

WHITE PAPER

Grid Computing with Oracle Database 11g

Sponsored by: Oracle Corporation

Carl W. Olofson

March 2008

IDC OPINION

The IT world is in the midst of a sea change in how computing resources are managed. The previously normative models based on program code assigned to fixed assets are being replaced by models based on flexible code deployment and resource assignment. A new model has emerged based on a shared infrastructure, where the underlying physical hardware is abstracted from the application and logical resources built by aggregating or disaggregating the physical resources that are used for application deployment. This shared infrastructure, also known as grid computing, provides many benefits to today's IT operations. The benefits of grid computing in general, and Oracle's exploitation of it in Oracle Database 11g, are:

- Ending the paralysis that results from having computer and storage systems locked down on specific applications, databases, and files
- Enabling operations and administrative staff to flexibly assign compute and storage resources as needed, enabling them to avoid locked-down and overprovisioned resources
- Providing a means for easily assigning and tuning resources to meet the business' service levels
- Reducing cost and enhancing business agility by more effective utilization of resources using commodity hardware

IN THIS WHITE PAPER

This paper discusses the evolution of the datacenter toward the adoption of grid computing and details the benefits derived from this development. It considers the implications for database systems and details the requirements of a database for successful deployment in a grid computing context. It then reviews the features of Oracle Database 11g and shows how that product meets those requirements.

SITUATION OVERVIEW

Most datacenters today feature a range of computer and storage systems, all dedicated to specific workloads or combinations of workloads, usually carefully segregated by application and database. Typically, the computer and storage systems are even named after the database or application they serve. Switching them around to achieve greater efficiency is simply not done; instead, they tend to be overprovisioned to ensure sufficient capacity for peak periods and for growth within budgeted periods.

If storage utilization grows faster than the expected outlay, a complicated and expensive project to reallocate and expand storage results. If computing power becomes insufficient to handle the load, a "forklift upgrade," consisting of replacing the computer systems in question, is usually required. These actions are generally accompanied by substantial staff efforts to load and configure the new systems and to convert or migrate the existing workload to the new systems.

One way to reduce the frequency with which this occurs is to allocate much more system capacity to given applications or databases, allowing them to "grow into" the underutilized capacity over some manageable period of time. In the interim, of course, a good deal of system and storage capacity goes unused.

What Is the Grid?

The grid represents an alternative to this expensive model of fixed IT assets. It derives its name from its resemblance to a power grid, which enables power producers (such as generating stations) to be added or removed as needed and likewise power consumers to be added and removed, without complicated configuration changes. An IT grid enables server and storage resources to be added to or removed from the system without requiring complicated configuration changes, thereby enabling flexible deployment that avoids the problems inherent in the fixed topology model.

Definition of Grid Computing

Grid computing is a term that has been applied to various architectures designed to deliver the benefits of an IT grid. It is an approach to computing that detaches the software functionality from the specifics of hardware deployment by blending system and storage resources into a continuum of resources that can be allocated to, and deallocated from, a particular function or functional locus; in this case, a database. Simply put, it enables administrators to assign computing tasks to computing resources, and it assigns data to storage resources in a way that enables such resources to be easily added or removed or tasks and data to be moved as needed. Various vendors have taken different approaches to delivering grid computing; these approaches tend to vary based on the type of workload involved and the requirements of that workload. They include infrastructure-based distributed components for high-throughput compute grids, redundant application server deployment for application grids, and various approaches to providing scalable yet flexible storage management through storage grids.

An IT grid enables server and storage resources to be added to or removed from the system without requiring complicated configuration changes, thereby enabling flexible deployment that avoids the problems inherent in the fixed topology model.

In the case of database workloads, grid computing contrasts with the classic model that involves dedicated servers associated with dedicated storage in that the servers and storage are fluid; that is, they can be assigned, added, and reassigned as necessary without upsetting the overall topology of the database server environment.

