on the SPARC T4-4 servers under a virtualized Solaris 11 instance, connecting over the InfiniBand fabric to Exadata Database Storage cells and providing numerous advantages over and above the normal benefits of the Oracle database. Features such as Automatic Storage Management (ASM) and Hybrid Columnar Compression (HCC), unique to the Exadata software, are present in this platform and can offer up to 90% compression for content management workloads running on the SPARC SuperCluster. The capacity of the Exadata Storage Cells depends on the disk drives that are purchased. Both high capacity and high performance disk options are supported. High-capacity disk drives can be housed in additional disk trays in an expansion cabinet.
- *Oracle's Sun ZFS 7320 Storage Appliance*. Providing 40 TB of disk capacity for shared file systems, Oracle's Sun ZFS Storage 7320 Appliance uses Flash-enabled Hybrid Storage Pools to dramatically decrease application response times. DTrace performance monitoring features provide fine-grained control and insight into every block, and all I/O activity accessing the storage array. The Oracle WebCenter Content software uses the appliance over an NFS mount for application tier binaries and configuration files, as well as for OS backups.

- *Oracle InfiniBand Fabric.* Oracle's Sun QDR InfiniBand Switch provides the high-speed low-latency InfiniBand fabric interconnecting the SPARC T4-4 servers, the Exadata Storage Cells, and the Oracle Sun ZFS Storage 7320 Appliance within the SuperCluster. This internal network provides for direct access to all components in the SuperCluster. The InfiniBand fabric itself is never exposed to the outside client network and is used for internal communication between the nodes exclusively, utilizing either Internet Protocol over InfiniBand (IPoIB) or Socket Direct Protocol (SDP) depending on specific component connection guidelines.

Integrated Enterprise Application Stack

The Oracle Optimized Solution for Oracle WebCenter Content provides a fully-integrated software stack that includes all of the software, middleware, and operating system components needed to deliver an effective unified content management solution.

- *Oracle WebCenter Content on Oracle WebLogic Server*. The Oracle WebCenter Content software runs on the Oracle WebLogic application server and is the main application software in the Solution. Oracle WebCenter Content acts as the arbiter for all incoming data into the content management system. As tested, the Oracle WebLogic application server was run with the Oracle Exalogic Elastic Cloud software enhancements enabled, providing a performance boost when running over InfiniBand networking.
- *Oracle Solaris 11*. The SPARC T4-4 servers run Oracle Solaris 11 which offers numerous benefits for running Oracle WebCenter Content. Oracle Solaris 11 is a high performance enterprise operating environment, with Oracle Solaris representing the number one deployment platform for the foundational Oracle Database worldwide. Included in Oracle Solaris 11, integrated virtualization technologies such as Oracle Solaris Zones and Oracle VM Server for SPARC are recognized Oracle license boundaries for purposes of determining license costs on the SPARC SuperCluster.

Solution Architecture

Access to content is critical in the modern enterprise. The Oracle Optimized Solution for Oracle WebCenter Content is keenly focused around several main tenants.

- *High availability.* The solution was designed for high availability, so it was critical that no single hardware failure cause a loss of access to any software service. Likewise, the system was designed so that the replacement of any failed hardware component never causes a loss of software services.
- *Performance.* Simply put, the Oracle Optimized Solution for WebCenter Content on SPARC SuperCluster is capable of performing at levels far higher than off-the-shelf deployments would offer. The Optimized Solution in this case offers insight into getting the best performance out of the combined hardware and software stack. For example, there are performance benefits are realized from tuning the architecture with regard to JVM size and heap allocation, virtualization and storage I/O connectivity, and database throughput.
- *Balance in design choices to favor future growth.* Allocating the correct resources to the correct software uses drives the highest levels of performance and high availability while allowing for future

growth. This capability exists in terms of both adding additional storage (through the Exadata database subsystem) or through the addition of more SPARC SuperCluster nodes (in half or full rack increments) as well as using different processor sets for the database and software allocation to meet specific customer needs.

All hardware components, including those pertaining to the network are pre-configured in a traditional content management reference architecture to provide redundancy and increase application availability. The Oracle Optimized Solution for WebCenter Content reference architecture is depicted in Figure 1.

