Oracle 10*g*: Infrastructure for Grid Computing

An Oracle White Paper
September 2003

Oracle 10*g*: Infrastructure for Grid Computing

EXECUTIVE OVERVIEW

What is all the hype about grid computing?

Today, islands of computing within organizations make inefficient use of resources. Systems are slow to change and expensive to maintain.

Grid computing addresses these problems by providing an adaptive software infrastructure that makes efficient use of low-cost servers and modular storage, which balances workloads more effectively and provides capacity on demand. By scaling out with small servers in small increments, you get performance and reliability at low-cost. New unified management allows you to manage everything cheaply and simply in the grid.

What is new about grid computing?

Grid computing removes the fixed connections between applications, servers, databases, machines, storage – every component of the grid. By treating everything in the grid as a virtualized service, intelligent systems can optimize resource utilization and responsiveness.

Grid computing is based on five fundamental attributes: virtualization, dynamic provisioning, resource pooling, self-adaptive systems, and unified management.

What is Oracle doing with grid computing?

Oracle 10*g* allows companies to begin evolving their IT toward a grid computing model. For storage, databases, application servers, and applications, Oracle's new technology meets the requirements of grid computing. Together, Oracle Database 10*g*, Oracle Application Server 10*g*, and Oracle Enterprise Manager 10*g* provide the first complete grid infrastructure software.

This paper introduces and defines grid computing for businesses and describes how Oracle 10*g* products implement a grid infrastructure.

INTRODUCTION

Enterprise IT is under continuous pressure to do more with less. Change is constant, and companies need to adapt quickly to stay competitive. Meanwhile, requirements for availability and performance increase, while budgets tighten. To deal with the unpredictability and immediacy of computing needs, companies typically size servers for peak loads and staff IT organizations to handle ad hoc requests.

A new model of computing is emerging to address these problems. The vision of the grid computing model is shared throughout the industry, although some industry leaders have coined new terms for it, for example: Computing on Demand, Adaptive Computing, N1, Utility Computing, Hosted Computing, Organic Computing, and Ubiquitous Computing. In recognition of the significant new capabilities required to power grid computing, Oracle has named its new technology products Oracle 10*g*. This is the first major name change since Oracle added internet capabilities to Oracle8*i*.

Vision of Grid Computing

The central idea of grid computing is that computing should be as reliable, pervasive, and transparent as a utility. It shouldn't matter where your data or application resides, or what computer processes your request. You should be able to request information or computation and have it delivered – as much as you want, whenever you want. This is analogous to the way electric utilities work in that you don't know where the generator is or how the electric grid is wired. You just ask for electricity and you get it. The goal is to make computing a utility – a ubiquitous commodity. Hence, it has the name, "grid."

Origins of Grid Computing

Grid computing was conceived in the academic and research communities. Much like internet computing, which grew from the communication needs of dispersed scientific researchers, grid computing originated from the needs of the scientific community's needs to:

- Create a dynamic computing environment for sharing resources and results
- Scale to accommodate petabytes of data, and teraflops of computing power
- And keep costs down

Scientific Grids

SETI@home, the Search for Extraterrestrial Intelligence, is one of the earliest examples of a scientific grid. Signals from telescopes, radio receivers, and other sources monitoring deep space are distributed to the PCs of individual science buffs via the internet. This loose network of small computers crunches numbers, looking for patterns that could suggest signs of intelligent life.

Although the idea of harnessing idle computers across the internet is intellectually interesting, businesses will never want their data or their computing distributed to random computers. But, just as businesses have brought the concepts of the public internet in-house to make *intranets*, enterprises can bring the concepts of the scientific grids in-house to make *enterprise grids*. With both public grids and enterprise grids, grid computing is about harnessing the work of many of small computers.

Benefits of Enterprise Grid Computing

The need for low-cost computing drove the SETI@home innovation. The primary benefit of grid computing to businesses is achieving high quality of service and flexibility at lower cost. Enterprise grid computing lowers costs by:

- Increasing hardware utilization and resource sharing
- Enabling companies to scale out incrementally with low-cost components
- Reducing management and administration requirements

WHAT IS ENTERPRISE GRID COMPUTING?

