
WHITE PAPER ON ORACLE 10g GRID

By: Jaya Wazirani Agarwal

Introduction

Grid computing is a new IT architecture that leads to more reliable, recoverable and lower cost enterprise information systems. By efficient use of grid computing, groups of independent, modular hardware and software components can be connected and rejoined to meet the changing needs of businesses.

The motto of grid computing is to resolve some common problems with enterprise IT: the problem of application clusters that lead to under utilized, dedicated hardware resources; the problem of large, unwieldy systems that are very expensive to maintain and difficult to change; and the problem of distributed, fragmented and disintegrated information that cannot be fully exploited by the enterprise as a whole.

For over a decade Oracle has been providing software that allows organizations to adopt distributed computing architectures as part of their business transactional applications, content management applications, and business intelligence applications. Grid or utility Computing is based upon a well-established foundation. A distributed computing architecture often utilized in the world of high-performance technical computing. This approach allows organizations to harness together the power of the newest generation of high-performance and low-cost systems. This paper represents how Oracle 10g allows organizations to deploy Grid Computing as a foundation for business-oriented transactional, content management, and business intelligence applications.

What is GRID COMPUTING?

The grid computing treats all IT resources as a single resource and exploits the distinct nature of individual resources within single pool

Grid computing achieves a balance between the benefits of collective resource management and flexible independent resource control. IT resources managed in a grid include:

- **Infrastructure** - the hardware and software that create a data storage and program execution environment;
- **Applications** - the program logic and flow that define specific business processes;
- **Information** – the meanings inherent in all different types of data used to conduct business.

Infrastructure Grid

Infrastructure grid resources include hardware resources such as storage, processors, memory, and networks, as well as software designed to manage this hardware, such as databases, storage management, system management, application servers, and operating systems.

Virtualisation and provisioning of infrastructure resources mean pooling resources together and allocating to the appropriate consumers based on policies. For example, one policy might be to dedicate enough processing power to a web server that it can always provide sub-second response time. That rule could be fulfilled in different ways by the provisioning software in order to balance the requests of all consumers.

Treating infrastructure resources as a single pool and allocating those resources on demand saves money by eliminating under utilized capacity and redundant capabilities. Managing hardware and software resources holistically reduces the cost of labour and the opportunity for human error.

Spreading computing capacity among many different computers and spreading storage capacity across multiple disks and disk groups removes single points of failure so that if any individual component fails, the system as a whole remains available. Grid computing also affords the option to use smaller individual hardware components, such as blade servers and low cost storage, which enables incremental scaling and reduces the cost of each individual component, giving companies more flexibility and lower cost.

Infrastructure is the dimension of grid computing that is most familiar and easy to understand, but the same concepts apply to applications and information.

Storage Virtualisation

Oracle Automatic Storage Management (ASM), a feature of Oracle Database 10g, provides a virtualisation layer between the database and storage so that multiple disks can be treated as a single disk group and disks can be dynamically added or removed while keeping databases online. Existing data will automatically be spread across available disks for performance and utilization optimisation. In Oracle Database 10g Release 2, ASM supports multiple databases, which could be at different software version levels, accessing the same storage pool.

Grid Management

Because grid computing pools together multiple servers and disks and allocates them to multiple purposes, it becomes more important that individual resources are largely self-managing and that other management functions are centralized.

The Grid Control feature of Oracle Enterprise Manager 10g provides a single console to manage multiple systems together as a logical group. Grid Control manages provisioning of nodes in the grid with the appropriate full stack of software and enables configurations and security settings to be maintained centrally for groups of systems.

Another aspect to grid management is managing user identities in a way that is both highly secure and easy to maintain. Oracle Identity Management 10g includes an LDAP-compliant directory with delegated administration and in Release 2, federated identity management, so that single sign-on capabilities can be securely shared across security domains. Oracle Identity Management 10g closely adheres to grid principles by utilizing a central point for applications to authenticate users – the single sign-on server – while distributing control of identities via delegation and federation to optimise maintainability and overall operation of the system.