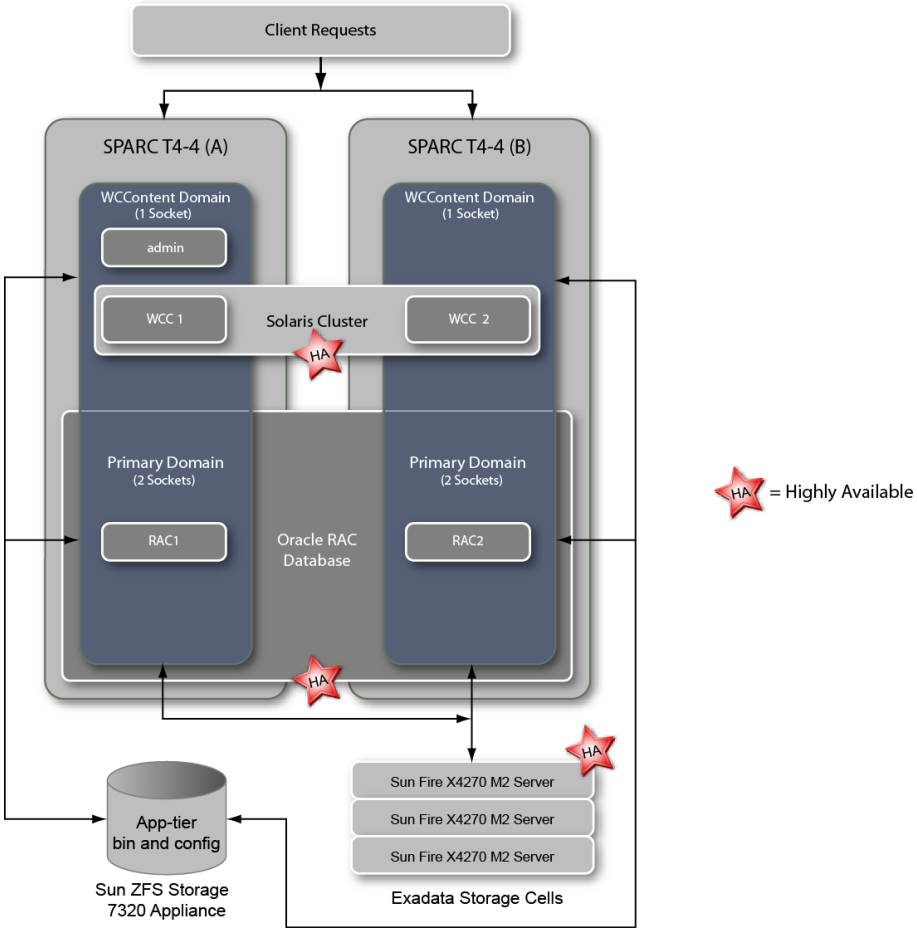

Figure 1. The Oracle WebCenter Content reference architecture features redundancy at all levels.

Redundancy and high availability is provided at all levels of the architecture.

- Solaris Cluster can provide redundancy between Oracle WebCenter Content instances.
- The Oracle RAC 11g database offers redundancy across multiple servers.
- Oracle Exadata Database server software provides redundancy across multiple Exadata storage cells.

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

In the reference architecture, Oracle Solaris Cluster would normally handle failover duties and no specific external failover protection would be needed outside of the SuperCluster. While Solaris Cluster would certainly work with Oracle WebCenter Content, many organizations employ external hardware load balancers for similar content management workloads. As a result, testing with Oracle WebCenter Content utilized an external load balancer to spread the load from incoming clients as well as to provide failover services for incoming connections.

Figure 2 provides detail on the configuration of the SPARC SuperCluster for the purposes of the testing described in this document. For testing, the FABAN toolkit (www.faban.org) was used to generate loads, deployed on Oracle x86 servers. All incoming data into the SuperCluster is directed through an external load balancer first, and then to the content management nodes which in turn communicate with the Exadata database servers. The external load balancer offers round-robin scheduling and external failover. A total of two WebCenter Content instances are in the tested configuration. In Oracle testing, WebLogic 11g for Exalogic was employed to provide communication between the Enterprise Content Management instances and for performance enhancements over the InfiniBand fabric.

After initial installation and testing, the SPARC SuperCluster is then loaded with the Oracle WebCenter Content software. This process includes loading Oracle WebLogic Server, as well as loading the Oracle Database 11g software with the Exadata enhancements for connection to the Exadata Storage Cells. Each Oracle SPARC T4-4 server was divided into two domains/LDOMs (through Oracle VM Server for SPARC) with two SPARC T4 processors (8 cores, 8 threads) on each system dedicated to running Oracle RAC database 11g. An additional SPARC T4 processor on each server was dedicated to running the Oracle WebCenter Content domain. As provided through Oracle VM Server for SPARC, each "domain" is a completely separate Oracle Solaris 11 instance, with inter-domain network connectivity and to external services all served by dedicated InfiniBand or 10 Gigabit Ethernet ports. No virtual networking was configured as is an option through Oracle VM Server for SPARC.

Again, the HA glyphs denote redundancy in the tested configuration. The external load balancer provides redundancy in terms of selecting the appropriate server in case of internal server failure, and will respond by moving application requests and load to the remaining Oracle WebCenter Content node upon failure. Multiple load balancers can also be deployed to eliminate a single point of failure.

