An Oracle White Paper
June 2009

Oracle Grid Computing

Introduction

Enterprise grid computing is an emerging IT architecture that delivers flexible enterprise information systems that are more resilient and less expensive than traditional legacy systems. In grid computing, groups of independent hardware and software components are pooled and provisioned—on demand—to meet changing business needs.

The accelerating adoption of grid technology is in direct response to the challenges facing information technology (IT) organizations. 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 (SLAs).

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 revolutionized IT economics by extending the life of existing systems and exploiting rapid advances in processing power, storage capacity, energy and space efficiency, and network bandwidth.

This white paper provides an overview of grid computing, highlights the benefits of using it, and describes key grid computing techniques that enable IT resource consolidation, agile IT operations, predictable high performance and scalability, and continuous availability.

What Is Grid Computing?

One way to think about grid computing is as the virtualization and pooling of IT resources—compute power, storage, network capacity, and so on—into a single set of shared services that can be provisioned or distributed, and then *redistributed* as needed. Just as an electric utility uses a grid to deal with wide variations in power demands without affecting customer service levels, grid computing provides IT resources with levels of control and adaptability that are transparent to end users, but that let IT professionals respond quickly to changing computing workloads.

The term *utility computing* is often used to describe the metered (or *pay-per-use*) IT services enabled by grid computing. *Cloud computing* (where dynamically scalable and often virtualized resources are provided as a service over the internet) is another term that describes how enterprises are using computing resources—on both private and public networks—over the internet. Because grid computing provides superior flexibility, it is the natural architectural foundation for both utility and cloud computing.

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

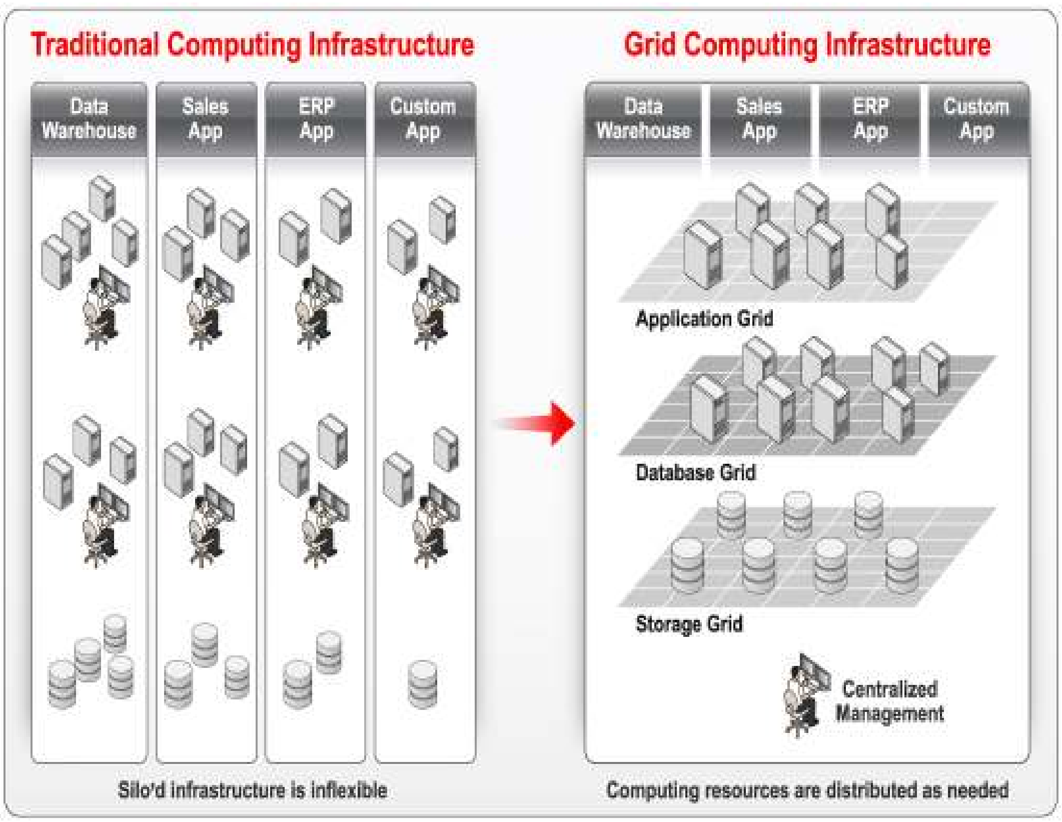

Compare the difference in infrastructures between traditional computing and grid computing.

