

Sun™ Oracle 10g Grid Reference Architecture

Optimizing Scalability and Performance

Highlights

- The Sun™ Oracle 10g Grid Reference Architecture offers a flexible and powerful solution with superb database performance and throughput.
- Building on the proven Sun Cluster software foundation, Sun Cluster Advanced Edition for Oracle RAC provides an end-to-end software stack for running a clustered Oracle database.
- The Solaris™ 10 Operating System (OS) is the most advanced operating system available, providing key technologies that help ensure scalable and reliable infrastructure.
- Sun Fire™ x64 rackmount servers and the Sun Blade™ 8000 modular system offer a wide range of scalable x64 platforms.
- Sun Fire servers based on CoolThreads™ technology, such as the Sun Fire T1000 and T2000 servers, provide dramatic scalability for Oracle 10g RAC while delivering breakthrough space and power efficiency.
- Sun StorageTek™ storage brings value, access, simplicity, and trust to data storage



Enterprise grids from Oracle provide scalability on demand, allowing organizations to rapidly scale out their databases using small increments of cost-effective compute power. By adding resources on demand, organizations gain agility as well as the ability to accommodate spikes in demand for application services. Most importantly, Oracle Database 10g Enterprise Edition supports highly available deployments of large numbers of concurrent users, letting applications scale quickly from tens to tens of thousands of online users as needs dictate.

Enterprise grid challenges

Even with the considerable advantages of enterprise grids, providing suitable and effective infrastructure for Oracle 10g Real Application Clusters (RAC) remains challenging. The move to low-cost x86/x64 hardware has often shifted many of the integration issues from vendors to in-house operational staff. The endless system, networking, and middleware configuration choices can be daunting and complex. Even with a highly available database such as Oracle 10g, failure to properly architect infrastructure can result in unplanned downtime that can be catastrophic for the business.

For many vendors, it is enough to ship rack-mount systems and say they work with Oracle. Unfortunately, this approach can create new complexities, leaving extensive engineering and integration tasks for the customer. For instance, while Oracle Database 10g RAC provides a highly available database, it assumes that highly available supporting network and storage infrastructure exists. Organizations deploying or consolidating Oracle 10g RAC face a number of challenges:

- Designing and configuring systems, networking, high-speed interconnects, and middleware for maximum performance and availability while minimizing complexity

- Ensuring that consolidation efforts result in systems that support business needs in terms of performance and scalability
- Providing highly available infrastructure for other related applications and middleware without duplication and undue complexity

With over 20 years of collaboration and alignment between Sun and Oracle, and a shared vision for grid and network computing, customers enjoy considerable peace of mind. Working closely with Oracle, Sun takes a comprehensive approach to enterprise grid infrastructure, providing the key components that strengthen Oracle Database 10g installations while reducing time to deployment and lowering risk.

To address the need for scalable and reliable grid infrastructure, Sun provides database solutions that carefully consider elements such as software, hardware, and network infrastructure. Sun complements the built-in availability features of Oracle Database 10g Enterprise Edition and RAC with the tested and proven Sun Oracle 10g Grid Reference Architecture. These tools, plus a broad range of capable and innovative x64 and UltraSPARC® systems, help take the risk out of grid deployments, letting organizations get the most from their infrastructure investments.

The Sun Oracle 10g Grid Reference Architecture

Sun Reference Architectures have been designed, tested, tuned, and documented so that customers can reduce the complexity, costs, and risks of deploying new technology in the enterprise. Reference Architectures include:

- A documented multitiered architecture
- Recommended technology products from Sun and other vendors
- Technical documents that guide architecture, sizing, and implementation

Before choosing to implement the Sun Oracle 10g Grid Reference Architecture, customers can also run a proof-of-concept at the Sun Solution Center for Oracle Competency. For more information on Sun Reference Architectures, see sun.com/service/refarch.

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 either AMD Opteron™ or UltraSPARC processors and running the Solaris Operating System (OS), the Sun Oracle 10g Grid Reference Architecture delivers a robust and 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. Key components of the Sun Oracle 10g Grid Reference Architecture are depicted in Figure 1.

