2013 Enterprise Service Availability Benchmark Survey

Published by

Executive Summary

This benchmark survey presents service availability metrics and related corporate practices that allow IT executives to compare their organization's performance and practices to their peers.

The results presented here are based on responses from 147 professionals from a wide range of industries and geographies collected through an online survey.

Some of the key findings of the survey include:

- **Most organizations have commendable service availability goals for mission-critical systems:** 90% have an availability goal of >99.76% (less than 22 hours of downtime a year), 68% have a goal of higher than 99.91% availability (less than 8 hours of downtime a year), and 27% have a goal of 99.99% availability (less than one hour of downtime a year).

- **However, 32% of the organizations surveyed did not meet their service availability goals for mission-critical systems in 2012.**

- The lowest downtime figures were reported by respondents from the telecommunication industry (3.7 hours for the year), followed by high-tech (4.0) and business services organizations (4.2). Highest downtime figures were reported by respondents from healthcare (14.1), financial services organizations (8.9) and public sector organizations (8.5).

- 71% of the respondents run at least some mission-critical applications in a private cloud environment, while 33% run such applications in the public cloud. At the same time, **less than half of these organizations (47%) have a disaster recovery solution for their cloud systems** and **40% have no process in place to test the reliability of their private cloud systems.**

Table of Contents

What Drives Service Availability Goals?

Ensuring customer satisfaction is the most common business driver for service availability goals, cited by 56% of the respondents.

Avoiding productivity loss was mentioned by just over a quarter (26%) of respondents.

56%
Ensuring customer satisfaction

26%
Avoiding productivity loss

10%
Regulatory compliance

2%
Keeping up with the competition

6%
Other
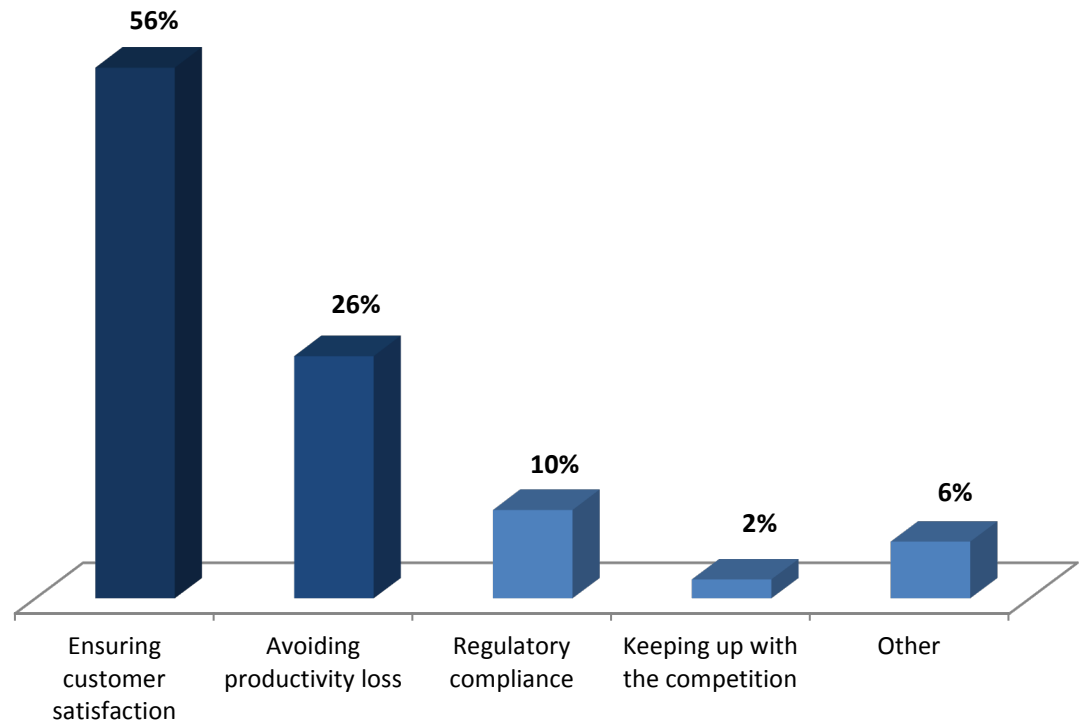
Strategies and Tools to Ensure Service Availability

The most common strategy for ensuring service availability is replication to a disaster recovery site, used by more than three-quarters (77%) of respondents.

