IDC
Analyze the Future

WHITE PAPER

Oracle's MAA Portfolio: Deploying High-Availability Solutions Across the Enterprise

Sponsored by: Oracle

Jean S. Bozman
June 2010

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

IDC OPINION

High availability (HA) of applications and databases is always an important consideration when workloads are deployed. But for business-critical and mission-critical workloads (those workloads that are essential for continued business operations), high availability is a "must." When these workloads are interrupted, business processes come to a halt — causing downtime and increasing operational costs.

There are many approaches to achieving high availability, but acknowledging that different types of workloads have different requirements for availability is a pragmatic way to ensure business continuity. Business processes must keep going, or there will be impacts on employee productivity, end-user access — and business revenue. Therefore, offering a spectrum of HA software, combined with reliability features in the hardware platform, is an efficient approach to "mapping" the right HA technology to the right workloads.

Oracle has taken a robust approach to ensuring high availability, with multiple products under the unified Oracle Maximum Availability Architecture (MAA) initiative. The range of these MAA products includes high-availability software (Oracle Real Application Clusters [Oracle RAC]) and Oracle Data Guard and support for alternate "paths" to connectivity in a heterogeneous computing environment, offered via Oracle GoldenGate.

SITUATION OVERVIEW

High-availability initiatives at customer sites are driven by organizations' need to avoid downtime — when systems go offline or become unavailable to end users accessing a database or applications. High-availability software preserves business continuity by maintaining the presence of the application online — even if that application is moved to alternate computing resources due to an outage (e.g., power outage, network outage, or outage caused by failure of a physical server).

However, because organizations, whether midsize or large organizations, have a range of workloads (applications and databases) running in their computing environment, there is also a need to "tap" different approaches to HA for different types of workloads. IDC calls this processing "mapping" — as in mapping different HA technologies to different types of workloads, as appropriate.

IDC Availability Spectrum

Enterprises have a broad array of applications running within their environment. Thus, it is important not only to "map" the appropriate HA software solution to specific workloads but also to offer a "spectrum" of HA solutions that address availability for a range of business requirements. To make this point, IDC developed the Availability Spectrum matrix, which associates different availability levels with specific qualitative results related to the type of downtime experienced. End users experience different amounts, or types, of interruption in their work, depending on what type of HA solution is applied to an outage. The degree of interruption, as experienced by end users, differentiates the availability levels shown in the IDC Availability Spectrum chart.

Table 1 shows this "spectrum" of availability levels (labeled 1–4). For example, the broadest range of high-availability solutions is classified in the AL2 and AL3 levels within the chart. Typically, clustering and failover software covers this broadest range within the Availability Spectrum. End users may experience a short, or very short, period of downtime, followed by a return to normal operations.

These two levels, AL2 and AL3, are the most popular and widely deployed approaches to ensuring business continuity by means of "failover" of key applications and databases to alternate server resources with access to shared data. Restarting workloads, when combined with access to production data, ensures continued access to applications and databases, even if unplanned downtime occurs. Top causes of unplanned downtime include power outages, network outages, outages caused by natural disasters, and outages caused by human error or IT operational errors. That means that outages will continue to occur — and that businesses must plan for them and work to reduce operational costs associated with downtime. In contrast, planned downtime is associated with routine maintenance of systems and the need to take physical servers offline for purposes of repairs, software updates, or security patches.

 ©2010 IDC

TABLE 1

IDC's Availability Spectrum

	Impact of Component Failure on Priority User	System Protection Features
Availability level 4 (AL4)	Transparent to user; no interruption of work; no transactions lost; no degradation in performance	100% component and functional redundancy
Availability level 3 (AL3)	Stays online; current transaction may need restarting; may experience performance degradation	Automatic failover transfers user session and workload to backup components; multiple system connections to disks
Availability level 2 (AL2)	User interrupted, but can quickly log on again; may need to rerun some transactions from journal file; may experience performance degradation	User work transferred to backup components; multiple system access paths to disks
Availability level 1 (AL1)	Work stops; uncontrolled shutdown; data integrity ensured	Disk mirroring or RAID and a log-based journal file system for identification and recovery of incomplete in-flight transactions

Source: IDC, 2009

In addition, IDC has classified fault-tolerant systems within the AL4 level; these are systems that show no perceptible interruption in data services to end users through the use of built-in hardware redundancy and lockstep software. These fault-tolerant approaches address a small segment of the total availability opportunity, while the failover, clustering, and workload-balancing solutions address the majority of HA installations at customer sites, based on IDC demand-side, customer-based research. Short periods of restart are acceptable to most businesses; long periods of downtime (minutes to hours) are not acceptable.

