

ORACLE DATABASE 10g – THE DATABASE FOR THE GRID

Brajesh Goyal, Oracle Corporation

EXECUTIVE SUMMARY

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. The basic idea of grid computing is the notion of computing as a utility, analogous to the electric power grid or the telephone network. As a client of the grid, you do not care where your data is or where your computation is done. You want to have your computation done and to have your information delivered to you when you want it. From the server-side, grid is about virtualization and provisioning. You pool all your resources together and provision these resources dynamically based on the needs of your business; thus achieving better resource utilization at the same time.

Trends in the IT industry are moving enterprises towards grid computing. This is largely driven by innovations in hardware, with low-cost blade servers and shared storage technologies. These technologies simplify virtualization and provisioning of hardware. Hardware vendors have also started offering management software that enables virtualization and provisioning of this hardware. Hardware innovations can only be useful when the software running on them can efficiently leverage those innovations. Oracle offers the software – both Oracle Database 10g and Oracle Application Server 10g – that efficiently leverages these innovations.

Admittedly, the industry has not yet converged on a specific, concrete definition of grid and grid standards are in active development. There are three fundamental attributes that a grid must exhibit, and enterprises can begin reaping the benefits of grid computing by enhancing existing IT infrastructure with the fundamental attributes of a grid:

- Virtualization between every layer of the computing stack
- Provisioning of work and resources based on policies and dynamic needs of the enterprise
- Pooling of resources to increase utilization

Oracle Database 10g leverages the grid-enabling hardware innovations and makes it easy for you to install and configure Oracle database on these low-cost standard modular hardware components. Oracle Database 10g virtualizes the way Oracle database uses the hardware components – storage and servers. It automatically provisions clustered storage and servers to different databases running in your grid. As a provider of data, Oracle Database 10g provides technologies with which Oracle database administrators can pool, virtualize, and provision data to grid users and applications. In addition, an enterprise running on the grid imposes stringent operational requirements in terms of security, high availability, self-reliance, and manageability. Oracle Database 10g offers unrivaled operational benefits that simplify the management and operations of your grid.

INTRODUCTION

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

This view of utility computing is, of course, a client side view. From the server side, or behind the scenes, grid

computing 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.

ENTERPRISE GRID COMPUTING

Grid computing standards that enable computing across enterprises are in active development. However, numerous trends in the IT industry are moving enterprises towards grid computing. Every major IT vendor is talking about grid computing and is promising to offer grid technologies for the enterprises. Although, vendors have their own interpretation of grid computing, these interpretations share certain fundamental attributes.

TRENDS PROMOTING GRID COMPUTING

Grid computing is poised to be the next big thing. There are a number of trends in the IT industry that are moving enterprises towards grid computing.

HARDWARE TRENDS

Much of what makes grid computing possible today are the innovations surrounding 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 processors used in high-end SMP servers.
- **Servers:** Every major hardware vendor is delivering blade servers today. Blade server technology provides much higher density computing at a very low-cost as compared to their SMP equivalents. These blade servers also come with remote management capabilities that make it easy for data center administrators to manage these systems. Data centers have started to leverage these technologies.
- **Networked storage:** Network Attached Storage (NAS) and Storage Area Networks (SANs) are popular in today's data centers. Storage is no longer tied to a single server.
- **Network Interconnects:** Gigabit Ethernet and Infiniband interconnect technologies are driving down the cost of connecting clusters of servers.

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.

Vendors such as Sun, HP, IBM, and Topspin are talking about or delivering hardware virtualization and provisioning technologies. These technologies allow you to dynamically group and network a set of servers and storage components. You can also dynamically move servers and storage components from one group to another. Some of these technologies also allow dynamic loading and starting OS and applications on these servers.

SOFTWARE 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 clusters of standard low-cost blades naturally go well with Linux, the standard inexpensive OS.

Hardware innovations can only be useful when the software running on them can leverage those innovations. Software has started to leverage these hardware innovations. One issue with software today is that it is designed to use the resources it is provisioned, but it is not designed to deprovision or give up resources it no longer needs. Oracle

provides software today – both Oracle Database and Oracle Application Server – that leverage these hardware innovations. Oracle Database and Oracle Application Server can utilize the resources they are provisioned and can easily renounce resources they no longer need.

