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

TABLE OF CONTENTS

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

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 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.

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Survey respondents hold a variety of job roles and represent a wide range of organization types and industry verticals. The majority (51%) hold the title of database administrator, followed by that of director or manager. One-third work for very large organizations with more than 10,000 employees. By industry sector, the majority of respondents come from government agencies, IT service providers, healthcare establishments, financial services, and manufacturing operations. (See Figures 39–41 at the end of this report for more detailed demographic information on job titles, company sizes, and industry groups.)

Key highlights and findings from the survey, which explores high availability issues and solutions, include the following:

- Downtime is adding up to not just hours, but days per year for many organizations. Twenty percent of respondents report having three or more days of unplanned downtime during the past year, and close to half report three or more days of planned downtime. Whether it's planned or unplanned, there is a cost to downtime, and most respondents report their businesses regularly feel the effects of such interruptions—both at a business and an IT operational level.
- Network outages were the leading cause of unplanned downtime over the past year, followed by human error, server failures, and storage failures. On average, these interruptions result in more than a day's worth of business interruptions, and point to the need for a multi-layered approach to high availability to address all sources of potential failure.
- The surge in volume of both structured and unstructured data types being managed—combined with increasing user demand for higher availability and faster performance—is creating challenges for IT to meet service-level expectations. Half of the respondents already are tasked with delivering data on a real-time basis, and many are managing more than a petabyte of data. However, only one-fourth of companies report they back up this data on a real-time basis.
- The shift to cloud services and architectures is adding a new dimension to database high availability. Close to half the organizations surveyed use or plan to use public or private cloud services and architectures, and are likely to seek the same types of service-level agreements seen in their traditional on-premise environments.
- Organizational inertia—combined with the surge in data growth—makes business continuity challenging. Business continuity and disaster recovery planning is essential to data availability, but most respondents do not have formal plans in place to address any potential outages.

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 1: Total Data Managed

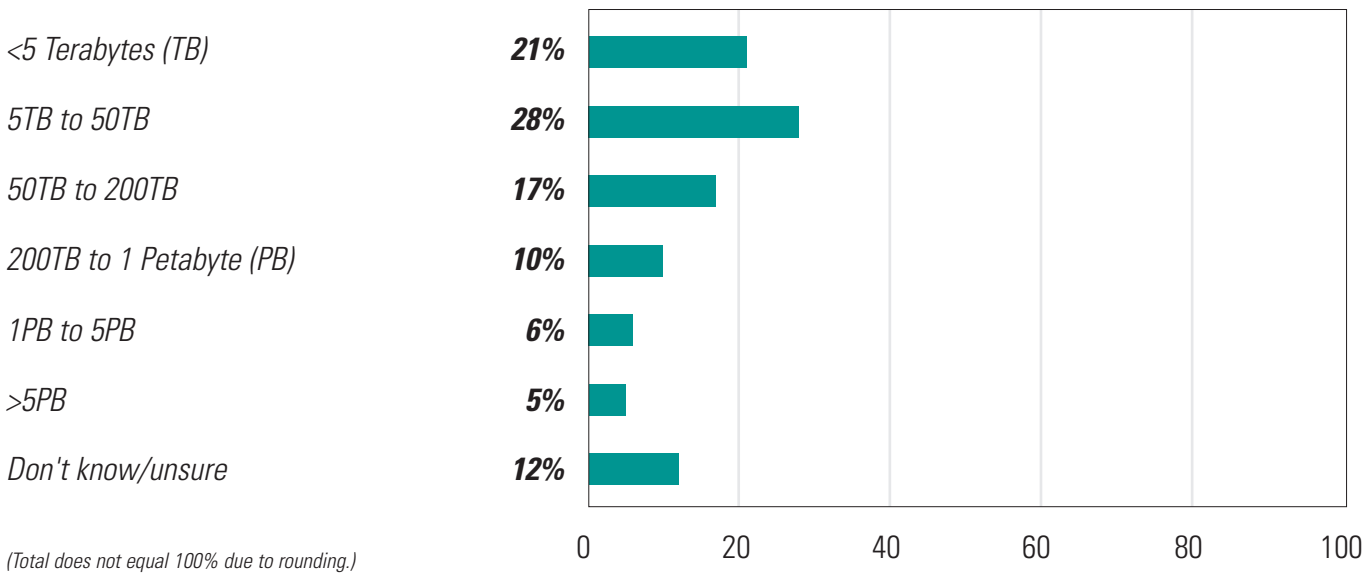

(Total does not equal 100% due to rounding.)

Figure 2: Total Data Managed—By Number of Employees

	<1,000	1,000 to 10,000	>10,000
<5TB	38%	11%	13%
5TB to 50TB	32%	29%	23%
50TB to 200TB	14%	25%	14%
200TB to 1PB	1%	14%	16%
1PB to 5PB	3%	4%	11%
>5PB	1%	5%	7%
Don't know/unsure	11%	11%	15%

(Totals may not equal 100% due to rounding.)

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 3: Oracle Availability Solutions

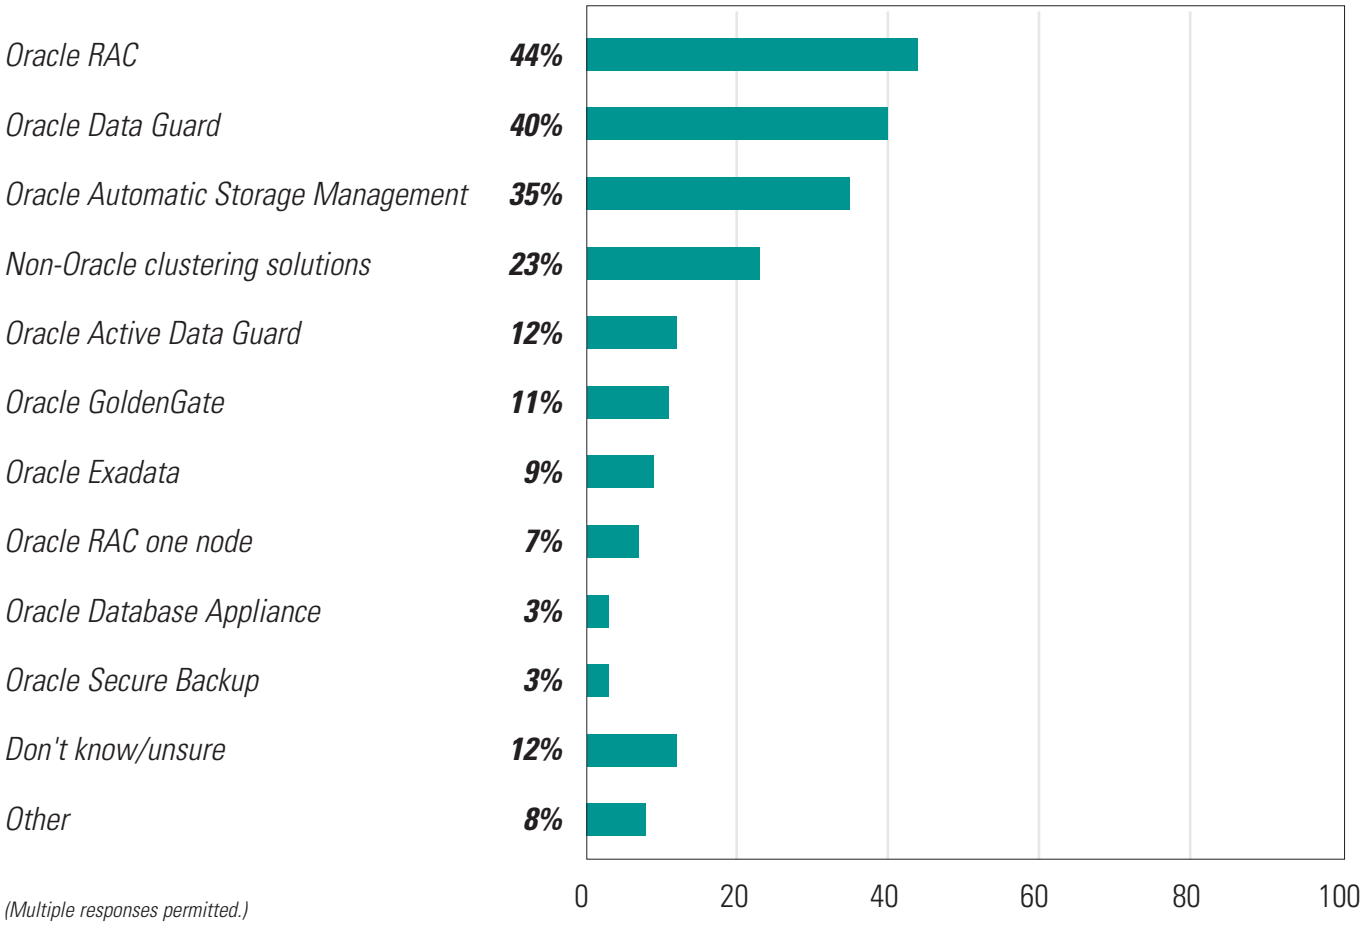

(Multiple responses permitted.)

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 4: Database Availability Infrastructures

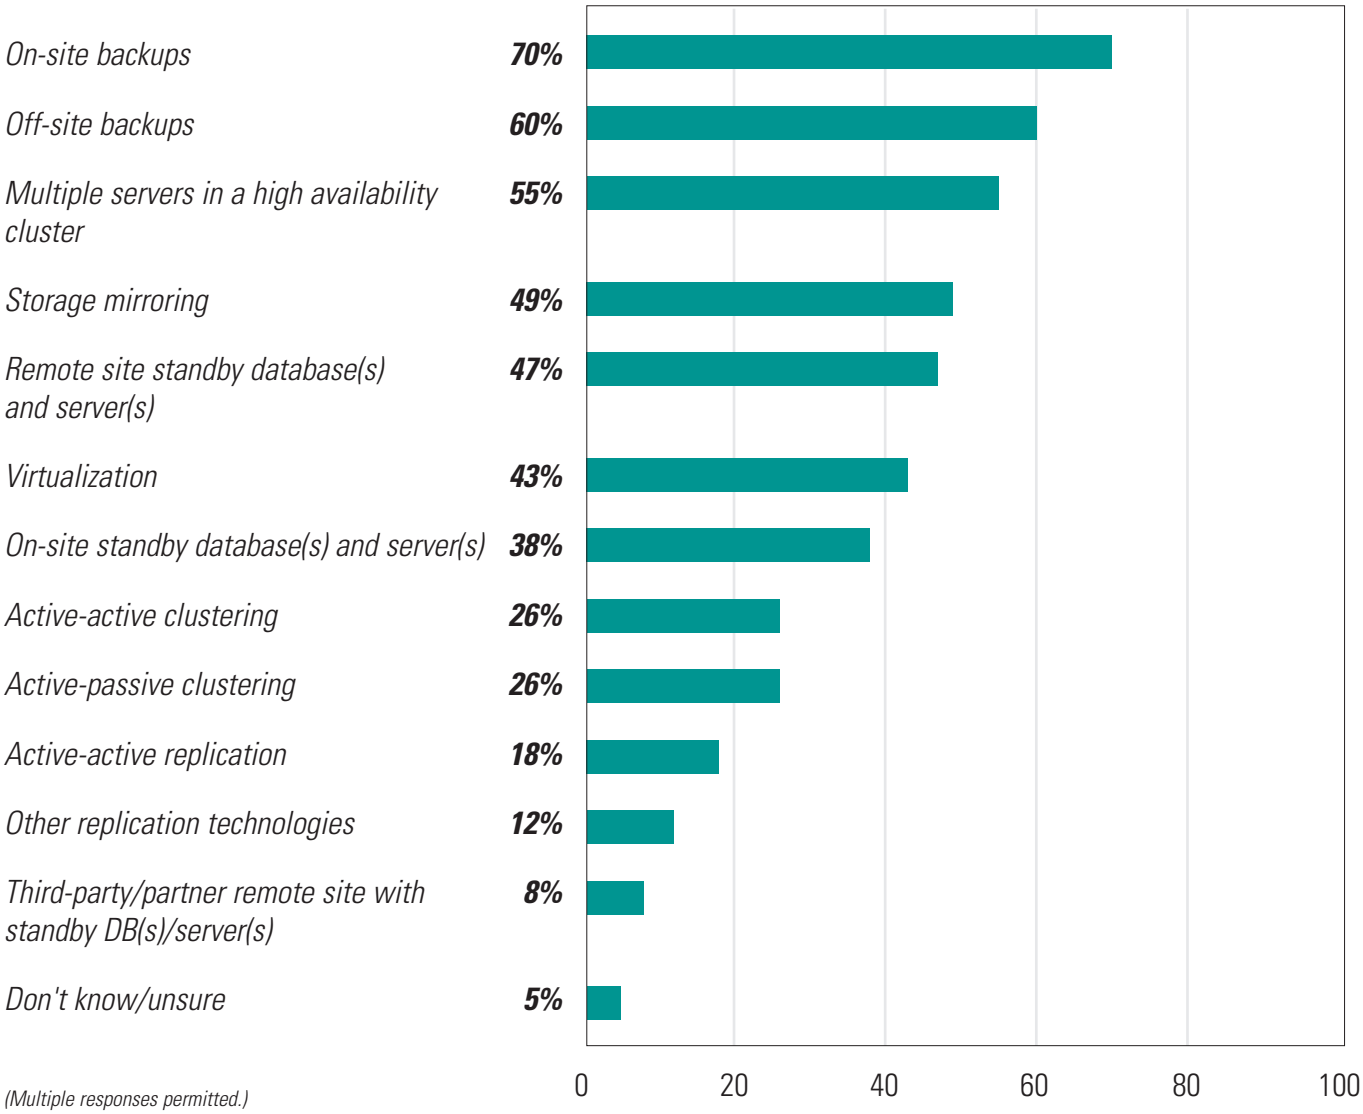

(Multiple responses permitted.)