Common Benefits

The key benefits of grid computing come in the form of resource flexibility, scalability, and optimization of operations through parallel processing. These benefits are expressed through an architecture that gives users the following capabilities:

- ☒ To **avoid** unnecessary hardware, power, and staffing costs of **overprovisioning** IT systems, commonly done to avoid capacity upgrades.
- ☒ When capacity upgrades are necessary, to **scale incrementally** by adding (or in some cases, redeploying) system and storage resources without expensive "forklift upgrades" or time-consuming and error-prone upgrade procedures.
- ☒ To **ensure continuous availability** through the provisioning of redundant resources, ensuring automatic failover when necessary.
- ☒ To **increase transaction throughput** through parallelization of tasks.

All these benefits combine to enable better business agility in responding to changes in load or business priorities.

Implementations of Grid Computing

Grid computing makes it possible to add and remove resources in relation to an operational domain such as an application, database, or storage environment as easily as producers and consumers of energy can be added to and removed from an electricity grid. The earliest forms of this approach were used to network large numbers of small computers together to do the job of large, expensive supercomputers at a fraction of the price. Part of the beauty of this approach was that computers could be added or removed without impacting the high-powered computations under way.

Later, this principle was applied to various other problem spaces, including more in the enterprise and/or commercial computing market segments. Examples of different types of grid computing actively used today include the following:

- ☒ **Compute grids.** These grids are used to achieve processing scalability on low-cost hardware by deploying compute tasks in parallel in a manner similar to the techniques used by supercomputers.
- ☒ **Application grids.** These grids are used to achieve better throughput and availability of applications by building them as components that are deployed redundantly on a number of application servers; adding power is done by adding systems. If a server fails, its work is taken up by another.

- ☒ **Storage grids.** Storage is made available on a range of disks, with data managed in an optimized way to ensure scalability and high availability while ensuring that disks may be added or removed in a manner that is transparent to the user and has no definitional effect on the application or database (that is, the application or database does not need to be changed to refer to different files on different volumes when disks are added or removed); the application or user is insulated from the exact nature of the physical storage layout.

The database grid, representing the approach taken with Oracle Database, is a hybrid that contains elements of the application and storage grid. Like an application grid, it deploys code redundantly on multiple servers (or nodes), which break up the workload based on an optimized scheme and execute tasks in parallel against common data resources. If any node fails, its work is taken up by a surviving node to ensure high availability. The storage for the database is managed in a manner consistent with a storage grid.

Grid Computing with Oracle Database 11g

Oracle Real Application Clusters as an Example of Grid Computing

As mentioned earlier, Oracle has employed a database grid to deliver database services that offer the benefits of grid computing. Oracle calls the feature that delivers the grid-based server cluster capability Real Application Clusters (RAC) and the feature that offers storage grid capability Automatic Storage Management (ASM). Together, these features provide:

- ☒ Scalability, including the ability to build clusters of database servers with low-cost hardware that do the kind of work once reserved for expensive symmetrical multiprocessing (SMP) systems
- ☒ High availability through automatic failover so that if one node in the cluster fails, then the other nodes can take up the work
- ☒ Flexibility, since resources, including both systems and storage, can be assigned to the database as needed, and removed if necessary, without reconfiguring the database on either the server or the storage so that there is no need for excess capacity and unused processor power or storage

In addition, manageability features have been built into both RAC and ASM to ensure that users can deploy and manage a database grid without performing complex and technically sophisticated system, storage, or network configuration tasks. This last capability is very significant because without it, deploying a database grid would be beyond the technical expertise of most IT shops. Without the built-in self-managing capabilities, IT shops would need to have advanced IT skill sets to tune and optimize this scale-out deployment of clustered servers running the database workload. These built-in manageability features for RAC represent several years of evolution of this technology at Oracle that was produced as a consequence of Oracle's long experience in supporting database on clustered server systems, having fielded database grid technology since 2000. These self-management features are augmented by Oracle Enterprise Manager's Grid Control, a component that consists of tools designed to make provisioning, administering, operating, and monitoring applications and systems with the grid simple and easy to do.

Oracle calls the feature that delivers the grid-based server cluster capability Real Application Clusters (RAC) and the feature that offers storage grid capability Automatic Storage Management (ASM).

