Oracle Grid Computing

An Oracle Business White Paper
October 2008

ORACLE

Oracle Grid Computing

INITRODUCTION

Enterprise grid computing is an emerging information technology (IT) architecture that delivers more flexible, resilient and lower cost enterprise information systems. With grid computing, groups of independent, modular hardware and software components can be pooled and provisioned on demand to meet the changing needs of businesses.

The accelerating adoption of grid technology is in direct response to the challenges that IT organizations face with today's rapidly changing and unpredictable business climate. IT departments are under increasing pressure to manage costs, increase operational agility and meet IT service level agreements.

Using enterprise grid computing technology, IT departments can adapt to rapid changes in the business environment while delivering high service levels. Enterprise grid computing has also revolutionized IT economics by both extending the life of existing systems and exploiting rapid advances in processing power, storage capacity, network bandwidth, as well as energy and space efficiency.

Oracle first introduced grid computing capabilities in the Oracle Database and Oracle Application Server in 2003 and continues to lead the software industry in its commitment to grid computing products and practices. Today over 10,000 customers have deployed Oracle grid computing solutions. With the current releases of the Oracle Database and Oracle Fusion Middleware products, including the Oracle WebLogic and Tuxedo product lines, Oracle has introduced a second generation of grid computing technology. This technology builds on Oracle's strong foundation of scalable, fault-tolerant database and middleware clusters, virtualized computing and storage resources, and highly automated end-to-end systems management.

This white paper provides an overview of grid computing, its benefits and key grid techniques for:

- IT resource consolidation
- Agile IT operations
- Predictable high performance and scalability
- Continuous availability

WHAT IS GRID COMPUTING?

The simplest way to think of grid computing is as the virtualization and pooling of IT resources, such as compute power, storage and network capacity, into a single set of shared services that can be provisioned or distributed *and re-distributed* as needed. Just as an electric utility deals with wide variations in power demands without affecting customer service levels, grid computing provides a level of control and adaptability to IT resources that can respond to changing computing workloads while being transparent to end users.

In fact, the term *utility computing* is often used to describe the metered IT services enabled by grid computing. *Cloud computing* is another related term often used to describe how enterprises are using computing resources over the Internet, both on private networks as well as public, multi-tenant clouds. Because it provides superior flexibility, grid computing is the natural architectural foundation for both utility computing and cloud computing.

As workloads fluctuate during the course of a month, week or even through a single day, grid computing infrastructure analyzes demand for resources in real-time and adjusts supply accordingly.

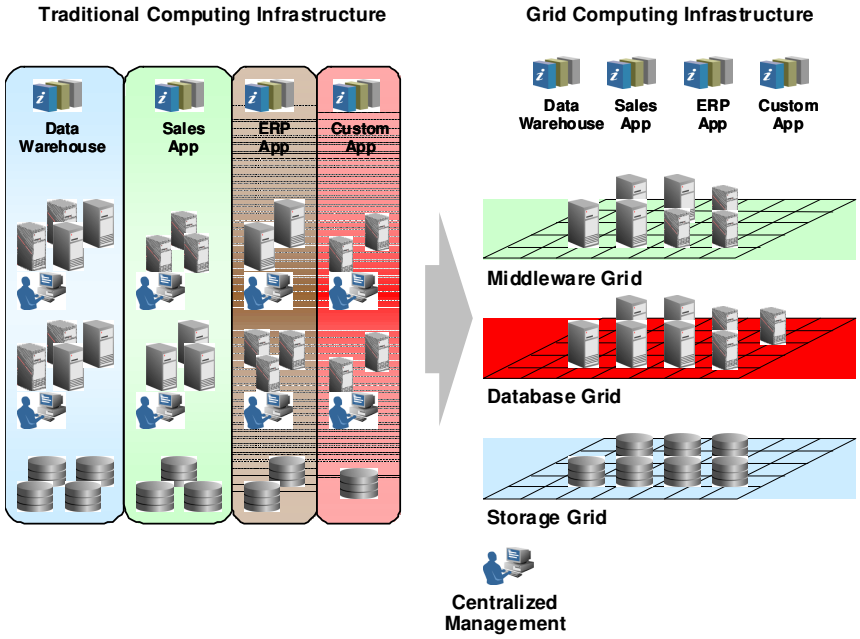

Grid computing operates on these basic technology principles:

- **Standardization –** IT departments have enjoyed much greater interoperability and reduced systems management overhead by standardizing operating systems, server and storage hardware, middleware components and network components in their procurement activities.

