

Oracle and the Grid

An Oracle White Paper
November 2002

Executive Overview	3
Introduction	4
Trends Promoting Grid.....	4
Hardware Trends	4
OS Trends.....	4
Virtualization	5
Grid Momentum.....	5
Grid: The Next Big Thing.....	5
Grid Taxonomy	5
Scavenging Resources	6
Sharing Resources.....	6
Dedicating Resources	7
Oracle Resonates with the Grid	8
Leveraging Commodity Components.....	8
Information Sharing	9
Scheduling and Resource Management	10
Standard Grid Environments.....	11
Portability.....	12
Operational Characteristics	12
Positioning for the Future.....	12
Portability Across Grid Phases	13
Right Architecture for Future Grid Technologies	13
Product Directions Aligned with Grid	13
Grid Standards Support	13
Conclusion.....	13

EXECUTIVE OVERVIEW

Grid computing is poised to drastically change the economics of computing. Grid computing can dramatically lower the cost of computing, extend the availability of computing resources, deliver higher productivity, and higher quality. Enterprises have already started to leverage the Grid, and the Oracle database has already gained popularity in Grids.

Oracle has thought about and studied the Grid for years, Oracle builds the Oracle9i Database product using a Grid. The Grid can be a confusing term, since it is used to describe several different styles of computing. To make the Grid easier to understand and utilize, this paper describes a taxonomy of Grid adoption. Understanding this taxonomy can help enterprises in gaining remarkable Grid benefits with their IT investments.

The basic idea of Grid computing is the notion of computing as a utility, analogous to the electric power grid or the telephone network. Utilities must be Unbreakable. If electric utilities weren't unbreakable, you would own a power generator. If the phone system weren't secure, you would not use it. The Grid must be unbreakable and highly available. Your data, your computation, and your resources in the Grid must be secure.

Oracle resonates with the Grid. The ideas of Grid computing are aligned with capabilities and technologies Oracle has been developing for years. Oracle provides substantial Grid computing technology, available today, that can help you capitalize on the Grid. Oracle9i possesses the key differentiating technologies – Oracle Real Application Clusters, Oracle Streams, and Oracle Transportable Tablespaces – for building the Grid. Only Oracle9i delivers the operational requirements – portability, RAS, security, scalability, and manageability – for the Grid.

Oracle9i has the right architecture for delivering future Grid computing technologies. Oracle is making its technology stack work the best in Grid environments. The Oracle Globus Development Kit (OGDK) enables Oracle technology to be used with the Grid standard Globus toolkit. Oracle plans to fully support emerging Grid standards. Oracle is working with the Global Grid Forum, to help develop Grid standards. Oracle will make it easy for you to take advantage of Grid computing. Your investments in Oracle9i are fully leveraged for you to utilize future Grid computing technologies.

INTRODUCTION

Today's mantra for computing is affordability. Enterprises are looking at ways of reducing costs and increasing efficiencies of their processes and systems., and Grid computing offers exactly that. Grid computing increases the efficiencies of enterprise resources. Grid computing is a way to consolidate your hardware – eliminating islands of underutilized computers. You can create centralized pools of computing and allocate computing resources to the priorities of your organization.

At the highest level, the central idea of Grid computing is computing as a utility. You shouldn't care where your data 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, and 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 commodity, and ubiquitous. Hence it has the name, the Grid.

This view of utility computing is, of course, a client side view. From the server side, or behind the scenes, the Grid is about resource allocation, information sharing, and high availability. Resource allocation ensures that all those that need or request resources are getting what they need. Resources are not standing idle while requests are going unserved. Information sharing makes sure that the information users and applications need is available where and when it is needed. High availability ensures that all the data and computation must always be there--just as a utility company must always provide electric power.

Grid computing transforms the way enterprises use resources; it virtualizes the use of enterprise resources. It virtualizes the use of enterprise resources across servers in a data center, across data centers in an enterprise, and across enterprises.