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.

FUNDAMENTAL ATTRIBUTES OF THE GRID

Admittedly, the industry has not yet converged on a specific, concrete definition of grid computing. The standards that will enable grid computing to become as pervasive as the Internet are still in active development. Despite the fact that the full promise of grid computing remains many years away, there are fundamental attributes that a grid must exhibit, and enterprises can begin reaping the benefits of grid computing by enhancing existing IT infrastructure with these attributes of a grid.

VIRTUALIZATION AT EVERY LAYER

Virtualization enables grid components such as storage, processors, databases, application servers, and applications to work together without creating rigidity and brittleness in the system. Rather than making static ties that determine where a database physically locates its data or which exact server the database runs on, 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.

POLICY-BASED PROVISIONING

Provisioning means allocating resources where they are needed. Once the resources are virtualized, resources need to be dynamically allocated for various enterprise tasks based on the changing business priorities. In the context of hardware resources, hardware resources such as storage and servers need to be allocated to databases and application servers. In the context of data, data needs to be allocated to where it is needed.

RESOURCE POOLING

Some definitions of grid computing stop with virtualization and provisioning. Oracle believes, however, that pooling resources is the third fundamental attribute that constitutes a grid. Consolidation and pooling of resources is required for grids to achieve better virtualization and provisioning and hence better utilization of resources. Once you have pooled individual disks into storage arrays and disk farms, it is easier to virtualize and provision storage. NAS and SAN technologies have already shown some benefits of pooling of storage. Similarly pooling individual servers together makes it easier to virtualize and provision these servers to various databases. Pooling data together makes it easy to virtualize and provision data.

ORACLE DATABASE 10g: THE DATABASE FOR THE GRID

Oracle Database 10g is the first database that is designed for grid computing. Oracle Database 10g leverages the grid-enabling hardware innovations and has made it easy for you to install and configure the Oracle database on these standardized modular hardware components. Oracle Database 10g, both as the consumer of grid resources – servers and storage – and as a provider of data, leverages the three fundamental attributes of grid computing. Oracle Database 10g virtualizes the way the Oracle database uses the hardware components – storage and servers. It automatically provisions clustered storage and servers to different databases running in your grid. As a provider of data, Oracle Database 10g provides technologies with which database administrators can pool, virtualize, and provision data to grid users and applications. In addition, an enterprise running on the grid imposes stringent operational requirements in terms of security, high availability, self-reliance, and manageability. Oracle Database 10g offers unrivaled operational

benefits that simplify the management and operations of your grid.

CONFIGURING AND INSTALLING ORACLE DATABASE 10g ON STANDARD-BASED COMPONENTS

Oracle Database 10g makes it easy for you to run your database on a grid running on standard low-cost modular hardware components — storage, blades, and interconnects.

AUTOMATIC STORAGE MANAGEMENT

Automatic storage management (ASM) virtualizes your storage and provides easy provisioning of your database storage. In addition, you can now use standard low-cost modular storage components for storing all your Oracle data. You can use single ASM to manage storage for multiple Oracle databases. Instead of managing many database files, ASM only requires you to manage a small number of disk groups. A disk group is a set of disk devices that ASM manages as a single, logical unit. You 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.

ASM also offers benefits of storage technologies such as mirroring or logical volume managers (LVMs). Similar to these technologies, ASM enables you to create a single disk group from a collection of individual disk devices. It balances I/O to the disk group across all of the devices in the disk group. It also implements striping and mirroring to improve I/O performance and data reliability. ASM automatically re-balances the database's storage whenever the storage configuration changes.

PORTABLE CLUSTERWARE

Clusterware is the software that provides clustering services for communication between servers in a cluster. Prior to Oracle Database 10g, you were dependent on hardware vendors for clusterware on their platforms except for Linux and Windows. Oracle provided clusterware only for Linux and Windows.

Oracle Database 10g now offers portable clusterware; that is, Oracle offers a clusterware that runs on all platforms. You don't need to depend on vendor-offered clusterware. This has many benefits:

- You have a single place of support. You don't need to depend on multiple vendors for your support issues. It makes it easy to diagnose and fix issues.
- You don't have any vendor-imposed limits. With portable clusterware, you can now grow the size of your cluster to any number you like.