Without the built-in self-managing capabilities, IT shops would need to have advanced IT skill sets to tune and optimize this scale-out deployment of clustered servers running the database workload.

Many IT organizations must dedicate servers and storage to specific applications and their databases, and the servers and storage must have the capacity to manage the application and database they serve as application and data demands and size increase. Typically, this means allocating much more in terms of server and storage resources to the application and database than they actually need. Thus we can see that in the traditional silo scenario, systems cannot be easily redeployed and management of the systems and storage is discrete and inherently complex, with the complexity growing as the number and size of applications and their databases increase.

Thus we can see that in the traditional silo, scenario systems cannot be easily redeployed and management of the systems and storage is discrete and inherently complex, with the complexity growing as the number and size of applications and their databases increase.

By contrast, by using the grid approach, application and database management may be concentrated in clusters of commodity servers, clusters can be expanded with additional servers as needed, and management of applications and databases is shared across the cluster. Similarly, storage units can be added to the storage grid as needed, and management of storage is spread across the units in the storage grid. As a result, overprovisioning of servers and storage is not needed, and staff time is optimized.

How Oracle Database 11g Delivers the Key Benefits of Grid Computing

A number of key benefits were enumerated earlier with respect to grid computing. Any vendor seeking to present a product or product set as offering grid computing must deliver these benefits. Oracle Database 11g delivers these benefits as follows:

- ☒ Oracle Database 11g delivers the ability to move system and storage resources around flexibly as IT needs change in two ways. To add system capacity, the user can use Enterprise Manager Grid Control to clone a server and automatically add it to a cluster. Removal can be done at any time; the cluster will automatically accommodate the change. Users can add or remove storage resources by using facilities of ASM that provide for the automatic formatting of disks and redistribution of data onto or off of storage, with appropriate restriping and mirroring adjustments as needed.
- ☒ Because of the preceding benefit, users can scale up by adding resources without expensive, time-consuming, and error-prone upgrade procedures and can avoid the unnecessary cost of overprovisioning IT systems by removing unnecessary resources even more easily.
- ☒ RAC and ASM both optimize process dispatch across systems and storage access across disks to maximize opportunities to achieve more processing and storage access efficiency through parallelization of tasks.

Growing Grid Momentum in the Oracle Customer Base

Although grid technology is becoming a salient feature of many various infrastructure software products, Oracle has been offering its version of grid technology since 2000, and product adoption has been steady since the middle of 2001. According to Oracle, the past few years have witnessed continued growing momentum for grid computing among Oracle's customers. Oracle reports that grids are being deployed all over the world, in every industry vertical and across all sizes of organizations on Linux, Windows, and Unix platforms. The company further notes that these grids are supporting both large-scale, purpose-built applications and the spectrum of horizontal applications including Oracle eBusiness Suite, PeopleSoft, Siebel, SAP, and many other "off-the-shelf" applications. In addition, Oracle reports that grids are supporting a wide range of vertical applications from stock exchange and billing systems to business intelligence and data warehousing applications to spatial and Web 2.0 applications.

Oracle reports that grids are being deployed all over the world, in every industry vertical and across all sizes of organizations on Linux, Windows, and Unix platforms.

