

A large, abstract, light gray graphic on the left side of the page, consisting of several overlapping, curved shapes that create a sense of depth and movement.

# **SUN ORACLE 10G GRID REFERENCE ARCHITECTURE GUIDE**

White Paper

August 2006

## Table of Contents

<b>Chapter 1: Executive Summary</b> .....	1
Scope .....	1
<b>Chapter 2: Introduction</b> .....	2
IT Challenges and Grid Computing .....	2
The Sun Oracle 10g Grid Reference Architecture .....	3
Reference Architecture Components .....	3
Key Benefits .....	4
<b>Chapter 3: The Sun Oracle 10g Grid Reference Architecture</b> .....	6
Reference Architecture Design .....	6
Reference Architecture Components .....	7
Sun Fire™ x64 Servers .....	7
Sun Fire™ V240 Server .....	8
Sun StorEdge™ 3510 Fibre Channel Array .....	8
Solaris™ 10 Operating System .....	8
Sun Cluster Advanced Edition for Oracle Real Application Clusters .....	8
Sun N1™ System Manager Software .....	8
Oracle Real Application Clusters 10g .....	9
Oracle Enterprise Manager 10g Grid Control .....	9
Cisco SFS 7000 Series Product Family Element Manager .....	9
Cisco InfiniBand Technology .....	9
Connecting the Network, Servers, Applications, and Storage .....	10
Application Logical Architecture .....	12
System Components .....	12
<b>Chapter 4: Summary</b> .....	14
The Sun Advantage .....	14
<b>Chapter 5: References</b> .....	15
Web Sites of Interest .....	15
Books, Papers, and Guides .....	16

## Chapter 1

# Executive Summary

In the unending quest to handle the challenges associated with doing more with less, IT professionals are continually exploring new avenues for planning ahead, getting ahead, and staying ahead of business demands. One of the newest tools in the IT arsenal is grid computing, and one of the newest uses for grid computing is enterprise database implementations.

As with any software, deploying grid and database software across the enterprise mandates a solid infrastructure, along with best-fit hardware and software. The endless choices can be daunting. To ease this effort, Sun provides the Sun Oracle 10g Grid Reference Architecture. Designed around proof-of-concept deployments and tested, validated configurations, the Sun Oracle 10g Grid Reference Architecture provides a high performance, highly available database infrastructure in a cost-effective grid environment. Deployed on the latest hardware and software from Sun, and incorporating hardware and software from Oracle and Cisco Systems, this reference architecture includes failover mechanisms and redundancy to eliminate single points of failure. By using these advanced products and technologies, the Sun Oracle 10g Grid Reference Architecture delivers a robust, scalable database platform that maintains or increases service levels and helps ensure high system availability.

Built on the industry-leading Solaris™ Operating System (OS), the Sun Oracle 10g Grid Reference Architecture leverages state-of-the-art features found in the Solaris™ 10 OS. Among those features are indisputable performance advantages for database services, world record benchmark results, and optimizations that yield faster performance without requiring changes to existing applications. Running on Sun's new extremely fast Sun Fire™ x64 servers, and coupled with Sun Cluster™ software and Sun StorageTek™ QFS software, this reference architecture provides an ideal platform for a powerful database grid solution.

A key strength of this architecture, and one of its greatest benefits, is its ability to integrate complementary elements from both Sun and third-party vendors. A hallmark of Sun infrastructures, vendor interoperability fosters a more adaptable enterprise, and enables computing resources and bandwidth to be re-provisioned to meet changing business demands. By combining low cost servers with other architecture components, the Sun Oracle 10g Grid Reference Architecture balances cost, performance, and availability. The outcome — a flexible, scalable, and powerful solution with superb database and throughput, that can help reduce costs throughout the solution life cycle.

### Scope

This document is intended for executives evaluating the use of a horizontally scalable grid for database application deployment, as well as project managers, system engineers, integration specialists, and other technical staff looking to design, implement, and tune database applications in a grid computing environment. It provides an overview of database deployment challenges and options, and describes how the Sun Oracle 10g Grid Reference Architecture delivers on these solutions. An architectural overview and descriptions of the hardware and software components used in the reference architecture are also included. Note that only the features and techniques directly relevant to the reference architecture are discussed here. For more complete information on Sun servers and storage, Oracle Real Application Clusters (RAC) 10g, Cisco Systems switches, or other components of this architecture, please refer to the product documentation from Sun, Oracle, and Cisco Systems.

## Chapter 2

### Introduction

#### IT Challenges and Grid Computing

IT organizations are grappling with the challenges posed by providing a greater variety of services with fewer resources. Pressure is mounting on IT executives to deploy reliable, high performance database solutions within severely constrained capital and expense budgets. Expectations are high for horizontally scalable solutions that can maintain or increase application availability, reliability, and performance while supporting service levels and simultaneously reducing total cost of ownership (TCO). These requirements demand integrated solutions that are often increasingly complex to design, implement, and maintain, and require specialized technical and business expertise.

Rapid changes in business cycles and increasing demand for compute power are prompting IT organizations to look for alternative solutions. Today, grid computing is shifting from a mere concept to a viable computing framework. Now, IT organizations are turning to the grid computing model for solutions that can adapt to changing business needs, deliver greater efficiency, minimize costs, and provide investment protection and rapid return on investment (ROI). Organizations have high expectations for these new IT solutions, and the outcome of implementations is critical. As a result, IT organizations should consider several questions when evaluating the use of a grid computing solution:

- Is a low cost grid architecture a viable option?
- What are the cost factors in scaling a grid horizontally?
- How do additional connections, software licensing fees, and implementation and management of more complex systems impact solution costs?
- Is a grid solution capable of providing a reliable, horizontally scalable platform for running enterprise-grade commercial applications at acceptable service levels?
- What are the components of a grid architecture?
- How are applications distributed across servers?
- Where are the potential bottlenecks?
- In a database cluster, how does the system recover from a component failure?
- With respect to interconnect technology, how do InfiniBand and Gigabit Ethernet compare in terms of latency, bandwidth, and overhead?
- How can Sun products add value to a grid solution?
- How do Sun products complement Oracle RAC deployments?