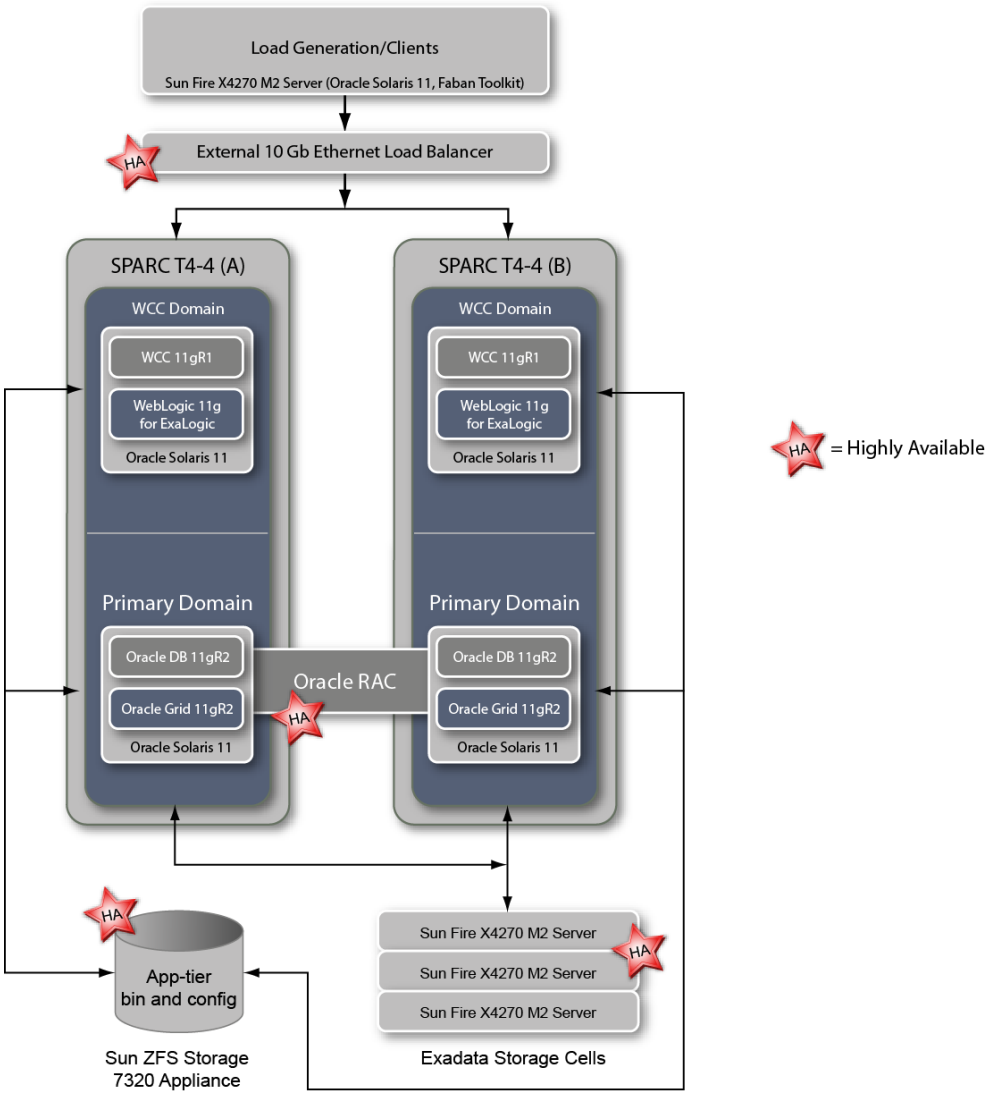

Figure 2. The software configuration of the Oracle Optimized Solution for WebCenter Content provides a single tuned and configured solution.

Analyzing Solution Characteristics

The Oracle Optimized Solution for WebCenter Content based on the SPARC SuperCluster employs numerous approaches that offer advantages in terms of networking, virtualization, and performance.

Network Considerations

Figure 3 illustrates the network from a logical perspective. Virtual LANs (VLANs) are employed to secure and separate communications between clients and the database itself, providing redundancy and a degree of isolation for the database. Clients communicate through the load balancer directly to the Oracle WebCenter Content server instances running on SPARC T4-4 servers. The server instances, in

turn, communicate with the underlying Oracle Exadata database which occupies its own VLAN. As a result, there is no direct connection between the clients and the database, only with the Oracle WebCenter Content server instances.