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 5: Database Infrastructures—By Number of Employees

	<1,000	1,000 to 10,000	>10,000
On-site backups	65%	68%	81%
Off-site backups	55%	64%	60%
Multiple servers in a high availability cluster	42%	54%	72%
Storage mirroring	36%	57%	56%
Remote site standby database(s) and server(s)	23%	51%	64%
Virtualization	34%	42%	53%
On-site standby database(s) and server(s)	23%	30%	64%

(Multiple responses permitted.)

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

COSTS OF DOWNTIME—UNPLANNED AND PLANNED

Downtime is adding up to not just hours, but days per year for many organizations. Twenty percent of respondents report having three or more days of unplanned downtime during the past year, and close to half report three or more days of planned downtime. Whether it's planned or unplanned, there is a cost to downtime, and most respondents report their businesses regularly feel the effects of these interruptions—both at a business and IT operational level. Network outages were the leading cause of unplanned downtime over the past year, followed by human error, server failures, and storage failures. On average, these interruptions result in more than a day's worth of business interruptions, and point to the need for a multi-layered approach to high availability to address all sources of potential failure.

Addressing downtime can be a challenge for any data center. Even the most well-designed and state-of-the-art operations suffer their share of glitches, errors, and outages. The key is being able to minimize and shield end users from these failures, but the survey finds that incidents that take data offline tend to be all too commonplace. A majority of respondents report they suffered both disruptions to business operations as well as disruptions to their database or IT operations in the last three years.

What were the leading causes of any unplanned downtime experienced within database environments during the past three years? Network outages topped the list, as cited by half of the respondents. Both people and machines were also to blame for additional types of interruptions—45% cite human errors as the leading cause of unplanned outages, while the same number of respondents cite server failures. Another 42% cite storage systems issues. (See Figure 6.)

Interestingly, human and application errors took a greater toll at the sites managing the largest volumes of data (more than 200TB maintained). Sites with smaller data stores suffered the greatest issues with infrastructure challenges—network and power outages. (See Figure 7.)

With data center disruptions, there is fallout. More than a third suffered reduced business user productivity, while a similar number say there was a negative impact on the reputation of their IT department. (See Figure 8.)

In total, how much accumulated unplanned downtime did respondents' data environments experience during the year 2011? About 40% say they experienced, in total, one day or more of unplanned last year, while 49% suffered less than a day. About one-fourth report their total downtime was minimal, totaling fewer than four hours for the entire year. (See Figure 9.) These amounts were also consistent across the various types of infrastructure solutions, suggesting that more holistic and supportive management approaches must accompany technology to mitigate this downtime. What is required is an enterprise management approach, in which the business drives data availability decisions. IT and business managers need to collaborate to identify those data sets that are mission-critical to the business—and target both technical and staff resources to enable this level of high availability.

Planned downtime—when systems and databases are purposely taken offline by IT departments for maintenance, upgrades, or migrations—is another common source of business interruptions. More than two-thirds of respondents, 67%, say this kind of downtime totaled two or more days. (See Figure 10.)

Most respondents are aware of the costs of downtime to their businesses over the past three years as well. More than one in eight, 13%, say these costs have exceeded half a million dollars. Another 35% report that there have been costs, but often not exceeding $500,000. (See Figure 11.) Downtime costs exceeding the $500,000 mark are most likely to have been borne by the largest organizations—18%, versus 8% of the smallest firms. (See Figure 12.)

Among the sites managing the largest volume of data in the survey (those with more than 200TB), the costs are even higher for downtime—24% of this group report cumulative costs exceeding half a million dollars. (See Figure 13.)

A sizable segment of respondents have a goal to reduce annual unplanned downtime over the next 24 months—close to half, 44%, have active efforts underway to mitigate their levels of unexpected downtime. About 15% report that their goal is to reduce unplanned downtime by 50% or more. (See Figure 14.) A majority, 58%, also report they have definitive timelines for bringing down their levels of unplanned downtime, mostly within a six-to-12-month timeframe. (See Figure 15.)

Many companies accept that planned downtime is a necessary evil, and only about one-third of respondents, 31%, say they are attempting to reduce the length of these types of interruptions, the survey finds. (See Figure 16.) Among this group, about 45% have an actual set timeline for achieving their goal, with documented plans to reduce downtime within the next six-to-12 months. (See Figure 17.)

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 6: Leading Causes of Downtime Over the Past 3 Years

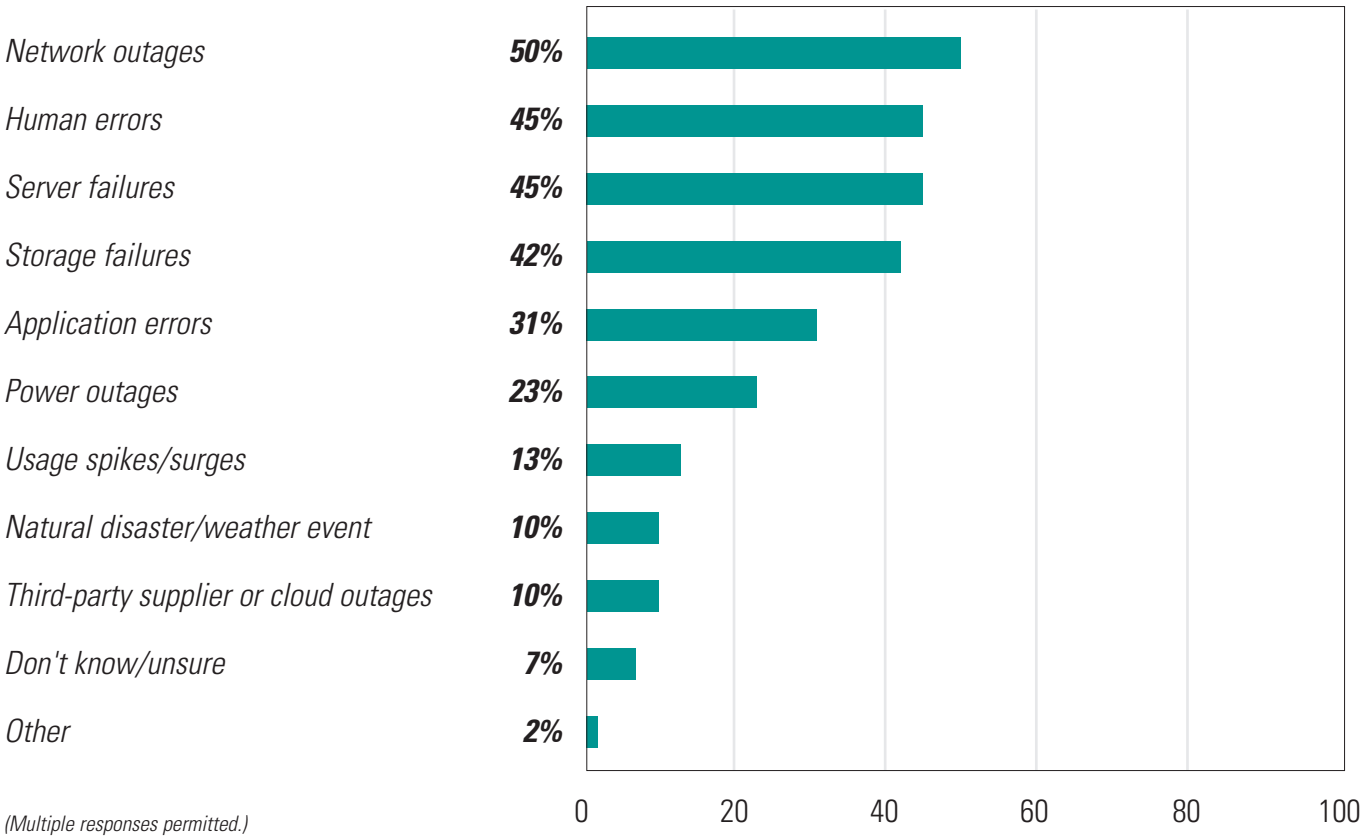

(Multiple responses permitted.)

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 7: Leading Causes of Downtime Over the Past 3 Years —By Data Volume

	"Large-Volume Data" sites	All others
Network outages	42%	53%
Human errors	58%	44%
Server failures	44%	46%
Storage failures	45%	44%
Application errors	37%	33%
Power outages	13%	28%
Usage spikes/surges	19%	15%

(Multiple responses permitted.)

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 8: Impact of Downtime in Past 3 Years

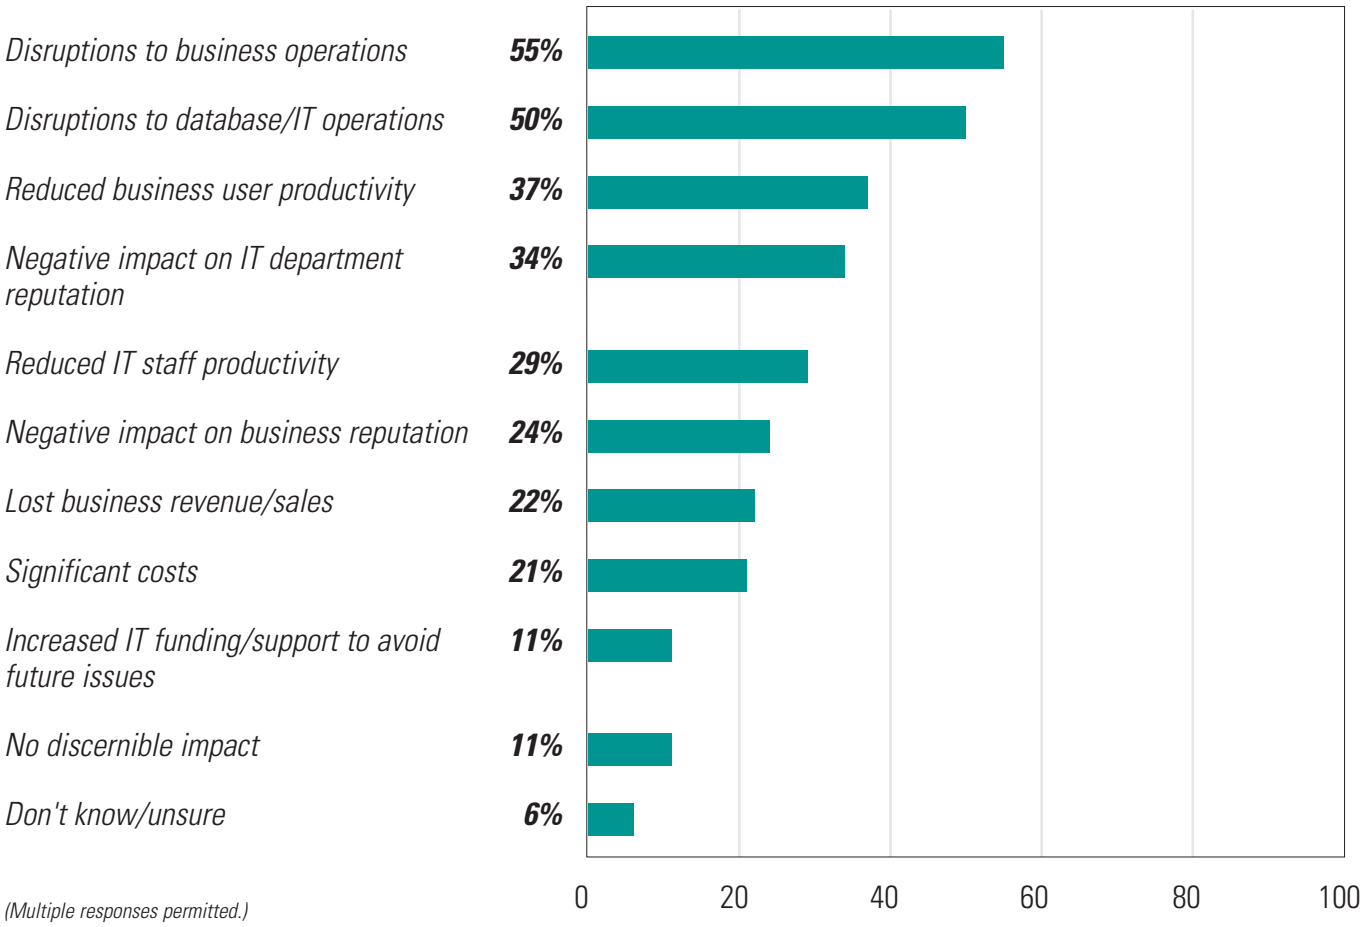

(Multiple responses permitted.)