In an effort to help customers address these questions, and to assist in evaluating a grid computing solution, Sun has invested tremendous resources in developing the Sun Oracle 10g Grid Reference Architecture. Sun Reference Architectures are integrated solutions consisting of specific combinations of hardware and software products that also incorporate third-party offerings from top tier software and networking equipment vendors. These solutions

include a set of recommended components that are sized, tested, tuned, and documented to work together, delivering specific functionality and reducing the complexity, costs, and risks of deploying new technology.

## The Sun Oracle 10g Grid Reference Architecture

The Sun Oracle 10g Grid Reference Architecture provides a high performance, highly available database infrastructure in a cost-effective grid environment. Deployed on Sun systems with AMD Opteron™ processors running the Solaris 10 Operating System (OS), the Sun Oracle 10g Grid Reference Architecture delivers a robust, scalable database platform that maintains or increases customer service levels with a low TCO. The architecture provides an optimal combination of cost savings, superb performance, and high availability with its ultimate value being the ability to synthesize these individual elements into an integrated architecture.

## Reference Architecture Components

Based on best-of-breed products from Sun and Oracle, the Sun Oracle 10g Reference Architecture consists of the following key components:

- Sun Fire x64 Servers — The Sun Fire x64 server line consists of high throughput, energy-efficient servers containing one to 16 AMD Opteron processors. Almost half the price of traditional servers, the Sun Fire x64 server line combines a low cost of entry with high performance, scalability, and versatility that helps to protect IT infrastructure investments.
- Solaris 10 Operating System — One of the most advanced operating systems available, the Solaris 10 OS is the latest version of Sun's industry-leading operating system. Interoperable with Linux and Windows, the Solaris OS also offers binary compatibility within each Sun server line, whether based on UltraSPARC®, AMD Opteron, or Intel Xeon processors. As a result, all Sun servers running the Solaris 10 OS provide powerful features that can help to reduce costs, complexity, and risk.
- Oracle RAC 10g — Designed for grid implementations, Oracle RAC 10g is Oracle's second-generation real application cluster database product. Providing database availability and flexibility, Oracle RAC 10g serves as a powerful foundation for enterprise database grids.
- Sun Cluster Advanced Edition for Oracle RAC — An end-to-end software stack comprised of Sun StorageTek QFS software, Sun Cluster RAC agent, and Solaris™ Volume Manager software, Sun Cluster Advanced Edition for Oracle RAC provides a complete solution for running a clustered Oracle database. Including the Sun Cluster Advanced Edition for Oracle RAC software in the reference architecture extends high availability for the Oracle RAC 10g database by building on the Sun Cluster software foundation.
- InfiniBand technology — InfiniBand is a high performance, switched fabric interconnect that fosters greater scalability and improved manageability. Providing increased bandwidth and low latency, InfiniBand technology helps accelerate database server speed and gives an added boost to clustered databases.

## Key Benefits

The Sun Oracle 10g Grid Reference Architecture provides a horizontally scalable solution that delivers high performance and service levels at low cost. Using best-of-breed products from Sun, Oracle, and Cisco, the Sun Oracle 10g Reference Architecture also offers the following benefits.

- *Utilize a proven solution portfolio*

Following the reference architecture assures a valid blueprint for deploying Oracle RAC 10g in a grid computing environment with the rigor of a Sun tested, certified, and supported configuration. The reference architecture also formulates best practices for an Oracle RAC 10g infrastructure with other best-of-breed third-party products within a grid context.

- *Maximize availability*

The Sun Oracle 10g Grid Reference Architecture maximizes database availability by leveraging Oracle RAC 10g and deploying redundant components to eliminate single points of failure. Using a clustered grid infrastructure, the Sun Oracle 10g Grid Reference Architecture employs built-in load balancing and failover mechanisms to minimize interruption of service levels due to individual component failure. Oracle RAC 10g enables databases to span across multiple physical servers, eliminating the server as a single point of failure and ensuring high availability for applications and users.

- *Maximize reliability*

Implementing Sun Cluster software enhances the reliability and availability of Oracle RAC 10g and helps improve the overall resiliency of the architecture. Sun Cluster software offers the advantages of a well-proven and solid cluster framework that is tightly integrated with the Solaris 10 OS kernel for the quick detection and handling of failures. Combined with an effective I/O fencing mechanism, full redundancy, and load-balancing of interconnect traffic, the Sun Cluster software helps ensure a secure and dependable Oracle environment on the Solaris 10 OS.

- *Maximize performance*

Sun's low cost, high performance Sun Fire x64 servers with AMD Opteron processors deliver the core computing power behind the reference architecture. These servers can be equipped with single- or dual-core AMD Opteron CPUs, providing extreme versatility and setting a new standard for server price/performance, availability, and total cost of ownership.

- *Provide flexible scalability*

Today's business needs can change in an instant, and these rapid changes often translate to increased system demands and higher service level agreements. System infrastructures must allow organizations to increase system capacity and throughput by *scaling up* or *scaling out* — deploying larger servers, adding a number of small servers, or both. The ability of Oracle RAC 10g to span the database over multiple physical servers provides the flexible scaling so vital to the architecture design.

- *Minimize costs*

IT organizations are increasingly pressured to do more with less. As a result, organizations are always looking for opportunities to reduce overall implementation costs, including costs of acquisition, deployment, and maintenance. The Sun Oracle 10g Reference Architecture aims to address budgetary concerns by utilizing low cost, robust commodity hardware and software.

- *Simplify management*

One of the challenges of grid computing is managing the additional complexity of a grid infrastructure. Comprehensive management framework solutions, including the Oracle Enterprise Manager 10g Grid Control and Sun N1™ System Manager software, are incorporated in the reference architecture. These tools provide click-of-a-button, automated, self-management capabilities that can be leveraged to manage grid complexity, simplify support, and drive down operating costs.