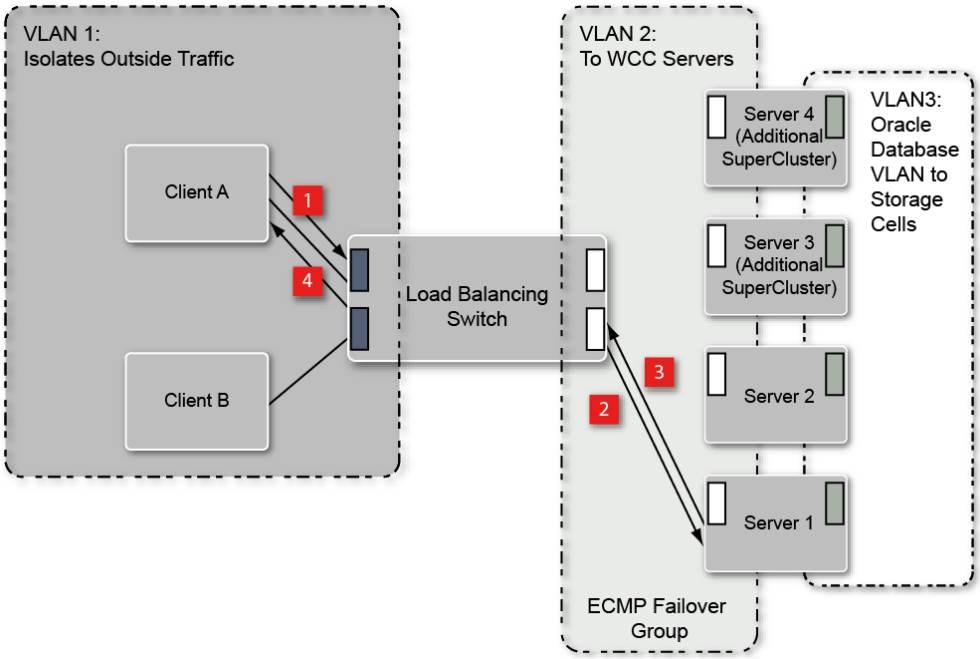

Figure 3. Through the use of VLANs, clients can only communicate with the Oracle Exadata database through Oracle WebCenter Content server instances.

Virtualization Considerations

When deploying software on the SPARC SuperCluster, some care needs to be taken to ensure that the right resources are allocated to the appropriate layers in the platform. Oracle Solaris running on the SPARC T4-4 servers provides two options for controlling resource allocation.

- *Oracle Solaris Zones* (formerly Oracle Solaris Containers) offer a convenient virtualization mechanism with a very fast (~5 second) reboot time. Any main Oracle Solaris 11 instance is referred to as a global zone, whether it is hosting other zones or not. Multiple additional zones can be run inside the global zone, with all sharing same kernel. Applications and patching can be performed either through software installed into a local zone or shared from the global zone. Resource allocation for CPU threads and memory is performed through Oracle Solaris Resource Manager.
- *Oracle VM Server for SPARC* represents a true hypervisor solution with a slightly longer boot time. Logical Domains (LDOMs) provide completely separate Oracle Solaris instances, requiring separate management and setup. Hypervisor support built into each SPARC T4 processor provides isolation between all domains. Virtual storage and networking devices can be created within the LDOM, with resource allocation performed through the hypervisor, allowing fine-grained I/O control. Oracle Solaris Zones can be deployed inside of LDOMs if desired for additional virtualization.

Oracle VM Server for SPARC was selected for testing. The Oracle WebCenter Content and Oracle Database 11g servers were each assigned to LDOMs based on the best results of performance testing. In addition, LDOMs were chosen because assignment of I/O paths and instance separation was more easily accomplished using LDOMs since physical Root Complexes (RC) can be assigned to specific LDOMs. Solaris Zones cannot have such physical paths assigned directly.

Obtaining High Performance Levels on SPARC SuperCluster

SPARC T4-4 based servers running Oracle Solaris 11 deliver very high performance for all workloads running on them. Not only do they excel at running the Exadata Database software (Oracle Database 11gR2), but they are ideally suited for Java-based software such as Oracle WebCenter Content running within Oracle WebLogic Server. There are some generalities that can be made about evaluating software performance on the SPARC T4-based systems.

- Workloads that ran on previous UltraSPARC servers such as Oracle's SPARC T3-based servers will see an immediate improvement by moving directly to the new SPARC T4 servers. Even without tuning, nearly twice the throughput and processing power can be expected, though individual testing needs to be performed to see how any given application will actually perform.
- Some specific tuning parameters should be taken into account to get the best performance on the SPARC T4-based systems.
- The higher clock speed of the SPARC T4 servers (3.0 GHz) allows for much improved single-threaded performance for those applications that are not optimized for parallel multiprocessor computing.
- Applications that run within the Oracle WebLogic application server will see inherent performance improvements since the application server is tuned to exploit the multithreaded design of the SPARC T4 processors.
- Though the SPARC T4-4 servers have excellent single-threaded performance, the true value of these systems and of the SPARC SuperCluster platform itself extends beyond single-threaded application performance. The platform has dramatic advantages that result in significant business value, including reliability, a single support organization for all issues, and an enterprise-class software suite.
- Enhancements to the cryptographic functions built in to the SPARC T4 processor allow for all cryptographic functions to be accessed in user space rather than kernel space. Previous generations of SPARC T-Series processors required at least one thread of any given core to be assigned to a domain in order to access the cryptographic acceleration aspects of that core. In the SPARC T4 processor, all available crypto cores are available for use no matter the core or thread assignments. Dedicated hardware accelerates cryptographic ciphers, including Advanced Encryption Standard (AES), Data Encryption Standard (DES), Camellia, CRC32c, KASUMI, MD5 Message-Digest Algorithm, mpmul, MonMul, MONTSQR, and secure hash algorithms (SHA-1, SHA-256, and SHA-512).

