

# RUNNING ON THE

At its Austin Data Center, Oracle and key partners have brought an enterprise computing grid to life.

Approaching Oracle's Austin Data Center (ADC) through the open hills and plains of central Texas, you might initially mistake the nondescript stucco-clad facility for a traditional office building or warehouse. As you get closer and pass through the perimeter fencing, however, you realize that this is no ordinary corporate structure. Embassy-grade gates monitor vehicles entering and leaving the facility. Armed guards patrol the premises, and dozens of cameras constantly record activity. Biometric hand and iris scanners restrict entry to authorized personnel, who must pass one by one through metal detectors and weight-sensitive traps.

Once inside, you gaze through protective security glass at an acre of raised floor space upon which glossy rows of computers, storage devices, and networking gear have been carefully configured. You've reached the heart of the Oracle On Demand Grid—an advanced computing infrastructure consist-

Below: Mitchell McGovern, Vice President of Oracle Global Data Center Operations (left), and Chris Pohto, Senior Director of Austin Data Center (right), in the command center of Oracle's Austin Data Center. Additional photos: A day in the life of Oracle's Austin Data Center, as staff regulate system performance

FRANK CURRY



# GRID

ing of more than 9,500 servers and 2.5 petabytes of storage capacity. The ADC has two primary functions: Hundreds of Oracle customers house their applications and data here as part of the Oracle On Demand program, and the ADC also hosts many of Oracle's internal production applications, such as global single-instance ERP, new development activities, global education services, and demo software. There are multiple grids within the ADC that serve these different groups.

"The Austin Data Center is what a power station of the future will look like," explains Benny Souder, vice president of distributed database development for Oracle. "Not only is it massive but it's also incredibly precise. Every computing resource is cabled, labeled, stacked, configured, and deployed in exactly the same way, over and over again."

This standardization allows systems administrators to quickly isolate and troubleshoot problems, eliminating downtime so that Oracle can provide its customers with the most-reliable computing services available. Because of primary, backup, and tertiary systems within its electrical, mechanical, and power infrastructure, the ADC is fully redundant.

"One of our requirements is one-hour time to recovery (TTR), 90 percent of the time," explains Chris Pohto, senior

## ADC Stats

- More than 9,500 operational servers
- 355 servers on average installed every month
- Largest Dell/Linux installation on earth
- 2.5 petabytes (2,500 terabytes) of disk storage, the largest NetApp installation on earth
- Certified by the EPA as a Green Power Partner, with 30 percent of the energy used coming from renewable sources such as wind and solar
- 1 acre of raised floor space, soon to be 2 acres
- Hundreds of mission-critical customers' apps hosted on the premises



director at the center. "Utilizing the grid, we always have excess capacity provisioned, with at least 30 systems and several terabytes' worth of storage just for this purpose. If a server fails, our first task is service restoration—fixing the box is secondary. So we simply move the affected instance to a free server, and the corresponding systems are back up in 15 or 20 minutes."

Continued uptime is promoted in other ways as well. Oracle is a U.S. Environmental Protection Agency Green Power Partner and Climate Leader, which means the company not only uses renewable energy sources such as a large wind farm in West Texas but also works to reduce greenhouse gasses. With its intensive focus on energy usage and mechanical efficiency, the ADC has not experienced any electrical or mechanical infrastructure outages since it opened, in 2002.

The ADC is rated by an independent engineering firm as equivalent to a Tier IV site, which means it has multiple active power and cooling distribution paths, redundant components, fault-tolerant computing devices, and 99.995 percent availability. According to the Computer Uptime Institute, this level of availability exceeds industry best practices. It is achieved by a focus on continual process improvement through the application of methodologies such as Six Sigma and capability maturity

## Enterprise Grid Alliance

In April 2004, Oracle and several other vendors formed the Enterprise Grid Alliance (EGA), to identify obstacles to the adoption of grid technologies and promote open, interoperable solutions to these problems. The alliance now numbers more than 30 organizations, including founding members EMC, Fujitsu-Siemens Computers, HP, Intel, NEC, Network Appliance, Oracle, and Sun Microsystems.