TRENDS PROMOTING GRID

The Grid is poised to be the next big thing. There are a number of trends which, taken together, will make Grids unstoppable.

Hardware Trends

In hardware, every vendor has announced or is delivering blades. Computer blades offer the lowest cost computing power, sometimes as much as 80% less than SMP. These blades can easily be assembled into blade farms, which are the most effective and scalable form of commodity computing. These blade farms are now being fitted with interconnects, making them hardware clusters. As such, they form the most cost effective form of commodity clusters--the most likely future architecture of computing.

OS Trends

In software, Linux continues to grow faster than any other OS. Today, Linux cannot scale to large SMP. But since blades are one to four cpus, Linux runs well

on them today. The economic advantage of blades over SMP will cause blades to dominate, and since Linux already works well for blades, this will accelerate Linux growth. Finally, Linux has a price advantage, which becomes more important as the number of blades grows, again accelerating Linux adoption. So commodity clusters naturally go well with Linux, the commodity OS.

Virtualization

In both the software and hardware industries, one of the big buzzwords at the moment is virtualization. But nothing is more virtual than a utility. A lot of vendors are trying to claim that their new strategy is virtualization, or utility computing – which is exactly what Grid computing is all about. Soon, these people will most likely come to embrace the Grid.

Grid Momentum

In the technology industry, grid momentum is building. Some major vendors, such as Oracle, are offering grid-enabling technology. Others, such as IBM, are planning to offer grid-enabling technology. The Grid standards body, GGF, is in place and has the support of all major technology vendors.

In IT organizations, Grid momentum is also building. Grid technologies promise increased utilization of existing hardware. Grids let you allocate your resources to meet the needs of your business, instead of having islands of computing that are idle or overloaded. And, as existing hardware needs to be replaced, blades offer the lowest cost. The economics are so compelling that enterprises have already started leveraging blade servers for Grid computing.

Grid: The Next Big Thing

If you look at the Web, it is really about presentation of information over the Internet or your intranet. After information presentation, the next logical step is information processing. Processing information over the Internet or your intranet is exactly what the Grid is all about. So one way to think of it is that the Grid is the next phase of the Internet, after the Web. In 1997 it was hard to see everything the Web would become, but you could tell it was going to be big. That's the state of the Grid today.

GRID TAXONOMY

Oracle has thought about and studied the Grid for many years. Oracle began adopting Grid computing over three years ago. We build the Oracle database product using a Grid. The Grid can be a confusing term, since it is used to describe several different styles of computing. To make the Grid easier to understand and utilize, this paper describes a taxonomy of Grid adoption. Understanding this taxonomy can help enterprises in gaining remarkable Grid benefits with their IT investments.

Most likely, enterprises will adopt the Grid by implementing one phase and then

gravitate towards more advanced phases to reap better returns and results. As you advance through the phases below, the efficiencies of resource utilization also increase.

Scavenging Resources

Scavenging idle resources is one implementation of the Grid. The thought of reclaiming unused resources from the thousands of PCs sitting idle every night is extremely enticing and often first interests organizations in the benefits of Grid computing. Such a solution has enormous potential, and at first glance, has almost no cost. This implementation is popular with researchers who have very little money, and have embarrassingly parallel problems to solve. (Embarrassingly parallel problems can be broken into pieces and processed independently, without need to exchange any information between the parallel processes.) Typical implementations involve a central server doling out work to large numbers of small nodes. The central server collects and assembles the results of the processing.

While the thought of access to massive amounts of processing power for little cost is appealing, this model of Grid computing has many limitations. The most problematic is that the resources, the unutilized PCs, are often outside the administrative control of those using the resources. The availability and reliability of the distributed resources is limited, making scheduling extremely difficult. Also, the weak or non-existent trust model of such Grids prohibits their use for processing any type of sensitive or proprietary data. Finally, the class of applications that are embarrassingly parallel is quite limited, reducing the usefulness of these grids with many organizations.