This helps to reduce operational complexity in the data center by simplifying application deployment, configuration and integration.

Standardize **hardware and software components to reduce incompatibility and simplify configuration and deployment.**

- **Virtualization** – virtualization abstracts underlying IT resources, enabling much greater flexibility in how they are used. Virtualized IT resources means that applications are not tied to specific server, storage and network components. Applications are able to use any virtualized IT resource. Virtualization is accomplished through a sophisticated software layer that hides the underlying complexity of IT resources and presents a simplified, coherent interface to be used by applications or other IT resources.

Virtualize **IT resources by pooling hardware and software into shared virtual resources.**

Automate **systems management, including resource provisioning and monitoring.**

- **Automation –** grid computing demands large scale automation of IT operations due to the potentially large number of components – virtual and physical – that make up the grid computing environment. Each of these components requires configuration management, on-demand provisioning, top-down monitoring and other management tasks. This scale and complexity means the grid management solution must provide a high degree of automation out-of-the-box to ensure that infrastructure cost savings do not evaporate as a result of additional staffing for managing the grid. Because grid computing and virtualized environments take advantage of highly sophisticated workload distribution to maximize IT efficiency and utilization, a top-down view from the end-user or application level is needed to allow IT to effectively measure service levels and proactively resolve problems. When these capabilities are combined in a single, automated, integrated solution for managing grids, customers will be assured of maximum ROI for their grid investment.

Benefits of Grid Computing

Grid computing provides the following benefits:

With a grid computing architecture you can quickly and easily create a large-scale computing infrastructure from inexpensive, off-the-shelf components like server blades and commodity storage.

- **Ability to respond to volatile business needs at high speed -** businesses today operate in an unpredictable, global environment. Predicting business demands, competitive threats, supply chain risks and regulatory requirements are increasingly challenging. Businesses want IT to be able to provide them the ability to 'turn on a dime' – for example, the ability to change a pricing model to beat a competitor, adjust the order management process for new regulatory requirements, or integrate acquired companies. An underlying grid infrastructure provides IT with the agility to do so.
- **Real-time responsiveness to dynamic workloads** - most applications today are tied to specific software and hardware silos that limit their ability to adapt to changing workloads. This can be a costly and inefficient use of IT resources because IT departments are forced to overprovision their hardware so that each application can handle the peak or worst-case workload scenario. Grid computing enables the allocation and de-allocation of IT resources in a dynamic, on-demand fashion, providing much greater responsiveness to changing workloads on a global scale.

- **Predictability managing IT service levels** - grid computing enables an organization to tie its business requirements, through service level agreements, to its IT architecture with demonstrable metrics and proactive monitoring and maintenance. This encourages a "shared service bureau" approach to IT with a focus on measuring and meeting higher service levels and better alignment between IT and business goals. In the end, high systems administration overhead, costly integration projects and runaway budgets can be eliminated. In addition, a grid-based architecture eliminates single sources of failure and provides powerful high-availability capabilities throughout the entire software stack, protecting valuable information assets and ensuring business continuity.
- **Cost savings through greater efficiencies and smarter capacity planning** – grid computing practices focus on operational efficiency and predictability. Easier grid workload management and resource provisioning puts more power in the hands of IT staff, enabling IT departments to maintain current staffing levels even as computing demands continue to skyrocket. A new generation of server virtualization and clustering capabilities from Oracle means that IT departments can avoid costs by eliminating the need to "overprovision" to meet worst-case scenarios during peak periods. Because computing resources can be applied incrementally when needed, customers enjoy much higher computing and storage capacity utilization. They can also use a more cost-effective scale-out or "pay as you grow" procurement strategy. Companies can avoid buying extra hardware or additional software licenses before they are actually needed. They can also take advantage of the price-performance benefits that come with the rapid growth in processing power and greater energy efficiency.