Grid computing operates on three basic technology principles: Standardize hardware and software components to reduce incompatibility and simplify configuration and deployment; virtualize IT resources by pooling hardware and software into shared resources; and automate systems management, including resource provisioning and monitoring.

Grid computing operates on these technology principles.

- **Standardization.** IT departments have enjoyed greater interoperability and reduced their systems management overhead by standardizing on operating systems, servers, storage hardware, middleware components, and network components. Standardizing also helps reduce operational complexity in the data center by simplifying application deployment, configuration, and integration.
- **Virtualization.** Virtualizing IT resources means that applications are not tied to specific server, storage, or network components and can use any virtualized IT resource. Virtualization occurs through a sophisticated software layer that hides the underlying complexity of IT resources and presents a simplified, coherent interface used by applications and other IT resources.
- **Automation.** Because of the potentially large number of components—both virtual and physical—grid computing demands large-scale automation of IT operations. Each component requires configuration management, on-demand provisioning, top-down monitoring, and other management tasks. A grid management solution must ensure that infrastructure cost savings do not evaporate as a result of hiring additional staff to manage the grid. IT administrators need a top-down view from the end-user or application level so they can effectively measure service levels and proactively resolve problems. Combining these capabilities into a single, automated, integrated solution for managing grids gives organizations a maximum return on their grid investment.

The Benefits of Grid Computing

Using a grid computing architecture, organizations can quickly and easily create a large-scale computing infrastructure from inexpensive, off-the-shelf components. Other benefits of grid computing include

- Quick response to volatile business needs
- Real-time responsiveness to dynamic workloads
- Predictable IT service levels
- Reduced costs as a result of improved efficiency and smarter capacity planning

Responding Quickly to Volatile Business Needs

Businesses today operate in an unpredictable, global environment. Staying on top of business demands, competitive threats, supply chain risks, and regulatory requirements is increasingly challenging. Businesses expect their IT groups to "turn on a dime" and be able to, for example, change a pricing model to beat a competitor, adjust the order management process to accommodate new regulatory requirements, and integrate acquired companies. With an underlying grid infrastructure, IT has the agility to respond quickly to these types of changing business needs.

Responding to Dynamic Workloads in Real Time

Most of today's applications are tied to specific software and hardware silos, limiting their ability to adapt to changing workloads. This costly and inefficient use of IT resources means that IT departments must overprovision their hardware so that each application can handle peak or worst-case workload scenarios. Grid computing lets IT professionals dynamically allocate and deallocate IT resources as needed, providing much better responsiveness to workloads that change on a global scale.

Providing Predictable Service Levels

Through the use of service-level agreements (SLAs), organizations can tie business requirements to IT architecture to get demonstrable metrics and proactive monitoring and maintenance. This encourages a shared-service-bureau approach to IT, with the focus on measuring and meeting higher service levels and better aligning IT and business goals. A grid-based architecture eliminates single sources of failure and provides powerful, high-availability capabilities throughout the entire software stack, protection for valuable information assets, and business continuity. It lets IT groups eliminate expensive systems administration overhead, costly integration projects, and runaway budgets.

Reducing Costs with Improved Efficiency 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 the IT staff, letting them maintain current staffing levels even as computing demands skyrocket. With a new generation of server virtualization and clustering capabilities from Oracle, IT departments no longer have to overprovision to meet worst-case scenarios during peak periods. And because computing resources can be applied incrementally when needed, organizations enjoy much higher computing and storage capacity utilization at a reduced cost.

With grid computing, organizations can choose a more cost-effective scaling with a "pay-as-you-grow" procurement strategy. Companies can avoid buying extra hardware or additional software licenses before they actually need them, and they can take advantage of the price/performance benefits that come with rapid growth in processing power and greater energy efficiency.

Using Grid Techniques to Modernize Data Centers

Grid computing IT architecture and methodology include these technologies and best practices.

- Consolidated IT resources
- Agile IT operations
- Predictable high performance and scalability
- Continuous availability

Not every IT department will adopt every grid computing technology or technique; however, many groups are already seeing dramatic benefits by using selected Oracle grid technologies and best practices.

Consolidating IT Resources

> Grid computing enables the sharing of IT resources and the consolidation of servers, storage, applications, and even data centers. Results include reduced costs and lower power, cooling, and space requirements.