Oracle has also made it easy for you install the portable clusterware. With a single install, you identify the nodes where you would like to install the portable clusterware and Oracle Universal Installer will install portable clusterware on all those nodes.

HIGH-SPEED INFINIBAND NETWORK SUPPORT

Oracle Database 10g has enhancements to provide you with better performance and scalability with upcoming high-speed interconnects such as Infiniband. You can use Infiniband across the board for your network communications. It offers many benefits:

- Infiniband offers multitudes of performance improvement over Gigabit Ethernet networks. The low latency and high-bandwidth of Infiniband makes it especially useful as a cluster interconnect.
- You use a single network infrastructure for your communication between different servers and between servers and storage. This simplifies the cabling requirements of your data center.
- With simplified network infrastructure, you use a single network backplane, which makes network provisioning easier.

With Oracle Database 10g, you can now use Infiniband for your application server to database server communication, for server-to-server communication in a clustered database, and for server to storage communication. This provides you with all round performance improvement and flexibility in your data center.

EASY CLIENT INSTALL

The easy client install feature simplifies deployment of applications in a grid. Clients of the database only need to download or copy a very small subset of Oracle client files and set an environment variable. These applications – OCI or JDBC applications – can access a database on your grid. You no longer need to go through the install process on the database client. This feature is especially useful for deployment of ISV applications. ISVs can include these Oracle client files in their install process and customers need not install Oracle client separately. Also in grid environments where client machines are dynamically identified and configured, this feature simplifies the installation and configuration of Oracle client software.

EASY ORACLE DATABASE INSTALL

Oracle Database 10g has simplified the installation of the Oracle database. You can install the Oracle database with a single CD. Oracle Universal Installer (OUI) can also perform multi-node install of the clustered Oracle database. During the install you identify the hostnames where you would like to install the Oracle database. OUI will then install the Oracle Database software on all those nodes. You can also decide to have either a single shared image of the software or a separate image on each host machine.

COMPUTE RESOURCE PROVISIONING

The tenet of grid computing is the ability to dynamically align resources to your changing priorities. Oracle Database 10g has numerous enhancements and new features that make it easy for you to align your computing resources with your business needs.

REAL APPLICATION CLUSTERS (RAC)

Oracle Real Application Clusters (RAC) enables high utilization of a cluster of standard low-cost modular servers such as blades. You can run a single Oracle database on a cluster of blade servers. 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.

Oracle Database 10g offers automatic workload management for services within a RAC database. RAC automatically load balances connections as they are made across instances hosting a service. In addition, using Resource Manager, you can specify policies for resource allocation to services running within a RAC database. To meet these policies, RAC will automatically provision database instances to these services.

Oracle Database 10g also offers a single-button addition and removal of servers to a cluster. With a push of a button, you can now add a server to your cluster and provision this server to the database. Oracle database will automatically install all the required software – Portable Clusterware and Oracle Database software – and start an Oracle database instance on it. Similarly, with the push of a button, you can remove a server.

RESONANCE

The most groundbreaking feature of Oracle Database 10g is resonance. You define your cluster. A cluster is a set of servers that are clustered together using Oracle Portable Clusterware. You run many databases on the same cluster and define your service policies for these databases. Resonance will dynamically grow and shrink the number of servers your individual databases are running on in this cluster to meet your service-level objectives. This is done automatically without requiring any user intervention.

Imagine you have a large cluster with many databases in it. If you had to manage it manually, you would have to constantly monitor the load on each of those databases and then you would have to manually shut down or bring up additional instances of those databases. Your shut down command might take a very long time, as there might be

active sessions on the database instance. Oracle Database 10g does this for you. It constantly monitors the load for you. When it needs to shut down an instance, it automatically migrates active sessions to other active instances of the database. Similarly, when it brings up an additional instance of a database, it automatically balances the workload across all the instances of this database.

ORACLE SCHEDULER

Oracle Database 10g introduces Oracle Scheduler, which provides you with many advanced capabilities to schedule and perform business and IT tasks in your grid. You can provision your workload within a database across time to get more efficient resource utilization. You define your jobs, which can be stored procedures or external jobs such as C or Java programs. You define your schedule. Then you assign jobs to schedules. You have the ability to define any arbitrarily complex schedule. You can also group jobs into job classes to simplify management and prioritization of jobs. Using Oracle Resource Manager, you can define your resource plans and assign these resource plans to these job classes. You can also change these resource plans across time. For example, you may consider the jobs to load a data warehouse to be critical jobs during non-peak hours but not during peak hours.