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 9: Accumulated Unplanned Downtime in 2011

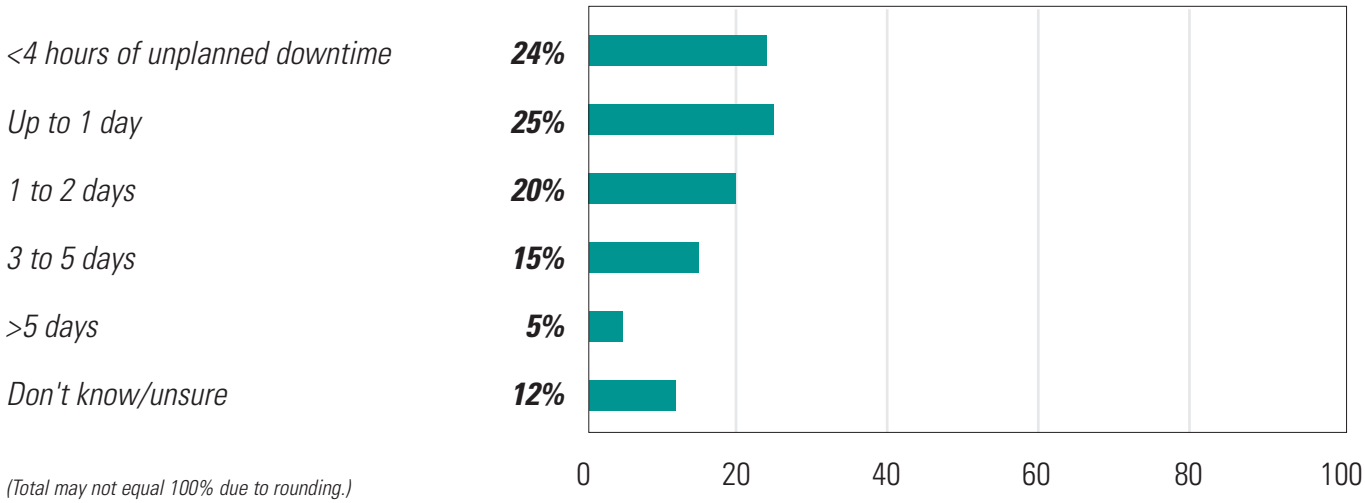

(Total may not equal 100% due to rounding.)

Figure 10: Accumulated Planned Downtime in 2011

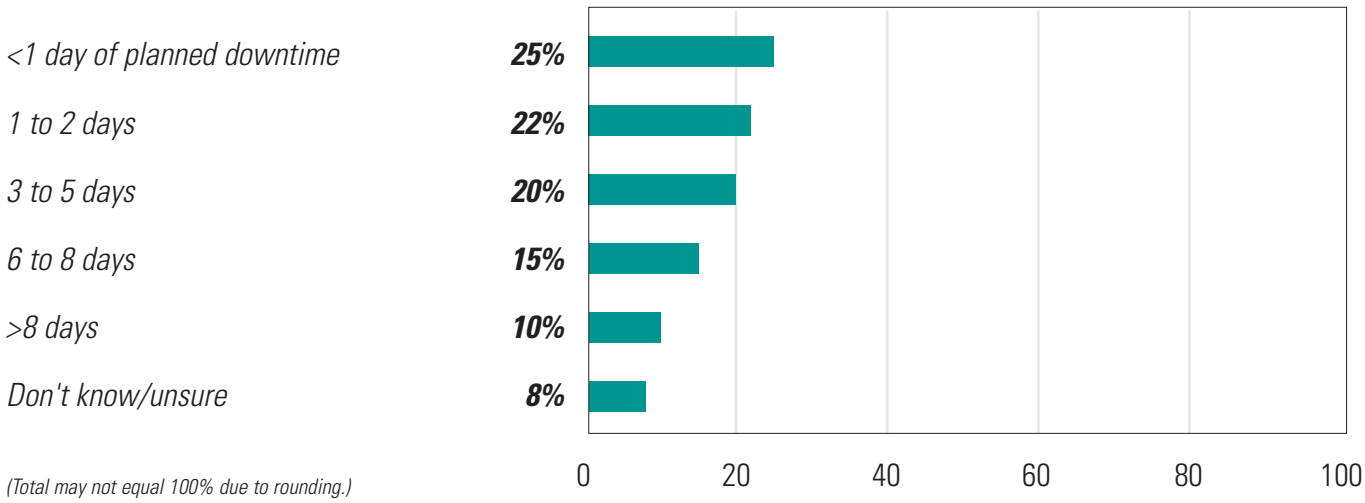

(Total may not equal 100% due to rounding.)

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 11: Total Cost of Downtime Over Past 3 Years

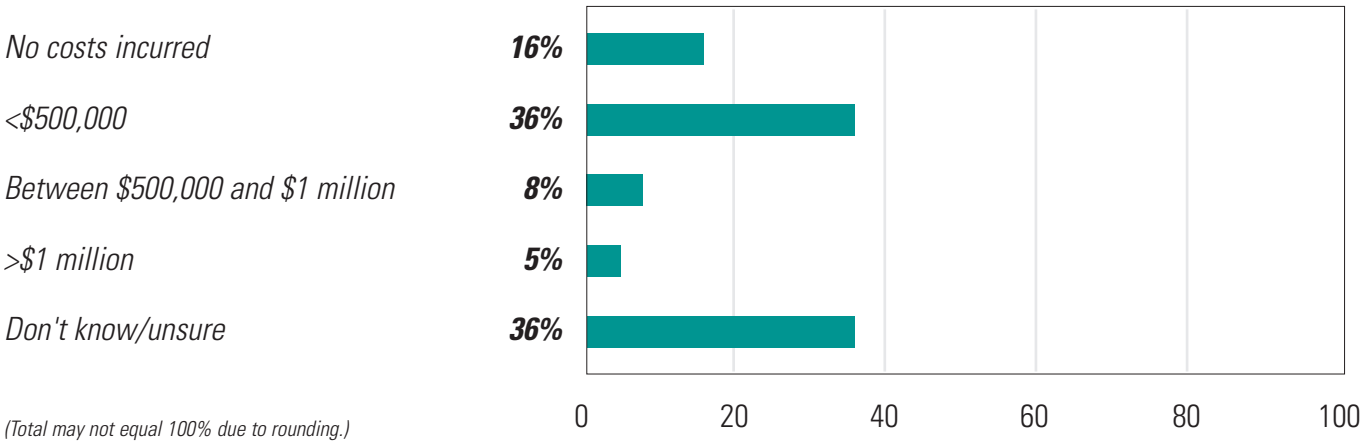

(Total may not equal 100% due to rounding.)

Figure 12: Total Cost of Downtime Over Past 3 Years —By Number of Employees

	<1,000	1,000 to 10,000	>10,000
No costs incurred	17%	20%	8%
<$500,000	35%	39%	29%
Between $500,000 and $1 million	5%	9%	8%
>$1 million	3%	3%	10%
Don't know/unsure	24%	29%	46%

(Total may not equal 100% due to rounding.)

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 13: Total Cost of Downtime Over Past 3 Years—By Data Volume

	"Large-Volume Data" sites	All others
No costs incurred	13%	20%
<$500,000	37%	41%
Between $500,000 and $1 million	8%	8%
>$1 million	16%	3%
Don't know/unsure	26%	29%

(Total may not equal 100% due to rounding.)

Figure 14: Goal to Reduce Annual Unplanned Downtime Over Next 24 Months

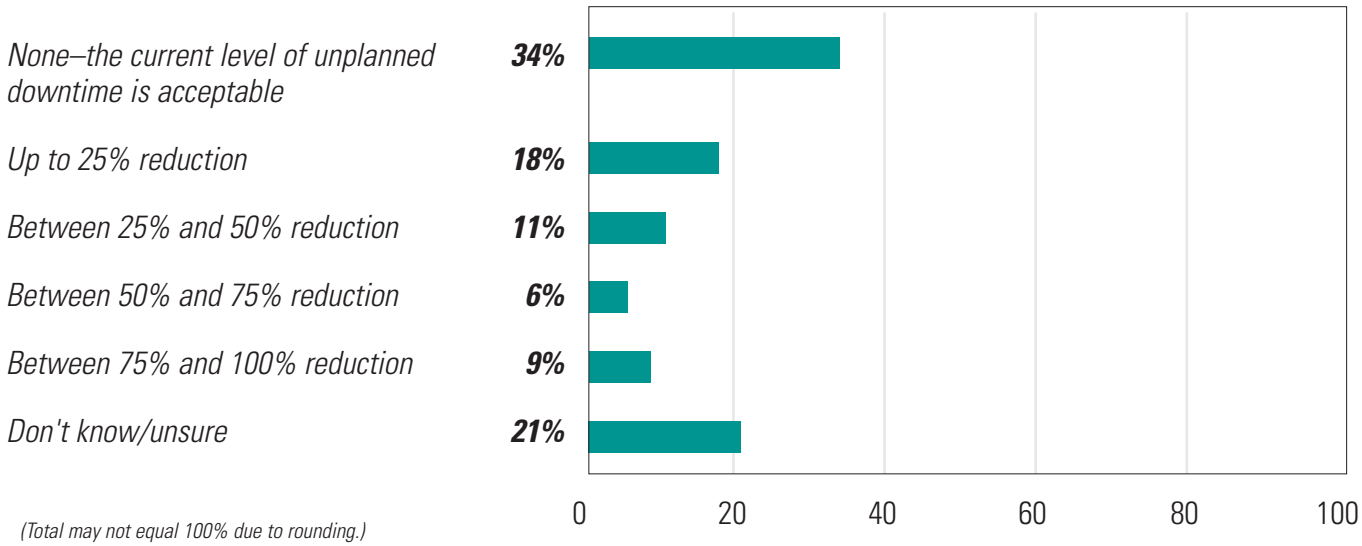

(Total may not equal 100% due to rounding.)

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 15: Timeframe for Reducing Annual Unplanned Downtime

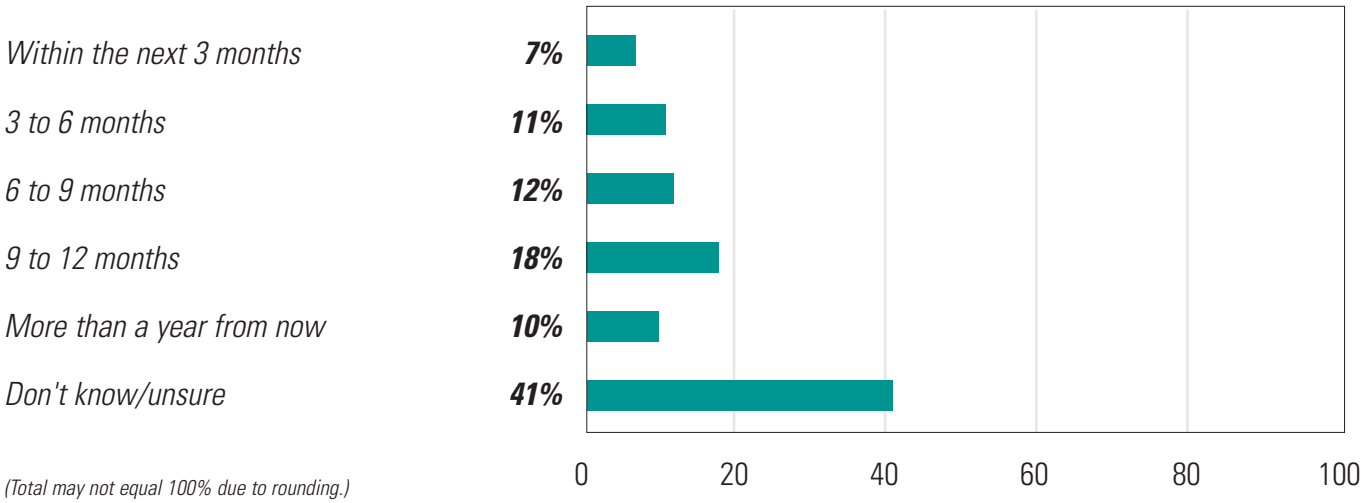

(Total may not equal 100% due to rounding.)