Consolidating IT resources can provide dramatic cost and energy savings. Forrester Research estimates that in data centers today, the average server utilization is only about 30 percent. Considering how many servers can be in an enterprise, that inefficiency is staggering. Because application usage varies greatly by certain times of the day or year, it provides an opportunity to apply grid techniques for a combination of better management, utilization, and overall efficiency.

IT departments are significant users of electricity, so they must also consider those energy costs in their data center operations. When metrics such as efficiency and operating margins are scrutinized, IDC estimates that power, cooling, and other management costs account for 70 percent of a server's lifetime cost. Many organizations 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, organizations can have a greener data center. As organizations centralize and consolidate servers and storage, their overall server and storage utilization increases. The IT staff no longer needs to overprovision hardware and can improve overall energy efficiency.

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

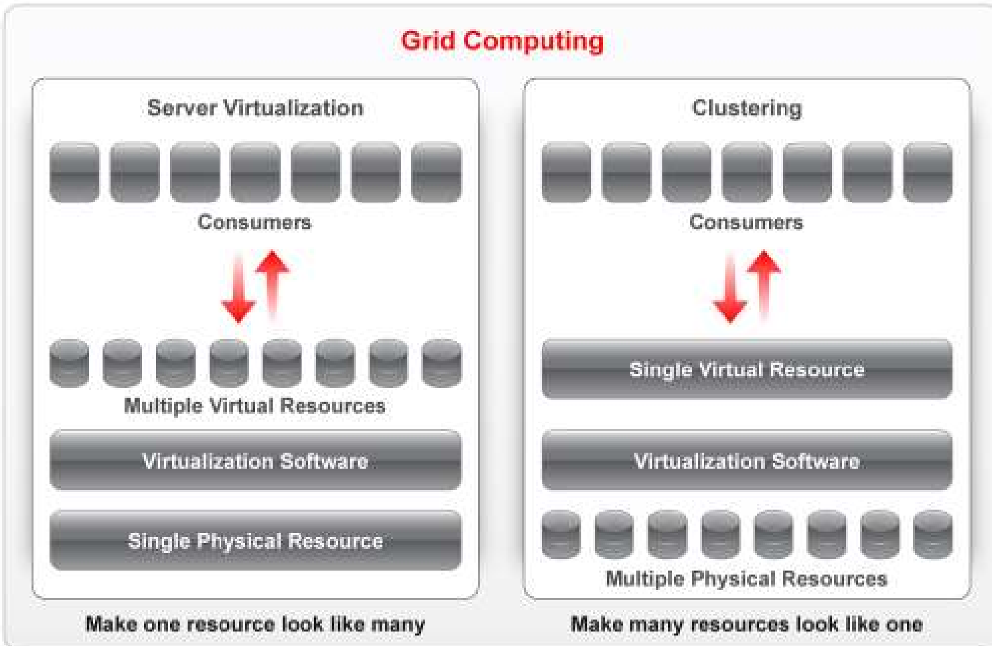

Server virtualization and clustering are two key grid computing techniques that can make one resource look like many and many resources look like one.

Consolidating with Virtual Servers

Many customers are using hypervisor-based virtual machines (VMs) to consolidate multiple applications onto fewer centrally managed shared servers. A virtual machine (or virtual server) is software that simulates the operation of computer hardware and allows an application to run on it as if it were a physical computer. The advantage of this approach is that many virtual machines can run on a single physical computer, allowing the consolidation of multiple small servers onto one larger server. This approach helps establish a standardized computing environment in 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 on a single server running Oracle Enterprise Linux. Oracle VM lets the IT staff quickly add or release more server resources to match spikes or lulls in demand. This increases server utilization and reduces energy costs. Using Oracle VM, IT professionals can create preconfigured virtual machine images for quick deployment, and then quickly perform live migration to other servers to maintain high levels of availability.

For example, the University of Massachusetts serves more than 60,000 students and 16,000 faculty and staff on five campuses.[1] Previously, the University had 500 servers running at 20 percent capacity, except during the peak period of fall registration. Using Oracle VM, the University of Massachusetts consolidated its 500 servers to fewer than 300 servers. The lower investments in hardware made possible by Oracle VM will reduce the University's operational expenses. As an added bonus, virtualization will improve availability and promote the University's ecofriendly initiatives, which will also save money. When compared to a more traditional IT environment, the University's will save close to US$100,000 a year in power and cooling costs.

Consolidating with Server and Storage Clustering

In addition to server virtualization, Oracle also offers server and storage clustering to allow the consolidation of even the largest application environments—including those 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. (For more information, see "Clustering," later in this paper.)