Workload Profiles and Test Results

In testing performed by Oracle, the load on the system was driven by a collection of smaller servers running the FABAN load generator. The sections that follow detail individual workload profiles.

- ***High Volume Content Ingestion (HVCI).*** High Volume Content Ingestion is designed to drive the SPARC SuperCluster to accept as many incoming transactions as possible, placing those documents in the database as the primary/only storage location. This workload acts similarly to use cases where documents are being generated at a high rate through customer interactions and have to be stored for future use. Specific example use cases include retail applications where receipts are created through customer purchases, or the creation of medical records and Explanation of Benefits (EoB) documents when a patient visits a doctor. Examples also abound in the rich media space where user-generated content needs to be stored and sent back out so others can view what has been created.
- ***Full Text Search.*** Full text search is designed to optimize the SPARC SuperCluster to search through the entire contents of all files in a database index in order to find a certain expression or query to return to the user or requesting application. While the HVCI use case focuses on performance, full text search has much applicability to other processes within the enterprise based on the specific query that is run (such as high volume analytics or compliance reporting batch jobs). Sometimes queries can be focused at finding a certain specific document, but full text search can also be used to determine a certain business trend based on the frequency of some dimension of data in the database.
- ***HVCI with SSL Encryption.*** SSL encryption tests were also run on the HVCI use case above. Encryption is increasingly necessary, and demonstrating acceptable performance while running encrypted workloads is of critical importance to many businesses. The SPARC SuperCluster platform is ideally suited for this application since it features a carefully co-engineered hardware and software design that dramatically accelerates encryption. The dedicated cryptographic acceleration provided by SPARC T4 processors makes running a fully-encrypted workload highly performant, and also very easy to implement. Specific support is provided in the Oracle WebLogic application server to enable SSL encrypted workloads with a minimum of effort. In Oracle testing, all SSL operations were terminated (endpoint) in the WebLogic Server instances themselves, and not at the load balancer level.

Table 1 displays the performance achieved in the testing on the Oracle SPARC SuperCluster platform. Even equipped with only two SPARC T4-4 servers in a half-rack SuperCluster configuration (with high performance disk option), performance for both the HVCI and FTS workloads was extremely good.

Performance with full SSL encryption was also remarkable. Many organizations decide to forego using security on certain parts of their dataflow due to performance issues, even when doing so would cause issues during audits if governing or regulatory bodies were to investigate. As mentioned, the cryptographic acceleration units are available via user-level libraries, no matter what core(s) those libraries are running on. As a result, the full cryptographic resources of the processor cores are available to all workloads that need them, regardless of core assignment.

TABLE 1. PERFORMANCE RESULTS OF ORACLE WEBCENTER CONTENT ON ORACLE SPARC SUPERCLUSTER.

TEST	RESULTING PERFORMANCE METRICS
High Volume Content Ingestion (HVCI)	3151 documents/second @ 100 KB docsize no SSL
Full Text Search	3717 concurrent searches per second at a database index size of 3.5 million documents requiring 80 milliseconds for each search operation to return 4313 concurrent searches per second at a database index size of 2 million document requiring 66 milliseconds for each search operation to return
SSL Encryption of HVCI test	2718 documents/second @ 100 KB docsize using **hardware** crypto acceleration through WebLogic Server Certicom functionality (16% performance hit) 2515 documents/second @ 100 KB docsize using **software** crypto acceleration through WebLogic Server Certicom functionality (25% performance hit)

Tuning for Best Workload Performance

To get the best performance out of the various components in the Oracle SPARC SuperCluster, some tuning was performed in several key areas. Though the tuning below was used in this configuration and for these use cases, these settings and environment variables should be viewed as a starting point for tuning, and not a comprehensive list of changes that should be made for all application profiles.