DATABASE RESOURCE MANAGER

Resource Manager provisions resources to database users, applications, or services within an Oracle database. It allows database administrators to limit the Oracle database resources allocated to grid users, applications, or services. This ensures each grid user, application, or service gets a fair share of the available computing resources. Administrators define resource plans that allocate resources to various consumer groups based on resource usage criteria such as CPU utilization or number of active sessions. Each consumer group comprises one or more database users. Oracle Database 10g provides additional mappings for consumer groups based on user host machine, application, OS username or service.

DATA PROVISIONING

Data provisioning means getting information delivered to users when they need it, regardless of where it resides on a grid. To process information on any available resource, a grid must efficiently make information available across distributed systems. There are essentially three ways to do this: consolidate, share, and federate information.

Consolidating information in a single database makes it easy to provision information. You use a single pool of resources where you run your database. This allows for more efficient resource utilization. Since the database resources are pooled together, it does not happen that one database is melting down while the other one is sitting idle. It provides easy provisioning. You can now use Oracle Resource Manager to efficiently provision Oracle database resources to various users and applications that are running on this database.

If you cannot pool your resources and cannot consolidate your information in a single database, you can share information to get more efficient resource utilization. Oracle Database 10g features such as Oracle Transportable Tablespaces and Oracle Streams makes it easy for you to efficiently share data between different databases to share processing of information across different resources.

For certain information, for example a terabyte database, that is infrequently accessed, it is more efficient to leave it in place. For this information, you can use Oracle Database 10g's federated features such as distributed SQL, gateways, and materialized views to get access to that data on demand.

ULTRA LARGE DATABASE SUPPORT

Oracle Database 10g now supports a single database with up to 8 exabytes (8 million terabytes) of data. This virtually removes a limit on how large your consolidated database can be. You can also store data in much larger files, thus decreasing the number of files in large databases. Additionally, the Bigfile Tablespace simplifies the management of datafiles in large databases, minimizes scalability issues related to having a large number of datafiles, and simplifies management of storage utilizing such features as Automatic Storage Management and Oracle Managed Files.

ORACLE TRANSPORTABLE TABLESPACES

It is not always possible to pool your hardware resources together. For example, you might have geographically distributed hardware resources that you cannot cluster efficiently. Or your data center limitations prevent you from pooling the hardware together. In these situations, Oracle Transportable Tablespaces offers a very efficient way of

sharing large subsets of data and then sharing processing on this data on different hardware resources.

Oracle Transportable Tablespaces offer grid users an extremely fast mechanism to move a subset of data from one Oracle database to another. It allows a set of tablespaces 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.

Oracle Database 10g now supports heterogeneous transportable tablespaces. This feature allows tablespaces to be unplugged, converted with RMAN if need be, and transported across different platforms for example, from Solaris or HP/UX to Linux. Many users are using this feature to migrate their databases to Linux.

As an example for a usage of the transportable tablespaces, if you look at a financial application at a typical enterprise, it is very light during normal time. It has a couple of inserts or updates every hour. But during quarter end, it needs lots of resources for reporting. What you can do is, during the normal time, you run this application on an appropriately powerful resource. At the quarter end, you use the transportable tablespace feature to move the data to a more powerful resource and perform your processing there.

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, and can keep two or more copies in sync as updates are applied. It also provides a unified framework for information sharing, combining message queuing, replication, events, data warehouse loading, notifications and publish/subscribe into a single technology.

If you take the example above, once you are done with reporting, either you can throw away the database or use Oracle Streams to keep the two databases in sync so that you can do next quarter reporting again using the second database. For this purpose, you can define a schedule that fits your needs for syncing. As an example, you can sync every midnight or sync as the changes are happening.

Oracle Database 10g has many enhancements for grid users. You can now offload processing from the primary database to a second database; that is, in place of capturing changes as they are happening on the primary database, you can set up a second database to capture those changes. This provides another way of sharing workload across your resources.