Server virtualization and clustering are complementary and can even work together on the same physical machine. For example, Oracle Database Real Application Clusters (Oracle RAC) can run within an Oracle VM environment to provide more flexibility and greater efficiency. With Oracle VM, capacity can be adjusted with a finer granularity, allowing development and test environments to share machines with production environments. Smaller applications can be made fault tolerant by running in VM environments on multiple host computers.

In another example, Burlington Coat Factory consolidated dozens of separate database servers into two, 18-node clusters, each hosting multiple Oracle RAC databases.[2] The result was significant cost savings due to the elimination of unnecessary hardware. In addition, the server centralization improved overall IT manageability and reduced maintenance time.

During the modernization process, Burlington Coat Factory also consolidated and virtualized its storage using Oracle Automatic Storage Management, a feature of Oracle Database. The company consolidated more than 1,000 logical storage volumes into fewer than 40 volumes,

[1] Alan Joch, "See Savings with Linux," *Oracle Magazine*, July 2009.

[2] IDC, *Grid Computing with Oracle Database 11g*, March 2008.

dramatically improving manageability and increasing storage utilization by 50 percent. This storage optimization also increased CPU (central processing unit) utilization from 50 to 97 percent.

Keeping IT Operations Agile

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

Managing a grid computing system requires a new breed of agile, efficient systems management software that can provide rapid resource provisioning, real-time visibility into end-user service levels, proactive monitoring, and diagnostics. As the administrative workload associated with grid computing grows and evolves, the new systems management software must be able to automate administrative tasks so that IT staff can manage the growing complexity. In addition, the systems management software must be able to work within the underlying complexity of hardware and software infrastructures, and be able to configure and modify those infrastructures to meet the dynamic needs of the business.

Manual management of grid computing infrastructures is not economically feasible. However, the comprehensive functionality and grid automation provided by Oracle Enterprise Manager make managing the grid possible. Oracle Enterprise Manager provides top-down, end-to-end application management with broad coverage across Oracle databases, middleware, and applications.

Oracle WebLogic Operations Control, the latest addition to the Oracle Enterprise Manager product family, automates and optimizes the dynamic allocation of resources across the application grid. It also ensures that individual applications have the memory, storage, and computing resources they need to meet their SLAs.

Automated provisioning of database and middleware servers is due to the server cloning or bare metal provisioning capabilities of Oracle Enterprise Manager. By cloning servers, IT departments can create reference copies of test, development, and production servers—including all patches, installed applications, and configuration data—and then rapidly deploy them to new server machines.

One example of server provisioning was implemented by Gas Natural, an international oil and gas utility based in Spain.[3] Gas Natural deployed a pool of Oracle RAC cluster nodes to power its data warehouses and business intelligence applications, electricity market applications, and other

[3] Mainstay Partners, *Gas Natural Strengthens Global Business with Enterprise Grid Computing*, August 2007.

internal applications. The utility provisioned Oracle RAC cluster nodes to serve production, development, and test servers, and achieved the following results.

- Significant savings in hardware costs due to the consolidated data warehouse
- Tremendous performance increases in business intelligence applications
- Low cost and high availability in remote locations
- A stable, scalable environment for both transactional and business intelligence applications

Delivering 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.

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 significantly more users than internal applications and more-challenging requirements for response time, scalability, and availability.
- Cloud computing service providers often have service-level commitments for performance and availability. They must be able to scale quickly as new clients are added and workload usage peaks.
- SOA enables disparate applications and data sources to be integrated into loosely coupled, composite applications. As more and more applications are exposed as Web services by SOA, other programs are quickly consuming those services and creating greater, often unpredictable workloads. This is especially true if Web services are exposed outside the organization's firewall. Such exposure can put a strain on Web services and create 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. Enterprise grid computing delivers consistent high performance as workloads increase from initial rollout to full-scale deployment. In addition, grid architecture provides the ability to scale all levels of the technology stack by using clustering and virtualization technologies to add storage, network, and computing capacity as they are needed. Capacity can be added in smaller, less-expensive increments at any time, so organizations can take advantage of lower pricing opportunities.

Clustering

Oracle has long been a leader and industry visionary in the area of *clustering* (making many computers work together as one and allowing individual computers to be added or removed from the cluster as demand requires). Oracle has clustering capabilities at all layers of the technology stack. One of the most significant innovations, and a key component of Oracle's grid computing solutions, is Oracle Real Application Clusters. Oracle RAC offers unmatched database scalability and availability by using clusters of database servers.