Applications Grid

Business Process Management

Business Process Execution Language (BPEL) provides a standard for orchestrating processes into complex business flows in service oriented architecture. The Oracle BPEL Process Manager 10g is the industry's first native implementation of BPEL for modelling, deploying, and managing business flows in a standards-compliant way. It comprises an easy-to-use BPEL modeller, a scalable native BPEL engine, an extensible WSDL binding framework, a monitoring console, and a set of built-in integration services.

Oracle Business Activity Monitoring 10g (BAM) enhances Oracle's support for managing business processes by providing real-time visualization of business activities, alerts and notifications based on business-level thresholds, metrics and key performance indicators, and overall monitoring of business events.

Standard Web Services Support

Oracle Application Server 10g provides a cohesive SOA (Service Oriented Architecture) platform consisting of four main components. First, Oracle Containers for J2EE (OC4J) is a comprehensive, J2EE-certified service oriented architecture platform to develop and deploy simple and composite Web Services. Oracle 10g Release 2 includes support for J2EE 1.4, which includes multiple standards for implementing Web Services with Java.

Second, Oracle JDeveloper is the integrated J2EE development environment for modelling, developing, debugging, optimising, and deploying Java applications and Web Services.

Third, the Oracle Application Development Framework (ADF), which is embodied by default in JDeveloper, provides an SOA framework based on the Model-View-Controller design pattern that will dramatically improve developer productivity.

Finally, Oracle TopLink provides an object-relational mapping solution and middle tier persistence layer that simplifies how J2EE applications map to and access relational and XML-oriented data.

Oracle Enterprise Manager 10g enhances Oracle's support for SOA by monitoring and managing Web Services and any other administrator-defined services, tracking end-to-end performance and performing root cause analysis of problems encountered.

Information Grid

Data Provisioning

Information starts with data, which must be provisioned wherever consumers need it. For example, users may be distributed geographically, and fast data access may be essentially important for these users than access to an identical resource. In these cases, data must be shared between systems, either in bulk or real time. Oracle's bulk data movement technologies include *Transportable Tablespaces* and *Data Pump*.

For more fine-grained data sharing, the feature of Oracle Database 10g captures database transaction changes and propagates them, keeping two or more database copies in sync as updates are applied. It also unifies traditionally distinct data sharing mechanisms – message queuing, replication, events, data warehouse loading, notifications, and publish/subscribe – all into a single technology.

Centralized Data Management

Oracle Database 10g manages all types of structured, semi-structured, and unstructured information, representing, maintaining and querying each in its own efficient way while providing common access to all via SQL and XML Query. Along with traditional relational database structures, Oracle natively implements OLAP cubes, standard XML structures, geographic spatial data, and unlimited sized file management, which virtualises information representation. Combining these information types enables connections between disparate types of information to be made as readily as new connections are made with traditional relational data.

Metadata Management

Oracle Warehouse Builder is more than a traditional batch ETL tool for creating warehouses. It enforces rules to achieve data quality, does fuzzy matching to automatically overcome data inconsistency, and uses statistical analysis to infer data profiles. With Oracle Database 10g Release 2, its metadata management capabilities are extended from scheduled data pulls to handle a transaction-time data push from an Oracle Database implementing the Oracle Streams feature.

Oracle's series of enterprise data hubs (e.g. Oracle Customer Data Hub) provides real-time synchronization of operational information sources so that companies can have a single source of truth while retaining separate systems and separate applications, which may include a combination of packaged, legacy, and custom applications. In addition to the data cleansing and scheduling mechanisms, Oracle also provides a well-formed schema, established from years of experience building enterprise applications, for certain common types of information, such as customer, financial, and product information.