AL1 category software generally addresses data replication and the preservation of the data that is generated by a production workload. The ability to move workloads from one server to another, when combined with the ability to access production data, provides a holistic solution for high availability that preserves business continuity via backup/restore and data replication and allows work to proceed normally — even when an unplanned outage occurs.

Unplanned Downtime and Planned Downtime

Planned downtime is the other broad category of downtime — and it is the counterpoint to unplanned downtime. Preparation for unplanned outages, which are inevitable everywhere and at all times of the day or night, is what allows organizations to avoid prolonged periods of unplanned downtime. However, physical servers must be maintained — and workloads must be shifted to alternate servers in order to carry out those maintenance activities. Planned downtime — during which physical servers are repaired or maintained or system software is updated — can also be accommodated, given thorough planning on the part of IT staff and the business units they support.

Rising adoption of virtualization in the x86 server world has led to an increase in attention regarding planned downtime and the use of live migration of virtual machines (VMs) for purposes of addressing planned or unplanned downtime. However, these approaches usually do not address issues of both planned and unplanned downtime for specific applications or databases. This is the aspect in which technologies that address downtime become necessary for business continuity. Fortunately, many software products for unplanned downtime work as easily with virtualized servers as with non-virtualized servers, making it easier for IT managers to move important workloads between physical and virtual servers, if needed.

The Worldwide Availability and Clustering Software Market

The availability and clustering software (ACS) market is driven by customers' need to avoid downtime and to preserve continuity for applications, databases, and production data. With high availability, businesses can quickly recover from the inevitable outages caused by natural disasters, power outages, network outages, and the like.

IDC tracks the size of the worldwide ACS market and publishes the results in an annual report. In addition, IDC provides a worldwide forecast, with regional breakouts, for the total market opportunity. It is important to note that while the worldwide server market declined in 2009, IDC research shows that demand for ACS software continued to increase, underscoring the fundamental need for highly available workloads (inclusive of applications and databases).

Figure 1 shows the growth of the ACS market from 2008 to 2013, as forecast by IDC's software group — based on supply-side software revenue data. Use of this ACS software, when combined with IT "best practices" learned over years of deploying highly available configurations of servers and storage, supports business continuity and, therefore, continued access to data by end users and end customers. IDC studies have shown that downtime is a leading driver of operational costs associated with IT staff time and reduced end-user productivity within the business. Across computing environments, regardless of hardware platform, chip type, or operating system, the drivers for high-availability software are the same.