How Customers Have Realized the Benefits of the Oracle Grid

- ☒ Starwood uses an Oracle Grid for its mission-critical reservations and preferred guest program systems. It runs Oracle RAC on a cluster of three HP nodes running HP-UX on an Itanium chip. Oracle's Grid architecture provides it with the scalability and availability critical to its business. Because Starwood supports properties all over the world, it is critical that its systems be online 24 x 7. Starwood is eager to adopt Oracle's latest database technology and has beta tested Oracle Database 11g since June 2006. Its Senior Director of Database Engineering, Arup Nanda, offered special praise for many new features in Oracle Database 11g, including the new partitioning features and table space encryption.
- ☒ ESPN uses Oracle Database to manage its video assets. Although most of the actual video is stored outside the database, the configuration of clips is constantly changing and new video is constantly added. System availability and performance are critical because ESPN's video programmers need constant access to video clips 24 hours a day. ESPN relies on an Oracle RAC solution to provide the performance and availability they require. Their production systems are IBM pSeries running AIX with Oracle RAC on clusters of two nodes. They have been very pleased with the experience they've had beta testing Oracle Database 11g and plan to upgrade to this new release soon. ESPN anticipates using the new release's connection pooling features to greatly enhance data throughput on the network and freeing up system RAM.
- ☒ Mike Prince, CTO at Burlington Coat Factory, reported that by consolidating dozens of application databases to two 18 node clusters hosting multiple RAC databases, the retail company has been able to manage its costs much better by avoiding overprovisioning. Since that move, the company's IT staff spends less time on system maintenance activities, which has freed IT staff to ratchet up service levels while eliminating time-consuming onerous IT task backlogs. ASM, with its ease of storage management, and improvements in Oracle Enterprise Manager functionality improve IT task turnaround and system performance. Looking forward, Prince says that implementing Oracle Database 11g's data compression feature will reduce storage costs and amplify performance even further.

- ☒ Stephen Taylor, director of IT infrastructure for Chesapeake Energy, said that the company has grown very rapidly over the past few years and that Oracle RAC has enabled IT to keep up with that growth. It has deployed a three-node cluster running Dell four-way 64-bit systems against 986GB in databases stored on Network Appliance storage, connected by fiber and managed by ASM. The company deploys a number of in-house and packaged applications and requires continuous availability for certain applications that track business-critical statistics on a continuous basis. The self-managing nature of Oracle Database, combined with the easy-to-use aspects of Oracle Enterprise Manager, has freed Taylor's DBAs from routine maintenance tasks so they can concentrate on more high-value development activities.

FUTURE OUTLOOK

Oracle is not the only vendor working on grid technology. Various system vendors have sought to optimize both hardware and operating environments using grid techniques. A number of other database management systems vendors have also deployed some level of grid technology, although Oracle may be said to have a more evolved product and strategy in this area, due to its early development and adoption of grid technology. Grid technologies are also being deployed in data and application integration, Web service message management, and content management.

Grid technology is a rare example of a fairly new IT development that makes both engineering sense and economic sense, and Oracle is at the forefront of this movement. It seems clear that this approach will come to dominate most areas of what used to be a rigidly configured datacenter.

CHALLENGES/OPPORTUNITIES

Oracle Database, with its RAC and ASM features, represents a comprehensive approach to the use of grid computing to provide flexible, scalable, and reliable database environments. Yet this is a new technology area; user reluctance to fully embrace this approach, combined with problems posed by the business implications of the grid approach for system and storage vendors, could inhibit adoption to some extent in the near term. Overall, however, the logic of this approach seems unavoidable, so adoption is likely to continue and even accelerate.

CONCLUSION

IT environments are more complex than ever before and becoming more so by the day. If the deployment of physical systems and storage matches the complexity of the applications and databases they support, the datacenter is likely to become a horrible tangle of racks, cables, drives, and confusion (many already are). Grid technology cannot necessarily make applications or databases less numerous or complex, but it can enable the deployment of multiple applications and databases in consolidated environments for more efficient execution and resource utilization and make the systems and storage used to run them easier to track and manage and more affordable.

Grid technology cannot necessarily make applications or databases less numerous or complex, but it can enable the deployment of multiple applications and databases in consolidated environments for more efficient execution and resource utilization and make the systems and storage used to run them easier to track and manage and more affordable.

The Oracle approach to grid computing in general, and its database grid in particular, as expressed in Oracle Database 11g with RAC and ASM is designed to deliver on the promises of grid computing:

- ☒ Easier, more flexible and efficient management of systems and storage
- ☒ Better performance through parallel processing of database operations
- ☒ More reliability through the failover features of servers and the managed redundancy of database storage
- ☒ The ability to provide high-throughput processing capability on low-cost systems and storage

Companies that are seeking a DBMS that delivers these capabilities ought to include Oracle Database 11g with RAC and ASM in their product evaluations.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2008 IDC. Reproduction without written permission is completely forbidden.