



For the Complete Technology & Database Professional

ENTERPRISE DATA AND THE COST OF DOWNTIME

2012 IOUG DATABASE AVAILABILITY SURVEY

By Joseph McKendrick, Research Analyst
Produced by Unisphere Research, a Division of Information Today, Inc.
July 2012

Sponsored by **ORACLE®**

Produced by  **UNISPHERE®**
RESEARCH
Delivering Certainty
Thomas J. Wilson, President

EXECUTIVE SUMMARY

For every organization in today's always-open global economy, success means ensuring access to enterprise data whenever it is needed. However, business operations are often hampered by interruptions in the flow of the vital data required to make decisions. When evaluating your organization's approach to high availability, you must consider a range of issues: For example, if a mission-critical database goes offline, how long will it take to get things back up and running? What are the business costs of the delays, and what can be done to mitigate or eliminate delays in delivering data?

These are some of the questions posed in a new survey of Independent Oracle Users Group (IOUG) members, which explored issues with planned and unplanned downtime, alongside database high availability and disaster recovery solutions. The survey, underwritten by Oracle Corporation and conducted by Unisphere Research, a division of Information Today, Inc., included input from 358 data managers and professionals.

At least half of the survey respondents report they are working to provide their enterprises with data on a real-time or near-real-time basis. The challenge of maintaining this level of data access is increasing as enterprise databases are tasked with the management of large volumes of enterprise data, both structured and unstructured. This increase in volume—as well as the velocity, and variety of data—presents some management and storage challenges, but also offers great opportunities for businesses to better serve customers and make more insightful decisions in a timely manner.

The survey finds that systems scaling into the hundreds of terabytes are commonplace, and that more than one out of 10 companies are managing more than a petabyte of data within their enterprises—taking into account all clones, snapshots, replicas and backups. (See Figure 1.) This varies by size: While only 4% of smaller organizations surveyed scale into the petabyte range, 18% of the largest now support petabyte-plus environments. (See Figure 2.)

Respondents were also asked to describe the types of systems implemented to support high availability expectations. A majority report that they support both on-site and off-site backup capabilities, stated by 70% and 60%, respectively. Most, 55%, also pursue clustering as a strategy, where servers are lashed together to automatically pick up each other's workloads in the event of a server failure. While about half also use storage mirroring, in which a copy of the file system is maintained at another location, there is even broader adoption of Oracle-aware technologies for data protection and availability. Respondents

are commonly deploying Oracle Database with Oracle Real Application Clusters, Oracle Active Data Guard and Oracle GoldenGate for higher availability of database systems. A closer look at technologies used to maintain synchronized database copies show that Oracle technologies have overtaken third-party storage mirroring solutions. (See Figure 3.)

Historically, IT infrastructure for data protection and availability (backup and disaster recovery) was based upon horizontal storage-based infrastructure solutions provided by numerous third parties. This has been changing as Oracle database users have implemented Oracle-aware capabilities for data protection and availability.

While storage technologies see an Oracle Database block as an opaque collection of bits with little to differentiate it from any other type of data on disk, Oracle sees a database block as an object having an Oracle-defined structure that can be checked at multiple levels for both physical and logical consistency. This deep integration with database internals has produced comprehensive solutions for data protection and availability included with Oracle Database. Evidence that the user community has recognized this value-added approach is reflected in these responses.

A high percentage of respondents maintain a standby database for the purpose of data protection and high availability. Various backup approaches—employing both Oracle and non-Oracle solutions—are shown in Figure 4.

These varied approaches are far more prevalent at the largest organizations in the survey. For example, 72% of the largest organizations have clusters of multiple servers to support their high availability needs, versus 42% of smaller firms. In addition, while 64% of the large sites have either remote or on-site standby databases and servers, only 23% of the smaller organizations surveyed claim these capabilities. (See Figure 5.)

Meeting database availability requirements on a consistent basis is a challenge for most organizations, as clearly not all service-level agreements (SLAs) are regularly met. If large portions of the system vary in the service level of availability required, or the requirements include the latest changes as opposed to being up to date within an hour, a day, a week, or a month, this can create dramatic savings in systems cost for the systems at lower requirement points. This enables additional budget to be spent on the systems with higher requirements, as well as providing key information so that systems can be grouped on physical servers by availability requirements and the need to deliver the latest information.