SELF-PROPELLED DATABASE

Oracle Database 10g offers a new self-propelled database feature. This feature utilizes Oracle Transportable Tablespaces and Oracle Streams and offers you an easy way to share processing across distributed hardware resources. In addition, it offers an efficient way for migrating your applications to the grid.

With a single command, you can snap a set of tablespaces from one database, ship the tablespace to another database, reformat the tablespace if the second database is on a different OS, and plug this tablespace into the second database. During this time, there might be changes happening at the first database. Oracle Streams would have already started capturing those changes and these changes are synced with the second database. All of this is done with a single command. If the second database is on a grid, what you have just done is migrate your application to a grid with a single command. You can later migrate all the applications running on the first database to the second database by simply repointing the connection string to the second database.

ORACLE DATA PUMP

With Oracle Database 10g, Oracle introduces a new data movement facility that greatly improves performance when getting data into and out of a database. Oracle Data Pump is a high-speed, parallel infrastructure that enables quick movement of data and metadata from one database to another. This technology is the basis for Oracle's new data movement utilities, Data Pump Export and Data Pump Import, which have greatly enhanced performance over Oracle's original Export and Import.

When moving a transportable tablespace, Data Pump Export and Import are used to handle the extraction and recreation of the metadata for that tablespace. Data Pump is a flexible and fast way to unload entire databases or subsets of databases, and reload them on the target platform.

DISTRIBUTED SQL, GATEWAYS, AND DISTRIBUTED TRANSACTIONS

It is not always possible to either consolidate or share information. The limitations of the data centers or geographical distribution of resources may prevent you from doing so. It may also be because of security issues. You may not want the entire data set to be visible to the users on the second database. Or it may not be efficient to move the data – for example, you might have a one-terabyte data set that is infrequently accessed. Oracle Database 10g offers an extremely strong federated technology that helps you in those situations. With these technologies, you can leave data in place and access this data on demand.

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

In addition, Oracle Database 10g also offers external tables and Bfiles features, that let you keep your data on the file system while providing access to grid users via Oracle database APIs. External tables provide you with SQL access to structured data in the files. Bfiles provides read-only access to the unstructured data in the files.

OPERATIONAL BENEFITS

The dynamic nature of the grid imposes stringent operational requirements on your grid infrastructure. You would want the grid infrastructure to be self-reliant – it should be able to tolerate system failures and adapt to changing business needs. You would like to be able to control your grid environment in a more holistic manner rather than on a component-by-component basis. Lastly security is paramount in grid environments – you would not like any unwanted exposure of your data and resources.

SELF-RELIANT DATABASE

A truly responsive enterprise requires grid infrastructure to self-manage, and to learn and adapt to changing circumstances. Your grid infrastructure needs to be able to tolerate the failures of individual components and provide high availability in all circumstances.

HIGH AVAILABILITY

Oracle Database 10g brings the highest levels of reliability and availability to the grid. You get the same levels of reliability and availability on standard, low-cost modular hardware – servers and storage. Acknowledging that failures will happen, Oracle Database 10g provides near instantaneous recovery from system faults, meeting the most stringent service level agreements. ASM provides the reliability and availability on low-cost standard storage. RAC provides the reliability and availability on low-cost standard servers.

Oracle Database 10g provides robust features to protect from data errors and disasters. The new flashback database feature provides the ability to recover a database to a specific time to recover from human error. The recovery time is proportional to the time duration it needs to go back to. With this flash backup feature, database administrators can now use low-cost standard disks for maintaining their backups.

Oracle Database 10g also includes tools to minimize planned downtime, critical for any interactions in a 24x7 environment. The new rolling upgrade feature enables online application of patches to the database software. You don't need to bring down the entire database to apply a patch. You can apply patches to the clustered database – one instance at a time--thus keeping the database online while applying the patch.

SELF-MANAGING

With the new self-managing features, Oracle Database 10g has taken a giant leap toward making the Oracle database self-reliant. Oracle Database 10g includes an intelligent database monitor that records data regarding all aspects of database performance. Using this information, Oracle Database Automatic Memory Management dynamically allocates memory to different components of the Oracle database. Automatic health management automatically generates alerts regarding various aspects of the database that simplify database monitoring for DBAs. Automatic storage management, covered previously in this paper, provides automatic rebalancing and provisioning of storage.