Tuning for the WebLogic Server

Oracle testing employed the 64-bit Oracle WebLogic server optimized by the Exalogic Elastic Cloud Software. To install the server, issue the command below. Exalogic optimization must also be enabled (or checked) either from the Oracle WebLogic Server admin console or in the `config.xml` file.

```
# java -d64 -jar wls1036_generic.jar
```

A number of UCM server UVM start options were provided as shown below. In particular, `-Dweblogic.SocketReaders`, `-Dweblogic.threadpool.MinPoolSize` and `-Dweblogic.threadpool.MaxPoolSize` need to be tuned for different use cases on different configurations. In the use cases tested by Oracle, 20 socket readers and a 200-thread pool were used for the full-text search while 40 socket readers and a 300-thread pool were used for the file ingestion case.

```
-server -Xms15g -Xmx15g -Xmn11g -Xss256k -XX:MaxPermSize=256m -XX:PermSize=256m -
XX:+UseNUMA -XX:+ForceNUMA -XX:+AggressiveOpts -XX:+UseParallelGC -
XX:+UseParallelOldGC -XX:ParallelGCThreads=32 -XX:+DisableExplicitGC -XX:-
ReduceInitialCardMarks -Xnoclassgc -XX:ReservedCodeCacheSize=64m -
XX:CICompilerCount=4 -XX:+PrintCommandLineFlags -XX:+AlwaysPreTouch -
XX:LargePageSizeInBytes=256m -XX:+UseCompressedOops -XX:-UseBiasedLocking -
XX:InitialSurvivorRatio=4 -XX:SurvivorRatio=4 -Dweblogic.SocketReaders=20 -
Dweblogic.threadpool.MinPoolSize=200 -Dweblogic.threadpool.MaxPoolSize=200 -
Dweblogic.GatheredWritesEnabled=true -Dweblogic.ScatteredReadsEnabled=true -
Dweblogic.replication.enableLazyDeserialization=true -Djps.auth.debug=false -
Djava.security.egd=file:/dev/./urandom
```

For JDBC connection pools, one multiple data source pool was created (`MultiDS`) with two generic data sources (`CSDS` and `CSDS2`). `CSDS` and `CSDS2` had the same parameters except for the URLs as shown below.

```
CSDS URL:
jdbc:oracle:thin:@(description=(sdu=32768)(tdu=32768)(address=(protocol=sdp)(host=ucm-
ib-vip-1.us.oracle.com)(port=1521))(connect_data=(sid=ucmDB1)))
CSDS2 URL:
jdbc:oracle:thin:@(description=(sdu=32768)(tdu=32768)(address=(protocol=sdp)(host=ucm-
ib-vip-2.us.oracle.com)(port=1521))(connect_data=(sid=ucmDB2)))
Driver class name: oracle.jdbc.driver.OracleDriver
Properties: sendStreamAsBlob=true
oracle.net.CONNECT_TIMEOUT may need to be adjusted to 20000 (20 seconds) from the
default 10000 (10 seconds) to avoid connection timed out problem.
Initial Capacity: 100
Max Capacity: 300
Minimum Capacity: 50
Statement Cache Type: FIXED
Statement Cache Size: 50
Test Connection On Reserve: Enabled (Checked)
Test Table Name: SQL SELECT 1 FROM DUAL
Wrap Data Type: Disabled (Unchecked)
```

The 64-bit `libumem.so` was loaded for all WebCenter Content server instances, by making a change in `<domain>/bin/startWebLogic.sh`. In addition, the following should be added to `JVM_OPTIONS` in the same file to enable SDP.

```
-Djava.net.preferIPv4Stack=true -Doracle.net.SDP=true
```

To make sure that Oracle WebLogic Server native operating system libraries can be loaded by WebCenter Content instances without issue, it may be necessary to add an execution bit to all dynamic library files in `$WLS_HOME/server/native/solaris` and its sub-directories. In addition, HTTP access logs should disabled, and WebLogic Server's "Running Mode" should be set to "Production" in the WebLogic Server config file.

Oracle WebCenter Content Application Tuning

A number of tuning parameters were specified for the various use cases. For the file ingestion use case, the following parameters were specified in the `<domain>/ucm/cs/config/config.cfg` file.

```
NumConnections=1000
MaxRequestThreadCount=1000
MaxSearchConnections=1000
DisableTotalItemsSearchQuery=true
SkipAnalyzeServerConfigClassesDir=true
#SearchIndexerEngineName=ORACLETEXTSEARCH
IndexerDatabaseProviderName=SystemDatabase
OracleTextIndexingParallelDegree=32
OracleTextDisableSearchSnippet=true
DoDocNameOrder=false
DisableAuthorizationTokenCheck=true
DisableRecentQueriesList=1
DisableDrillDownParsingWithSAX=1
useFastCheckin=1
ServiceAllowRetry=true
```