## Chapter 3

# The Sun Oracle 10g Grid Reference Architecture

### Reference Architecture Design

Throughout the process of designing the Sun Oracle 10g Grid Reference Architecture, certain requirements have been paramount. These requirements include reliability, high availability, resiliency, cost-effectiveness, a high degree of scalability, and use of best-of-breed products. By adhering to these fundamental criteria, and investing extensive effort in development, integration, testing, and tuning the architecture, Sun technologists have produced a database grid ideally suited for enterprise applications. Figure 3-1 depicts the reference architecture implementation.

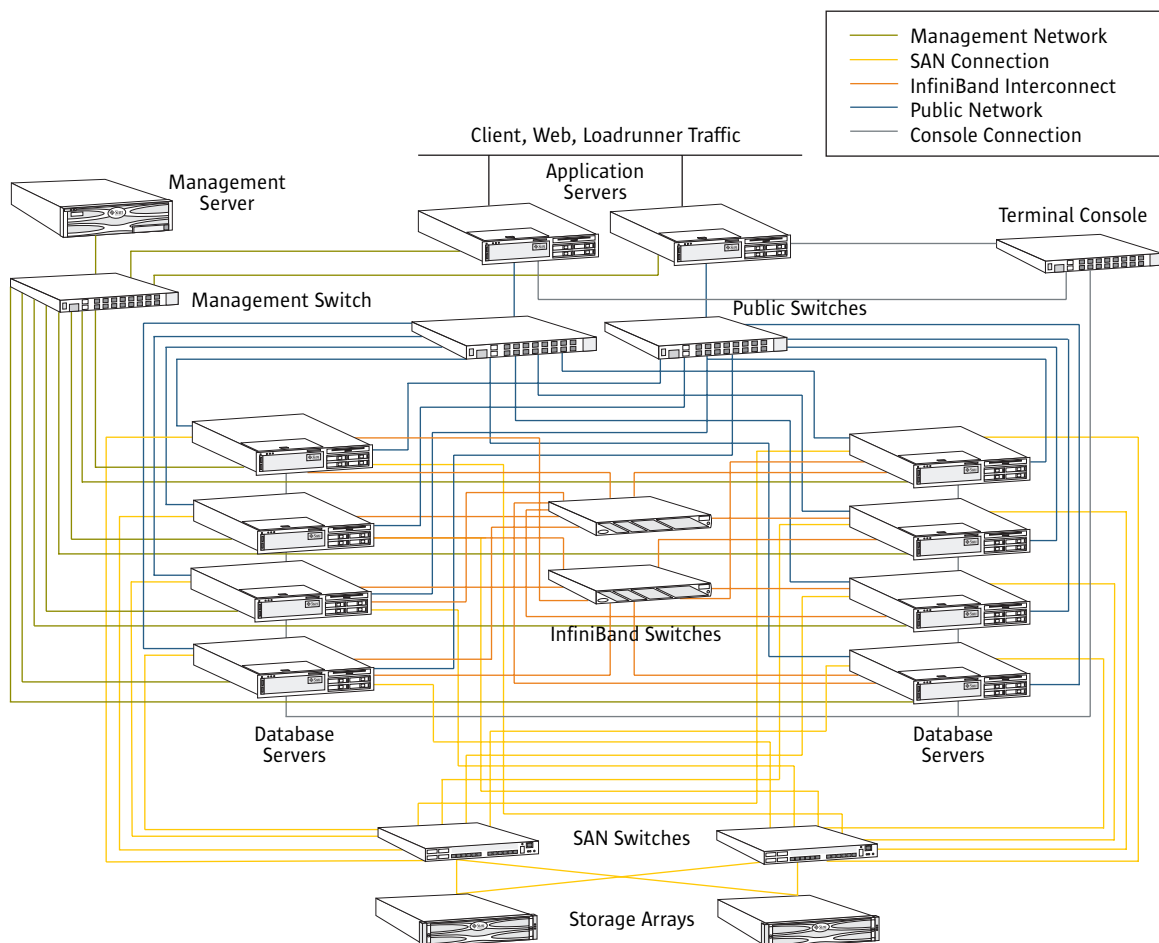


Figure 3-1. The Sun Oracle 10g Grid Reference Architecture implementation



## Reference Architecture Components

The Sun Oracle 10g Reference Architecture combines quality products from top-tier vendors (Table 3-1). By utilizing fully redundant components, the reference architecture delivers a highly robust, available, and reliable cluster grid infrastructure at a low acquisition cost. With ultra-dense, rack-optimized Sun Fire x64 servers that incorporate next-generation AMD Opteron x64 processors, as well as industry-standard, low latency, high throughput InfiniBand interconnect technology, the reference architecture enables a high performance, scalable database platform.

*Table 3-1. Hardware and software components used in the Sun Oracle 10g Reference Architecture*

Hardware Components	Software Components
Sun Fire x64 servers with AMD Opteron processors (Database and application tier)	Solaris 10 Operating System
Sun Fire™ V240 server (Management server)	Oracle Real Application Cluster 10g R2
Sun StorEdge™ 3510 Fibre Channel Array (Configured with dual RAID controllers)	Sun Cluster Advanced Edition for Oracle RAC software (Sun StorageTek QFS shared file system, Sun Cluster RAC agent, Solaris Volume Manager)
Cisco InfiniBand server switches	Oracle Enterprise Manager 10g Grid Control
Sun Storage Area Network (SAN) switches	Sun N1 System Manager software
Gigabit Ethernet switches	Cisco SFS 7000 Series Product Family Element Manager
Terminal console	

### Sun Fire™ x64 Servers

Sun Fire x64 servers deliver the next leap in enterprise-wide IT return on investment. Combining extreme performance, reliability, serviceability, and flexibility with the ability to run 32-bit and 64-bit operating systems and applications, Sun Fire x64 servers help protect IT infrastructure investments while enabling the migration to next-generation 64-bit computing. Additional features include high rack density, redundant power and cooling, RAID storage to help ensure continued server operation, and Lights Out Management for remote system monitoring, control, and easy firmware upgrades. With up to sixteen dual-core or single-core AMD Opteron processors, the Sun Fire x64 servers deliver exceptional versatility in a small footprint, making them an ideal platform for Oracle RAC deployments.