DATA CENTER MODERNIZATION USING GRID TECHNIQUES

Grid computing is an IT architecture and methodology comprised of both technology and best practices. Not every IT department will adopt every grid computing technology or technique. However, many IT departments are successfully using specific Oracle grid technologies and best practices with dramatic benefits.

IT Resource Consolidation

Grid computing enables sharing of IT resources and consolidation of servers, storage and even data centers, resulting in cost savings, as well as lower power, cooling and space requirements.

Consolidation of IT resources such as servers, storage, applications and even data centers can provide dramatic cost and energy savings.

Forrester Research estimates that average server utilization in data centers today is only about 30%. With hundreds or thousands of servers around the enterprise, the inefficiency is staggering. While application usage varies greatly by certain times of the day or year, it is also being seen as an opportunity to apply grid techniques for a combination of better management, utilization and overall efficiency.

Also, as a significant user of electricity, IT departments must consider energy costs in their data center operations. The comparison of grid computing to an electrical

utility has been made many times, and like any utility, metrics like efficiency and operating margins are scrutinized. IDC estimates that power, cooling and other management costs account for 70% of a server's lifetime cost. Many companies and government entities are now putting energy efficiency and "green computing" initiatives into their buying criteria for technology components.

With the power and space optimization that comes from consolidating resources into a grid infrastructure, enterprises can have a greener data center. Generally speaking, by centralizing and consolidating servers and storage, overall server and storage utilization should increase, thereby avoiding overprovisioning hardware and achieving improved energy efficiency.

Two key grid computing technologies, *server virtualization* and *clustering*, enable sharing of IT resources and consolidation of servers, storage and entire data centers.

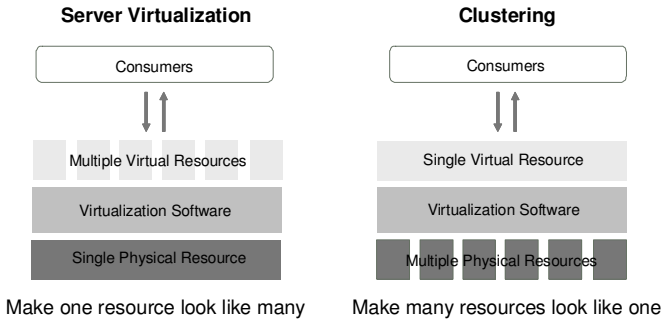

Many customers are now turning to hypervisor-based *virtual machines* (VMs) to consolidate multiple applications on to a smaller number of shared, centrally managed servers. A virtual machine, or virtual server, is software that simulates the operations of computer hardware, enabling an application to run on the virtual machine just as it would on a physical computer.

The advantage of this approach is that many virtual machines can run on a single physical computer, thus enabling the consolidation of many small servers onto one larger server. This approach helps establish a standardized computing environment on which to run applications, middleware servers, database servers and storage servers.

Oracle's server virtualization product, Oracle VM, provides a highly efficient way to run multiple Oracle and non-Oracle databases, middleware and application environments in a single server running Oracle Enterprise Linux. Oracle VM enables IT to quickly add or release more server resources for spikes or lulls in demand, increasing utilization and reducing energy costs. Oracle VM also provides the ability to create pre-configured virtual machine images for quick deployment. It also features live migration to other servers so that high levels of availability can be maintained.

The University of Massachusetts serves 75,000 students, faculty and staff on five campuses. Using Oracle VM, University of Massachusetts consolidated 500

servers running at 20% capacity except during the peak period during Fall registration down to 175 servers, reducing IT overhead by 30-60%.