Restore from backup and virtualization HA were mentioned by just shy of two-thirds of the respondents (64% and 63% respectively).

Strategy	Percentage
Replication to disaster recovery site	77%
Restore from backup	64%
Virtualization HA	63%
High availability clusters – Others	54%
High availability clusters – Unix	50%
VMware SRM	41%
DR as a service	26%
Other	2%
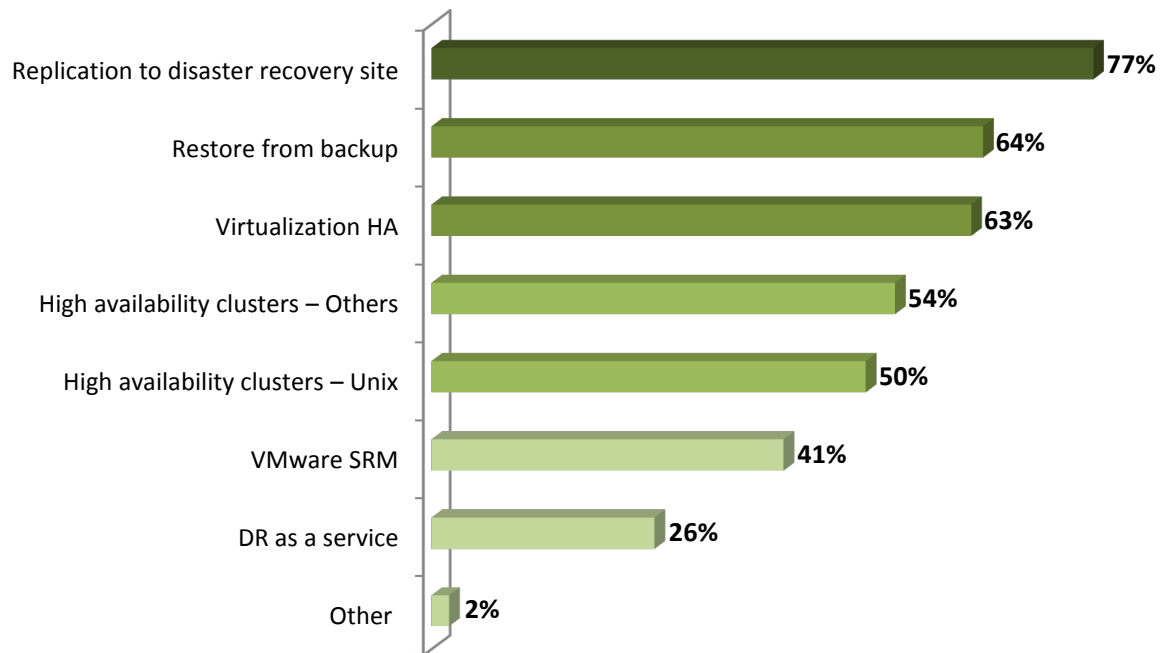
Respondents could select more than one option, results add up to more than 100%.

How Effective Are these Strategies and Tools?

According the survey results, replication is also the most effective strategy for ensuring service availability.

Restore from backup, although a popular strategy, is viewed as less effective than high availability clusters and virtualization.