Sun now offers the following x64 servers:

- The Sun Fire™ X2100, Sun Fire™ X2200, Sun Fire™ X4100, Sun Fire™ X4200, and Sun Fire™ X4600 servers, ranging from entry-level to enterprise-class rackmount servers with superior scalability and minimal complexity.
- The Sun Fire™ X4500 server, offering massive data storage with extremely high throughput rates at nearly half the cost of traditional solutions.
- The Sun Blade 8000 Modular System, providing an innovative modular approach to data center requirements that delivers up to three times the throughput of typical blade computers while using 20 percent less power and 30 to 50 percent less space than standard rack mount servers.

### **Sun Fire™ V240 Server**

Built on the Solaris platform and SPARC® architecture, the Sun Fire V240 server is a rack-optimized server designed for high availability. It maximizes compute power and data center space, boosts high speed network connectivity, and provides a secure, flexible platform for expansion. With embedded Sun Advanced Lights Out Manager (ALOM) capabilities, the Sun Fire V240 can be managed remotely in a *lights out* environment. These features, combined with redundant configuration capabilities, make the Sun Fire V240 ideal for running grid management framework software.

### **Sun StorEdge™ 3510 Fibre Channel Array**

The powerful Sun StorEdge 3510 Fibre Channel array uses a modular, building-block approach to help reduce costs. The array offers affordable enterprise-class features and functionality, such as dual hot-swap power and cooling, hot-swap redundant RAID controllers with mirrored cache, hot-swap disk drives, global and local hot sparing, dynamic LUN expansion, dynamic capacity expansion, non-disruptive firmware code loading, and remote status monitoring. The highly intuitive management interface makes the array extremely easy to deploy, configure, manage, and monitor. In addition, utilizing the Sun StorEdge 3510 Fibre-Channel arrays in the reference architecture offers reliability and impressive I/O performance in a small footprint.

### **Solaris™ 10 Operating System**

The latest version of Sun's industry-leading operating system, the Solaris 10 OS provides powerful, leading-edge features that can help reduce costs, complexity, and risk. It offers over 600 exciting new features — including innovative functionality such as Solaris™ Containers technology, Solaris™ Dynamic Tracing (DTrace), Solaris™ ZFS, and Predictive Self-Healing — that support the relentless availability, extreme performance, and unparalleled security today's enterprises demand. Capable of vertical or horizontal scaling, the Solaris 10 OS can help businesses leverage the benefits of low cost hardware, yet standardize on a single operating system. In addition, the Solaris 10 OS offers broad hardware platform support and expanding interoperability with Linux and Windows, making it the ideal operating system for grid deployments.

### **Sun Cluster Advanced Edition for Oracle Real Application Clusters**

Sun Cluster Advanced Edition for Oracle RAC provides enterprises with additional high availability for Oracle Real Application Clusters on the Solaris 10 OS. In addition to providing I/O fencing to help guarantee data integrity, this product suite leverages two key components of the industry-leading Sun Cluster software to deliver high performance. The Sun StorageTek QFS shared file system software helps eliminate administrative overhead and provides maximum data management and throughput for the most data-intensive applications. The Solaris Volume Manager software provides partitioning and automatic device relocation for data mirroring and high availability. Together, these components result in a high performance clustered file system for traditional file system management and administration. Furthermore, the Sun Cluster Advanced Edition for Oracle RAC is designed for tighter integration with the Solaris OS to help minimize application downtime.

### **Sun N1™ System Manager Software**

Provided at no cost as part of the Solaris Enterprise System, the Sun N1 System Manager software addresses the challenges associated with managing system infrastructures. It enables administrators to discover, provision, monitor, update, and manage hundreds of Sun Fire x64 and UltraSPARC servers from a single management console

anywhere on the Web. The Sun N1 System Manager offers an innovative and user friendly hybrid graphical and command line interface (CLI) to manage each step of the infrastructure life cycle, providing increased visibility into the operations performed across selected systems.

### **Oracle Real Application Clusters 10g**

Oracle Real Application Clusters (RAC) 10g is the first database designed for grid implementations. A proven technology that allows multiple, low cost servers to perform like a single large server, Oracle RAC forms a key foundation for enterprise database grids. With Oracle RAC, even very high end systems can be constructed of small, very low cost clusters made from standard, commodity parts. Running Oracle RAC on a cluster provides the highest level of database availability along with flexibility in scaling. If a node in the cluster fails, the Oracle software continues running on the remaining nodes. If more processing power is necessary, new nodes can easily be added to the cluster. As a result, organizations can deploy low cost solutions while reaping the benefits of high availability and adaptability. For more information, see <http://www.oracle.com/database/>.

### **Oracle Enterprise Manager 10g Grid Control**

The Oracle Enterprise Manager 10g Grid Control software enables administrators to easily manage business applications, end user services, and the entire grid infrastructure. Containing improved tools such as new service modeling, broader support for service protocols, and comprehensive policies and templates, Oracle Enterprise Manager 10g Grid Control enables management of grids as a single entity. Graphical service topologies and service dashboards provide the right information to high level decision makers and line of business managers for more effective planning, while automated provisioning and patching functionality give IT administrators the power to make efficient, error free changes.

### **Cisco SFS 7000 Series Product Family Element Manager**

The Cisco SFS 7000 Series Product Family Element Manager provides a graphical user interface (GUI) that enables users to monitor, configure, and maintain individual switches, as well as handle fault management, troubleshooting, and upgrades. The Sun Oracle 10g Grid Reference Architecture incorporates the Element Manager as a key part of the management network for the grid infrastructure.

### **Cisco InfiniBand Technology**

InfiniBand is an interconnect and I/O technology designed to address the needs of computing platforms for increased performance, higher scalability, and improved manageability. An open industry standard, InfiniBand is capable of remaking mainstream business computing by unleashing these capabilities and delivering on the promise of a single, unifying I/O fabric for the data center.