Figure 16: Goal to Reduce Annual Planned Downtime Over Next 24 Months

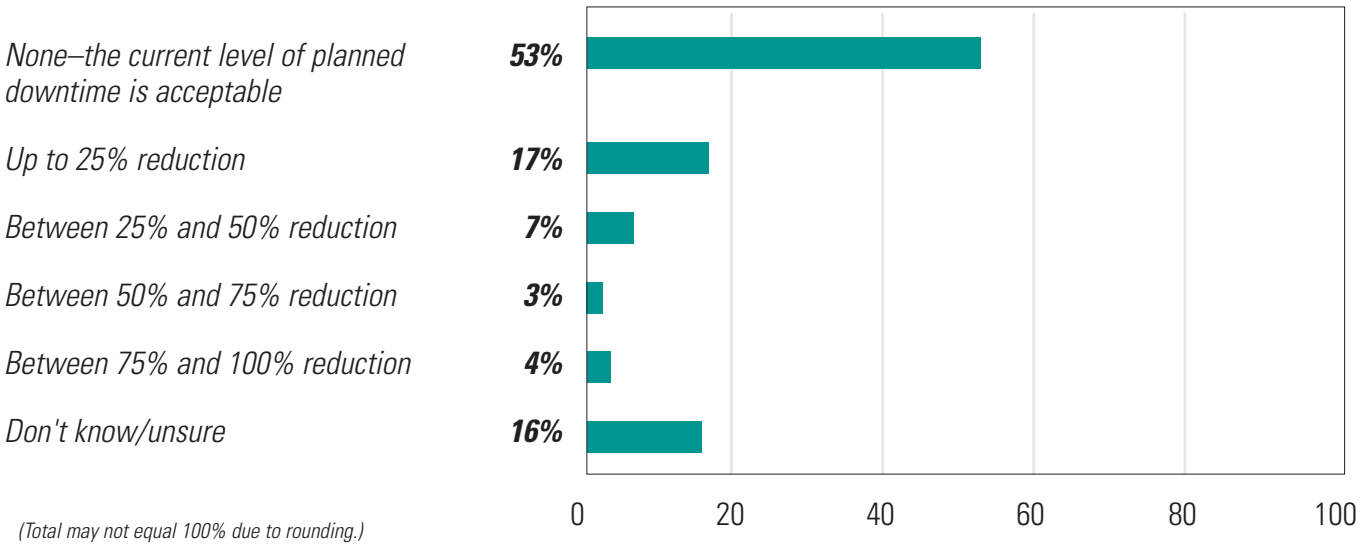

(Total may not equal 100% due to rounding.)

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 17: Timeframe for Reducing Annual Planned Downtime

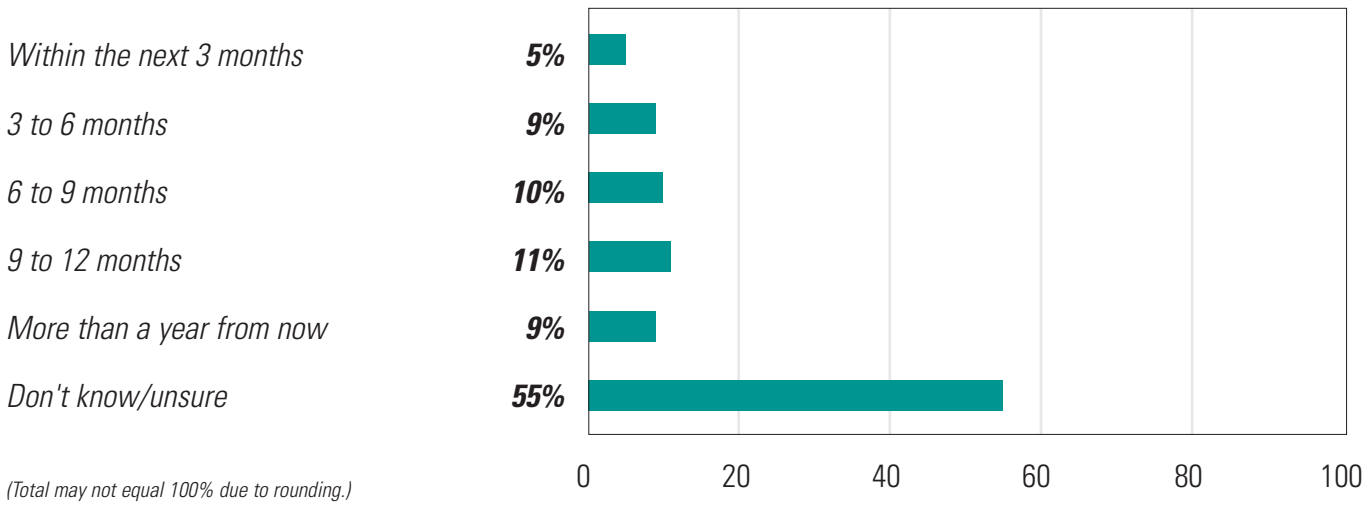

(Total may not equal 100% due to rounding.)

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

ACHIEVING DATABASE HIGH AVAILABILITY

The surge in volume of both structured and unstructured data types being managed—combined with increasing user demands for higher availability and faster query performance—is creating challenges for IT to meet service-level expectations. Half of the respondents already are tasked with delivering data on a real-time basis, and many are managing more than a petabyte of data. However, only one-fourth of companies report they back up this data on a real-time basis.

Today's database environments need to be available to provide information on demand to decision makers, unencumbered by technical issues that may arise in the background. This calls for high availability at data centers that operate at peak performance, support rapid failover to backup systems, and are capable of managing sudden spikes in workloads. The goal is to eliminate or avoid any extended wait times that may follow unexpected systems outages.

Respondents report a number of challenges in maintaining data availability in their current environments. Database performance and managing large volumes of business data are the two leading factors on the list of concerns, with each cited by half of respondents. Two-fifths also say increasing user populations are a challenge that affects availability—resulting in increasing query loads and usage spikes. Testing and quality assurance is a concern shared by more than one-fourth of respondents, particularly since users are demanding delivery of more real-time and near-real-time data. The ability to replicate this data is another challenge, particularly as data volumes grow, requiring larger and larger backup sites. (See Figure 18.)

Not surprisingly, these challenges are more acute among the large-volume data sites managing more than 200TB. About 69% of respondents at sites managing large data volumes worried about the impact of growing data volumes, versus 46% of small-data sites. Testing and quality assurance is also an issue more frequently seen at larger sites, while small data sites report more challenges with database performance. Interestingly, however, both large and small data sites were closer in terms of responding to increasing query loads and usage spikes, suggesting that sites of all sizes face increasing user demands. (See Figure 19.)

Meeting data availability requirements on a consistent basis is a challenge for most organizations, as not all database availability service-level agreements (SLAs) are regularly met. About one-fourth of respondents say these objectives are always met, but for a majority of respondents, 64%, capabilities are lukewarm —SLA requirements are met "most" of the time. (See Figure 20.)

Large-volume data sites in the survey report greater difficulty meeting service levels across their enterprise—only 13% are able to consistently meet their objectives, versus 32% of sites with smaller data stores. (See Figure 21.)

Real-time data is here, the survey finds. Half of respondents back up at least a portion of their organizations' data on a real-time or near-real-time (within one hour) basis. (See Figure 22.) This is a capability most often seen among the largest organizations in the survey. (See Figure 23.)

If respondents back up data on a real-time or near-real-time basis, how much of their organizations' data is protected this way? Only about one-fourth of respondents say a majority of their data is protected on a real-time basis. (See Figure 24.) This is less likely to be the case among organizations managing large data volumes —only 22% of companies with large, multi-hundred-terabyte stores back up most of this data in real time, versus 28% of smaller data sites. (See Figure 25.)

Respondents were asked to report their organizations' approaches to high availability, or the techniques and methodologies employed to recover malfunctioning systems as quickly as possible. While a majority of respondents report having a common strategy in place for ensuring high availability across their database services, its notable that this is still an area lacking for many companies. More than a third, 36%, say they either have no strategy, are still working on one, or simply don't know if such a strategy exists. (See Figure 26.)

How important are high availability architectures to achieving SLAs in respondents' environments? This is a critical approach for just about all respondents, and close to two-thirds rate it as "very important." (See Figure 27.)

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 18: Challenges in Maintaining Data Availability

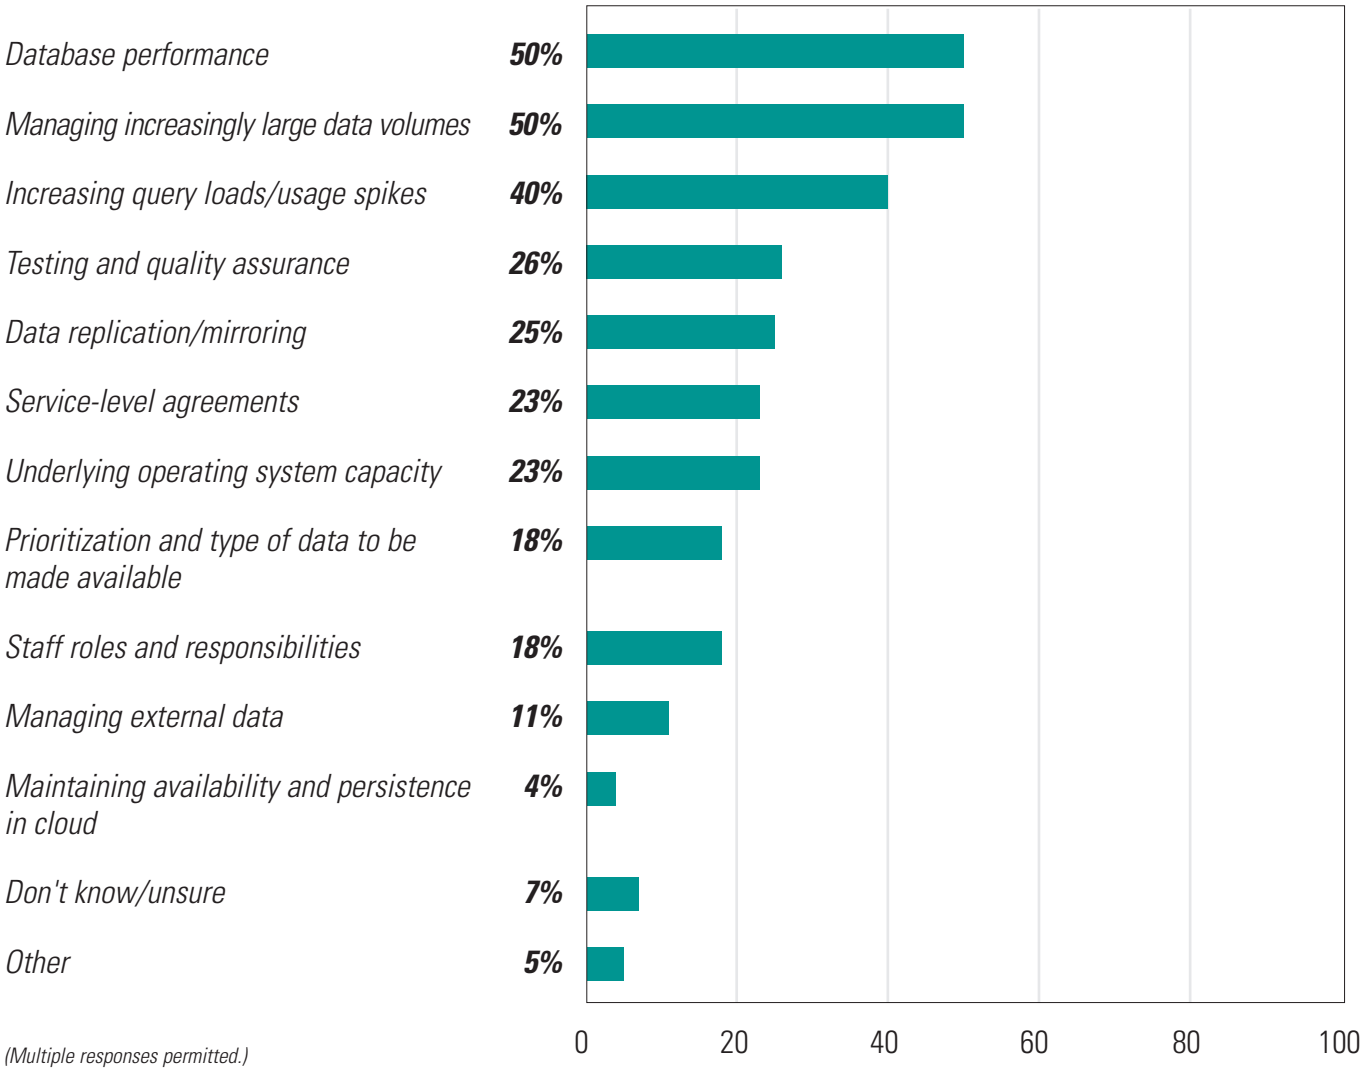

(Multiple responses permitted.)