FIGURE 1

Availability and Clustering Software Revenue by Region, 2008 and 2013

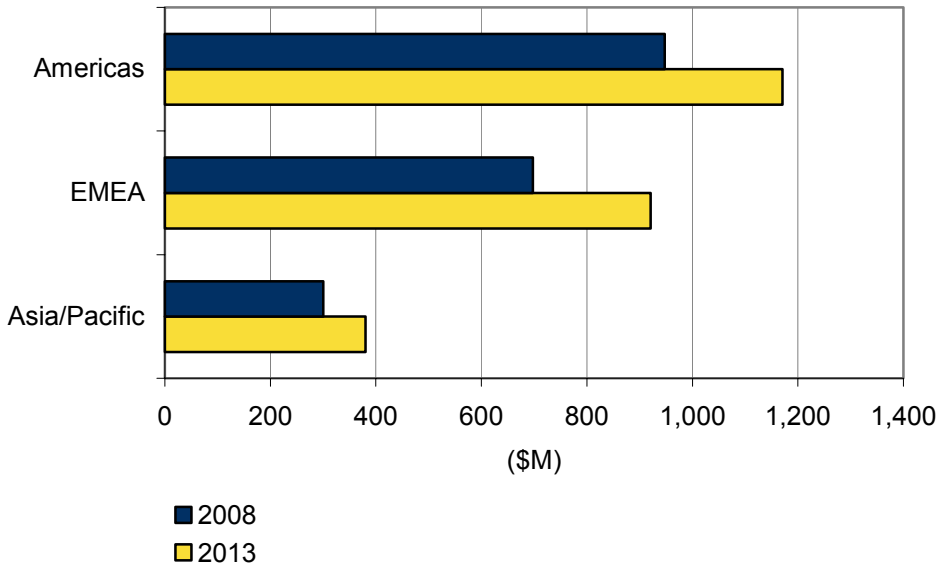

Source: IDC, 2010

Top drivers for high-availability software deployments have been remarkably consistent in recent years, based on IDC's ACS research. As Figure 2 shows, concerns about downtime, the need to support 24 x 7 x 365 business operations, the prospect of revenue loss, and the prospect of reduced employee productivity are among the top drivers for deployment of high-availability software. These drivers for high-availability software deployment exist, even when virtualized server infrastructure is considered. IDC demand-side data shows that customers expect to acquire more availability data, including ACS products that will be deployed in virtualized computing environments.

FIGURE 2

High-Availability Solutions Drivers: All Servers, All Applications

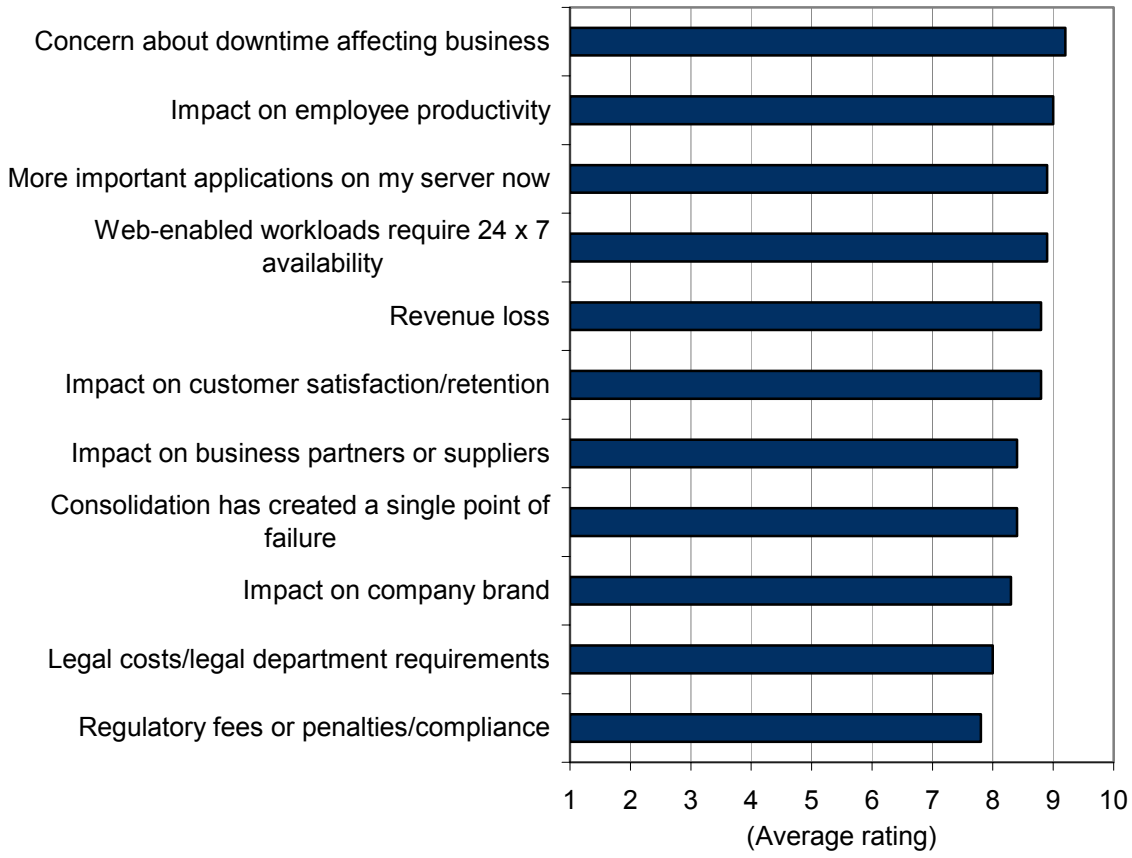

Note: Respondents were asked to rate the drivers for deployment of HA software on a scale from 1 to 10, with 1 being not at all important and 10 being very important.