In addition to server virtualization, Oracle also offers server and storage clustering that enables consolidation of even the largest application environments that may span multiple servers. Clustering makes a group of physical computers operate as one, providing improved performance, availability, scalability and cost. Oracle offers comprehensive clustering capabilities at the middleware, database and storage layers. Clustering is discussed in more detail later in this paper.

Server virtualization and clustering are complementary and even work together on the same physical machine. For example, Oracle Database Real Application Clusters (RAC) can run within an Oracle VM environment. This provides great flexibility and efficiency, since capacity can be changed with finer granularity, development and test environments can share machines with production environments, and smaller applications can be made fault tolerant by running in VM environments on multiple host computers.

Burlington Coat Factory[1] consolidated dozens of separate database servers into two 18-node clusters, each hosting multiple Oracle RAC databases. The result achieves significant cost savings by eliminating unnecessary hardware. The centralization of the solution also improves overall IT manageability and reduces maintenance time.

During their modernization process, Burlington Coat Factory also consolidated and virtualized their storage using Oracle Automatic Storage Management (ASM), a feature of Oracle Database. They consolidated over 1,000 logical storage volumes to under 40, dramatically improving manageability and increasing storage utilization by 50%. This storage optimization achieved 97% CPU utilization compared to 50% previously.

Agile IT Operations

Grid computing relies on agile and efficient systems management: rapid resource provisioning, proactive monitoring and automation of administrative tasks.

Grid computing requires agile and efficient systems management. Managing a grid requires a new breed of systems management, one that provides rapid resource provisioning, real-time visibility into end-user service levels, proactive monitoring, diagnostics and automation of administrative tasks.

With the administrative workload constantly growing and evolving, systems management software must step in to provide proactive monitoring, management automation and resource provisioning to enable IT staff to manage the growing complexity.

Server virtualization and clustering are able to simplify and abstract away the underlying complexity of the software and hardware infrastructure. Grid systems management software must be able to comprehend the underlying complexity and configure and modify the infrastructure to meet dynamic business needs.

Oracle Enterprise Manager provides top-down, end-to-end application management with broad coverage across Oracle databases, middleware and applications. Without the comprehensive functionality and grid automation

provided by Oracle Enterprise Manager, it would not be economically viable to manually manage the grid computing infrastructure.

The latest addition to Oracle Enterprise Manager is WebLogic Operations Control. WebLogic Operations Control automates and optimizes the dynamic allocation of resources at the application grid level, ensuring that individual applications have the compute, memory, and storage resources necessary to meet their respective SLAs and also ensuring that resources are allocated in an optimal way across the entire application grid.

Provisioning database and middleware servers is automated using the server cloning or bare metal provisioning capabilities of Oracle Enterprise Manager. Cloning servers enables an IT department to create reference copies of test, development or production servers, including all patches, installed applications and configuration data and then rapidly deploy them to new server machines.

A good example of server provisioning has been implemented by gasNatural[2], an international oil and gas utility based in Spain. They have deployed a pool of Oracle RAC cluster nodes that power their data warehouses, their business intelligence applications, electricity market applications and other internal applications. They are also able to provision RAC cluster nodes so that they can serve as production, development or test servers. The result is:

- A significant savings in hardware costs from their consolidated data warehouse;
- Tremendous performance increase in their BI applications;
- Low cost high-availability in remote locations;
- A scalable, stable environment for both transactional and business intelligence applications.

Predictable High Performance and Scalability

Grid computing leverages clustering and virtualization technologies at all layers of the technology stack to deliver predictable high performance and scalability.

Enterprise grid computing delivers consistent high performance as workload scales from initial rollout to full-scale deployments. A grid architecture provides the ability to "scale out" by adding computing, storage and network capacity on demand by using clustering and virtualization technologies at all layers of the technology stack. Capacity can be added in smaller, less expensive increments, and those purchases can be made at a later point in time, taking advantage of advancements in price-performance.