The Cisco InfiniBand switching solution is a new network infrastructure expressly designed for clustered databases. This intelligent fabric accelerates database server speed to 10 Gbps and slashes latency by an order of magnitude, resulting in improved performance at lower cost while enabling plug-and-play interoperability with legacy systems and databases. When used as the database cluster interconnect, InfiniBand technology can improve database performance by providing greater bandwidth and lower latency than traditional Gigabit Ethernet technology. This is achieved by bypassing operating system and transmission control protocol (TCP) overhead, speeding up inter-node communication for activities such as synchronizing a distributed lock manager or scanning tables in parallel, freeing up significant CPU cycles across the network.

## Connecting the Network, Servers, Applications, and Storage

Interconnect latency and throughput are essential elements to the scalability of a database grid cluster. Inter-node communication and cache fusion traffic can benefit from low latency, high throughput, low system overhead, such as that offered by InfiniBand technology. Deploying InfiniBand further aids integration of the Sun Oracle 10g Grid Reference Architecture elements since it leverages the high performance Reliable Datagram Sockets (RDS) protocol, a very low latency remote direct memory access (RDMA) based protocol driver. This RDS protocol can be used to create a unified fabric for all types of networking traffic, including inter-process communication (IPC), network, and storage traffic. The use of a single InfiniBand host channel adapter (HCA) simplifies network cabling and administration, and minimizes the associated management challenges of dealing with separate connections for IPC, local area networks (LAN), and storage area networks (SAN).

---

Note – At the time of the reference architecture development, storage connections were maintained as direct Fibre Channel Arbitrated Loop (FC-AL) connections to the storage subsystems without going through the InfiniBand fabric. However, it is anticipated that the Sun Oracle 10g Grid Reference Architecture can be extended and enhanced to take advantage of improvements and additions like the RDS protocol as they become available.

---

Figure 3-2 illustrates the two key network connection areas: the private interconnect and public network. Using high throughput InfiniBand technology for the private interconnects between the database servers speeds cache fusion traffic and Oracle RAC inter-node communication, and results in very low latency. The public network connects and enables traffic between the application and database servers.

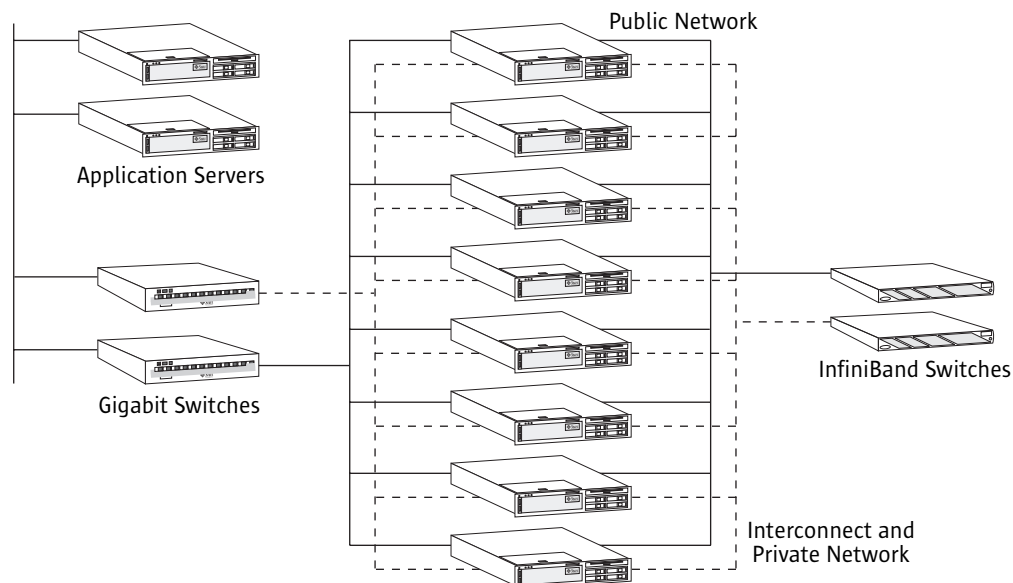


Figure 3-2. The private interconnect and the public network

The management network is reserved for telnet access, basic health and monitoring of components, connection to the Service Processor module, and monitoring and communication by the Oracle Enterprise Manager 10g Grid Control and Sun N1 System Manager software (Figure 3-3). The terminal console is connected to the built-in management port of the servers to enable console access.