There are, however, some successful examples of organizations scavenging resources on a Grid to solve sophisticated problems. SETI@home, a project that uses idle PCs on the Internet to mine radio telescope data for signs of extraterrestrial intelligence, is an example of such a Grid implementation.

Sharing Resources

Sharing resources is another implementation of the Grid. In a Shared Resource Grid, organizations take advantage of underutilized resources within or across organizations. Using the Grid, participants can easily share compute and data resources by moving applications and data around the grid to leverage available resources. The result is a moderate number of medium to large nodes, which may fall under a single administrative domain, or span multiple administrative domains. Schedulers on the Grid track resource availability, and assign resources accordingly. Participants benefit through access to additional shared resources with little investment. Such grids are built using existing resources, and can leverage underutilized legacy hardware.

As in the case of a scavenging resource Grid, the multiple administrative domains can provide a challenge to organizations. Not all systems within the Grid may be trusted for sensitive data, and moving large amounts of data to the available

resource may not be practical. Political issues can also derail resource sharing projects, as many organizations are loath to lose control over the resources they own.

Many customers are using Oracle technology to implement this type of Grid. CERN is currently in the process of selecting technology partners for the LHC Computing Grid, which will allow thousands of physicists across the world to analyze petabytes of distributed particle physics data. Oracle9i Database and Oracle9i Application Server are candidate solutions to meet the LHC Computing Grid's technology requirements. Oracle9i Database offers the scalability, information sharing, and VLDB features necessary for thousands of users to share and access data within the LHC Computing Grid. Oracle Transportable Tablespaces offers CERN an extremely fast mechanism to share large amounts of data across many sites.

Dedicating Resources

Sharing resources across multiple administrative domains, political barriers, trust issues, and the excess data movement needed to access Grid resources, can make sharing resources impractical in many organizations. Dedicating resources to Grid computing addresses these issues by creating centralized pools of computing that can be shared by all participating applications. The resulting small number of very large computing pools eliminates the need to ship large amounts of possibly sensitive data outside the trusted and controlled administrative domain.

At first glance, simply pooling existing resources appears a likely method for implementing a dedicated resource grid. However, in practice, many organizations find it difficult to pool existing resources. Beyond the logistical issues of finding data center space for pooled resources, political issues can again derail such an implementation. A more likely scenario is implementing a dedicated resource pool as new resources are purchased.

Farms of low cost commodity server blades will most likely dominate these new dedicated resource pools. These blades provide the best price performance and are designed for high-density installations, thus reducing the need to build new, expensive data centers. Farms of blade servers provide a heterogeneous resource that can be easily allocated to applications as needed. The granularity of resource is small, facilitating adding and removing resources allocated to applications as loads grow and shrink.

Many organizations have already successfully implemented dedicated resource Grids. The police force of Nordrhein-Westfalen has consolidated all their budgeting, warehousing, procurement, inventory and stock management, controlling, and cost and service accounting applications on Oracle Real Application Clusters on a 3-node linux cluster N800 server.

Electronic Arts performed an ROI study comparing a 4-node dual-CPU versus a 2-node 4-CPU Linux cluster with RAC. They chose the smaller commodity dual-

CPU cluster as it provided better scalability and price performance.

ORACLE RESONATES WITH THE GRID

The ideas of Grid computing are aligned with capabilities and technologies Oracle has been developing for years. Oracle provides substantial Grid computing technology, available today, that can help you capitalize on the Grid. Oracle9i is the best on commodity clusters, the hardware for the Grid. It possesses the key technology differentiators for the Grid. Only Oracle9i delivers the operational characteristics necessary for the Grid. Oracle is making it easy for you to utilize Grid technologies.

This section describes the key Grid problems and lists the Oracle technology that helps you in addressing these problems.

Leveraging Commodity Components

Run real applications on commodity hardware and OS: Commodity hardware, such as blade servers and commodity OS, such as Linux offer the most compelling cost benefits for the Grids. Grids must leverage these commodity components to achieve maximum utilization.