Recent trends toward more customer-facing Web applications, cloud computing and Service Oriented Architecture (SOA) are driving the need for predictable high performance and scalability.

- Customer-facing Web applications have orders of magnitude more users than internal applications used by employees. These often have challenging response time, scalability and availability requirements.

- Cloud computing service providers often have service level commitments for performance and availability, and they must scale out quickly as new clients come on board and workload peaks happen.
- SOA enables disparate applications and data sources to be integrated into loosely-coupled, composite applications. As more and more existing and new applications are exposed as Web services in a SOA, more programs are consuming these services, creating greater, often unpredictable workloads, especially if the Web services are exposed outside the firewall. This may strain these Web services, creating spikes in response time and bottlenecks in throughput.

Grid computing leverages clustering and virtualization technologies at all layers—middleware, database and storage—to deliver predictable high performance and scalability for applications, especially customer-facing, cloud computing and SOA applications discussed above

Clustering

Oracle has long been a leader and industry visionary in the area of clustering, making many computers work together as one and enabling individual computers to be added or removed from the cluster on-demand. Oracle has proven clustering capabilities at all layers of the technology stack.

For Oracle Database, one of Oracle's most significant innovations is Oracle Real Application Clusters (RAC). RAC is a key piece of Oracle's grid computing products and offers unmatched database scalability and availability using clusters of database servers. In addition, Oracle Database features Automatic Storage Management (ASM), which enables a shared pool of storage for both clustered and non-clustered database environments. ASM provides dynamic provisioning of storage, as well as simplified and automated storage administration. The Oracle Exadata Storage Server and HP Oracle Database Machine are revolutionary new products providing a smart storage grid connected to a database grid via fast InfiniBand connections. These products are a combination of hardware and optimized software, resulting in query performance gains of 10X or more and unlimited scalability by simply adding more physical units.

For middle-tier servers, Oracle's *application grid* provides clustering capabilities, including significant features in Oracle Web Cache, Oracle Coherence, Oracle WebLogic Server, Oracle Tuxedo and more. In addition, Oracle TimesTen In-Memory Database is a relational database running in the middle-tier, delivering extremely fast and predictable response times and throughput, and often serving as a transparent read/write cache for Oracle Database with RAC. Oracle JRockit Real Time is the industry's fastest Java Virtual Machine (JVM) and delivers predictable, low latency response times with deterministic Java garbage collection pauses.

Oracle Enterprise Manager provides IT managers with the ability to provision, monitor and manage all of Oracle's clustered databases, storage and middleware.

Pacific Gas and Electric Company (PG&E) provides power to 15 million people across 70,000 square miles of Northern and Central California. In 2005, as PG&E rolled out its SmartMeter program to provide customers with detailed rate and usage information to help them understand, manage and reduce their power consumption, PG&E's meter reading capabilities had to scale from once-a-month to once-an-hour, a factor of 720. They replaced an old IBM DB2 database running on a mainframe with a highly-scalable cluster of eight Oracle Database RAC nodes running on smaller, more energy-efficient systems. The resulting savings amounted to $5 million annually.

Mercado Libre[3] is the largest online auction house in Latin American. Their rapid rate of growth quickly outpaced their mid-sized SMP server. In 2004, they replaced the server with a 4-node RAC Cluster, which cost $500,000 less than the big SMP box alternative. Today, their grid footprint has grown to a 16-node cluster as their business has grown by a factor of four.

According to IDC, which documented Mercado Libre's grid transformation, their savings, totaling approximately US$5.1 million over five years, came from a combination of avoided hardware and software costs (US$1.18 million), increased uptime (US$1.97 million), increased search speed (US$1.40 million) and improved fraud prevention (US$0.5 million). Minus capital and operating expenses, IDC is projecting net-present benefits of US$2.4 million over five years, equating to a return on investment of 452%.

In-Memory Data Grids

With the increasing importance of middleware in modern architectures, virtualization and clustering in the middle tier are also critical for the continuous operation, predictable high performance and agile scale-out of enterprise applications.