Strategy	Rating
Replication to disaster recovery site	4.28
Restore from backup	3.38
Virtualization HA	4.13
High availability clusters - Others	4.09
High availability clusters - Unix	4.25
VMware SRM	3.91
DR as a service	3.63
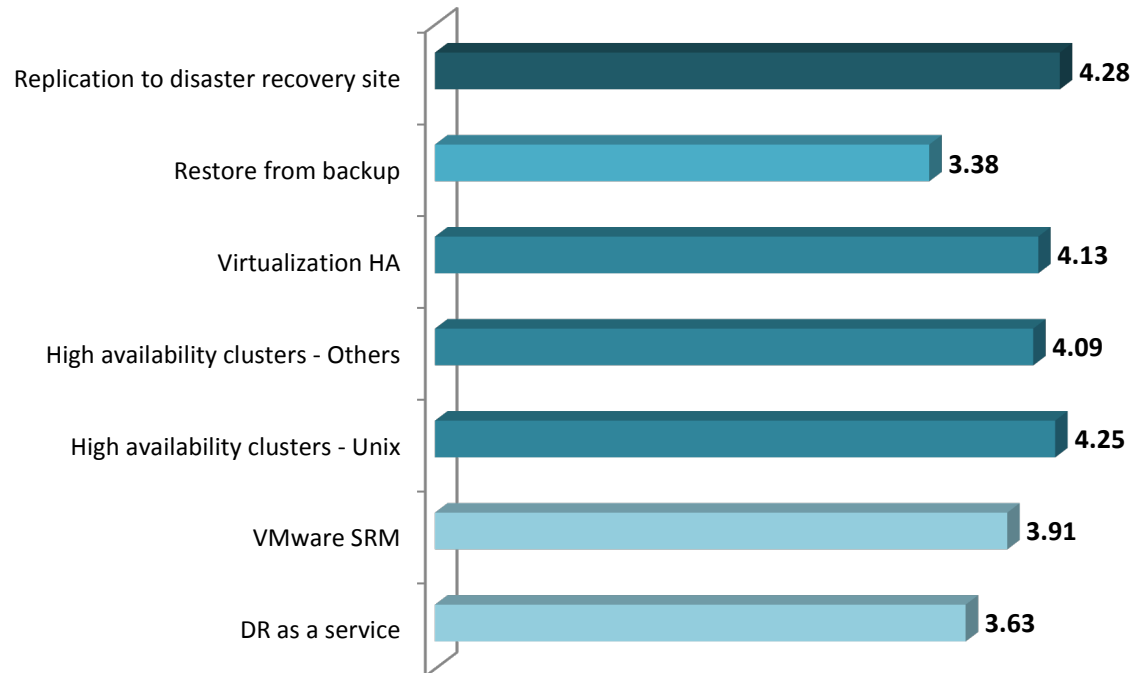
Top Challenges in Ensuring Service Availability
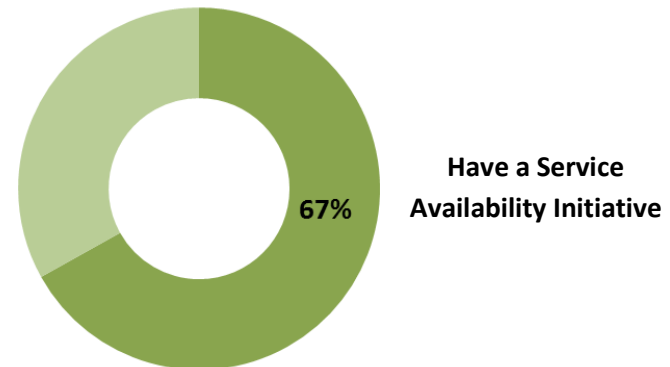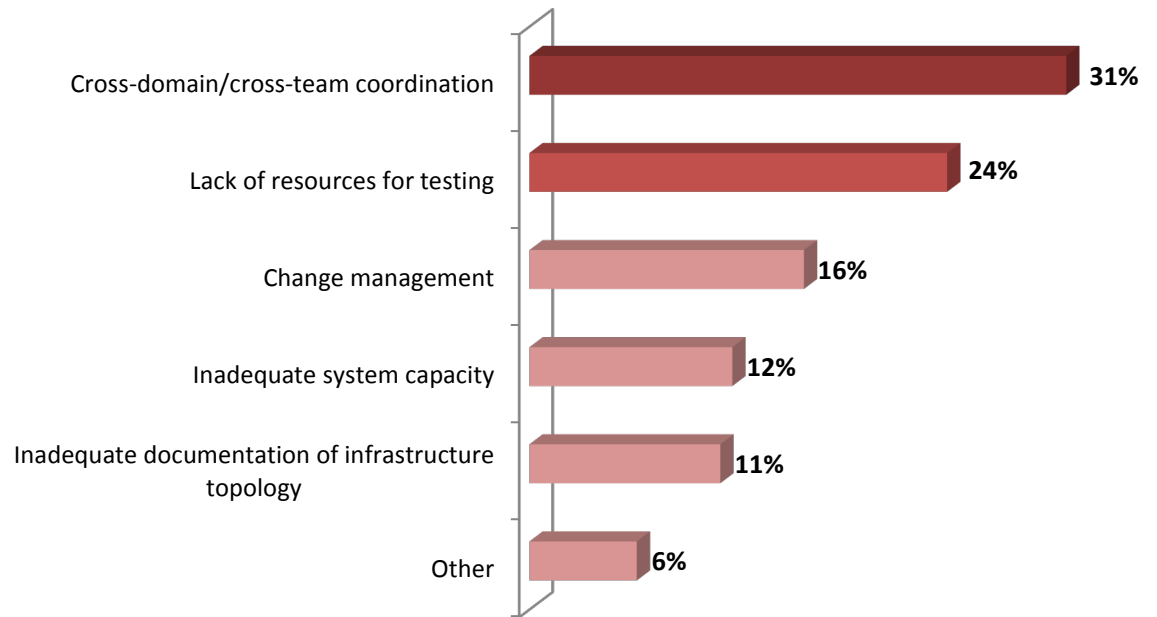
When asked to rate the top challenge in ensuring service availability, **close to a third of the respondents (31%) mentioned cross-domain and cross-team coordination**. Close to a quarter (24%) cited lack of resources for testing.