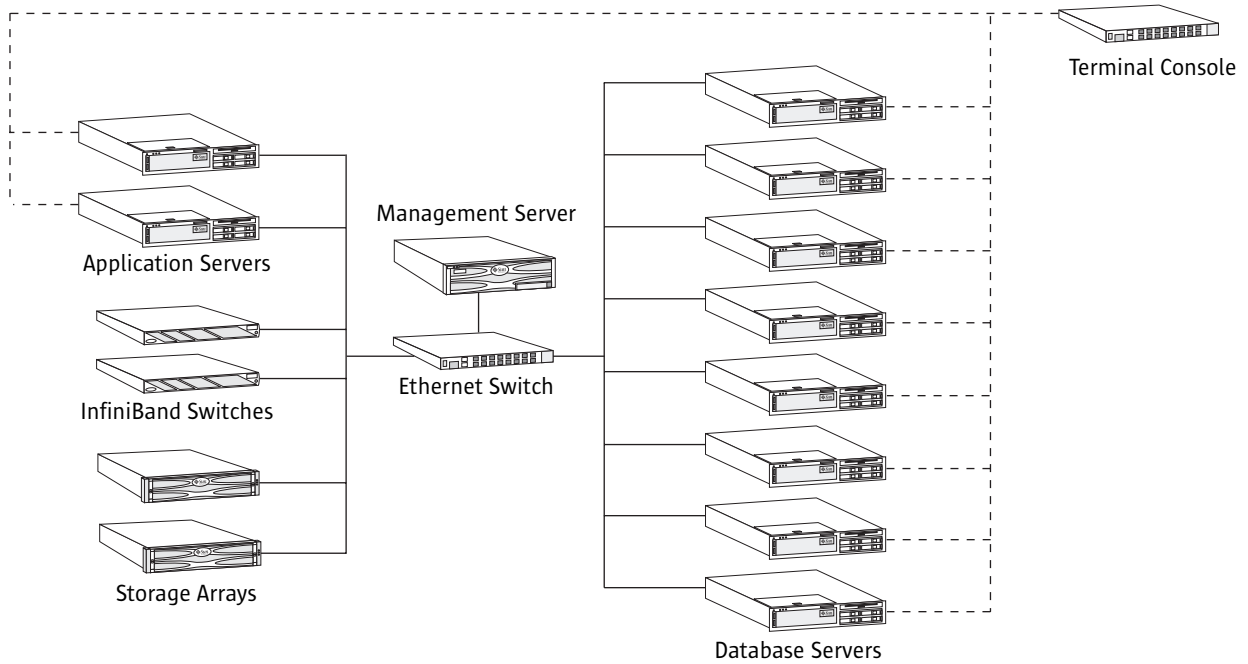


Figure 3-3. The management network

All eight database nodes share a pair of Sun StorEdge 3510 FC arrays via two SAN switches (Figure 3-4). This configuration, together with multiplexed I/O (MPxIO), helps provide full redundancy in the event of failures.

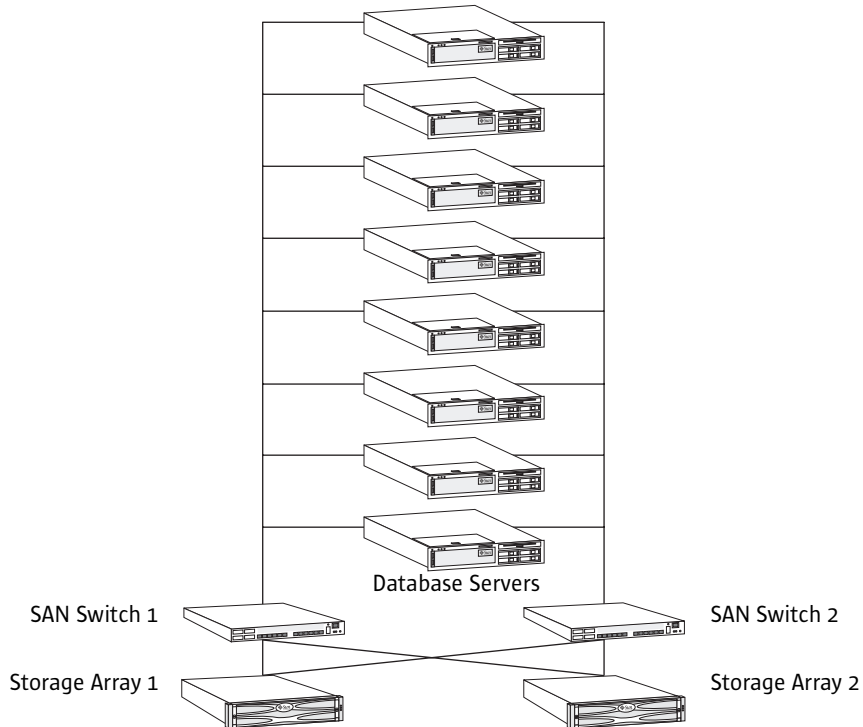


Figure 3-4. The database servers share the two storage arrays through the SAN switches

### Application Logical Architecture

Figure 3-5 represents the logical structure of the Sun Oracle 10g Grid Reference Architecture. The management software stack meets the critical need for a grid management framework by leveraging software from the Oracle Enterprise Manager 10g Grid Control, the Sun N1 System Manager, and the Cisco SFS 7000 Series Product Family Element Manager. In addition, other key components form the foundation of the logical architecture, including the Solaris 10 OS, Sun Cluster Advanced Edition for Oracle RAC software, and Oracle RAC 10g. Oracle Clusterware provides the database cluster framework, and is enhanced by the use of the Sun Cluster software, a robust and proven cluster framework that is now part of the Solaris OS kernel. Oracle datafile storage options can be raw devices through Oracle Automatic Storage Management (ASM), the Sun Cluster file system with Solaris Volume Manager and Sun StorageTek QFS shared file system software, or both.