Source: IDC, 2010

Oracle's High-Availability Solutions

Overall, Oracle's approach to high availability is seen in the Oracle MAA portfolio of offerings. Oracle provides a range of availability solutions, applying specific technologies at multiple layers of the "software stack." By providing HA functionality in multiple layers of the software stack, Oracle technologies address many causes of downtime and provide support for continued processing on other servers and storage devices within the enterprise.

Importantly, the Oracle solutions for availability address most of the availability levels in the IDC Availability Spectrum. With the exception of the fault-tolerant AL4, which is not a design point for most enterprise workloads, Oracle addresses all of the levels: AL2 and AL3 with active/active failover clustering (Oracle RAC) and workload balancing and grid computing — and AL1 through support for data replication. Given the ability to restart applications and the availability of production data that has been copied or replicated to alternate resources, business can resume following unplanned

downtime. It is also possible to use Oracle RAC to move workloads from one server to another for purposes of supporting planned downtime. This, too, preserves uptime for production workloads and supports IT staff productivity, which helps to control operational costs for IT and for the business. One way to view this range of offerings is to say that they cover many use-case scenarios for production deployments of enterprise workloads (applications and databases) across the enterprise.

Oracle's Maximum Availability Architecture

Oracle MAA supports a range of availability solutions, precisely because the variation in availability requirements within the datacenter means that multiple approaches to HA will be needed to address the full complement of workloads in the organization. For example, high availability in a Web server environment, given deployment of high numbers of small servers supporting Web requests, will be very different from high availability for banking applications, which requires that each database update be verified and accounted for. Similarly, enterprise applications will need higher levels of availability than workloads supporting simple deployments of IT infrastructure.

The major components of the MAA "umbrella" architecture for high availability are as follows:

- **Oracle Database 11g.** The foundation of Oracle's MAA is Oracle Database 11g, which is designed to provide extremely high levels of availability, scalability, reliability, robustness, and performance. With respect to availability, over time, more HA functionality has been embedded within the database, integrating that functionality for ease of use and consistency of deployment, compared with the need to install additional layers of software for providing those functions. The database supports the ability to take "snapshots" of the data for later playback, allowing database administrators to apply changes to the database in a non-disruptive way so that production work can continue.
- **Oracle Real Application Clusters.** Oracle RAC provides high-availability clustering that can be deployed across all major hardware platforms, including x86, RISC, and EPIC architectures. It provides extremely high levels of availability for databases by running databases concurrently on multiple servers. In the event that a failure occurs, the surviving nodes would continue to provide service. Unlike traditional HA failover solutions, there is no need to restart the database on a surviving server before service can resume because the database is already running. RAC can be deployed within a scalable server, or it can be deployed across a grid of servers running Oracle software. This means that RAC can support scale-up or scale-out computing. Importantly, RAC can run on physical servers or virtual servers (VMs).
- **Oracle Automatic Storage Management (ASM).** ASM presents multiple storage devices as a pool of storage for Oracle Database files. With support for online addition and removal of devices from the pool, and automatic management of database files, it simplifies many management tasks. ASM automatically stripes all files across all devices in the pool, ensuring there are no hot spots or I/O bottlenecks in the storage subsystem. It also provides integrated mirroring to protect against failures of individual storage devices. In the latest version, Oracle has extended ASM to support file system data as well as database files.

- **Oracle Data Guard and Oracle Active Data Guard.** Oracle Data Guard is an Oracle Database feature that supports safe storage of production data on standby servers, preserving that data for disaster recovery failover. Data Guard ensures that the standby databases are isolated from any data corruptions on the production database, ensuring data integrity. Oracle Active Data Guard extends Oracle Data Guard by enabling read-only access of the standby system. With Active Data Guard, IT managers have the option of using an active/active approach to disaster recovery, which eliminates the need to maintain inactive standby systems. This avoids over-provisioning in the datacenter, which, in turn, reduces the capital expenditures (capex) associated with standby devices.

- **Oracle Flashback Technologies.** To minimize downtime from human errors, Oracle Database includes the Flashback suite of features, which can rewind the state of the database to a known, safe point in time and so reverse (or undo) the effects of human errors without requiring long system outages. Flashback technologies provide an integrated Continuous Data Protection (CDP) solution for the Oracle database.