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 19: Challenges in Maintaining Data Availability —By Data Volume

	"Large-Volume Data" sites	All others
Database performance	44%	54%
Managing increasingly large data volumes	69%	46%
Increasing query loads/usage spikes	47%	40%
Testing and quality assurance	40%	23%
Data replication/mirroring	27%	25%

(Multiple responses permitted.)

Figure 20: Consistency of Meeting Database Availability Service-Level Objectives

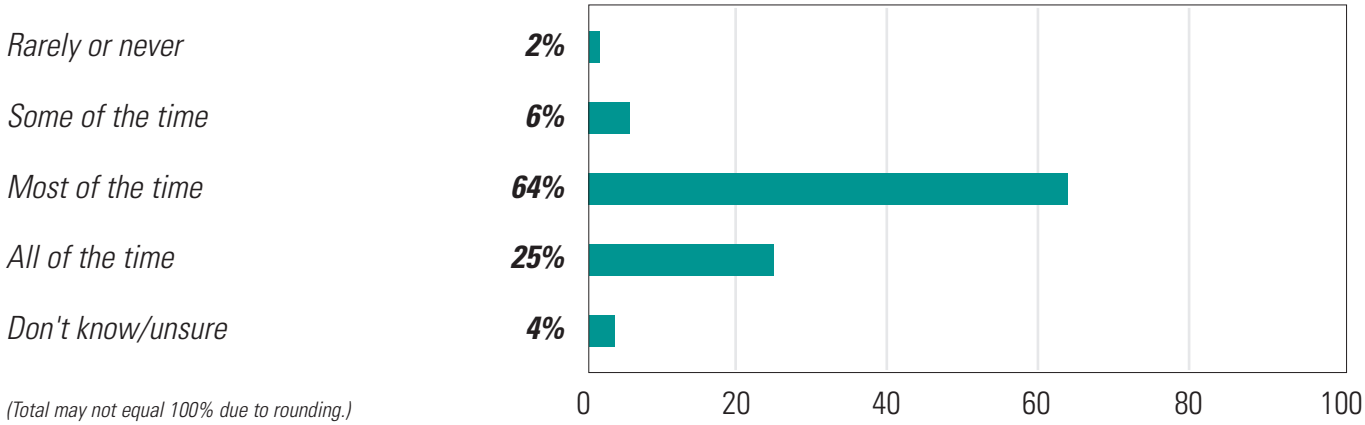

(Total may not equal 100% due to rounding.)

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 21: Consistency of Meeting Database Availability Service-Level Objectives—By Data Volume

	"Large-Volume Data" sites	All others
Rarely or never	3%	2%
Some of the time	3%	6%
Most of the time	79%	59%
All of the time	13%	32%
Don't know/unsure	2%	2%

(Total may not equal 100% due to rounding.)

Figure 22: Back Up Data Real-Time (Within 1 Hour)?

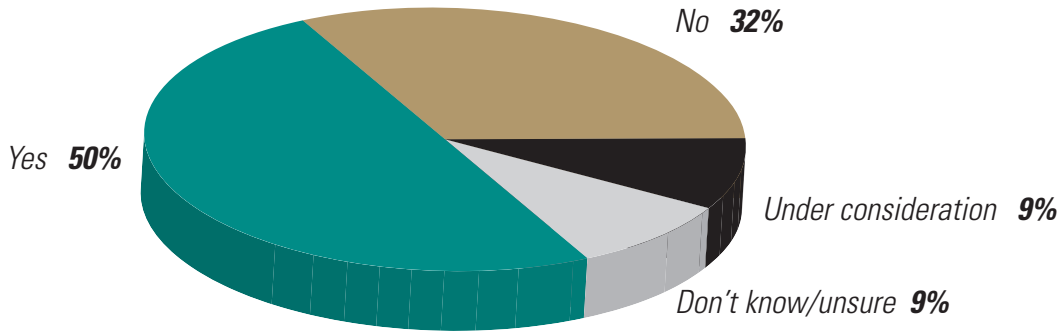

Figure 23: Back Up Data Real-Time (Within 1 Hour) —By Number of Employees

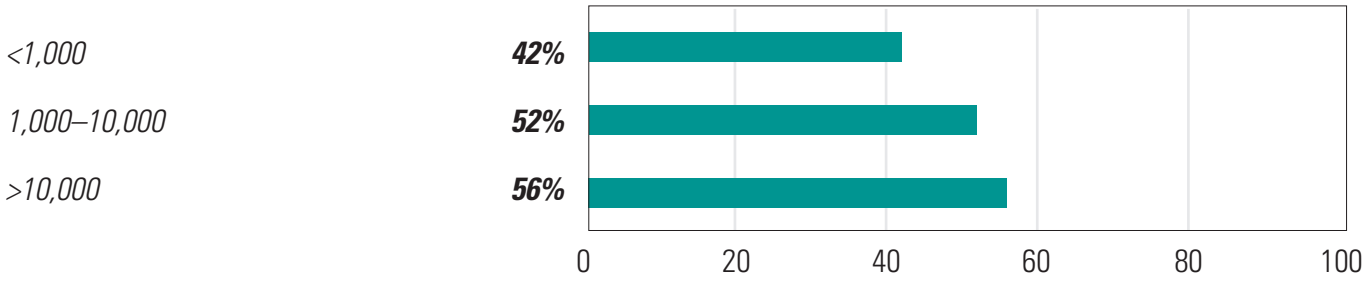

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 24: Percentage of Data Backed Up Real-Time or Near-Real-Time

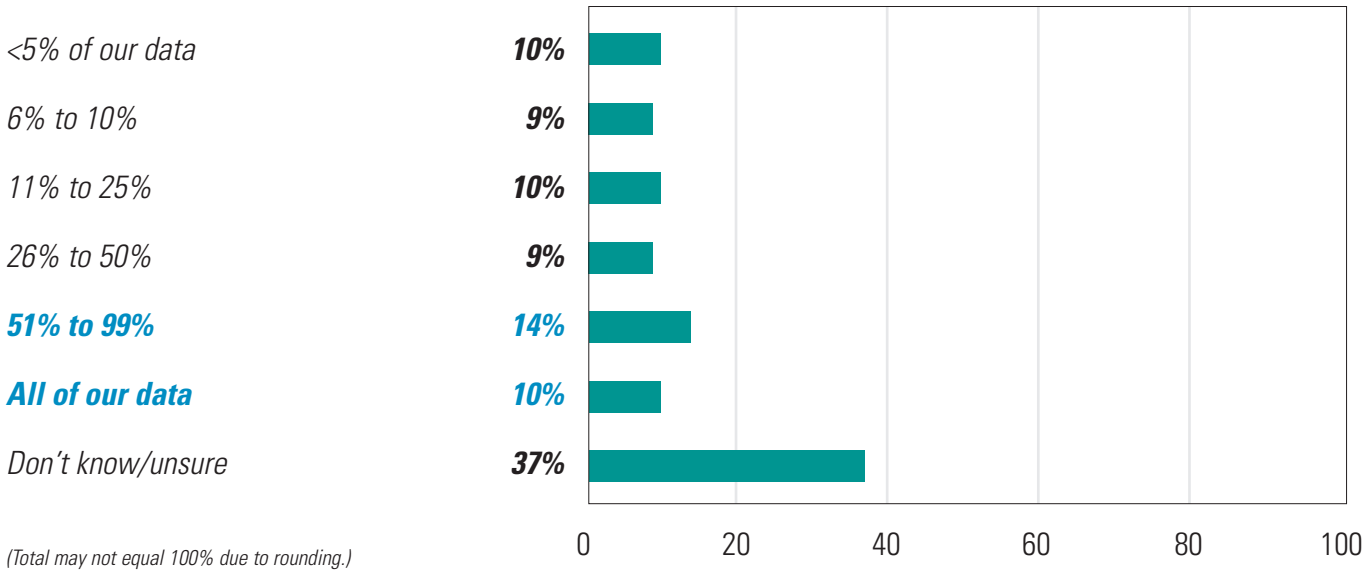

(Total may not equal 100% due to rounding.)

Figure 25: Percentage of Data Backed Up Real-Time or Near-Real-Time—By Data Volume

	"Large-Volume Data" sites	All others
<5% of our data	9%	13%
6% to 10%	14%	10%
11% to 25%	14%	10%
26% to 50%	5%	12%
51% to 99%	12%	17%
All of our data	10%	11%
Don't know/unsure	36%	27%

(Total may not equal 100% due to rounding.)

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 26: Have Common Strategy for Ensuring High Availability Across Databases?

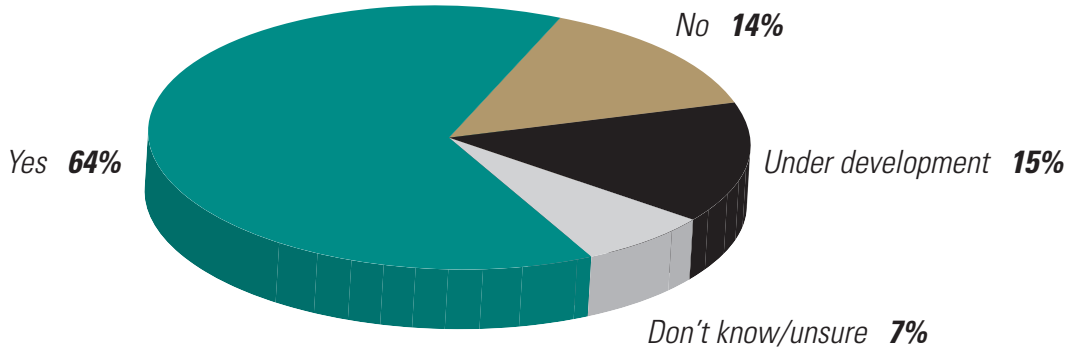

Figure 27: Importance of High Availability Architectures to Achieving SLAs

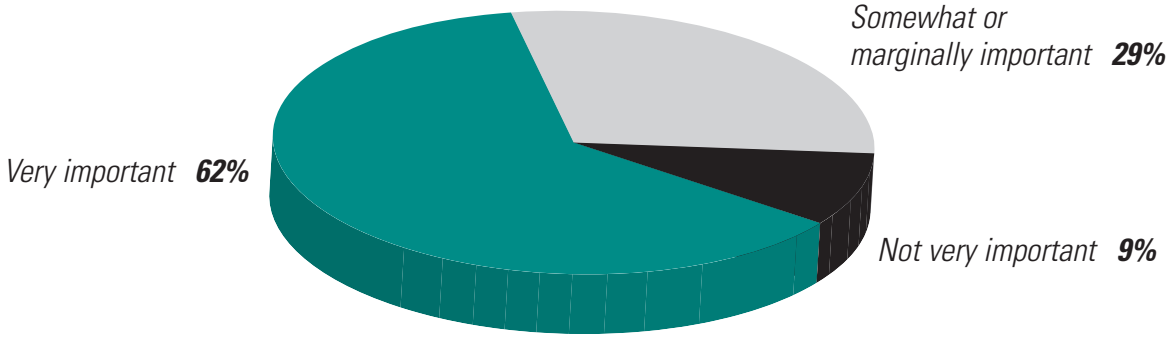

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

CLOUD COMPUTING'S IMPACT ON DATABASE HIGH AVAILABILITY

The shift to cloud services and architectures is adding a new dimension to database high availability. Close to half the organizations surveyed use or plan to use public or private cloud services and architectures, and are likely to seek the same types of service-level agreements seen in their traditional on-premise environments

Do respondents currently use, or plan to use within the next 12 months, any cloud services within their organizations? Close to half, 46%, indicate that they do use or plan to use at least one type of cloud computing approach, led by private cloud, which is being adopted by more than a third. Hybrid clouds—which involve a mix of both private cloud and public cloud services—are the choice of 13% of respondents. Thirteen percent also are employing services from public cloud providers. (See Figure 28.)