Sun Cluster Advanced Edition for Oracle RAC

Sun Cluster Advanced Edition for Oracle RAC provides a complete solution for running a clustered Oracle database, saving organizations the complicated task of integrating essential high-availability systems, networking, and storage. Including the Sun Cluster Advanced

Edition for Oracle RAC in the reference architecture extends high-availability for the Oracle RAC 10g database and also provides key high availability technology for related application components. Elements include:

- Sun Cluster RAC agent: The Sun Cluster RAC agent provides for high availability of the underlying redundant systems and interconnects used by Oracle 10g RAC, for example, providing redundancy of key interconnects via Sun Cluster private interfaces as well as application striping.
- Sun StorageTek™ QFS software: Sun StorageTek QFS shared file system software improves the quality of service of SAN infrastructure, delivering maximum scalability, data management, and throughput for the most data-intensive applications. With near raw access to information and data consolidation for read/write file sharing, Sun StorageTek QFS helps to boost productivity and eliminate administrator overhead.
- Solaris Volume Manager software: Solaris Volume Manager lets organizations manage large numbers of disks and their data. The

software helps increase storage capacity and data availability while easing administration of large storage devices.

The Solaris 10 Operating System

The Solaris 10 OS, the latest version of Sun's industry-leading operating system, is the most advanced operating system on the planet. Interoperable with Linux and Windows, the Solaris OS also offers binary compatibility within each server line, and provides features such as Solaris Containers and Predictive Self Healing that help to reduce costs, complexity, and risks.

Architecture design

Throughout the process of designing the Sun Oracle 10g Grid Reference Architecture, certain requirements such as reliability, high-availability, resiliency, cost-effectiveness, a high degree of scalability, and use of best-of-breed products have been viewed as critical to success. 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.

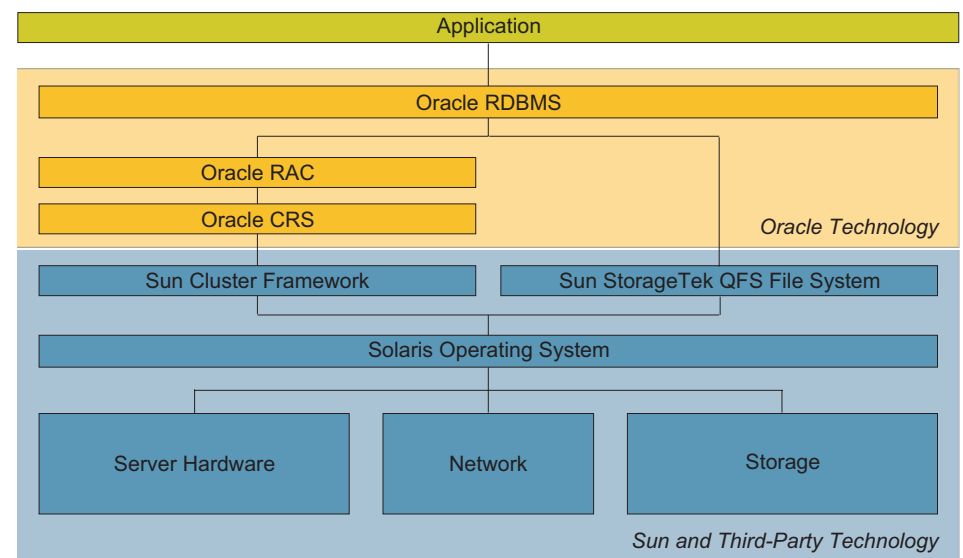


Figure 1: Sun technology provides an integrated foundation for the Sun Oracle 10g Grid Reference Architecture

Figure 2 depicts the reference architecture implementation. Multiple redundant application servers and database servers are provided for scalability and high availability. Database servers can be either x64 or UltraSPARC processor-based servers, depending on application requirements. Redundant internal and management networks help ensure continuity in the event of network failure. High-throughput, low-latency InfiniBand technology connects database nodes, providing essential scalability for enterprise grid deployments. SAN switches and powerful storage arrays provide redundant access to storage.