Enterprise grid computing builds a critical software infrastructure that can run on large numbers of small, networked computers, by combining two related concepts:

- *Implement One from Many.* Grid computing coordinates the use of clusters of machines to create a single logical entity, such as a database or an application server. By distributing work across many servers, grid computing exhibits benefits of availability, scalability, and performance using low-cost components. Because a single logical entity is implemented across many machines, companies can add or remove capacity in small increments, online. With the capability to add capacity on demand to a particular function, companies get more flexibility for adapting to peak loads, thus achieving better hardware utilization and better business responsiveness.
- *Manage Many as One.* Grid computing allows you to manage and administer groups of machines, groups of database instances, and groups of application servers at low-cost. Grid computing first removes many of the administrative costs of managing a single system by making each database and each application server adaptive to changing circumstances. Then, the model makes managing many systems simple, by allowing them to be managed as a single logical entity.

New Trends in Hardware

Much of what makes grid computing possible today are the innovations in hardware. For example,

- **Processors**. New low-cost, high volume Intel Itanium 2, Sun SPARC, and IBM PowerPC 64-bit processors now deliver performance equal to or better than exotic processors used in high-end SMP servers.
- **Blade servers.** Blade server technology reduces the cost of hardware and increases the density of servers, which further reduces expensive data center real estate requirements.
- **Networked storage**. Disk storage costs continue to plummet even faster than processor costs. Network storage technologies such as Network Attached Storage (NAS) and Storage Area Networks (SANs) further reduce these costs by enabling sharing of storage across systems.
- **Network interconnects**. Gigabit Ethernet and Infiniband interconnect technologies are driving down the cost of connecting servers into clusters.

Power Comes from Software

Although the *newness* of grid computing comes primarily from hardware, the *power* of the grid infrastructure must be embodied in software. The capability of a database, for example, to store and retrieve data through an abstract interface without knowing much about the underlying location or structure of that data requires software intelligence. The capability of an application server to begin distributing work to newly added blade servers without going offline can only be accomplished with software. By providing software to leverage and control new grid hardware, Oracle supplies the grid infrastructure, and powers enterprise grids.

Grid Computing Aligns with Oracle's Strengths

The ultimate vision of grid computing is a new way to build data centers and a new way to think about IT. The foundations of grid computing, however, have been in the making for many years, and every one of Oracle's strengths and focus areas aligns with and contributes to grid computing:

- **Focus on Consolidation**. For years Oracle has recommended that companies consolidate geographically dispersed data centers, consolidate all company data into single databases, and consolidate all applications onto application servers. With the advent of grid computing, consolidation means using pools of relatively small machines to run consolidated databases and application servers.
- **Support for Standards.** The technology industry can be divided into two camps: those that promote a proprietary single-vendor platform, and those that support open standards for interoperability. Although Oracle provides a complete, integrated infrastructure platform and

extols the benefits of software suites, it has been a fervent supporter of open standards. Open standards must solidify for grid computing to reach its full potential, and Oracle continues to work toward this goal.

- **Expertise on Linux.** Oracle has both the leading application server and the leading database on Linux. Oracle brings its strengths in security, reliability, and performance to Linux to make it strong enough for enterprise deployments. Because Linux runs very well on small computers (one to four CPUs) and provides the best price for performance, Linux is the perfect operating system for the large number of small computers that constitute the hardware in an ideal grid environment.

GRID COMPUTING ATTRIBUTES

The requirements for grid computing infrastructure can be described by the following attributes:

- **Virtualization** at every layer of the computing stack
- **Provisioning** of work and resources based on policies and dynamic requirements
- **Pooling** of resources to increase utilization
- **Self-adaptive software** that largely tunes and fixes itself
- **Unified management** and provisioning

Virtualization at Every Layer

Virtualization is the abstraction into a service of every physical and logical entity in a grid. Virtualization is important because it enables grid components (such as storage, processors, databases, application servers, and applications) to integrate tightly without creating rigidity and brittleness in the system. Rather than making fixed ties that determine which application server node will handle requests from a particular application, for example, or where a database physically locates its data, virtualization enables each component of the grid to react to changing circumstances more quickly and to adapt to component failures without compromising performance of the system as a whole.