Oracle Real Application Clusters

Oracle9i makes the commodity components – commodity hardware and commodity OS – unbreakable. Oracle Real Application Clusters (RAC) runs real applications on the commodity components. Oracle9i Database has proven to be the best database on commodity components. Oracle9i has unlimited scalability on Linux clusters. This has been proven with 32 nodes; although there is no architectural limit. Oracle9i has extended Linux with clustered file systems and other HA extensions. Oracle9i Database has the leading TPCC Linux benchmarks with 138,362 tpmC at \$17.21/tpmC on 8 node 4-CPU clusters.

RAC enables high utilization on commodity blade farms. Applications running on RAC can dynamically leverage more blades provisioned to them. Similarly, these applications can easily relinquish these blades when they no longer need them. Conversely, commodity databases have remarkably low utilization on commodity components. On the commodity databases, you need to allocate for peak loads and allocate spares. You cannot add and remove blades to the commodity databases without bringing down the entire system.

RAC, based on shared disk architecture, can grow and shrink on demand. This is not possible with databases from other vendors as they are based on shared nothing architecture, which does not offer this flexibility. With shared nothing, data is partitioned artificially. When more blades are added, all the data needs to be repartitioned to allocate data to the new blades. Similarly, when blades need to be taken off, data needs to be repartitioned before taking off the blades.

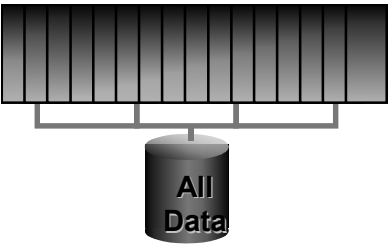
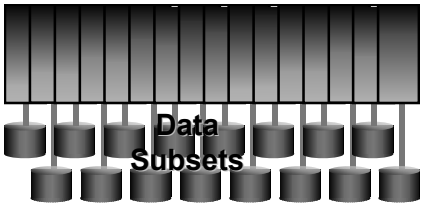
Shared Disk	Shared Nothing
Runs real applications	Runs Benchmarks
Natural data partitioning	Artificial data partitioning
Add blade	Add blade and repartition
Remove blade	Remove blade and repartition
	

Table 1: Shared disk offers greater flexibility.

Information Sharing

Get information delivered to you when you need it, regardless of where it resides: To process information on any available resource, Grid must efficiently share information across multiple systems. Grid must also provide access to data residing on heterogeneous systems -- database systems from multiple vendors and file systems.

Oracle Transportable Tablespaces

Oracle Transportable Tablespaces offers Grid users an extremely fast mechanism to move a subset of data from one Oracle database to another. Transportable Tablespaces allows Oracle data files to be unplugged from a database, moved or copied to another location, and then plugged into another database. Unplugging or plugging a data file involves only reading or loading a small amount of metadata. Transportable Tablespaces also supports simultaneous mounting of read-only tablespaces by two or more databases.




Figure 1: Oracle Transportable Tablespaces

Oracle Streams

Some data needs to be shared as it is created or changed, rather than occasionally shared in bulk. Oracle Streams can stream data between databases, nodes or blade farms in a Grid. It provides a unified framework for information sharing, combining message queuing, replication, events, data warehouse loading, notifications and publish/subscribe into a single technology. Oracle Streams can keep two or more copies in sync as updates are applied. Streams automatically captures database changes, propagates the changes to subscribing nodes, applies changes, and detects and resolves any conflicts. Streams can also be used directly by applications as a message queuing feature, enabling communications between applications in the Grid.

Distributed SQL and Distributed Transactions

Oracle Distributed SQL allows Grid users to efficiently access and integrate data stored in multiple Oracle and non-Oracle databases. Transparent remote data access with Distributed SQL allows Grid users to run their applications against any other database without making any code change to the applications. While integrating data across multiple data stores, the Oracle database intelligently optimizes the execution plans to access data in the most efficient manner. The Oracle database also automatically manages transactions across multiple data stores. Oracle XA capability allows Grid users to coordinate distributed transactions across multiple resources such as legacy applications and third-party application systems.