Sun x64 systems for horizontal, vertical, and diagonal scalability

As equipment ages and loads increase, many organizations are considering ways to replace their aging enterprise grid hardware. Most are also keenly interested in consolidation options that yield greater scalability and flexibility along with savings on real estate and power.

A choice of powerful rackmount X64 servers

With Sun, organizations have a broad choice in the types of Sun x64 systems that they can deploy for Oracle databases and enterprise grids. Sun offers a full line of rackmount servers based on next-generation AMD Opteron processors. Some organizations prefer to deploy larger numbers or horizontally scalable systems while others choose vertically scalable servers that support larger numbers of processors. Some choose a middle road (diagonal scalability) that combines numbers of larger multiprocessor servers. Sun Fire X4100, X4200, and X4600 M2 servers provide enterprise class features offering from two-way to 16-way operation and large memory support, yielding a wide range of choices to power unique database challenges.

Consistent with Sun's strong track record for providing single-instance database scalability with the Solaris OS, the Sun Fire X4600 M2

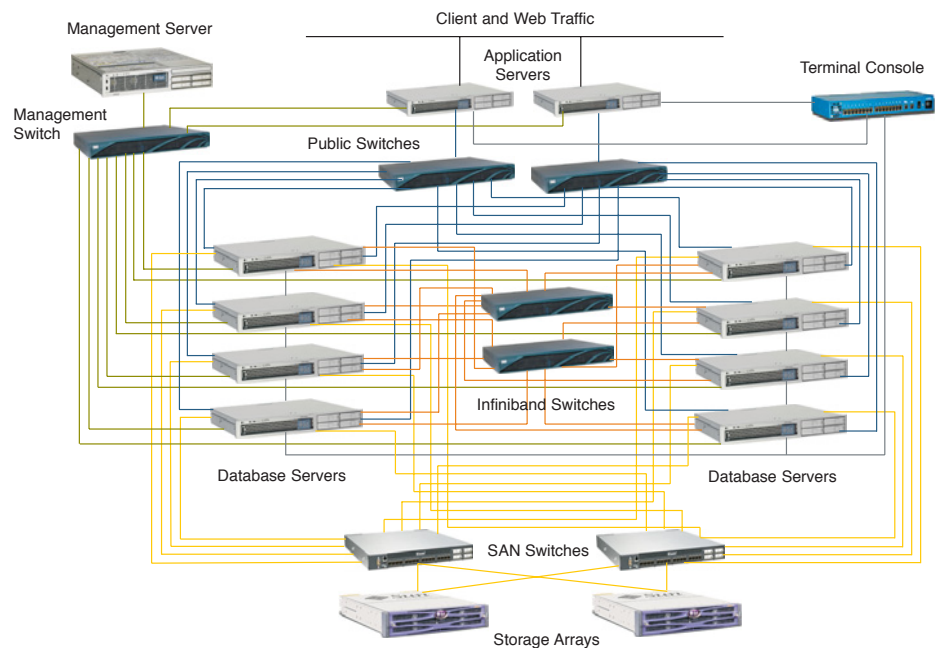


Figure 2: The Sun Oracle 10g Grid Reference Architecture provides redundant high speed connectivity, a choice of x64 or UltraSPARC processor-based systems, and a choice of storage.

server supports up to 16-way operation and up to 128 GB of memory. It can provide both single-system vertical scalability, and can also be combined in enterprise grid configurations for even greater levels of diagonal scalability.

x64-based modular computing for scalability and consolidation of enterprise grids

Given the need for consolidation and simplification, many have long viewed blade servers as potentially attractive for hosting enterprise grids. Unfortunately, most traditional blade platforms have proven disappointing for deploying larger databases, lacking both high-performance processors and memory capacity. In addition, most traditional blade platforms lack the truly extensive and expandable I/O functionality needed to successfully implement the redundant and highly available network infrastructure that enterprise grids require.