Dynamic Provisioning

Provisioning simply means distributing supplies where they are needed. In the context of the grid, "supplies" may mean server requests that need to be handled, data that needs to be accessed and used, or computations that need to be performed. Provisioning in the grid environment means a grid service broker that knows the resource requirements of one element of the grid and the resource availability of another element links the two together automatically and dynamically to make efficient use of resources. Then it adjusts the associations

as circumstances change. Policies, such as response time thresholds or anticipated peak demands, can be used to further optimize the associations of resource-requestors to resource providers.

Resource Pooling

Consolidation and pooling of resources is required for grids to achieve better utilization of resources, a key contributor to lower costs. By pooling individual disks into storage arrays and individual servers into blade farms, the grid runtime processes that dynamically couple service consumers to service providers have more flexibility to optimize the associations. Resource sharing also happens purely in software. Web services provide the model for applications to expose re-usable functionality for discovery and invocation by unrelated applications.

Self-Adaptive Software

With labor being the most significant portion of IT costs, savings due to better hardware utilization or more responsive systems become irrelevant if the everyday tasks of administrators are not automated and simplified. A grid infrastructure would be unworkable if every node required constant manual tuning and intervention. A critical grid infrastructure requirement is systems that automate the bulk of maintenance and tuning tasks traditionally performed by IT staff. More of the tasks that used to be performed by administrators must now be handled by the systems themselves.

Unified Management

Even with self-managing systems, human beings will always be involved in managing an enterprise grid, but the management tasks required by humans should be simplified with a single tool that can provision, monitor, and administer every element in the grid. Such a tool should evaluate availability and performance from the perspective of the user, such that any bottleneck in the system or any unavailable component raises alerts. Most importantly, with a grid infrastructure, IT professionals must be able to treat groups of systems as a single logical entity so that tasks can be performed once and executed on multiple machines.

- *Implement One from Many.* Together, the attributes of virtualization, dynamic provisioning, and resource pooling form the requirements for software that implements a single logical entity using many services running on multiple servers and crossing multiple disks—an entity which delivers high quality of service from low-cost components.
- *Manage Many as One.* Together, the attributes of self-adaptive software and a unified management model form the requirements for dramatically

lowering management costs by viewing the entire enterprise grid as one simple whole.

ORACLE 10*g*: GRID INFRASTRUCTURE

Oracle 10*g* provides the first complete, integrated software infrastructure to power grid computing. Oracle 10*g* takes the fundamental attributes of grid computing…

Implement One from Many

- Virtualization at every layer
- Dynamic provisioning
- Resource pooling

Manage Many as One

- Self-adaptive software
- Unified management

…and implements them throughout every element of the grid: storage, databases, application servers, and applications.

The diagram below visually depicts the way Oracle 10*g* products and features map to grid computing requirements.