Private cloud adoption is more prevalent among larger companies in the survey. At least 38% of the largest organizations (those with more than 10,000 employees) say they have adopted or plan adoption of private cloud architectures. There is stronger movement toward public cloud among the smallest firms versus medium-size companies (16% versus 8%), but large organizations are embracing public cloud as well (14%). (See Figure 29.)

A number of respondents expect service-level agreements forged for their cloud service arrangements to be comparable to existing SLAs for more traditional IT services—such as those with the IT department. Most respondents, 60%, aren't sure how cloud SLAs compare, but among those that are aware, the tendency is to establish SLAs that are on an equal footing with traditional IT environments. (See Figure 30.)

Figure 28: Cloud Services Employed Within Enterprises

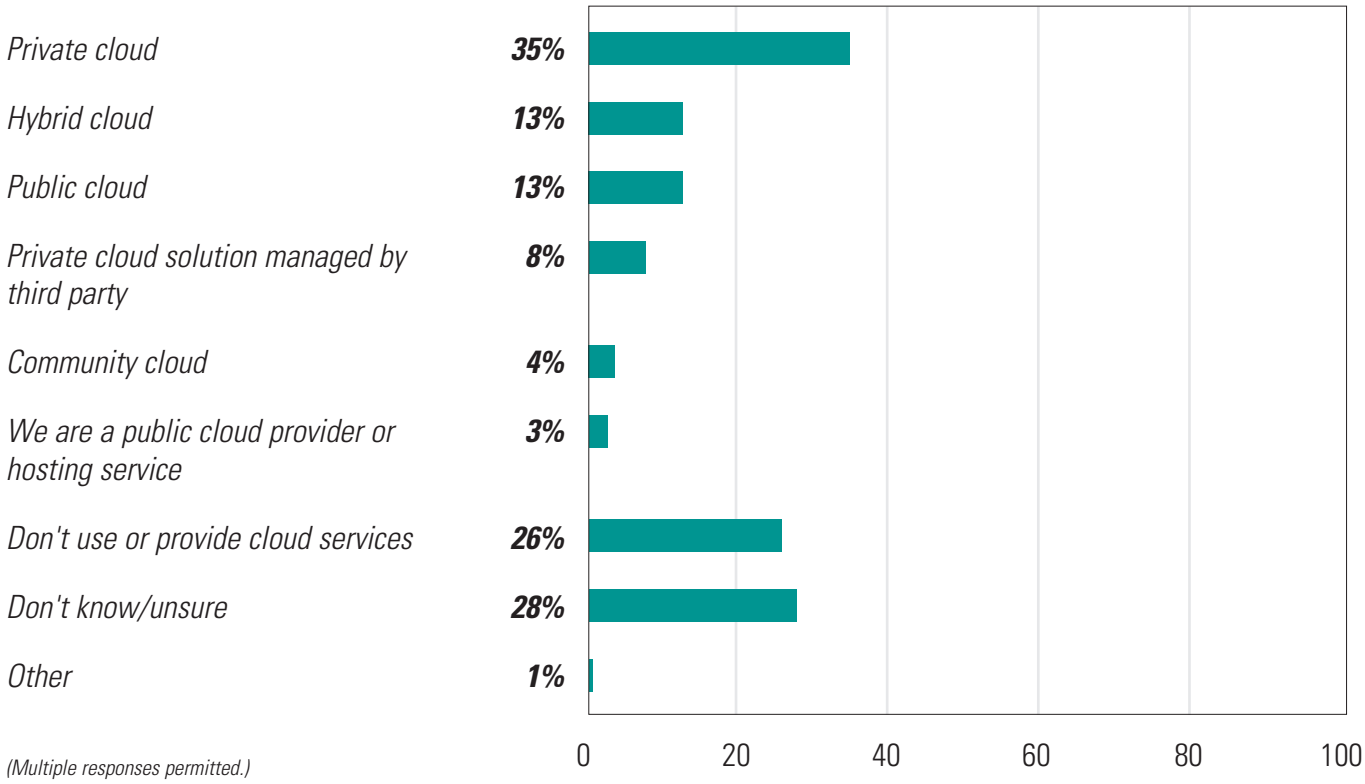

(Multiple responses permitted.)

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 29: Cloud Services Employed Within Enterprises —By Number of Employees

	<1,000	1,000 to 10,000	>10,000
Private cloud	28%	39%	38%
Hybrid cloud	11%	13%	19%
Public cloud	16%	8%	14%
Private cloud solution managed by third party	11%	8%	5%
Community cloud	3%	4%	3%
We are a public cloud provider or hosting service	1%	3%	4%
Don't use or provide cloud services	36%	24%	23%
Don't know/unsure	23%	27%	30%
Other	1%	0%	1%

(Multiple responses permitted.)

Figure 30: Service-Level Agreements Differ for Cloud?

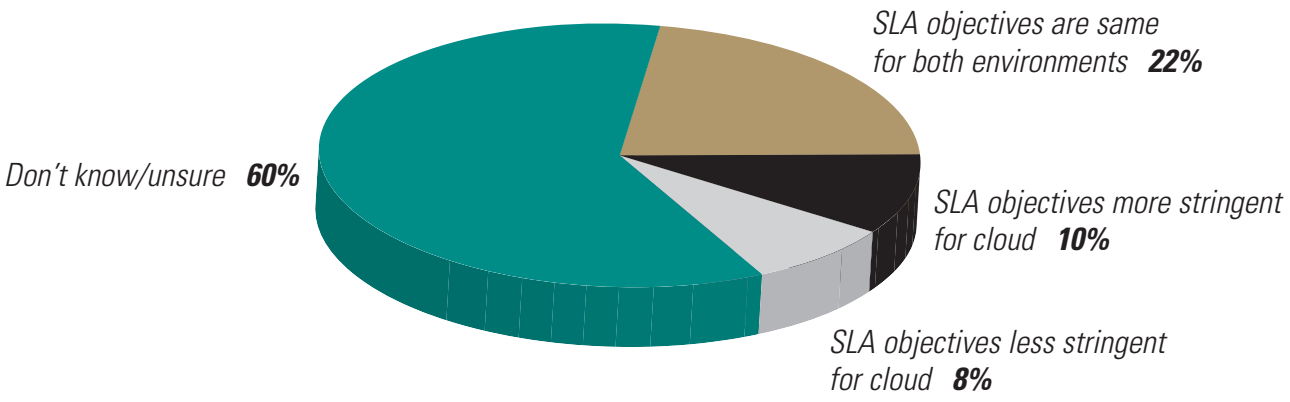

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

ACHIEVING DISASTER RECOVERY

Organizational inertia—combined with the surge in data growth—makes disaster recovery challenging. A solid disaster recovery plan is essential to database high availability, but most respondents do not have formal plans in place to address any potential outages.

Survey respondents were asked to provide their input and experiences regarding recovery, encompassing best practices and platforms that ensure data can be recovered following major outages. While not in the majority, a large segment of respondents' organizations have formal disaster recovery plans in place. Just under half, 47%, say they have plans in place that cover all enterprise data. (See Figure 31.) Coverage of data is most pronounced among the largest organizations in the survey, with plans covering all enterprise data, cited by 53% of respondents in this group. By contrast, only 36% of the smallest firms provide such coverage. (See Figure 32.)

How often are disaster recovery plans reviewed and tested at respondents' sites? For the most part, this testing takes place about once a year. A significant portion of respondents, 18%, report that their organizations conduct no such testing at all. (See Figure 33.)

What are respondents' greatest challenges in planning and managing a recovery strategy for their data assets at this time? Both organizational and technical issues are leading factors on the list. Forty-three percent cite funding and budget constraints, while 42% say they consider managing data growth to be their most vexing challenge. Network bandwidth is still an issue for many respondents, cited by more than a third. (See Figure 34.)

The challenges of disaster recovery planning are the same, no matter how much data an organization maintains—with one exception. For large-volume data sites, backing up and managing increasingly large quantities of data represents the biggest hurdle. Sixty percent cite this as their greatest challenge, compared with 40% of smaller sites. (See Figure 35.)

The distance to a remote data center or backup site is also another important consideration in data availability strategies. More than four-fifths of respondents, 83%, have remote data centers, and 43% indicate these sites are more than 100 miles from their primary data centers. (See Figure 36.)

In the event of an unplanned outage, how long would it take for all systems and data to be restored and available to respondents' organizations' end users? More than a third of respondents, 34%, estimate that they would need more than a business day (exceeding eight hours) to have everything up and running. Another one-third say things would be up and running within the day, and 15% even say they could do a full restore within one hour's time. (See Figure 37.) Large-volume data sites may be slightly more delayed in getting things back up and running than their counterparts with smaller data stores—40% of large-volume data sites would take more than a day, versus 35% of small data sites. (See Figure 38.)

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 31: Have Formal Business Continuity/Disaster Recovery Plan in Place?

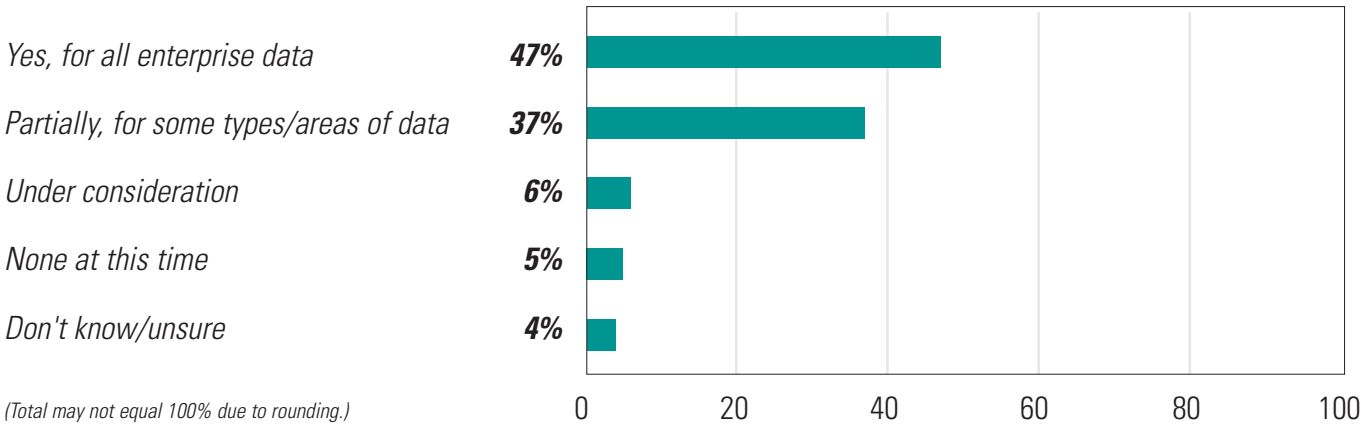

(Total may not equal 100% due to rounding.)

Figure 32: Have Formal Business Continuity/Disaster Recovery Plan in Place—By Number of Employees

	<1,000	1,000 to 10,000	>10,000
Yes, for all enterprise data	**36%**	**50%**	**53%**
Partially, for some types/areas of data	**42%**	**36%**	**34%**

Figure 33: How Often Business Continuity/Disaster Recovery Plan Reviewed and Tested?