With its industry-leading I/O and powerful four-socket Sun Blade X8400 server modules, each supporting up to 64 GB of memory, the Sun Blade 8000 modular system finally provides a

compelling and scalable high-performance database solution with extremely high density. With up to 10 Sun Blade X8400 server modules in a single 19U chassis, and up to two chassis per rack, the Sun Blade 8000 modular system offers up to 160 AMD Opteron cores and up to 1280 GB of memory in single rack. Not only can these systems consolidate growing enterprise grids based on Oracle 10g, but they can consolidate entire multitiered applications into a single chassis, offering massive simplification of infrastructure and administration. For example, most of the Sun Oracle 10g Grid Reference Architecture can be deployed in a single Sun Blade 8000 chassis.

In fact, recent internal Sun testing has demonstrated that four Sun Blade x8400 server modules easily matched the performance of an eight-node cluster of Sun Fire V40z servers. This consolidation provided slightly greater transactions per second and similar response times in fewer rack units (19U vs. 24U) while lowering CPU utilization. Most importantly, not only was there additional headroom on the four Sun

“The Sun Fire T2000 has certainly proven its worth [for Oracle]. With increased performance and higher efficiency, the T2000 will definitely be in our future. The lower energy consumption and lower heat output would be a welcome change in our already taxed datacenter and the compact 2U size should be an easy fit in any rack.”¹

Jon Emmons

Database Administrator, Plymouth State University

Blade X8400 server modules, there was also room for an additional six server modules in the chassis that could be used for other applications.

In addition to system consolidation, middleware consolidation too can greatly simplify multitiered applications and administrative operations. For example, the wide range of available Sun Cluster agents can provide high availability for applications as diverse as Oracle, Apache, and IBM WebSphere, saving organizations from having to engineer (and operate) diverse availability solutions for all of their application components. The result is simpler and cleaner software infrastructure that is easier to administer and less prone to complexity-induced failure.

CoolThreads™ technology for maximum throughput and efficiency

While horizontal scaling offers huge benefits, it often comes at a cost due to increases in power consumption, server sprawl, and complexity.



Figure 3: A wealth of Sun servers are available based on both AMD Opteron and UltraSPARC processors, ideal for Oracle 10g RAC deployments and enterprise grids.

Sun Fire servers based on CoolThreads technology offer a compelling solution by combining the low price point, small footprint, and capabilities of a horizontally scalable system with low power and space consumption. Based on the innovative UltraSPARC T1 processor with support for up to 32 threads and up to 32 or 64 GB² of memory, these servers are ideal for OLTP database workloads. Fully binary compatible with other UltraSPARC platforms, Sun Fire T1000 and T2000 servers offer 5-10 times the server efficiency of a Xeon-based server on a range of industry-standard benchmarks (sun.com/servers/coolthreads/t1000/benchmarks.jsp).

With support for large numbers of threads and technology such as Solaris Containers, Sun Fire CoolThreads servers let organizations run multiple database instances on a single system, removing some of the drawbacks of horizontal scalability. In this respect, servers based on CoolThreads technology essentially combine the best aspects of horizontal and vertical scalability in a single platform. To aid with characterizing performance in concert with power and density, Sun has developed the

SWaP metric, which stands for Space, Watts, and Performance. See sun.com/swap for the latest benchmarks and SWaP analysis.

Putting enterprise grids to work on Sun

With the Sun Oracle 10g Grid Reference Architecture, Sun can help design and deploy effective enterprise grid infrastructure tailored to a wide range of diverse needs. This carefully designed approach lets customers get enterprise grid architecture right the first time to support their most critical business needs with lower levels of risk.

Learn More

To learn more about the Sun Oracle 10g Grid Reference Architecture, visit sun.com/service/refarch or contact your Sun representative. You can also take advantage of Sun's latest server technology — at no charge for 60 days — through the Sun Try-and-Buy Program. Visit sun.com/tryandbuy for more information.



1. lifeaftercoffee.com/2006/05/08/sun-fire-t2000-vs-sun-fire-v440
2. 64 GB capacity target date November 2006

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA Phone 1-650-960-1300 or 1-800-555-9SUN Web sun.com

© 2006 Sun Microsystems, Inc. All rights reserved. Sun, Sun Microsystems, the Sun logo, CoolThreads, Solaris, StorageTek, Sun Blade, and Sun Fire are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States 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. AMD, Opteron, and the Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. Information subject to change without notice. Printed in USA