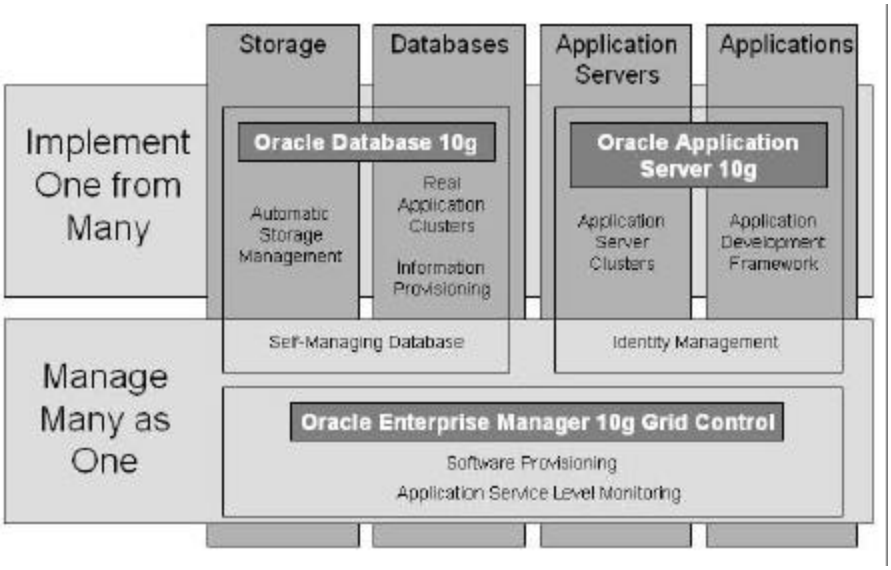

The following sections describe how grid computing attributes are embodied in Oracle's three grid infrastructure products:

- Oracle Database 10*g*
- Oracle Application Server 10*g*
- Oracle Enterprise Manager 10*g* Grid Control

Oracle Database 10*g*

Oracle Database 10*g* builds on the success of Oracle9*i* Database, and adds many new grid-specific capabilities. Other vendors implement certain portions of a grid infrastructure, for example pools of virtualized storage are becoming common, but no one else can provide a true grid database. Oracle Database 10*g* is based on Real Application Clusters, introduced in Oracle9*i*. There are more than 500 production customers running Oracle's clustering technology, helping to prove the validity of Oracle's grid infrastructure.

Real Application Clusters

Oracle Real Application Clusters enables a single database to run across multiple clustered nodes in a grid, pooling the processing resources of several standard machines. Oracle is uniquely flexible in its ability to provision workload across machines because it is the only database technology that does not require data to be partitioned and distributed along with the work.

In Oracle 10*g*, the database can immediately begin balancing workload across a new node with new processing capacity as it gets re-provisioned from one database to another, and can relinquish a machine when it is no longer needed—this is capacity on demand. Other databases cannot grow and shrink while running and, therefore, cannot utilize hardware as efficiently.

New integrated clusterware in Oracle 10*g* makes clustering easy by eliminating the need to purchase, install, configure, and support third-party clusterware. Servers can be easily added and dropped to an Oracle cluster with no downtime. Oracle has the only database technology to include clusterware for all operating systems, which dramatically reduces the opportunities for failure in a clustered environment.

Automatic Storage Management

Automatic Storage Management simplifies storage management for Oracle Databases. By abstracting the details of storage management, Oracle improves data access performance through sophisticated data provisioning, without requiring additional work from DBAs. Instead of managing many database files, Oracle DBAs manage only a small number of disk groups. A disk group is a set of disk devices that Oracle manages as a single, logical unit. An administrator can define a particular disk group as the default disk group for a database, and Oracle automatically allocates storage for and creates or deletes the files associated with the database object.

Automatic Storage Management also offers the benefits of storage technologies such as RAID or Logical Volume Managers (LVMs). Oracle can balance I/O from multiple databases across all of the devices in a disk group, and it implements striping and mirroring to improve I/O performance and data reliability. In addition, Oracle can reassign disks from node to node and cluster

to cluster, automatically reconfiguring the group. Because Automatic Storage Management is written to work exclusively with Oracle, it achieves better performance than generalized storage virtualization solutions.

Information Provisioning

In addition to the provisioning of work across multiple nodes and the provisioning of data across multiple disks, another type of provisioning happens within Oracle Database 10*g* — the provisioning of information itself. Depending on the volume of information and the frequency of access, it may be necessary to move data from where it currently resides or to share data across multiple databases. Oracle 10*g* includes various facilities to provide access to information when and where it's needed, matching information providers and information requestors. The most fine-grained and real-time of these facilities is Oracle Streams, which can migrate data from one database to another while both are online. Bulk data transfers are more suitable in some circumstances, for which Oracle provides Data Pump and Transportable Tablespaces. In Oracle 10*g*, all information provisioning facilities can move data to databases running on different operating systems, which is particularly useful for migrating databases to a grid environment, for example, blade servers running Linux.