Oracle Automatic Storage Management, a feature of Oracle Database, lets IT systems use a shared pool of storage for both clustered and nonclustered database environments. Oracle Automatic Storage Management provides dynamic provisioning of storage, along with simplified, automated storage administration.

For middle-tier servers, Oracle application grid technology provides clustering capabilities with significant features in Oracle Web Cache, Oracle Coherence, Oracle WebLogic Server, Oracle Tuxedo, and more.

- Oracle TimesTen In-Memory Database is a relational database running in the middle tier that delivers extremely fast, predictable response times and throughput. Oracle TimesTen In-Memory Database often serves as a transparent read/write cache for Oracle Database customers using Oracle RAC.
- Oracle JRockit Real Time is the industry's fastest Java Virtual Machine (JVM). It 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.

The HP Oracle Exadata Storage Server and HP Oracle Database Machine are revolutionary new products that provide a smart storage grid connected to a database grid with fast InfiniBand connections. These products combine hardware and optimized software, and offer unlimited scalability (by adding more physical units) and faster query performance that is at least 10 times faster than traditional data storage warehouses.

Flexible, Cost-Efficient Applications

Pacific Gas and Electric Company (PG&E) provides power to 15 million people across 70,000 square miles of Northern and Central California.[4] In 2005, PG&E rolled out its SmartMeter program, which gave customers detailed rate and usage information so they could manage and reduce their power consumption. As a result of the program, PG&E's meter-reading capabilities had to scale from once per month to once per hour—a factor of 720.

[4] Marta Bright, "Growing Green," *Oracle Magazine*, July 2008.

PG&E replaced its IBM DB2 database running on a mainframe with a highly scalable cluster of eight Oracle RAC nodes running on smaller, more-energy-efficient systems. The savings amounted to US$5 million annually.

Mercadolibre, Inc. is the largest online auction house in Latin America.[5] Its rapid rate of growth quickly outpaced a midsize symmetric multiprocessing (SMP) server. In 2004, the company replaced the server with a 4-node Oracle RAC cluster, which cost US$500,000 less than the big SMP box. Today, Mercadolibre's grid footprint has grown to a 16-node cluster—and the business has grown by a factor of four.

According to IDC, which documented the grid transformation, Mercadolibre's savings totaled approximately US$5.1 million over five years and came from a combination of these factors.

- Avoidance of hardware and software costs (US$1.18 million)
- Increased uptime (US$1.97 million)
- Increased search speed (US$1.40 million)
- Improved fraud prevention (US$500,000)

After subtracting Mercadolibre's capital and operating expenses, IDC projects net-present benefits of US$2.4 million over five years—equal to a return on investment (ROI) of 452 percent.

In-Memory Data Grids

For continuous operation, predictable high performance, and agile scalability, enterprise applications need virtualization and clustering capabilities in the middle tier. They also need in-memory data grids. Oracle Coherence—an in-memory data grid application—lets Java and .NET applications quickly access objects that are stored in memory spanning multiple physical machines in the middle tier.

Oracle Coherence provides horizontal scalability, predictable high performance, and high availability. Its architecture scales linearly as additional nodes are added (or removed) on the fly. Oracle Coherence achieves high performance by storing data in memory on middle-tier computers to avoid disk access and network delays. It ensures high availability by storing copies of the data on different servers in the data grid, thereby avoiding a single point of failure in case a node crashes or is taken offline for maintenance.

---

[5] IDC/Oracle, *Business Benefits Series—Mercadolibre*, 2005.

AbeBooks.com is the world's largest online marketplace for books, listing more than 110 million new, used, rare, and out-of-print books from 13,500 booksellers.[6] AbeBooks.com needed to enhance its online customer checkout experience and provide fast, reliable access to frequently used information such as publication inventory data, while increasing Web site performance and scalability. To achieve this, AbeBooks.com implemented Oracle Coherence across all six AbeBooks international Web sites, to support the company's growing customer base and extreme processing requirements. With Oracle Coherence, AbeBooks reduced transactional strain on system resources, seamlessly managed millions of daily updates from its seller community as well as 30,000 customer orders, and accelerated by threefold the delivery of information to the company's shopping basket application.

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. It is critical that IT departments be able to quantify and validate the services they provide as information technology becomes more tightly linked with business operations. As part of an organization's utility computing strategy, grid computing depends on accurate service-level management to ensure that the physical capacity of the data center is providing adequate service levels to end users.