"The goal of the EGA is to help enterprises understand how they can build and manage a computing grid in their own data centers today," explains Donald Deutsch, president, Enterprise Grid Alliance and vice president, standards strategy and architecture, Oracle Corporation. "To that end, the alliance is developing certification procedures, compliance programs, and new specifications that companies will leverage as they move to a grid infrastructure."

Because the EGA focuses exclusively on the needs of enterprise users, it will enable businesses to realize the many benefits of grid computing, according to Bernd Kosch, vice president of strategic alliances at Fujitsu Siemens Computers and an EGA board member. "We're putting together the technical prerequisites for implementing on-demand computing models by combining virtualization and automation technologies, and we're applying this knowledge not just to the physical infrastructure but to the application layer as well."

EGA vendors also focus on IT efficiency in the context of grid computing. For example, at most companies today, server resources are used at only 20 to 40 percent of capacity. That's because systems are dimensioned to handle peak loads; dedicated servers are usually installed for each application; and redundant systems are configured for testing, development, and backup purposes. As a result, systems keep getting bigger and system management keeps getting more complicated. But key technologies for grid computing of interest to EGA vendors, such as virtualization, dynamic provisioning, and automation of administration and management functions, will combat complexity.



Benny Souder, Vice President of Distributed Database Development, Oracle

model integration. In the unlikely event of a primary site failure, a backup site can pick up the load in under two hours, with all data and applications synchronized.

"We're proving that you can run fundamental business operations on an outsourced basis," says Mitchell McGovern, Oracle's vice president of global data center operations. "The grid is designed to simplify allocating capacity on demand. It allows us to run our customer environments very securely, very reliably, and with very high availability—often for less than what these customers would pay to run these environments themselves."

### GET A PIECE OF THE GRID

Pohto's team has responsibility for Oracle On Demand operations, including provisioning new customers and expanding their environments. When a customer initiates the Oracle On Demand service, the team provisions a "slice" of the grid for that customer's needs. As the implementation progresses and the applications are moved into production mode, system administrators establish grid control rules that allow the environment to dynamically respond to the need for additional processors or storage.

"Grid control is all about managing a large number of systems as if they were a single system," explains Pohto. "When you have a whole lot of computers, the ability to manage on an exception basis—to look across all of them and immediately see if there's a problem—is incredibly powerful."

Just as the local power company doesn't allocate electrical capacity for an individual house or business but, rather, creates distribution plans for a large set of consumers, Oracle sizes the grid infrastructure to service the aggregate demand. "When you walk into a room and throw the light switch, you don't think about who manufactured the switch or where the power came from," says McGovern. "All you care about is one thing—that the light goes on. That's the kind of turnkey capacity and reliability we supply at the ADC."

### BRAINS OF THE GRID

Grid computing is built around the concept of virtual software environments. In its entirety, the grid appears as a single

computing resource. But look a little closer, and you will find a pool of shared resources that is both self-correcting and self-managing. With its Grid Control features, Oracle Enterprise Manager 10g enables IT professionals to group multiple hardware nodes, databases, application servers, and other targets into single logical entities—making it easy to support distinct segments of the grid for individual customers. “By executing jobs, enforcing standard policies, monitoring performance, and automating tasks across a group of targets instead of on many systems individually, Oracle Grid Control enables us to scale smoothly,” says Souder. “Thanks to these advanced management features, the existence of many small computers in a grid infrastructure does not increase complexity.”

Additionally, Oracle Application Server 10g and Oracle Database 10g are designed to run very well in clusters. Oracle Application Server 10g takes advantage of grid technology features such as policy-based resource management, metric-based workload management, and centralized user provisioning, dramatically simplifying maintenance in multinode environments. Oracle Database 10g offers data provisioning capabilities such as Oracle Streams and transportable tablespaces, which allow DBAs to detach part of a database and attach it to another database without unloading and reloading it. And Oracle Real Application Clusters (RAC) technology enables a single database to be partitioned across multiple clustered nodes, pooling the processing resources of several standard machines.

