

# Oracle Grid Computing

*An Oracle Business White Paper*  
*May 2008*

## **INTRODUCTION**

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 cycles. IT departments are under pressure to increase operational agility, to establish and meet IT service levels and to control costs.

Using enterprise grid computing technology IT departments can adapt to rapid changes in the business environment while meeting higher 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 Application Server in the year 2003 and continues to lead the software industry in its commitment to grid computing products and practices. With the current releases of the Oracle Database and Oracle Fusion Middleware products, Oracle has introduced a second generation of grid computing capabilities that build on Oracle's strong foundation of scalable, fault-tolerant database and middleware clusters, virtualized computing and storage resources, and highly automated end-to-end application and business process management.

## 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 type of IT operations Enterprise grid computing enables.

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.
- **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.
- **On-demand Provisioning** - IT resources must be easily *provisioned*, meaning allocated, configured and maintained by grid management tools. As different parts of the system require additional computing power, such as when many new users are added, IT professionals need the ability to quickly and accurately establish user accounts and security privileges and allocate storage and computing capacity. In grid computing, powerful provisioning and resource management software determines how to meet the specific needs of users, while optimizing performance of the system as a whole.
- **Automation** – virtualization and provisioning can only be accomplished with large scale automation of IT operations such as system installation, patching, server cloning, workload management, user account creation and so on. In years past, IT staff created some of this automation through custom programs and scripts, but many have discovered that this does not scale effectively. Out-of-the-box management automation from system

**Standardize hardware and software components to reduce the risk of costly incompatibility and integration issues**

**Virtualize Infrastructure Resources: Pool hardware and systems software into a single virtual resource.**

**Provision Infrastructure Resources: Allocate capacity on demand based on policies to meet individual needs and optimize the system as a whole.**

providers such as Oracle can significantly boost productivity of system administrators.

- **Real-time and Predictive Monitoring** – with the growing scale and complexity of data center implementations, IT departments can no longer afford to work reactively to potential problems as they arise. IT professionals need increasingly sophisticated tools to monitor a vast number of systems in real time and predict problems before they occur. Grid computing relies on policy-based monitoring and management of quality-of-service thresholds and top down applications management. This enables IT staff to quickly identify the root cause of a problem or potential problem from the lowest level hardware issues through the database, middleware and user interface tiers.

## BENEFITS OF GRID COMPUTING

Grid computing provides the following benefits:

- **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 in 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 architecture eliminates single sources of failure and provides powerful high-availability capabilities through 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-

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.

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 can be described as 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**

Consolidation of IT resources such as servers, storage, applications and even data centers can provide dramatic cost savings. It is also the most common first step in modernizing IT operations.

Forrester Research estimates that average server utilization in data centers today is 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.

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, web servers, middleware servers and databases.

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. Oracle VM's ability to quickly add or release more server resources for spikes or lulls in demand provides the same type of opportunity for management and energy efficiency as storage virtualization. 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.

Oracle also offers server and storage clustering that provides additional scalability to consolidate even the largest application environments that may need to span multiple servers. The combination of server virtualization and clustering provides the ultimate consolidation environment for any IT requirement.

Burlington Coat Factory<sup>1</sup> consolidated dozens of separate databases 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 the Oracle Automatic Storage Management (ASM), a feature of Oracle Database 11g. 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**

Providing agility and predictable IT service levels in the data center requires real-time visibility into IT operations, proactive monitoring and diagnostics and large scale 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.

Grid techniques such as server and storage virtualization and clustering are able to simplify and mask the underlying complexity of the software infrastructure. Grid systems management software must be able to comprehend the underlying complexity and configure and modify the infrastructure to meet dynamic business needs.

Provisioning at the server level is often done using the server cloning or bare metal provisioning capabilities of Oracle Enterprise Manager for both database and middleware servers. 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<sup>2</sup>, an international oil and gas utility based in Spain. They have developed 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**

Enterprise grid computing delivers maximum scalability through the ability to add computing, storage and network capacity on demand. The ability to “scale out” comes from clustering standard hardware and software components and virtualizing them to effectively create one large, virtual computer.

Like grid computing, Service Oriented Architecture (SOA) applications often consume services from widely varied sources, greatly reducing silos of disconnected information and application logic. However, SOA applications can also introduce more unpredictability in the computing workload as newer, more powerful Web Services are introduced. As these Web services become increasingly popular, more and more programmers (and thus programs) will consume them. This may strain these Web services beyond the initial intention of the developer.

This is similar to the “success crisis” that many Web sites faced early on during the Internet revolution. As these sites became more popular and traffic increased, the underlying infrastructure could not handle it. Response times spiked, and some sites stopped working altogether. SOA Web services can also become the victim of their own success, where increased adoption by more and more applications can rapidly outpace the original scope of the Web services.

This is why grid computing is ideally suited as the underlying software infrastructure for SOA applications. New services introduced in a SOA environment require dynamic allocation of computing power in order to perform and scale predictably. With a grid computing infrastructure, these services can get access to a virtualized pool of compute power and storage on an as needed basis. This can provide significant cost savings due to reduced server hardware and software licenses and improved application uptime.

### **Server Clustering**

Oracle has long been a leader and industry visionary in the area of database and middleware clustering.

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 physical servers.

For middle-tier servers, Oracle provides load-balancing and failover capabilities for clustered Oracle Web Cache servers, Oracle Application Servers including clustered OC4J, Java Messaging Service (JMS), Oracle Internet Directory (LDAP) servers and more. Having a broad selection of clustering capabilities in the middle tier provides a higher level of agility for IT managers.

Mercado Libre<sup>3</sup> is the largest online auction house in Latin America. 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 is also critical for the continuous operation, predictable high performance and scale-out of enterprise applications.

Oracle Coherence establishes in-memory “data grids” for Java and .Net applications to access objects in memory that are distributed across multiple physical machines in the middle tier. This enhances the processing capability of middle-tier application servers and provides horizontal scalability, high-availability and predictable, high performance. Coherence provides this high performance because in-memory processing in the middle tier reduces network overhead and minimizes reading and writing of data to disk. The Coherence architecture has been shown to scale linearly as additional nodes are added. 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 system crashes or is taken offline for maintenance.

Oracle Coherence's grid architecture enables the addition of application instances to be started on the fly. Oracle Coherence is designed for lights-out management, which provides the ability to expand and contract Oracle Coherence almost instantaneously in response to changing demand.

For Allegient Systems<sup>4</sup>, 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 greater scrutiny in terms of budgeting and accountability 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 infrastructure is providing real services to end users.

For Petro Canada<sup>5</sup>, 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**

Oracle has been at the forefront of developing high-availability products and practices for its entire history. From server failover with RAC and Oracle Application Server, to data replication and Oracle Data Guard, Oracle provides IT organizations with a rich 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, a common drawback of setting up a standby server is the expense incurred paying for hardware and software that are only utilized when disaster strikes. Grid 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 several sever failover options. Some, such as Oracle Real Application Clusters 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.

## **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, the premiere disaster protection product in the industry, provides this capability. Not only does it help IT departments survive data center disasters, but it does so without have expensive back-up servers sitting idle or unused during normal operations. Oracle Data Guard uses Oracle Streams technology to provision data to a failover site.

Fidelity National Financial<sup>6</sup> (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, Automatic Storage Management, Oracle 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 at the time. 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 11g, 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.

---

## 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  
May 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](http://www.oracle.com)

Copyright © 2008, Oracle. All rights reserved.

This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission. Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.