- **Recovery Manager (RMAN).** RMAN provides an integrated backup, restore, and recovery infrastructure within the database, and it is key to restoring business processes in the event of unplanned downtime. In a database production environment, backup and recovery procedures must be scheduled, in a procedural way, so that the backup data is up to date, allowing faster recovery to the "state" of computing as it was prior to the unplanned outage event. RMAN enables optimization of the backup and recovery process for the Oracle Database through its integrated disk-based backup capability (Fast Recovery Area) and through its integration with media management products, such as Oracle Secure Backup. RMAN has also been extended to support cloud-based backup and recovery of Oracle Databases to Amazon Web Services (AWS) S3 storage for archiving.

- **Oracle Secure Backup (OSB).** Accurate backup of production data is vital to maintaining high availability in an enterprise datacenter environment — and across a corporate network. Within the hierarchy of storage devices, tape devices support long-term archiving of production data for later retrieval, while continuing access to nearline data after it has been stored to the secondary tier of storage devices is key to restoring production data in the event of unplanned downtime. Oracle Secure Backup is Oracle's centralized tape backup management product, part of Oracle's portfolio of an integrated data protection solution suite. OSB, which is also able to back up regular file system data, is tightly integrated with RMAN for use with the Oracle Database 11g. For this reason, it offers significant performance advantages and efficiency levels for database backups.

- **Oracle Streams.** The Streams feature within the database enables a flexible way to capture, stage, propagate, and apply changes between multiple Oracle databases. Streams has been used for a variety of purposes, including those requiring data transformation or distribution of data subsets — or any custom processing of the data in the stream. It ensures near-continuous access to production data that is in the process of being replicated to multiple storage devices.

- **Oracle GoldenGate.** GoldenGate, which was acquired by Oracle in 2009, supports real-time access to multiple data sources in a mixed environment that has multiple servers, multiple storage devices, and multiple types of database products. GoldenGate works by leveraging database log-based change capture and subsequent delivery of those changes across different types of databases and platforms while ensuring continuous access to the platforms and preserving business continuity. Oracle has stated that GoldenGate is Oracle's strategic logical replication and data integration product, which ultimately will address deployments currently supported by Oracle Streams.

A Robust Approach to Availability Solutions

Oracle offers a range of availability solutions, applying specific technologies to the "building blocks" of an end-to-end approach to high availability. By addressing restart of applications, access to production data, and support for heterogeneous, or mixed, computing environments — with multiple server types and operating systems — Oracle high-availability technologies manage many of the everyday realities of the IT organization and the enterprise datacenter.

Heterogeneous Computing in the Datacenter

This is a practical approach to addressing the uptime needs of enterprises and SMBs for which Web-enabled access to applications and data is a highly important business requirement. Some systems will be more mission critical than others — and these systems will require more protection, with restart of applications on alternative computing resources and access to production data that has been replicated prior to any downtime or outage. The combination of careful planning and adopting best practices in terms of IT deployments will maximize the business benefits of protecting against unplanned downtime. For planned downtime, best practices regarding maintenance for repair of physical servers and movement of workloads to alternate servers will also improve overall uptime for end users and end customers.

"Rip and replace" is not a practical solution for the in-house IT infrastructure of most businesses. Rather, careful consideration of what types of systems are installed — and what workloads are running on them — will determine what types of HA solutions should be installed to protect the business and the applications and databases that support that business. The economic downturn has limited IT budgets and caused budgetary constraints that will lengthen server life cycles within the datacenter, deferring previously planned replacement and buildout projects. That is why replacements will take place on a prioritized basis. That is why protection of business-critical and mission-critical workloads must take place — and that is why IT spend for availability software is continuing, even as the IT spend for servers declined in 2009.

Scale Up and Scale Out: Choices for Customers

Oracle products support both scale-up and scale-out computing, maximizing customer choice for server deployments. Other dimensions of customer choice involve the diversity of hardware platforms, chip types, and operating systems — and all of these choices are supported by Oracle high-availability software solutions.

Customer preferences and IT skill sets will determine which types of deployments work best in each individual case or usage scenario.

Oracle software supports a "grid" approach to scale-out computing, meaning that Oracle technologies run on collections of servers inside the datacenter, across the datacenter, or across the network. In this way, customers can deploy groups of small, or volume, servers and access networked storage to create a "grid computing" resource built on low-cost server hardware. Larger servers can also be deployed in clustered configurations if more capacity is needed on each individual server "node."