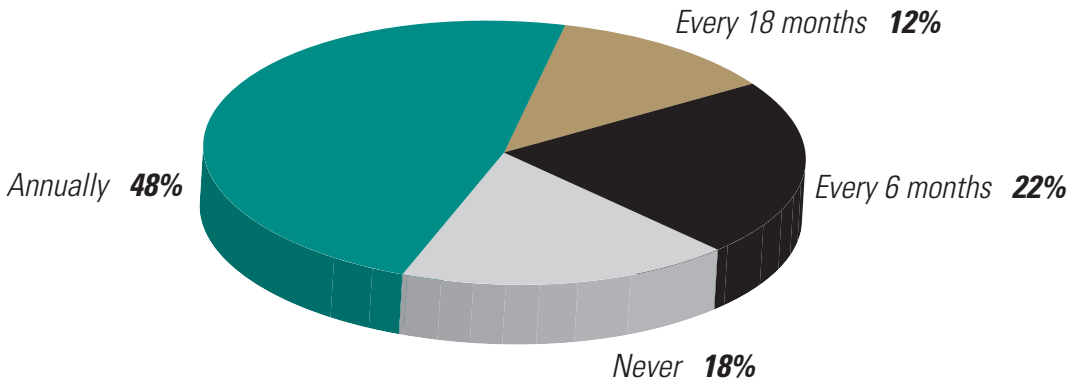

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 34: Greatest Challenges to Planning and Managing Business Continuity/Disaster Recovery

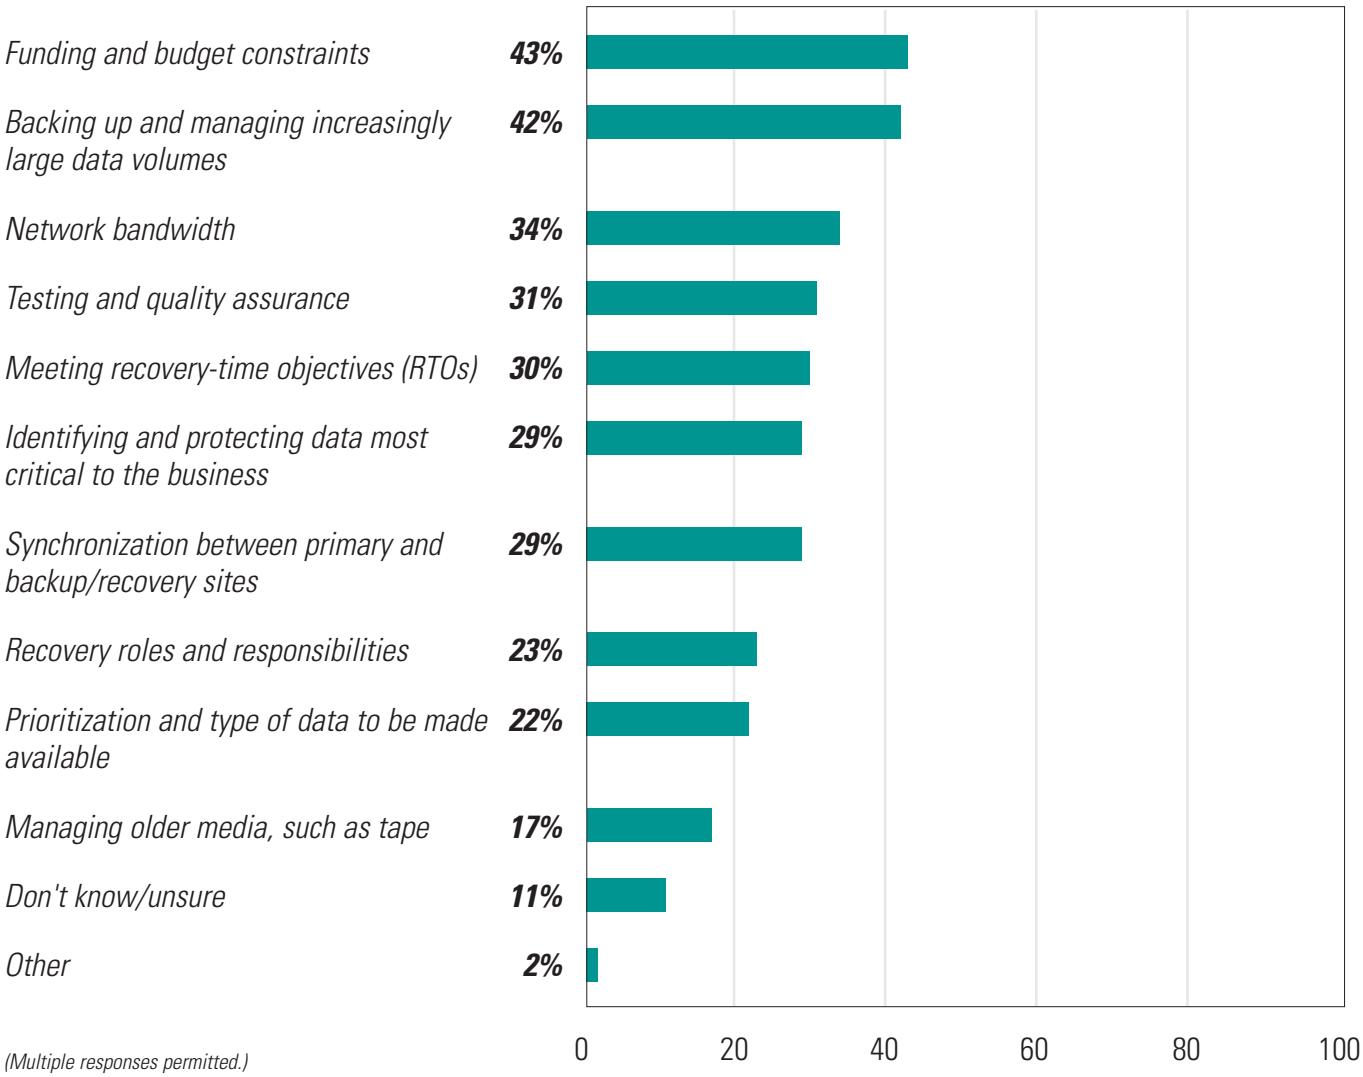

(Multiple responses permitted.)

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 35: Greatest Challenges to Planning and Managing Business Continuity/Disaster Recovery—By Data Volume

	"Large-Volume Data" sites	All others
Funding and budget constraints	42%	43%
Backing up and managing increasingly large data volumes	60%	40%
Network bandwidth	37%	37%
Testing and quality assurance	33%	30%
Meeting recovery-time objectives (RTOs)	27%	33%

(Multiple responses permitted.)

Figure 36: Distance Between Primary and Remote Data Centers

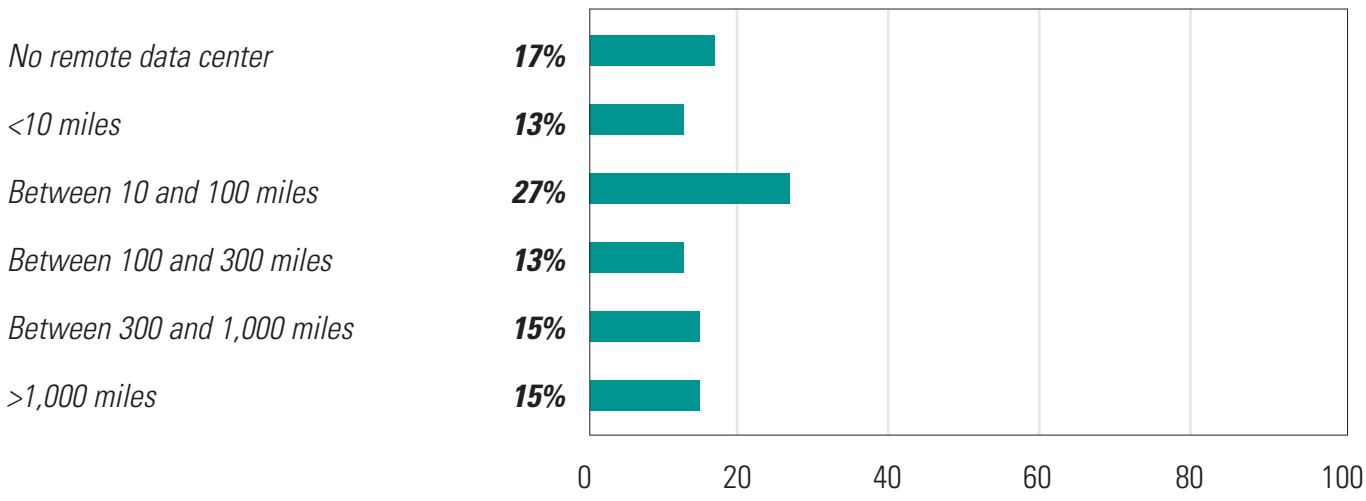

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 37: How Long Until All Systems and Data Restored from an Unplanned Outage?

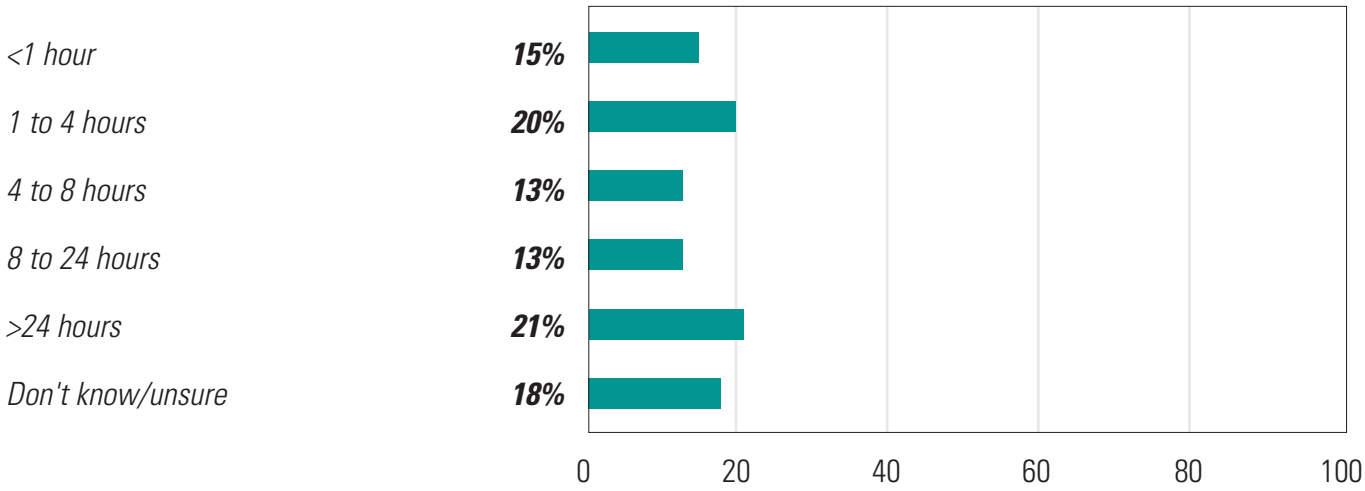

Figure 38: How Long Until All Systems and Data Restored from an Unplanned Outage—By Data Volume

	"Large-Volume Data" sites	All others
<1 hour	22%	14%
1 to 4 hours	18%	21%
4 to 8 hours	8%	17%
8 to 4 hours	15%	14%
>24 hours	25%	21%
Don't know/unsure	12%	13%

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

IOUG RECOMMENDS

Today's 24x7x365 business environment demands that data be available and accessible on a moment's notice. Organizations competing in this environment need to address the length of time it takes to get things back up and running after a systems disruption, as well as the business costs of these delays. IOUG recommends the following approaches to meet the burgeoning needs for availability in fast-growing and increasingly complex data environments:

Measure the length and costs of downtime. Both planned and unplanned downtime have costs, and it's important to understand the extent to which downtime impacts the business. In addition, efforts to address and reduce downtime depend on the full awareness of where the bottlenecks are.

Develop an integrated information management lifecycle strategy. Some data needs to reside online and be quickly accessible to end users, while other forms of data can be less accessible and stored in archive or backup systems. Data should be moved to various tiered storage systems as part of an organization's entire information lifecycle management, ensuring a strategy for final retirement or archiving of data.

Explore new data availability options. A range of technologies and approaches are available to help address downtime, ranging from server clustering to storage mirroring to virtualization and cloud-based solutions.

Get business buy-in and support. Data availability solutions only deliver value if they have the support and input of the business. The business needs to determine what data needs to be made highly availabile, and what data isn't as essential.

Make sure the organization is ready. While preparing data to be available is essential to business continuity, it is far from the total picture. People need to be ready to use the data and assume responsibilities for lost data for business continuity to be achieved.

The findings from this latest IOUG member survey show that at about half of companies are working to provide their enterprises with data on a real-time or near-real-time basis. The challenge of maintaining this level of online data access is increasing as Oracle and other database sites are being inundated with large volumes of enterprise data, both structured and unstructured. This increase in volume, velocity, and variety of data not only presents management and storage challenges, but also offers great opportunities for businesses to expand their capability to innovate, better serve customers, and make more insightful decisions.