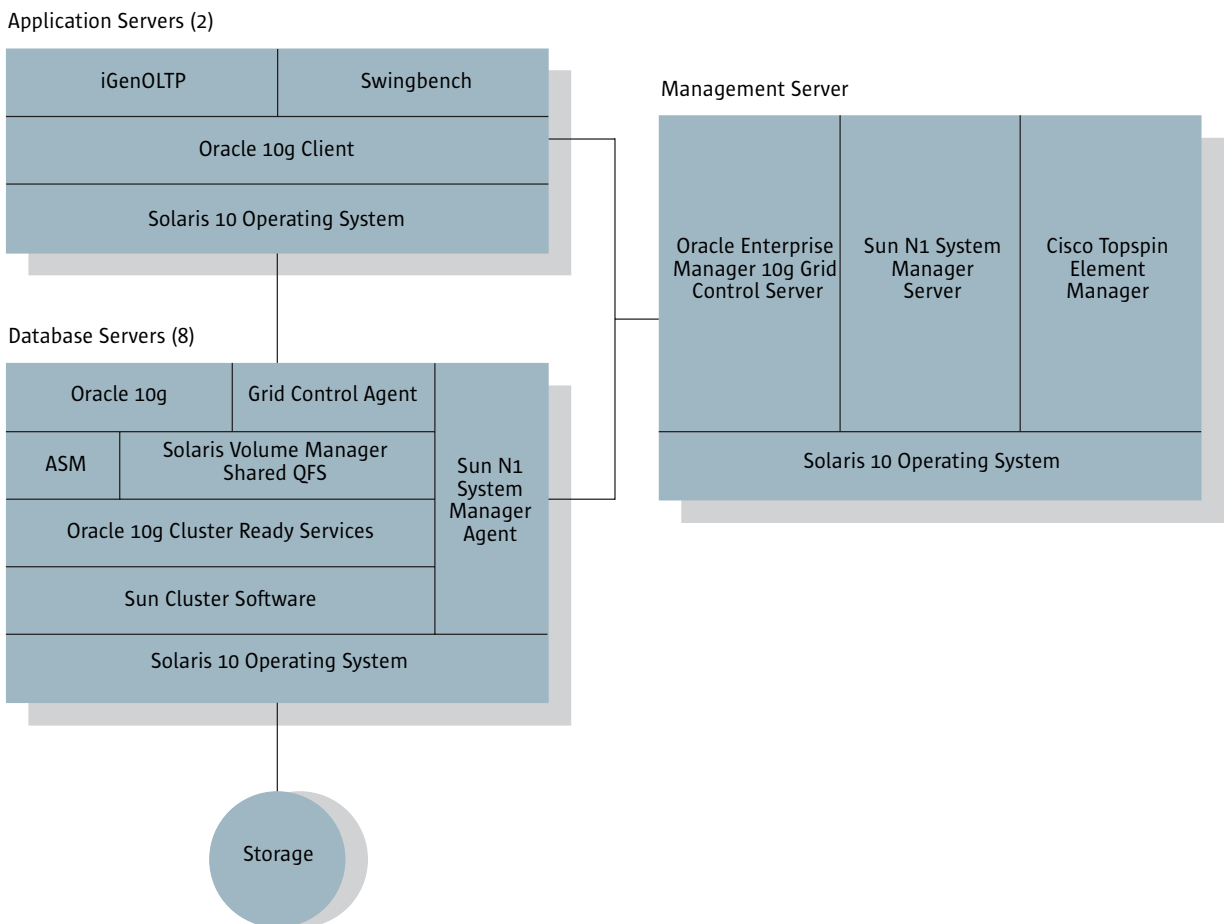


Figure 3-5. The Sun Oracle 10g Grid Reference Architecture logical structure

### System Components

Table 3-2 details the reference architecture hardware and software components deployed, tested, and tuned in the Authorized Sun<sup>SM</sup> Solution Center by Sun Competency Center for Oracle staff. Hardware and software specifications

are for models and releases available at the time the reference architecture was developed. Note, however, that any Sun Fire x64 server can be used for the application and database servers. For updates on available products, please refer to individual vendors' Web sites.

Sun and Oracle technologists at the Sun Competency Center for Oracle in Menlo Park, California designed, integrated, tested, and tuned the reference architecture for arbitrary workloads generated with the iGenOLTP and Swingbench tools. Developed at Sun, iGenOLTP is a benchmark suite that mimics an online customer order entry application. Swingbench is a free Oracle database load generation tool used in demonstrating load balancing and session failover.

*Table 3-2. Specific reference architecture components*

<b>Servers</b>	<b>Management Server</b>	Sun Fire V240 (2 x UltraSPARC® IIIi CPU, 6 GB RAM, 2 x 36 GB Drives)
	<b>Application Servers</b>	2 x Sun Fire™ x64 servers (4 x AMD Opteron CPU, 8 GB RAM, 2 x 72 GB Drives)
	<b>Database Servers</b>	8 x Sun Fire x64 servers (4 x AMD Opteron CPU, 8 GB RAM, 2 x 72 GB Drives)
<b>Storage</b>	<b>Storage Arrays</b>	2 x Sun StorEdge 3510 FC arrays
<b>Network</b>	<b>Interconnect Switches</b>	2 x Cisco Topspin 90 12-port InfiniBand switches
	<b>Public Switches</b>	2 x 24-port gigabit switches
	<b>Management Switch</b>	24-port gigabit switch
	<b>SAN Switches</b>	2 x 16-port Sun SAN switch
	<b>Adapter Cards (for Database Servers)</b>	8 x Dual-port gigabit cards 8 x Dual-port Fibre-Channel cards 8 x Dual-port Sun InfiniBand cards
	<b>Terminal Console</b>	16-port terminal console
	<b>Software</b>	<b>Operating System</b>
<b>Cluster Software</b>		Oracle 10g R2 (10.2.0.2) Cluster Ready Services Sun Cluster Advanced Edition for Oracle RAC (SC3.1U4)
<b>Database Software</b>		Oracle 10g R2 (10.2.0.2) RDBMS
<b>Management Software</b>		Oracle Enterprise Manager 10g R2 Grid Control Sun N1 System Manager 1.2 Cisco SFS 7000 Series Product Family Element Manager
<b>Volume Manager Software</b>		Oracle Automatic Storage Management Solaris Volume Manager and Sun StorageTek QFS Shared File System 4.5
<b>Workloads</b>		iGenOLTP Swingbench 2.2
	<b>Miscellaneous</b>	Oracle 10g R2 Client Oracle 10g R1 Grid Control Agent

Note – Cluster configurations typically require the use of two individual networking cards for redundancy. The Sun Cluster architecture supports the deployment of dual-ported cards, providing needed redundancy at reduced cost.

## Chapter 4

# Summary

### The Sun Advantage

IT managers are always seeking fresh ways of delivering new and enhanced services — even while being forced to maintain or reduce budget levels. Because database computing is a technology that constantly demands more compute power, expansion capabilities, performance, and availability, IT managers are turning to emerging technologies like grid computing to help deliver the capabilities needed while meeting budget constraints.

Today, grid computing increasingly is being adopted as a solution for handling growth, performance, and availability concerns. However, issues related to planning, implementing, and managing grid environments are emerging. The potential for disparate, difficult to manage elements is enough to deter any IT executive. Coordinating hardware, software, and networking equipment is a daunting task. In order to minimize the uncertainty of implementing a grid computing solution, Sun has taken on the effort of evaluating best-of-breed hardware and software and creating an optimized grid computing solution for Oracle databases. The result of this effort is the Sun Oracle 10g Grid Reference Architecture.