End-to-End High-Availability Solutions for the Enterprise

Availability software for business-critical and mission-critical workloads is part of a business' planning for business continuity — inclusive of servers, storage, software, and services. This is a key concept, as many sites are deciding to consolidate workloads onto fewer, but more powerful servers and to group small servers together via a grid or clustering deployment for scale-out configurations.

This "concentration" of computing onto fewer server footprints makes high availability for those servers more important. Any disruption to those servers will spell downtime for the entire business — and reduced productivity for employees and end customers accessing those servers on a regular basis. Oracle's MAA portfolio of HA solutions was designed to address a range of computing styles rather than focus on one or two use cases. The full spectrum of availability, from data replication to restart of applications, is supported by MAA software.

CUSTOMER SNAPSHOTS

State of Connecticut

The state of Connecticut is running line-of-business (LOB) mission-critical workloads on Oracle RAC, with a cluster of six RAC nodes. These extremely demanding workloads, which are accessed by thousands of state employees on a daily basis, are protected by a number of Oracle MAA products, including Oracle Database 11g and Oracle RAC. "They're considered mission-critical applications at the highest level of importance for the state," said Angelo J. Romano, IT Manager, Core-CT Technical Support for the state of Connecticut.

The state is running the PeopleSoft HR, Financials, Data Warehouse, and Enterprise Portal applications on its RAC deployments, which are based on production data stored in Oracle Database images. "Availability for the PeopleSoft solutions has been very good — and there is an alternate computing site, with another RAC deployment and mirrored data, for disaster recovery purposes," Romano said. The state uses Oracle Data Guard to replicate production data to this alternate site, where another Oracle RAC cluster is housed. "We're running Data Guard to keep our production databases synched up," Romano said. The state also regularly tests its ability to "fail over" applications to the alternate site as part of its ongoing disaster recovery planning.

There is a platform transformation story here as well. The state's financial applications, including the Oracle PeopleSoft modules, originally ran on scalable SMP Unix servers. However, as changes were made in the datacenter, the decision to move these enterprise applications from large SMP servers to a cluster, or grid, of smaller servers was made. As a first step, the PeopleSoft workloads were moved to Oracle RAC running on a cluster of midrange Unix servers in 2002 — and then they were transitioned to Oracle RAC running on a cluster of Linux x86 server nodes in 2007. Each server "node" within the RAC deployment is a dual-socket x86 server, and each socket supports four, or more, processor cores.

This multiphase approach allowed the IT move to scale-out x86 servers to happen over time while preserving high availability for the applications. The final switchover to the Linux x86 server platforms happened over a weekend. By transforming the server infrastructure for databases and applications, the state achieved two business goals — ensuring high availability and saving substantial money from moving from monolithic Unix servers to scale-out RAC deployments. Importantly, IT flexibility is another effect of moving to the RAC clustered servers, Romano said: "You can add more nodes into the RAC as you need more capacity."

Intermap Technologies

Intermap Technologies is based in Englewood, Colorado, and offers uniform high-resolution 3D digital models of the earth's surface. The company has proactively remapped entire countries and built uniform national databases called NEXTMap. The detailed data, which is stored within Oracle Database 11g using the Spatial functionality, can be accessed on a 24 x 7 x 365 basis — so downtime is not an option. Because of the way in which it is architected, Intermap provides digital elevation data at primary and alternate sites so that production data can always be accessed at one site or the other, even if unplanned downtime were to affect one of the two sites.

Customers use these high-resolution 3D digital models of countries worldwide to plan for a variety of projects: government services, engineering projects, oil and gas, hydrology, renewable energy, environmental planning, wireless communications, aviation, automotive, and insurance risk assessment applications relating to geographic location (e.g., flooding), among others. Intermap uses Oracle 11g for storing the geographic information system (GIS) images — and for high availability of those images — to ensure access by customers. "We are maintaining a very high level of accuracy and spatial referencing," said Sue Merrigan, Director of Intermap's Product & Business Solutions group. Many kinds of spatial data, including elevation and topography, are included with the stored images.

The company maintains a second site that mirrors data from its primary datacenter. "[Oracle] Active Data Guard is mirroring the data to a mirrored read-only architecture at the colocation facility," Merrigan said. The two centers are located about 30 miles apart from each other to ensure that any localized outages (network, power, natural disaster) at one site would not affect operations at the alternate site.