One of Oracle's application grid products is Oracle Coherence in-memory data grid, which allows Java and .NET applications to very quickly access objects in memory that span multiple physical machines in the middle tier. This provides horizontal scalability, high availability and predictable, high performance. The Coherence architecture scales linearly as additional nodes are added, and nodes can be added or removed on the fly. High-availability is achieved by storing copies of the data on different servers in the data grid to avoid a single point of failure in case an individual middle-tier node crashes or is taken offline for maintenance. High performance is achieved by storing data in memory on middle-tier computers, which avoids disk access and network delays.

For Allegient Systems[4], an Internet-based litigation management solution, application performance and scalability are directly linked to revenue and operating margins. Oracle Coherence enabled them to improve service levels to customers by 53%, while transaction volume grew by 27%. Efficiency gains with Coherence enabled them to process 30% more invoices per day while overall database load by 58 percent. This enabled Allegient to save $300,000 in hardware, software and maintenance fees in 12 months and potentially millions more in years to come.

Service Level Management

IT departments are under great pressure to deliver on their service level commitments and be accountable to the lines of businesses they serve within an organization. As information technology is more tightly linked with business operations than ever before, it is critical that IT departments can quantify and validate the services they provide. Grid computing, as part of a utility computing strategy, depends on accurate service level management to ensure that the physical capacity of the data center is providing adequate service levels to end users.

For Petro Canada[5], one of Canada's largest oil and gas companies, management of service levels is critical as they rapidly grow their work force and adopt new applications and systems. At the same time, they need to maintain current IT staffing levels, so the ability to automate the monitoring and management of IT service levels is also critical.

By adopting Oracle Service Level Management and other Oracle Enterprise Manager capabilities, Petro Canada has been able to manage their growth, improve security and gain a better understanding of their application availability and end-user service levels.

Continuous Availability

Grid computing enables continuous availability with replication, automatic failover and disaster protection.

Oracle has been at the forefront of developing high availability products and practices for its entire history. From server failover with Oracle RAC and Oracle Application Server, to data replication and Oracle Active Data Guard, Oracle provides IT organizations with a comprehensive portfolio of solutions to keep the data center, and thus the business, running smoothly.

Server Failover

Server failover has been available for many years from both hardware and software manufacturers. The protections that are afforded from a successful failover of a server are often critically important. Yet from a business standpoint, the drawback of setting up a standby server is the cost of hardware and software that are only used when disaster strikes. Grid computing enables standby resources to be used as active resources, resulting in higher utilization.

Another consideration of server failover is that of failover time. Some applications are so critical to the business that they can not afford to be down even for a few minutes, while others can tolerate some interruption of operation.

Oracle has pioneered many techniques of server failover that provide IT departments with automatic failover capabilities for several server types. For example, Oracle Database RAC, Oracle WebLogic Server, Oracle Tuxedo and Oracle Coherence clusters can withstand failures of several servers within a cluster and still remain in operation. IT departments can simply remove failed servers from service and repair or replace them and add them back to the server grid. Load balancing, work load management and overload protection, and automatic migration and failover of services or whole servers ensure applications stay up and running.

Disaster Protection

Even clusters cannot survive a complete data center failure from natural disasters, fires, floods and so on. In these cases, failover to a remote location is required. An enterprise grid can be designed to encompass multiple locations, dynamically shifting workloads across those locations for the highest reliability.

Oracle Active Data Guard provides the ability to create up-to-date replicas of the production Oracle Database for standby and disaster recovery. In addition, these replicas can also be used for resource-intensive read-only operations such as queries, reporting or backup.

Fidelity National Financial[6] (FNF) is a leading provider of outsourced products and services that includes technology solutions, financial and insurance services, claims management and more. FNF published a detailed case study that followed the progress of migrating a critical document management application (EDoc) to an Oracle grid implementation. EDoc manages over 50 million documents and impacts numerous systems that are critical to FNF's service offerings. FNF has standardized on Oracle Database with RAC, ASM, Oracle Active Data Guard and other high-availability components.