Metadata Inference

Joining the Oracle 10g software family in 2005 is the new Oracle Enterprise Search product. Oracle Enterprise Search 10g crawls all information sources in the enterprise, whether public or secure, including email servers, document management servers, file systems, web sites, databases, and applications, then returns information from all of the most relevant sources for a given search query. This crawl and index process uses a series of heuristics specific to each data source to infer metadata about all enterprise information that is used to return the most relevant results to any query.

Benefits of Grid Computing

Compared to other models of computing, IT systems designed and implemented in the grid style deliver a higher quality of service, at a lower cost, with greater flexibility. Higher quality of service results from having no single point of failure, a powerful security infrastructure, and centralized, policy-driven management. Lower costs derive from increasing the utilization of resources and dramatically reducing management and maintenance costs. Rather than dedicating a stack of software and hardware to a specific task, all resources are pooled and allocated on demand, which eliminates under-utilised capacity and redundant capabilities. Grid computing also enables the use of smaller individual hardware components, which reduces the cost of each individual component and providing more flexibility to devote resources in accordance with changing needs.

The motivation for the development of grid computing is to reduce the need to have dedicated resources sized for peak capacity. Research has shown that the average CPU usage is only 15-20% and storage usage is only 50%. Having dedicated resources for each major application also means that there can be a large number of systems to maintain. The solution to this is to create a grid - a pool of low-cost servers + storage that can be allocated to applications to meet peak loads.

Other enhancements for Oracle 10g are aimed at reducing costs and improving the quality of service by making the database easier to deploy and manage including backup and recovery and data warehousing enhancements.

1. Ease of Management

In Oracle 10g, a new Automated Storage Management (ASM) sub-system has been introduced. This feature removes the provision to have a file system or a volume manager for managing the database files and eliminates the complexity of balancing data across disks and controllers

manually. Instead, the ASM stripes data across the disks automatically, thereby maximizing throughput and reducing the risk of data loss. The ASM also maintains the data striping as disks are added/removed and re-balances the I/O load dynamically.

Oracle 10g has also been given a self-management infrastructure which captures information about the running of the database and stores it in the Automatic Workload Repository. This information is used by tools such as the new Automatic Database Diagnostic Monitor (ADDM) which analyses data such as the top SQL statements and passes this onto the SQL Tuning Advisor which provides recommendation for improving performance. These recommendations can be implemented transparently to the application making it a useful feature for 3rd-party applications.

Other major enhancements in this area include:- enhancing Enterprise Manager to be able to manage the whole grid and the complete stack of resources and simplifying installation and configuration by reducing the number of initialisation parameters - for example there are now just two memory size parameters, one each for the SGA and the PGA.

2. Availability

Real Application Clusters have been enhanced to provide Automatic Service Provisioning - servers are automatically allocated to work loads and clients are automatically assigned to the server with the least load. Also on failure of a server, the surviving servers are automatically re-allocated to work loads.

Enhancements in this area also include:

- Flashback available at the row, transaction, table or database level;
- Recovery area on disk that is maintained automatically by the database and contains only those blocks changed since the last backup - thereby enabling faster recovery from media failure;
- data guard (standby database) has been enhanced to enable compression and encryption of log traffic from the master database to the standby system;
- Tables can now be re-defined without invalidating stored procedures;
- Support for rolling upgrades of the hardware, operating system and the database to reduce planned down time.

3. Data Warehousing

Data Warehousing enhancements include:

- An increase in the size limits of the database to support ultra-large databases totalling millions of terabytes with ultra-large files of terabytes in size. The 4GB restriction on LOBs has been raised to 128 terabytes.
- Improvements to Real Application Clusters (RAC) enable resources to be allocated automatically that means that operational data can be used immediately without the need to copy it to another database.
- Enhancements to OLAP analytics, a data-mining GUI and a new SQL model allow query results to be treated as sets of multi-dimensional arrays on which complex inter-dependent operations - such as forecasting - can be run without the need to extract data to spreadsheets or perform complex joins and unions on the data.