The Sun Oracle 10g Grid Reference Architecture is a portfolio of Sun and third-party products designed to reduce the risks, uncertainty, and costs associated with implementing an Oracle database within a grid computing environment. Designed, tested, implemented, and tuned at the Sun Competency Center for Oracle in Menlo Park, California, the reference architecture consists of recommended, integrated hardware and software stacks for a proven grid database solution. Along with this architecture, Sun provides a best practices framework for obtaining the highest possible performance, availability, and resiliency.

The most compelling benefits of the Sun Oracle 10g Reference Architecture come not from the individual components, but from the ability to integrate complementary elements — whether they come from Sun or another vendor. The combination of Sun Fire x64 servers with AMD Opteron processors and Cisco InfiniBand technology ensures higher database throughput, while the adoption of Sun Cluster Advanced Edition for RAC enhances the reliability and availability of Oracle RAC 10g. The result is a solution that can help deliver the highest service levels and satisfy the most demanding business requirements of a database grid.

For more information on applying this technology solution to an Oracle database environment, contact a local Sun sales representative.



## Chapter 5

# References

### Web Sites of Interest

	Web Site URL	Description
<b>Hardware</b>	<a href="http://sun.com/x64">http://sun.com/x64</a>	Sun Fire x64 Servers
	<a href="http://sun.com/servers/entry/v240/">http://sun.com/servers/entry/v240/</a>	Sun Fire V240 Server
	<a href="http://sun.com/x64/resources.jsp">http://sun.com/x64/resources.jsp</a>	64-Bit Computing
<b>Clustering</b>	<a href="http://sun.com/software/cluster/ds/ds-cluster31">http://sun.com/software/cluster/ds/ds-cluster31</a>	Sun Cluster
	<a href="http://sun.com/software/cluster/3reasons.xml">http://sun.com/software/cluster/3reasons.xml</a>	Sun Cluster and Oracle
	<a href="http://sun.com/software/cluster/osp/">http://sun.com/software/cluster/osp/</a>	Sun Cluster Open Storage Program
<b>Software</b>	<a href="http://sun.com/solaris">http://sun.com/solaris</a>	Solaris Operating System
	<a href="http://oracle.com/database">http://oracle.com/database</a>	Oracle Database 10g
	<a href="http://oracle.com/clusters">http://oracle.com/clusters</a>	Oracle Real Application Clusters
	<a href="http://sun.com/software/products/system_manager/">http://sun.com/software/products/system_manager/</a>	Sun N1 System Manager
	<a href="http://sun.com/third-party/global/oracle/grid/">http://sun.com/third-party/global/oracle/grid/</a>	Sun and Oracle Joint Grid Initiative
<b>Storage</b>	<a href="http://sun.com/storage">http://sun.com/storage</a>	Sun Storage
	<a href="http://sun.com/storagetek/disk_systems/workgroup/3510/">http://sun.com/storagetek/disk_systems/workgroup/3510/</a>	Sun StorageTek 3510 FC Array
<b>Partners</b>	<a href="http://sun.com/third-party/global/oracle/">http://sun.com/third-party/global/oracle/</a>	Sun and Oracle Global Partnership
<b>Benchmarks</b>	<a href="http://sun.com/third-party/global/oracle/collateral/index.html#benchmarks">http://sun.com/third-party/global/oracle/collateral/index.html#benchmarks</a>	Sun and Oracle Benchmarks
	<a href="http://sun.com/x64/benchmarks">http://sun.com/x64/benchmarks</a>	Sun x64 Systems — Benchmarks

## Books, Papers, and Guides

	Title	Web Site URL or ISBN
<b>x64 Computing</b>	Server White Papers	<a href="http://sun.com/servers/wp.jsp">http://sun.com/servers/wp.jsp</a>
	Sun Fire™ X4500 Server Architecture	<a href="http://sun.com/servers/x64/x4500/arch-wp.pdf">http://sun.com/servers/x64/x4500/arch-wp.pdf</a>
	Sun Fire™ X4100 and Sun Fire™ X4200 Server Architectures	<a href="http://sun.com/servers/entry/x4100/x64-whitepaper.pdf">http://sun.com/servers/entry/x4100/x64-whitepaper.pdf</a>
<b>Clustering</b>	Designing Enterprise Solutions with Sun Cluster 3.0	ISBN 0-13-008458-1
	Developing Agents for Applications Running on Sun Cluster Software	<a href="http://sun.com/software/whitepapers/wp-clusterapi/Cluster3.0API_final-wp.pdf">http://sun.com/software/whitepapers/wp-clusterapi/Cluster3.0API_final-wp.pdf</a>
<b>Management</b>	Sun N1 System Manager Documentation	<a href="http://sun.com/software/n1gridsystem/docs.xml">http://sun.com/software/n1gridsystem/docs.xml</a>
<b>Solaris OS</b>	Solaris 10 Whitepapers	<a href="http://sun.com/software/solaris/whitepapers.xml">http://sun.com/software/solaris/whitepapers.xml</a>
<b>Storage</b>	Storage White Papers	<a href="http://sun.com/storage/white-papers/">http://sun.com/storage/white-papers/</a>
	Improving Backup and Recovery Strategies	<a href="http://sun.com/storage/white-papers/improving_backup_recovery.pdf">http://sun.com/storage/white-papers/improving_backup_recovery.pdf</a>
	Storage Consolidation	<a href="http://sun.com/storage/white-papers/cost_risk_complexity.pdf">http://sun.com/storage/white-papers/cost_risk_complexity.pdf</a>
<b>Sun and Oracle</b>	White Papers on Sun and Oracle	<a href="http://sun.com/third-party/global/oracle/collateral/index.html#whitepapers">http://sun.com/third-party/global/oracle/collateral/index.html#whitepapers</a>

Copyright © 2006 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A.

All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Sun, Sun Microsystems, the Sun logo, Sun Fire, Sun StorEdge, Sun Cluster, UltraSPARC, StorageTek, N1, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon architecture developed by Sun Microsystems, Inc.

The AMD logo and Opteron are trademarks or registered trademarks of Advanced Micro Devices, Inc.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a). DOCUMENTATION IS PROVIDED AS IS AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS HELD TO BE LEGALLY INVALID.