“Oracle 10g software can be scaled simply by addition of more servers to the RAC configuration,” Pohto notes. “The whole data center possesses the same kind of architecture. Eventually there will be no big servers anywhere but, rather, lots of little boxes, dynamically allocated as needed.”

In general, servers consist of a combination of two- and four-processor Dell PowerEdge servers running Red Hat Enterprise Linux AS v. 2.1 and 3. For data storage, the ADC uses clustered storage systems from Network Appliance. This infrastructure supports hundreds of Oracle On Demand customers running 1,600 independent database environments and application stacks.

“Preprovisioned equipment is always ready to go in certified configurations,” adds Pohto. “From the time a customer signs an Oracle On Demand contract, we have 21 days to turn over the entire environment. Right now we’re averaging 11 to 14 days.”

### COMPARING COSTS

When customers contract with Oracle to place their applications in the ADC, they get to take advantage of Oracle’s rigorous management and administrative practices. This is partly what attracted Commercial Net Lease Realty (CNLR) to adopt Oracle E-Business Suite On Demand. “We like being able to leverage Oracle’s internal consulting team for implementation, upgrades, and services, and we enjoy having real-time support available if problems occur,” reports David Lachicotte, vice president of

## Project MegaGrid

In late 2004, Dell Inc., EMC Corporation, Intel Corporation, and Oracle announced Project MegaGrid, a group project in which the four companies combine certain core technologies and technical resources to ease the burden of integration for their customers and to develop a complete enterprise grid computing solution that will challenge the performance of traditional SMP offerings at a fraction of the cost—in short, to develop an integrated approach to building and deploying an enterprise grid.

The initial phase of Project MegaGrid is focusing on designing, testing, and documenting best practices for building effective enterprise grid computing infrastructures, taking into account cost and performance requirements. These practices also include conducting a series of tests for scalability, performance, and manageability of a comprehensive database, server, and networked storage configuration. The working infrastructure is designed, configured, and validated at Oracle’s Austin Data Center.

The next phase will tackle more-complex scalability challenges, such as provisioning across the midtier and the database and storage tiers. It will include operating multiple midtiers and databases across the large cluster.

information technology at the Orlando, Florida-based firm.

As a real estate investment trust, CNLR acquires, owns, manages, and indirectly develops net-leased, single-tenant properties nationwide. According to Lachicotte, Oracle’s On Demand environment offers significant savings and faster resolution for potential technical problems, allowing his team to focus on internal support issues and on bolstering the business processes that give his company a competitive edge.

Before adopting Oracle E-Business Suite On Demand, CNLR conducted a comprehensive cost analysis that included hardware, software, professional services, and upgrades. The company figured that it would spend about US\$3.5 million over a three-year period to deploy, maintain, and upgrade the infrastructure itself, or just over US\$2 million if it let Oracle do the job.

### ALWAYS ONLINE

As Oracle pushes the boundaries of grid computing, the Austin Data Center remains a unique test bed for the industry as a whole. But the real beneficiaries are Oracle customers,

who end up with better products as a result of Oracle’s experience. “You can see some features in Oracle 10g that came about as a result of our experiences with grid computing in Austin,” says Souder. “Moving forward, I think you’ll see more and more of the knowledge we learn and the experience we gain in this grid environment coming back into the products.” ■

*David Baum (david@dbaumcomm.com) is a freelance business writer based in Santa Barbara, California.*

### SNAPSHOT

#### Commercial Net Lease Realty, Inc.

**Location:** Orlando, Florida

**Industry:** Professional services

**Total assets:** US\$1.3 billion

**Oracle products & services:** Oracle E-Business Suite On Demand, including Assets, Cash Management, Financial Analyzer, Financials, Internet Expenses, Payables, Procurement, Projects, Property Manager, Purchasing, Receivables, and Sales Analyzer; and Oracle Consulting Services