Where does grid computing fit:

Generation	Features
First(Host Based Computing)	<ul style="list-style-type: none"> • Dumb Terminal • Single Server • Bunch of Complex applications
Second(Remote Access)	<ul style="list-style-type: none"> • Single Server • Single Client
Third (Client Server)	<ul style="list-style-type: none"> • Multiple Clients(But Limits) • Upto Two servers
Fourth (Multitier)	<ul style="list-style-type: none"> • Multiple Client • More than two servers
Fifth(Grid Computing)	<ul style="list-style-type: none"> • Virtual environment where all systems are treated as single pool of resources • N tier • Service Oriented Architecture

Entering the Fifth Generation. Grid Computing

Grid Computing is the result of several trends coming together. Some of these are as follows:

- High-performance microprocessors have become available, making it possible to deploy large applications on a number of low-cost systems rather than a single mid-range system.
- New standards for object-to-object communications making it easier to build multivendor, multiapplication networks.
- High-speed networking technology is becoming both less costly and readily available, offering higher levels of performance when deploying distributed application architectures.

As these trends combine, applications are likely to be segmented by function or instance of a function. This approach will allow each function to be hosted on the most cost-effective platform. In some cases, both types of segmentation will be used. In the end, an organization's systems can be considered a pool of shared resources that adapt automatically to changing conditions and failures based upon rules of the organization's choosing.

Grid Resources Work Well Independently and Best Together

By managing any single IT resource – infrastructure, applications, or information – using grid computing, regardless of how the other resources are treated, enterprises can realize higher quality, more flexibility, and lower costs. It's also possible to deploy an applications grid, or an SOA, without changing the way information is managed or the way hardware is configured.

It's possible, however, to derive even greater benefit by using grid computing for *all* resources. For example, the applications grid becomes even more valuable when you can set policies regarding resource requirements at the level of individual services and have execution of different services in the same composite application handled differently by the infrastructure – something that can only be done by an application grid in combination with an infrastructure grid. In addition, building an information grid by integrating more information into a single source of truth becomes

tenable only when the infrastructure is configured as a grid, so it can scale beyond the boundary of a single computer.

Oracle's Vision for Grid Computing

Infrastructure resources managed in a grid will progress to the point that computing and storage capacity are delivered on demand like a utility. Applications in a grid will advance so that business and application logic are as massively connected and referenced as static web pages are on the Internet today, enabling frictionless, automated, global business between trading partners. Eventually, a global information grid will impart to every bit of digitally-represented information anywhere the same values we take for granted with relational databases; it will be as if all information resides in a single virtual database. All inherent relationships between information will be revealed, and anyone with appropriate authorization will have instantaneous access to all relevant information regardless of representation, location, or access method.

CONCLUSION

Grid computing is the next generation model for enterprise computing based on the core tenets of virtualisation and provisioning of every resource in IT. Grid computing delivers benefits of increased utilization and greater flexibility for infrastructure, applications, and information resources. Oracle 10g is the family of software products that supports grid computing and the software foundation that is enabling many companies and institutions to make the promise of grid computing a reality.

USA	India		
Corporate Office	Development Center	Regional Offices	
7 Lincoln Highway, Suite 205, Edison, NJ 08820 Phone 732.548.9268 Fax 732.548.8913	Plot No. 38, Electronics city Sector 18, Gurgaon -122 015, India Phone: 91-124- 2397660-62, 5017660 Fax: 91-124- 2397655, 5019955 http://www.binarysemantics.com	20, 2nd Floor, Arihant Complex 1st Cross, CKC Garden (Off: Mission Road) Bangalore - 560 027 Phone: 91-80- 22240222 / 22241222 Fax: 91-80- 22277867	Basera - Plot No.48, 3rd Floor, Santhawadi Lane Opp. Jain Upasarai, Jaiprakash Road, Andheri (W) Mumbai - 400 058 Phone: 91-22- 26705762 /26286748/62 TeleFax: 91-22-26705763