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

DEMOGRAPHICS

Figure 39: Respondents' Job Titles

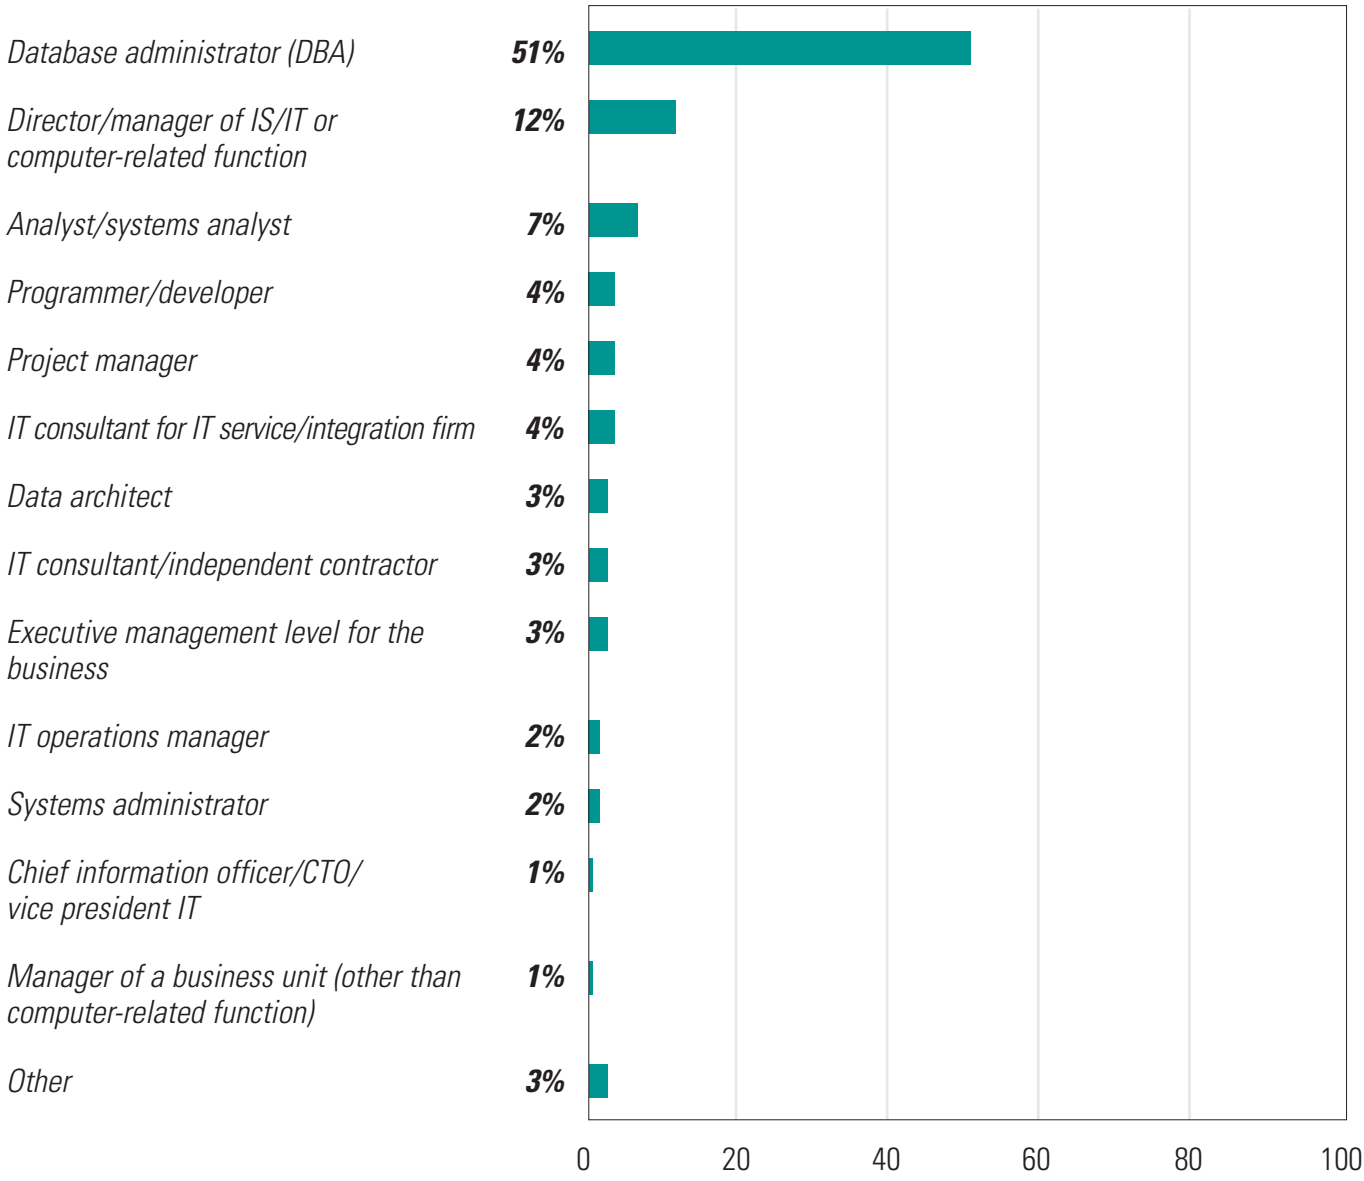

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 40: Respondents' Company Size—By Number of Employees

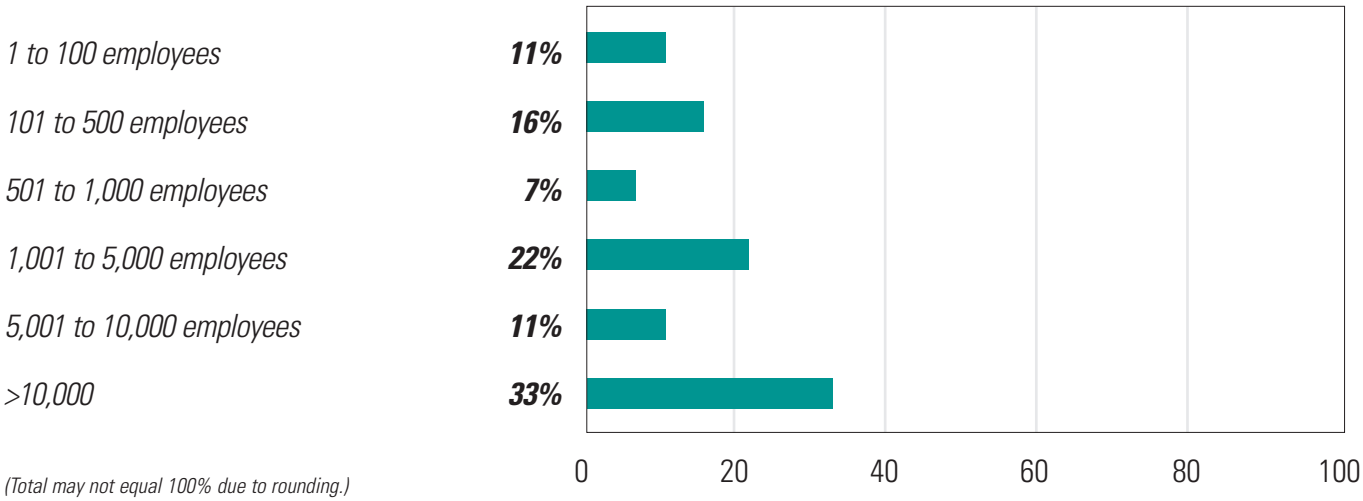

(Total may not equal 100% due to rounding.)

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.

Figure 41: Respondents' Primary Industries

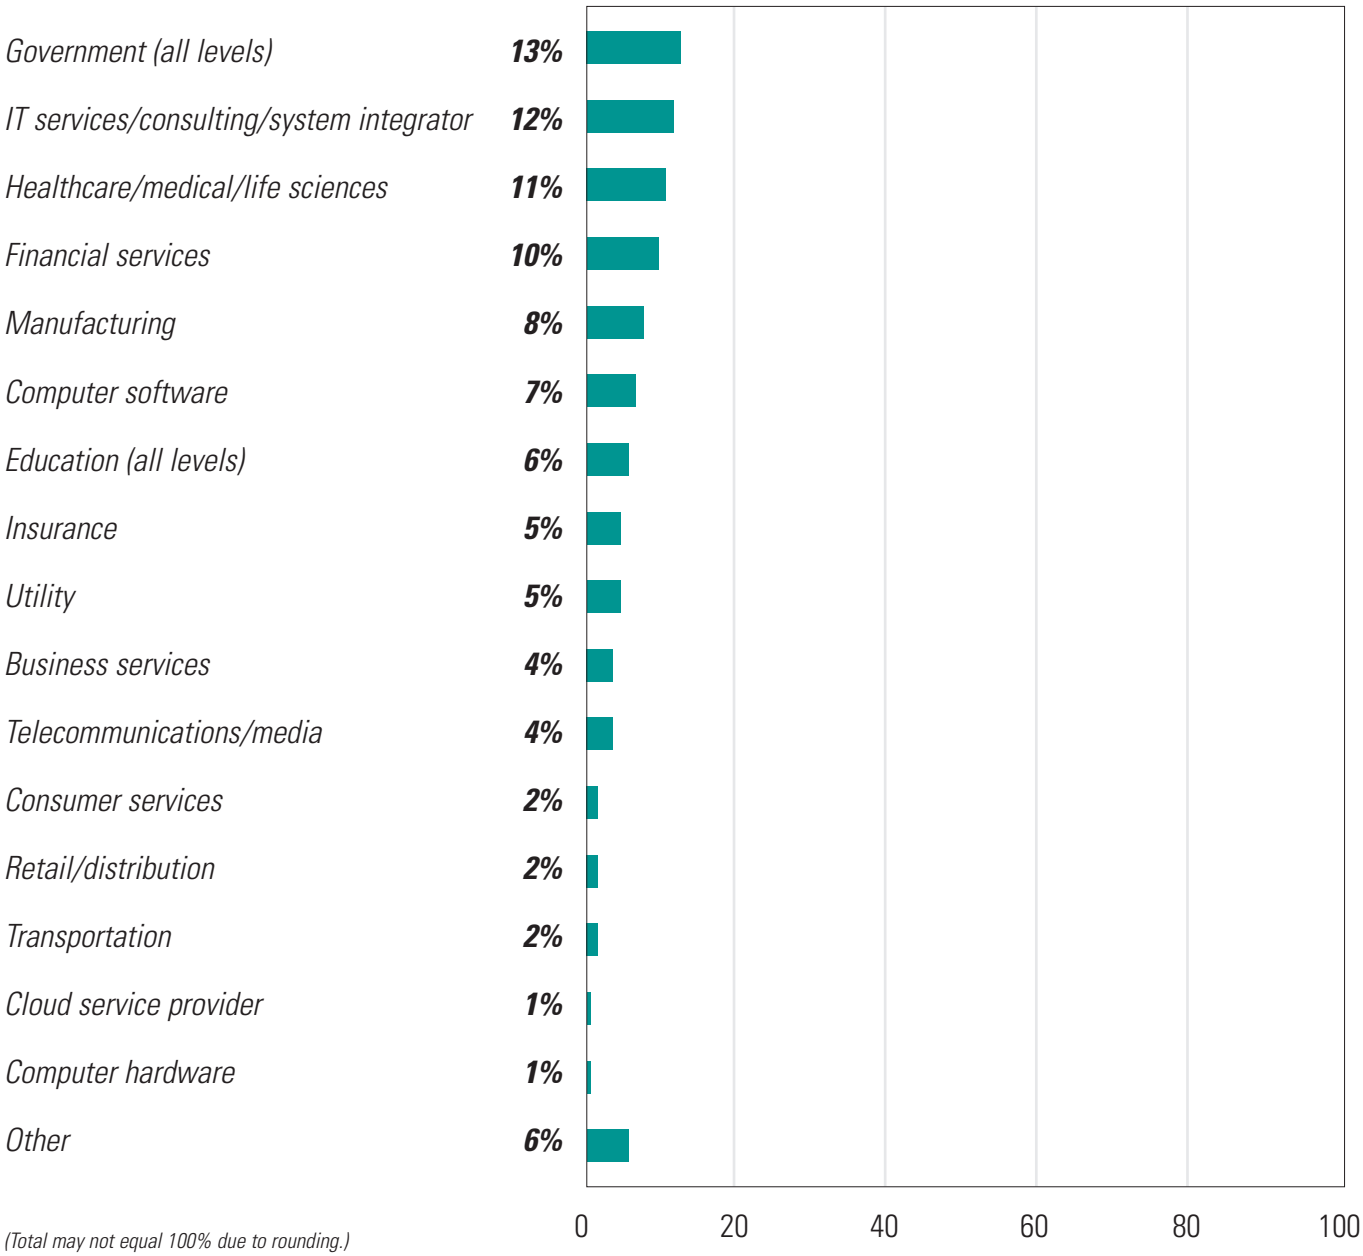

(Total may not equal 100% due to rounding.)

ENTERPRISE DATA AND THE COST OF DOWNTIME, 2012 IOUG Database Availability Survey *was produced by Unisphere Research and sponsored by Oracle. Unisphere Research is the market research unit of Unisphere Media, a division of Information Today, Inc., publishers of* Database Trends and Applications *magazine and the* 5 Minute Briefing *newsletters. To review abstracts of our past reports, visit www.dbta.com/About_Us#Unisphere. Unisphere Media, 630 Central Avenue, Murray Hill, New Providence, NJ 07974; 908-795-3701, Email: TWilson@infotoday.com; Web: www.dbta.com.*

Data collection and analysis performed with SurveyMethods.