The full-text search use case, set the `<domain>/ucm/cs/config/config.cfg` file as follows.

```
NumConnections=1000
MaxRequestThreadCount=1000
MaxSearchConnections=1000
DisableTotalItemsSearchQuery=true
SkipAnalyzeServerConfigClassesDir=true
SearchIndexerEngineName=ORACLETEXTSEARCH
IndexerDatabaseProviderName=SystemDatabase
OracleTextIndexingParallelDegree=32
OracleTextDisableSearchSnippet=true
DoDocNameOrder=false
DisableAuthorizationTokenCheck=true
DisableRecentQueriesList=1
DisableDrillDownParsingWithSAX=1
#useFastCheckin=1
ServiceAllowRetry=true
```

The following commands were added to `<domain>/ucm/cs/bin/intradoc.cfg`.

```
DisableDevChecks=true
TempDir=/tmp
```

Automatic indexing is not desirable for the file ingestion use case while it is useful for the full-text search use case. The automatic indexing cycle was disabled and re-enabled as appropriate by accessing the `UCM portal->Admin Applet->Repository Manager`.

JDBC Webless secured file storage was used for performance to use the database exclusively and not any other web or caching services, and only the following nine WebCenter Content system components were enabled.

```
CheckSCSHealth, ContentAccess-solaris, ContentAccess-solaris-x86,
FileStoreProvider, JpsUserProvider, NativeOsUtils, OracleQueryOptimizer,
ServletPlugin, YahooUserInterfaceLibrary
```

Network and OS Tuning

A number of TCP tunings were applied to the WebCenter Content nodes and the Oracle database nodes as follows.

```
# ndd -set /dev/tcp tcp_conn_req_max_q 16384
# ndd -set /dev/tcp tcp_conn_req_max_q0 16384
# ndd -set /dev/tcp tcp_time_wait_interval 20000
# ndd -set /dev/tcp tcp_max_buf 10485760
# ndd -set /dev/tcp tcp_cwnd_max 5242880
# ndd -set /dev/tcp tcp_recv_hiwat 655360
# ndd -set /dev/tcp tcp_xmit_hiwat 655360
# ndd -set /dev/ip ip_reass_queue_bytes 1310720
```

InfiniBand interface tuning was applied to both the WebCenter Content nodes and the nodes running the Oracle database as follows.

```
# dladm set linkprop <ib_int> -p _ibd_rc_num_swqe=4095
# dladm set linkprop <ib_int> -p _ibd_rc_num_rwqe=8000
# dladm set linkprop <ib_int> -p _ibd_rc_rx_copy_thresh=32768
```

The following tuning parameters were added to `/etc/system` on both the WebCenter Content nodes and the Oracle database nodes.

```
set rlim_fd_max=63536
set rlim_fd_cur=32768
set autoup=300
set zfs:zfs_arc_max=0x1680000000
# 256M pages
set max_uheap_lpsize=0x10000000
set mmu_ism_pagesize=0x10000000
set pg_contig_disable=1
# IB workaround to CR 7070731
set mac:mac_rx_fanout_inline_max=0
set shmsys:shminfo_shmmax=4294967295
```

The following NFS mount options were used for shares accessed from the Oracle Sun ZFS 7320 appliance.

```
vers=3,proto=tcp,rsize=131072,wsize=131072,rw
```

Oracle DatabaseTuning

A variety of tuning changes were made to the Oracle database running on the database nodes. Adequate shared memory is critical so the limit for Oracle users was increased.

```
project.max-shm-memory=(priv,268435456000,deny)
```

Common initialization parameters were used for both single-instance and Oracle RAC databases.

```
"_shared_io_pool_size"=512M
"_shared_iop_max_size"=512M
cursor_sharing=FORCE
db_block_size=8192
db_cache_size=16G
db_writer_processes=32
memory_target=48G
open_cursors=300
pga_aggregate_target=2G
session_cached_cursors=5
sga_target=0
shared_pool_size=8G
```

The following parameters were specific to the Oracle RAC database.

```
"_use_zero_copy_io"=FALSE
commit_write="IMMEDIATE, NOWAIT"
```

The following parameters were specific to the single-instance Oracle database.

```
commit_write="BATCH, NOWAIT"
```

Hot Tables and General Performance Guidelines

A number of guidelines can be provided for hot tables related to general performance.

- Several tables should be hash partitioned, along with the indexes created on those tables, including: `FILESTORAGE, DOCUMENTS, DOCUMENTHISTORY, DOCMETADATA, REVISIONS` and `REVCLASSES.`
- Adjust the `PCTFREE` and `INITRANS` table attributes as necessary to reduce contention and to allow multiple database sessions to access each of those tables simultaneously.
- It was found that having a large number of unused indexes on these `INSERT/UPDATE` heavy tables hurt the application performance. It is a recommended practice to monitor the index usage and remove the unused indexes from the database.