Heterogeneous Access with Generic Connectivity and Transparent Gateways

The heterogeneous access functionality of the Oracle database allows Grid users to run their Oracle database applications against non-Oracle databases transparently, without making any code changes. The applications access the non-Oracle database as if they were accessing the Oracle database. With generic connectivity, Oracle accesses the non-Oracle database using industry standard ODBC and OLEDB drivers. With transparent gateways, Oracle knows the capabilities of the non-Oracle database and does the necessary translation and optimization to access it in the most efficient manner. The supported database systems include DB2, Microsoft SQL Server, Informix, Ingres, and Sybase.

External Tables and BFILES

With Oracle external tables, Grid users get relational access data stored on a file system. They can join this data with data in any other Oracle table or perform any other Oracle database operation as if the data were stored in a regular Oracle database table. BFILES functionality allows Grid users to manage meta-data and to access data stored in files. BFILES are large binary data objects stored in operating system files outside database tablespaces.

Scheduling and Resource Management

Controlled resource access and allocation: Grid resources are accessed by a dynamic collection of

intra-enterprise or inter-enterprise users. In such an environment, tasks need to be scheduled and resources need to be appropriately allocated.

Oracle Scheduling Capabilities

Grid users can use the scheduling capabilities of the Oracle9i Database for submission of Oracle database tasks. The Oracle job queue functionality offers a flexible mechanism to schedule database tasks. You can run the jobs at a specific time and at specified intervals. You can also specify the maximum number of job queue processes that can be used to run these jobs. In an Oracle Real Application Cluster environment, you can either run the jobs at any available Oracle instance or at a specific instance.

Oracle Resource Manager

Oracle Resource Manager allows resource administrators to limit the Oracle database resources allocated to Grid users. Managing resources ensures each grid user gets a fair share of the available computing resources.

Standard Grid Environments

Support for standard Grid environments: The Globus toolkit has emerged as the standard for building Grids. Grid users must be able to leverage Oracle Grid technologies in a standard Grid environment.

Oracle Globus Development Kit

The Globus toolkit has a set of useful components that can be used to build Grid applications and programming tools. The Globus toolkit has become the de facto standard for doing Grid computing. The Oracle Globus Development Kit (OGDK) offers scripts to use Oracle technology efficiently with the Globus toolkit.

Grid Access to Oracle9i Utilities

The Globus Resource Allocation Manager (GRAM) provides resource allocation and process creation, monitoring, and management services. GRAM implementations map requests expressed in a Resource Specification Language (RSL) into commands understood by local schedulers and computers. Using this feature, Grid users can remotely invoke, using GRAM, Oracle Utilities such as export, import, and sqlplus to perform required actions in an Oracle9i Database. This is a powerful mechanism that can be used to set up remote databases for performing Grid tasks.

Grid Access to Oracle9i Database

This feature of the OGDK provides a Globus GRAM plug-in to invoke PL/SQL routines or SQL commands specified in Globus RSL. This can also be used for submitting jobs to the Oracle scheduler via the scheduler's PL/SQL API. This allows Grid users to perform, schedule, and monitor Grid tasks at the remote databases.

Grid Resource Information Service (GRIS) for Oracle Database

The Grid Resource Information Service (GRIS) component of the Globus toolkit provides information for monitoring and discovering a Grid resource. This feature allows Grid users to discover and monitor Oracle Databases. GRIS for the Oracle database exposes database attributes and properties useful to Grid users.

Portability

Utilize any available resource – hardware and operating system – for the task: Applications in the Grid environment must be portable. To use the available resource, the applications must be able to run on any hardware or operating system without any modification.