**Two third of the respondents (67%) indicated they have initiatives for improving service availability management in 2013.**

Cross-domain/cross-team coordination — 31%
Lack of resources for testing — 24%
Change management — 16%
Inadequate system capacity — 12%
Inadequate documentation of infrastructure topology — 11%
Other — 6%

67% — Have a Service Availability Initiative

Service Availability Goals

The vast majority of the organizations (90%) have a service availability goal of over 99.76% (less than 22 hours of downtime a year) for mission critical systems.

Over two thirds of the organizations (68%) have a goal of higher than 99.91% availability (less than 8 hours of downtime a year), and **27% have a goal of 99.99% availability (less than one hour of downtime a year).**

41%	
27%	
	22%
	7%
	4%

> 99.99% (less than one hour of downtime a year) | 99.91-99.99% (less than 8 hours of downtime a year) | 99.76-99.90% (less than 22 hours of downtime a year) | 99.50-99.75% (less than 44 hours of downtime a year) | 99.50>
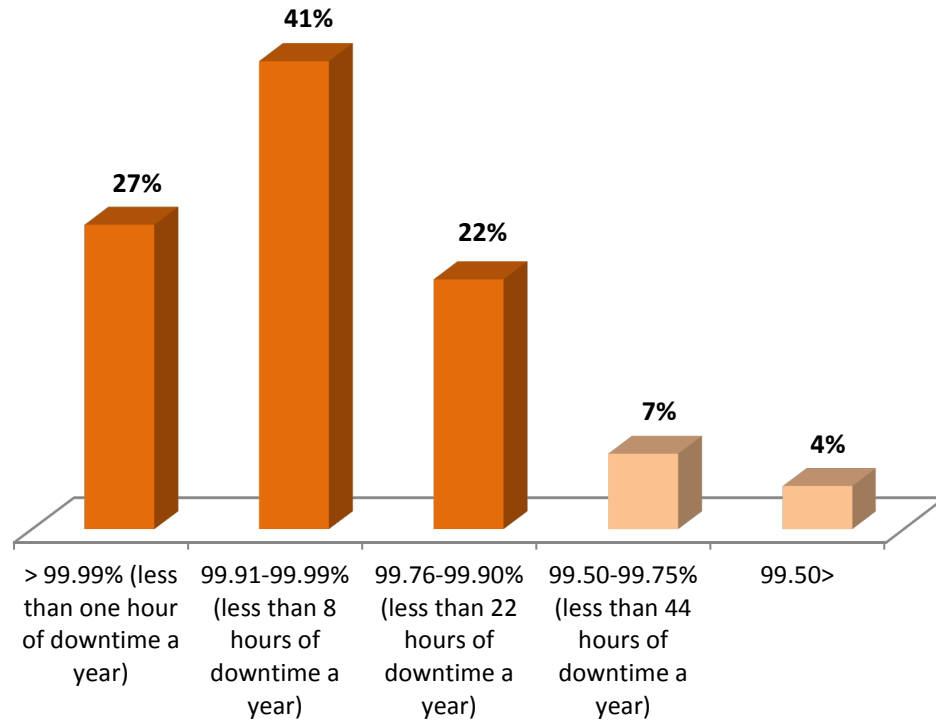
Service Availability: The Reality