Self-Managing Database

The first step toward manageability in a grid environment is making each individual system require less human attention. Oracle 10*g*, with the new self-managing database, reduces the maintenance and tuning tasks required by administrators. Oracle Database 10*g* includes an intelligent database infrastructure that takes snapshots of vital statistics and workload data to be analyzed for self-tuning and for advising administrators. The self-managing database automatically diagnoses problems such as poor connection management, lock contention, and poorly performing SQL. Oracle Database 10*g* fixes certain diagnosed problems and advises DBAs about simple corrective measure in other cases. Oracle's self-managing database enables DBAs to concentrate on more value-added work and dramatically reduces administration costs of databases.

Oracle Application Server 10*g*

Oracle Application Server 10*g* provides a complete infrastructure platform for developing and deploying enterprise applications, integrating many functions including a J2EE and Web services runtime environment, an enterprise portal, an enterprise integration broker, business intelligence, web caching, and identity management services.

Oracle Application Server 10*g* adds new grid computing features, building on the success of Oracle9*i* Application Server, which has hundreds of customers running production enterprise applications.

Application Server Clusters

Oracle Application Server 10*g* run-time services can be pooled and virtualized via application server clusters. Every service within the Oracle Application Server – HTTP, J2EE, Web cache, Web Services, LDAP, portal and others – can be distributed across multiple machines in a grid. New features in Oracle 10*g* enable performance thresholds to be defined beyond which new application server instances can automatically be added and started (or relinquished) to process additional work on new nodes of a grid, delivering capacity on demand.

With Oracle 10*g*, an administrator can define a set of policies or business rules that affect how individual work is provisioned across multiple machines. Specifically, workload allocation can be influenced by resource consumption metrics, such as CPU or memory usage, or application-specific metrics, such as transaction throughput or JDBC connections, or workload can be provisioned based on schedules, such as peak times of day or end of quarter. Oracle Application Server 10*g* provides out-of-the-box instrumentation that captures these various metrics and creates advisories based on historical and real-time information to help administrators make the best policy choices.

Oracle Application Server 10*g* also provides several availability enhancements. Because Oracle 10*g* includes clustering of every service within the application server, there is no single point of failure. Both planned and unplanned downtime of an individual instance will simply cause requests to be routed to another node. Because Application Server 10*g* includes efficient session replication, any type of failure (even that of a J2EE application holding state) will remain transparent to the user. Application Server 10*g* further improves application reliability through its interaction with Oracle Real Application Clusters. If an instance in the back-end database goes down, Application Server 10*g* is notified to reconnect. Without notification from a failed instance, an application server would wait for an IP time out, which takes several minutes, but the multi-tier failover notification feature reduces recovery time in such cases to mere seconds, and both failure and recovery remain transparent to the user.

Identity Management

Centralized application user administration becomes even more important in a grid environment. Identity management features within Application Server 10*g* simplify and centralize account creation, suspension, and deletion and privilege modification, all of which lower administration costs and reduce security vulnerabilities.

Oracle provides centralized user provisioning and single sign-on for users across all applications deployed to the Oracle Application Server. Access privileges for all applications can be created and revoked through a single interface. Identities can be managed through Oracle Internet Directory, a standards-based LDAP directory that benefits from the availability and scalability of being built on the Oracle Database.

Application Development Framework

Tightly integrated with Oracle Application Server 10*g* are the development tools that enable companies to quickly develop custom internet applications, and then easily deploy those applications to Oracle Application Server.

Applications for scientific grids, such as SETI@home, must be designed explicitly to run on loosely connected grids. In contrast, enterprise applications do not need to be re-designed to exploit the availability, scalability, and performance benefits of enterprise grids. When applications are deployed to an application server in a grid, those applications benefit immediately from the transparent workload distribution, load balancing, and scheduling necessary to efficiently coordinate work across multiple servers.