ORACLE ENTERPRISE MANAGER GRID CONTROL

Even with the self-managing Oracle Database 10g, there are aspects of the enterprise grid that administrators will want to manage and control. Oracle Enterprise Manager Grid Control provides a single tool that can monitor and manage not only every Oracle software element – Oracle Application Server 10g and Oracle Database 10g – in your grid but also web applications via APM (Application Performance Management), hosts, storage devices, and server load balancers. It is also extensible via an SDK so customers can use it to monitor additional components that are not supported out-of-box. Grid Control views the availability and performance of the grid infrastructure as a unified whole rather than as isolated storage units, databases, and application servers. IT staff can group hardware nodes, databases, and application servers into single logical entities and manage a group of targets as one unit.

Grid control provides a simplified, centralized management framework for managing enterprise resources and analyzing a grid's performance. With Grid Control, grid administrators can manage the complete grid environment through a Web browser throughout the whole system's software lifecycle, front-to-back, from any location on the network.

With Grid control, administrators can launch and execute any number of integrated Oracle Database features such as Data Pump, Resource Manager, Scheduler, Transportable Tablespaces, etc. Administrators can also monitor, diagnose, modify, and tune multiple databases throughout the grid. Administrators can also schedule tasks on multiple systems at varying time intervals, share tasks with other administrators, and group related services together to facilitate administration.

MANAGING SECURITY IN THE GRID

The dynamic nature of the grid makes security extremely important. Enterprises need to ensure that their data is secure. Exactly the right set of users must have access to the right set of data. At the same time, they need an easy way to manage security throughout their enterprise. Oracle Database 10g makes it easy for enterprises to manage their security needs in the grid.

ENTERPRISE USER SECURITY

Enterprise User Security centralizes the management of user credentials and privileges in a directory. This avoids the need to create the same user in multiple databases across a grid. A directory-based user can authenticate and access all the databases that are within an enterprise domain based on the credentials and privileges specified in the directory.

With Oracle Database 10g, grid users have the ability to store an SSL Certificate in a smart card, for roaming access to the grid. Oracle Database 10g also comes with Oracle Certificate Authority that simplifies provisioning of certificates to grid users.

VIRTUAL PRIVATE DATABASE (VPD)

VPD provides server-enforced, fine-grained access control, and a secure application context that can be used within a grid setting to enable multiple customers, partners, or departments utilizing the same database to have secure access to mission-critical data. VPD enables per-user and per-customer data access within a single database, with the assurance of physical data separation. VPD is enabled by associating one or more security policies with tables and views and now, in Oracle Database 10g, also with table columns.

ORACLE LABEL SECURITY

Oracle Label Security gives administrators an out-of-the-box row—and now column-level—security solution for controlling access to data based on its sensitivity, eliminating the need to manually write such policies. Using the GUI tool Oracle Policy Manager, administrators have the ability to quickly create and assign Oracle Label Security policies to rows and columns within application tables. Moreover, Oracle Database 10g adds integration of Oracle Label Security with Oracle Internet Directory (OID), allowing policies to be managed centrally within a dynamically changing grid setting.

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.

PRODUCT DIRECTIONS ALIGNED WITH GRID

Oracle product directions are aligned with the grid. Oracle Database 10g is the first database designed for the grid. Oracle already supports more grid computing technology than any of its competitors, as described in the previous sections. 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 standards, such as J2EE, Web Services, Xquery, and SQL, Oracle intends to fully support grid standards. Often, the work of various standards bodies overlap or various standards bodies develop complementary technologies. By investing in Oracle Database 10g, you can leverage complementary standards and interoperate with various standards and technologies.

CONCLUSION

Grid computing is poised to change the economics of computing. Even though, there is no single definition of the grid, there are fundamental attributes of grid computing that enterprises can utilize to reap benefits of the grid today. New innovations and new economics in hardware, now, make grid computing possible and sensible at the hardware layer. Only Oracle Database 10g leverages these hardware innovations and implements the fundamental attributes of grid computing. Only Oracle Database 10g with its strong security, self-reliance, and manageability offerings addresses the stringent operational needs of enterprise grids. With Oracle Database 10g, you can realize grid benefits today and leverage your investments in Oracle for future grid computing technologies.