Enabling Continuous Availability

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

Oracle continues to be at the forefront of developing high availability products and practices. From server failover with Oracle RAC and Oracle Application Server, to Oracle Active Data Guard and data replication, Oracle provides IT organizations with a comprehensive portfolio of solutions to keep the data center—and the business—running smoothly.

Server Failover

Server failover has been available for many years from both hardware and software manufacturers. The protections offered by a successful failover of a server are critical, and yet the hardware and software costs of setting up a standby server—used only if disaster strikes—can be prohibitive. With grid computing, standby resources can be used as active resources, resulting in higher utilization and lower costs.

[6] Oracle Success Story, *AbeBooks Inc. Accelerates Information Delivery with Flexible, Extreme Transaction Processing Solution,* August 2008.

Another consideration of server failover is failover time. Some applications are so critical that the business cannot afford to be without them even for a few minutes, while other applications can tolerate some interruption of operation.

As a pioneer in server-failover techniques, Oracle offers automatic failover capabilities for several server types. For example, Oracle Database with Oracle 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 simply remove the failed servers from service, repair or replace them, and then add them back to the server grid. Automatic migration and failover of services (or whole servers)—along with load balancing, workload management, and overload protection—ensure that mission-critical applications stay up and running.

Disaster Protection

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

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

Fidelity National Financial (FNF) is a leading provider of outsourced products and services, including technology solutions, financial and insurance services, and claims management. FNF published a detailed case study[7] that followed the progress of migrating a critical document management application (called EDoc) to an Oracle grid implementation. EDoc manages more than 50 million documents and affects numerous systems that are critical to FNF's service offerings. FNF has standardized on Oracle Database with Oracle RAC, Oracle Automatic Storage Management, Oracle Active Data Guard, and other high-availability components.

FNF faced heightened uptime requirements for EDoc because its customers were increasingly using online 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. FNF decided to migrate to Oracle RAC so it could scale the application to meet the company's evolving growth and quality-of-service requirements. FNF determined that an Oracle RAC implementation of EDoc would provide increased CPU horsepower, higher utilization of resource levels, and load-balancing capabilities.

According to the FNF case study, "[C]ollectively, 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." 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."

[7] Fidelity National Information Centers, *Oracle Maximum Availability Architecture—Architecture Case Study*, 2006.

Conclusion

Oracle first introduced enterprise grid computing in 2003. The state-of-the-art technologies and new database and middleware capabilities helped change the way IT departments operate. At the time, data center projects—including server consolidation, SOA development, space and power optimization, and large-scale implementations of rack-mounted Linux servers—seemed unrelated. We can now see that—taken together—these techniques can be described as a "grid computing approach to data center modernization."

The customer examples in this white paper illustrate how benefits derived from these techniques can be compounded. For example, server and storage consolidation increases utilization levels, which allows 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. With the current releases of Oracle Database and Oracle Fusion Middleware products—including the Oracle WebLogic and Oracle Tuxedo product lines—Oracle introduces a second generation of grid computing technology that builds on its strong foundation of scalable, fault-tolerant database and middleware clusters; virtualized computing and storage resources; and highly automated, end-to-end systems management. Today, with more than 10,000 Oracle customers deploying some level of grid computing technology, Oracle continues to lead the software industry in its commitment to grid computing solutions, products, and practices.

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	• Oracle Real Application Clusters enables database clustering so a single database can be deployed across a cluster of servers • Oracle Automatic Storage Management enables a shared pool of storage for the database • Oracle Active Data Guard creates read-only replicas for standby databases and disaster protection • Oracle In-Memory Database Cache delivers predictable low latency, high throughput, and high availability • Oracle Advanced Compression enables disk compression for cost and power savings and better performance • Oracle 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 delivers in-memory clustered data grids for Java and .NET objects • Oracle Tuxedo offers a clustered transaction processing monitor for C, C++, and COBOL applications • Oracle JRockit Real Time delivers a JVM with deterministic response times
HP Oracle Exadata Storage Server and HP Oracle Database Machine	• Oracle Database, Oracle Real Application Clusters, Oracle Automatic Storage Management, and Oracle Partitioning with database clusters connected to storage clusters by high-speed InfiniBands
Oracle Enterprise Manager	• Top-down management • Server cloning and bare metal provisioning • Automated administration • Real-time and predictive monitoring

Oracle Grid Computing
June 2009

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com

Oracle is committed to developing practices and products that help protect the environment

0109