Oracle Portability

Oracle portability across operating systems and hardware means Oracle9i can support heterogeneous grids. Heterogeneous grids allow you to use all of your existing hardware, not just one vendor's hardware. Oracle9i is portable into and out of Grid configurations, so you can take an application developed on SMP and easily move it to Grid infrastructure. Oracle9i runs on all major operating systems, and has the same consistent features and capabilities on all of them, since it uses a single code base. Oracle portability means that whatever direction the Grid takes, whatever operating system dominates the Grid, Oracle9i will work there.

Operational Characteristics

Grid must be unbreakable: You can't break into the Grid. Your data, your computation, and your resources in the Grid must be secure. You can't break the Grid. The Grid must be highly available.

Reliability, Availability, and Serviceability

Oracle9i brings the highest levels of reliability and availability to the Grid. Acknowledging that failures will happen, Oracle9i provides near instantaneous recovery from system faults, meeting the most stringent service level agreements. It provides robust features to protect from data errors and disasters. Oracle9i also includes tools to minimize planned downtime, critical for any interactions with global business partners in a 24x7 environment.

Oracle Advanced Security

Oracle9i provides industry-proven security for the Grid. Oracle9i security mechanisms have been validated by 15 independent security evaluations. Oracle9i offers comprehensive functionality for addressing security needs throughout the enterprise. It offers the most complete functionality for securing transport, authentication, authorization, and auditing. Row-level security and data encryption are provided to protect your most sensitive data.

POSITIONING FOR THE FUTURE

Grid technologies are evolving rapidly. Oracle assures the cost conscious

enterprises that their investments in Oracle today will be leveraged for future Grid technologies. Oracle possesses the right architecture and has its product directions fully aligned to deliver future Grid computing technologies.

Portability Across Grid Phases

Oracle9i is portable across Grid phases. You can utilize your existing resources better with shared resource Grids. Later when you consolidate your resources on blade farms or cluster hardware, Oracle Real Application Clusters will enable your applications to run on the consolidated hardware without any application code change.

Right Architecture for Future Grid Technologies

Oracle9i possesses the right architecture to expand its Grid technology stack. Oracle9i's hardware and OS portability guarantee the same application behavior on different platforms. Oracle's shared everything architecture is best suited to exploit growing blade server technology. Oracle Streams provides the most comprehensive information sharing offering for autonomous communication between Grid entities. Oracle Transportable Tablespaces offer an extremely fast mechanism to mount and unmount data.

Product Directions Aligned with Grid

Oracle product directions are aligned with the Grid. Oracle already supports more Grid computing technology than any of its competitors, as described in the previous section. Your investments in Oracle are well leveraged--you can incrementally adopt additional Grid computing technology as Oracle expands its technology stack.

Grid Standards Support

Oracle is committed to supporting industry standards. Oracle is working with the Global Grid Forum to help define Grid standards. Just as Oracle has supported in its products, and is helping other standard, such as J2EE, Web Services, Xquery, and SQL, Oracle intends to fully support Grid standards. Often, the work of various standards bodies overlaps or various standards bodies develop complementary technologies. By investing in Oracle9i, you can leverage complementary standards and interoperate with various standards and technologies.

CONCLUSION

Grid computing is poised to change the economics of computing. Grid computing enables efficient utilization of enterprise resources and thus drastically reduces enterprise computing costs. Fortune 100 companies have already started reaping Grid benefits. Enterprises can use Oracle technology, available today, to benefit from the Grid. Only Oracle9i possesses the key technology differentiators that can help you capitalize the most on the Grid. Oracle is committed to supporting Grid standards and making it easier for you to utilize the Grid. Oracle's product

directions are aligned with the Grid. Investments you make in Oracle9i today are leveraged to reap additional Grid benefits when Oracle enhances its Grid technology stack.



Oracle and the Grid
November 2002
Author: Brajesh Goyal
Contributing Authors: Benny Souder

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 is a registered trademark of Oracle Corporation. Various product and service names referenced herein may be trademarks of Oracle Corporation. All other product and service names mentioned may be trademarks of their respective owners.

Copyright © 2002 Oracle Corporation
All rights reserved.