FNF faced heightened uptime requirements for EDoc because their customers were increasingly using on-line documents for faster real-estate transactions. FNF also wanted to lower operating costs by reducing manual operations and promoting a more efficient e-business model.

The previous implementation of EDoc was continually running at full server capacity. So FNF migrated to Oracle RAC so that they could scale the application to meet their evolving growth and quality-of-service requirements. FNF determined that an Oracle RAC implementation of EDoc could provide increased CPU horsepower, higher resource utilization levels and load balancing capabilities.

According to the FNF case study "collectively, the utilization of Oracle High-Availability Features and Oracle Maximum Availability Architecture best practices has enabled FNF to meet service level agreements at the lowest cost." Furthermore, the new implementation has "proven to be more reliable by eliminating single points of failure and more scalable by providing the ability to add capacity on demand."

CONCLUSION

Oracle's introduction of enterprise grid computing in 2003 introduced a state-of-the-art methodology and a set of new database and middleware capabilities that have helped evolve the way IT departments operate.

At that time, data center projects such as server consolidation, SOA development, space and power optimization and large-scale implementations of rack-mounted Linux servers seemed unrelated. We can now see that these techniques taken together are in fact interrelated and can be described as a grid computing approach to data center modernization.

We've also seen, through several customer examples, how the benefits derived from these techniques are compounded. For instance, server and storage consolidation increases utilization levels. This in turn enables IT departments to save energy, reduce systems management costs, and get a better return on their hardware investments.

The widespread adoption of open standards, IT resource virtualization, on-demand provisioning, highly automated systems management and real-time monitoring has created a new generation of data center best practices. Oracle Database, Oracle Fusion Middleware and Oracle Enterprise Manager were designed with this next generation data center in mind.

By using these grid computing techniques and these Oracle products, IT professionals will find that the data centers they are building for the next decade of business challenges can also provide immediate benefits in terms of cost savings, sustainability and operational agility.

APPENDIX A: ORACLE GRID COMPUTING PRODUCTS

Oracle Product	Grid Computing Features
Oracle VM	• Server virtualization enables multiple software instances to run on a single computer
Oracle Database	• Real Application Clusters (RAC) enables database clustering, so a single database can be deployed across a cluster of servers • Automatic Storage Management (ASM) enables a shared pool of storage for the database • Active Data Guard creates read-only replicas for standby databases and disaster protection • In-Memory Database Cache delivers predictable low latency, high throughput and high availability • Advanced Compression enables disk compression for cost and power savings and better performance • Real Application Testing supports database replay and SQL performance analysis to test all changes before putting them into production
Oracle Application Grid	• Oracle WebLogic Server provides a clustered Java EE application server • Oracle Coherence: in-memory clustered data grid for Java and .NET objects • Oracle Tuxedo: clustered transaction processing monitor for C, C++ and Cobol applications • Oracle JRockit Real Time: Java Virtual Machine (JVM) with deterministic response times
Oracle Exadata Storage Server and HP Oracle Database Machine	• Oracle Database, RAC, ASM, Partitioning: database cluster connected to storage cluster with high speed InfiniBand
Oracle Enterprise Manager	• Top-down management • Server cloning and bare metal provisioning • Automated administration • Real-time and predictive monitoring

REFERENCES

[1] IDC, Grid Computing with Oracle Database 11g, March 2008

[2] Mainstay Partners, GasNatural Strengthens Global Business with Enterprise Grid Computing, August 2007

[3] IDC/Oracle, Business Benefits Series – Mercado Libre, 2005

[4] Oracle White Paper, Calculating the Return on Investment for Clustered Caching and Data Grid Solutions, May 2007

[5] Oracle Customer Snapshot, Petro-Canada, April 2006

[6] Fidelity National Information Centers, Oracle Maximum Availability Architecture – Architecture Case Study, 2006

Oracle Grid Computing
October 2008
Author: George F. Demarest

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
www.oracle.com