To gain additional benefits from grid computing, however, enterprise applications can expose their behavior to other applications and to management tools through standardized interfaces in a service-oriented architecture. Oracle Developer Suite 10*g*, which includes JDeveloper 10*g*, enables developers to create dynamic Web sites, J2EE applications, and Web services and to make these services accessible through enterprise portals and wireless devices. Applications designed to a service-oriented architecture can leverage a set of standards-based internet protocols to communicate with other applications and heterogeneous resources across a grid. Designing to a service-oriented architecture enables companies to reduce development time and integration costs.

Oracle Enterprise Manager 10*g* Grid Control

Oracle Enterprise Manager 10*g* Grid Control is the complete, integrated, central management console and underlying framework that automates administrative tasks across sets of systems in a grid environment. Grid Control helps reduce administration costs through automation and policy-based standardization. With Oracle Grid Control, IT professionals can group multiple hardware nodes, databases, application servers, and other targets into single logical entities. By executing jobs, enforcing standard policies, monitoring performance and automating many other tasks across a group of targets instead of on many systems individually, Grid Control enables IT staff to scale with a growing grid. Because of this feature, the existence of many small computers in a grid infrastructure does not increase management complexity.

Software Provisioning

Because of the potentially large number of physical nodes, it's especially important in a grid environment that installation and configuration of the software running on those nodes is fast and requires no human intervention. Manually installing software on hundreds of nodes would be time consuming and cumbersome. Administrators would certainly find ways to work around a manual installation, but the workarounds could lead to unsupportable upgrade situations and lost information about the configuration of the system.

With Grid Control, Oracle 10*g* automates installation, configuration, and cloning of Application Server 10*g* and Database 10*g* across multiples nodes. Oracle Enterprise Manager provides a common framework for software provisioning and management, allowing administrators to create, configure, deploy, and utilize new servers with new instances of the application server and database as they are needed. This framework is used not only to provision new systems but also to apply patches and upgrade existing systems.

In Oracle Application Server 10*g*, applications can be deployed once to a single application server instance, registered with the central repository, then automatically deployed to all relevant nodes in the grid. As changes are made to the application and as new nodes are added to the grid, nodes can be kept in sync.

Application Service Level Monitoring

Oracle Grid Control views the availability and performance of the grid infrastructure as a unified whole, as a user would experience it, rather than as isolated storage units, processing boxes, databases, and application servers. An administrator can trace a performance or availability problem as experienced by a user from end to end – from the user visible Web page, through external and internal networks, to application code, application server, and database access. Grid Control then allows an administrator to trace the root cause of the problem down to the individual Java class, for example, or the individual system configuration parameter.

OPEN GRID STANDARDS

With Oracle 10*g*, companies can begin implementing grid computing today, but the open standards that will make grid computing as pervasive as the internet are still under development, primarily by the Global Grid Forum (GGF). Oracle is a GGF sponsor and participates in working groups, chairing the Data Access and Integration (DAI) group. The Open Grid Services Architecture (OGSA) is a specification in active development within the GGF to define the general services-based approach to grid computing. Other working groups, such as Open Grid Services Infrastructure (OGSI) and OGSA-DAI, endeavor to define the common interfaces and protocols for various grid services. Oracle plans to actively support all grid-related open standards as they emerge.

CONCLUSION

Many phrases have been coined to describe new computing models created by the IT industry. Grid computing is the emerging standard, and grid computing is Oracle's approach to lowering costs while improving quality. The benefits of grid computing to businesses are real: increasingly flexible systems that can largely self-manage; better availability, performance and scalability at lower cost; and the opportunity for incremental investment and immediate return.

Grid computing will not radically change enterprise data centers, and it does not require throwing out existing investments and best practices. However, grid computing is also not just a passing fad. Enterprise grid computing, based on the Oracle 10*g* infrastructure, will be the foundation of information technology for the future, resulting in more cost effective computing for running more nimble, data-driven businesses.

Oracle 10*g*: Infrastructure for Grid Computing
September 2003
Author: Miranda Nash
Contributing Authors:
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

Oracle Corporation provides the software that powers the internet.