In addition to the Oracle Database 11g with GIS, Intermap has deployed Oracle Active Data Guard for data protection purposes. "The data is open to our clients for read-only access, even at times when our data management staff is supplying

changes to it." The site also uses the automated storage management (ASM) feature of 11g. "It makes storage management much simpler for our back-end storage," Merrigan said. "Being able to dynamically add LUNs, and [being able to] reboot them, without much manual intervention is very helpful. It allows our DBAs to focus more on administration and less on configuration."

MorphoTrak Inc. (SAFRAN Group)

MorphoTrak is a software company that applies a software-based matching technology based on advanced algorithms to biometric and identity management solutions. The company's MorphoTrak BIS (Biometric Identification Solution) system is used by police and government agencies for crime solving, border control, and civil identification. The company was formed from the merger of Sagem Morpho Inc. and Motorola's biometric division and is a unit of SAFRAN of France. MorphoTrak is headquartered near Washington, D.C., and has 450 employees.

Unique identifying characteristics of fingerprints are stored in very large databases — in excess of 10TB — inside each customer's MorphoTrak system. The MorphoTrak BIS safeguards all of the images stored within the customers' large databases, which is important because data loss and downtime are not acceptable options in the public safety arena.

The company leverages a number of products that are key elements of the Oracle MAA portfolio. "With over 100 production customers, each having different availability requirements and operational considerations, MorphoTrak has had the opportunity to deploy every Oracle availability technology, including Active Data Guard, RAC, and Multi-Master Replication," said Aris Prassinos, Chief Engineer at MorphoTrak.

MorphoTrak also needs to address the variation in IT skill sets that exists at customer sites worldwide. "We deploy our systems as turnkey fixed-cost solutions to customers who don't have onsite DBAs, and we have to be very cost-effective, both in terms of development and operational cost, as we operate in a very competitive market," Prassinos said.

Oracle Database 11g with its self-managing and tuning capabilities and ASM for simplifying storage administration make it easier for MorphoTrak customers to deploy and to use the MorphoTrak system, without having onsite database administrators, and at the same time meet their system availability requirements.

CHALLENGES/OPPORTUNITIES

Most IT organizations have inherited a mix of products and architectures installed throughout their organizations. A series of decisions, made over many years, has resulted in a "mixed" computing environment with a broad range of HA requirements. Some systems need minimal protection because they support "stateless" computing, such as the serving up of static Web pages, or they support file/print services. Others, supporting enterprise applications and enterprise databases, need high levels of availability to ensure business continuity. Without high availability for these systems, downtime will impact business revenue and employee productivity, causing operational costs to rise and resulting in business losses.

Oracle MAA is an umbrella initiative that provides a broad array of high-availability software solutions, which run on a wide variety of server and storage systems. In this way, Oracle MAA addresses the key pain points associated with unplanned downtime, providing choices for IT organizations in terms of functionality, price, and mode of deployment. Oracle is a leading provider of HA software, partnering with many system OEMs to deliver comprehensive solutions that minimize downtime and to provide alternative resources for continued data processing.

The worldwide software market is a competitive one, and there are many HA software solutions in the marketplace. While Oracle offers a comprehensive portfolio of availability products, it will continue to compete for customers' business with large systems vendors and other ISVs that also offer HA solutions — along with dozens of smaller software companies that sell ACS and HA software products worldwide.

CONCLUSION

There is a wide range of software solutions that improve availability for enterprise systems. The use of availability software to support combinations of servers and storage, and to achieve highly available data and applications, provides broad solution sets for customers who would otherwise need to integrate these solutions together by themselves, leveraging IT skill sets in their sites.

To the extent that vendors can provide a variety of HA solutions, aimed at a number of usage scenarios, and to the extent that they can make production deployments easier and faster, they can be said to be addressing customers' operational costs in the form of IT staff time and productivity. For customers in large enterprises, where mixed platforms and many operating systems are the norm, availability solutions must anticipate these infrastructure realities and current business conditions.

Oracle's MAA portfolio is a comprehensive set of product solutions supporting high availability and business continuity for many use-case scenarios. As such, its products address a number of IT pain points and can be used, often in combination, to improve overall availability for enterprise workloads and for the business processes they support, running across multiple computing tiers. Importantly, MAA products allow servers to continue running production workloads — even during test and development and software update activities — by supporting active/active workloads on those servers. Because of this, MAA is designed to reduce operational costs associated with ongoing maintenance and IT staff time devoted to maintenance.

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 2010 IDC. Reproduction without written permission is completely forbidden.