Generic Securefile Large Object (LOB) Guidelines

The following guidelines apply to dealing with Securefile Large Objects (LOBs):

- *BLOBs.* All the checked-in documents are stored as secured file BLOBs (Binary Large OBjects) in a UCM application table called `FILESTORAGE`. BLOBs should be created in a different tablespace than the rest of the database.
- *Block size.* A larger block size should be considered if the expected size of the checked-in documents is large. The default block size is 8 KB.
- *Chunk size.* A larger chunk size should be considered if larger documents are expected to be checked-in and retrieved. The default chunk size is 8 KB.
- *Inline or Out-of-line:* Choose "`DISABLE STORAGE IN ROW`" (out-of-line) if the average document size is expected to be greater than 4 KB. The default inlining is fine for smaller documents.
- `CACHE` *or* `NOCACHE.` Consider bypassing the database buffer cache (`NOCACHE`) if the checked-in documents are large in volume and not expected to be retrieved frequently.

- `COMPRESS` *or* `NOCOMPRESS`. Choose the `COMPRESS` option if storage capacity is a concern and a constraint as the option saves disk space at the expense of some performance overhead. In an Oracle RAC database environment, it is highly recommended to compress the checked-in documents (LOBs) to reduce the interconnect traffic.
- *De-duplication.* By default, duplicate documents are stored as a separate copy in the database. Choosing the `DEDUPLICATE` option enables sharing of the same data blocks for similar files, thus reducing storage overhead and simplifying storage management.
- *Partitioning.* Partition the parent table to maximize application performance.
- *Zero-Copy I/O protocol*. The Zero-Copy I/O protocol is enabled by default. Turning it off in an Oracle RAC database environment could be beneficial. Set the initialization parameter `_use_zero_copy_io=FALSE` to disable the Zero-Copy I/O protocol.
- *Shared I/O pool*. The database uses the shared I/O pool to perform large I/O operations on securefile LOBs. The pool is created using shared memory segments. If this pool is not large enough or if there is not enough memory available in the pool for a securefile LOB I/O operation, the database uses a portion of PGA until there is sufficient memory available in shared I/O pool. As a result, the shared I/O pool should be sized appropriately by monitoring the database during peak activity. Relevant initialization parameters include: `_shared_io_pool_size` and `_shared_iop_max_size.`

Oracle Exadata Specific Tuning Guidelines

The following tuning advice pertains to use of the Exadata database specifically.

- *Compression*. If the application database is hosted on an Exadata storage server, data can be compressed using the "Exadata Hybrid Columnar Compression" (EHCC) feature in the Oracle Database 11g. The tradeoff for compression is the slight performance overhead for the disk space.
- *Flash Cache.* The frequently referenced database objects in database smart flash cache should be pinned by setting the database object's storage class attribute `CELL_FLASH_CACHE to KEEP.`

Database NUMA Optimizations

There is no general recommendation for running the database with NUMA optimizations turned on or off. Those deploying a similar solution should experiment with and without database NUMA optimizations and decide whether they make sense. The relevant initialization parameter is `_enable_NUMA_optimization`. A value of "true" enables the optimization where as "false" disables optimization.

Database Patches

It is important to apply the database patch that contains the fix for the following bug"

- 12797420 - MUTEX X WAITS WITH TWO JVMS.

The bug fix eases the Mutex X contention in the library cache.

Conclusion

Oracle testing has verified the value of the Oracle Optimized Solution for Oracle WebCenter Content based on the SPARC SuperCluster. Key content management use cases such as high volume content ingest and full text search demonstrated substantial levels of performance. In addition, the minimal impact of no-added-cost encryption acceleration provided by the SPARC T4 processors lets organizations secure their data without costly or cumbersome processes. Perhaps most important, the SPARC SuperCluster provides fully tested and integrated infrastructure for Enterprise Content Management, demonstrating considerable head-room for future growth and scalability.

For More Information

For more information on the Oracle Optimized Solution for Oracle WebCenter Content on SPARC SuperCluster, see the references listed in Table 2.

TABLE 2. REFERENCES FOR MORE INFORMATION

Oracle Optimized Solutions	http://www.oracle.com/goto/optimizedsolutions
Oracle Optimized Solution for Oracle WebCenter Content on SPARC SuperCluster	http://www.oracle.com/us/solutions/solutions-webcenter-content-496763.html
Oracle Solaris 11	http://www.oracle.com/solaris
Oracle's Sun ZFS Storage Appliance	http://www.oracle.com/us/products/servers-storage/storage/unified-storage/sun-storage-7000-uss-103104.html
Oracle Technology Network: SPARC product technical information	http://www.oracle.com/technetwork/server-storage/sun-sparc-enterprise/overview/index.html

Oracle Optimized Solution for Oracle WebCenter Content on Oracle's SPARC SuperCluster Technical Whitepaper
October 2011
Author: Nick Kloski

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