By their own account, close to a third of the organizations surveyed (32%) did not meet their service availability goals for mission-critical systems in 2012.

Organizations with higher services availability goals had lower rates of achieving their goals.

Met goal ■ Did not meet goal

Did Not Meet Goal
Met Goal

Service Availability Goals (Mission Critical Systems)	Met Goal	Did Not Meet Goal
> 99.99%	46%	54%
99.91-99.99%	71%	29%
99.76-99.90%	82%	18%
99.50-99.75%	86%	14%

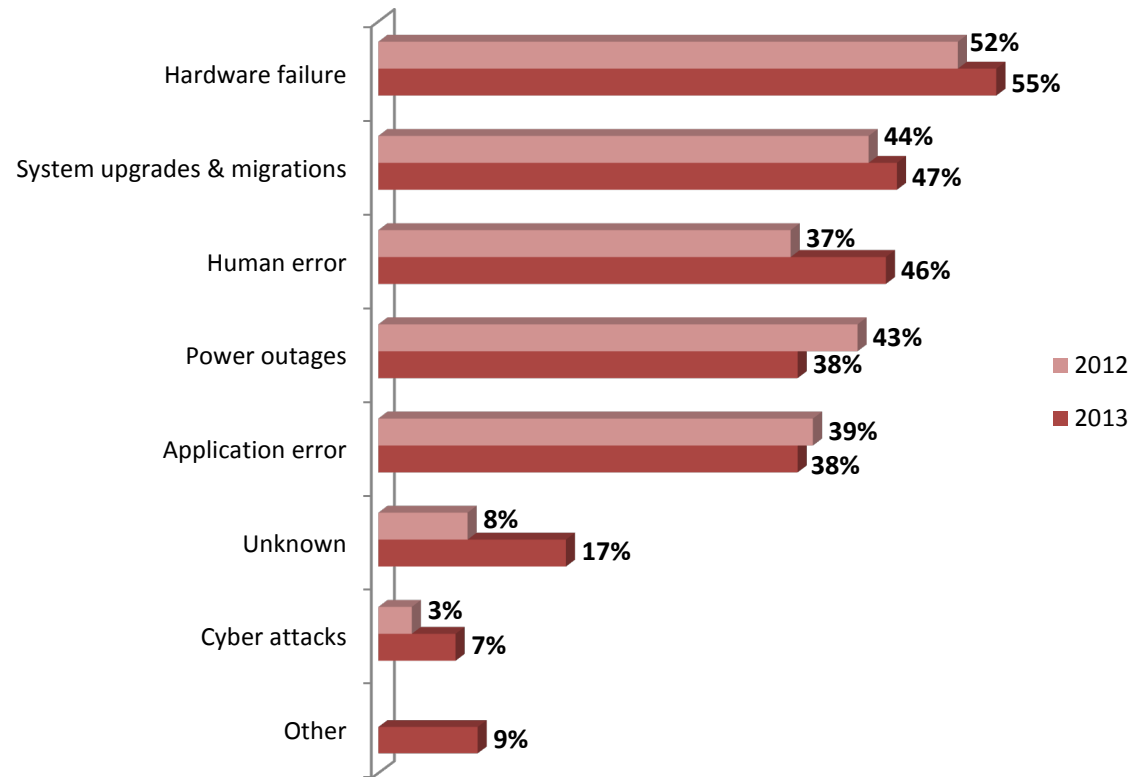**Service Availability Goals (Mission Critical Systems)**
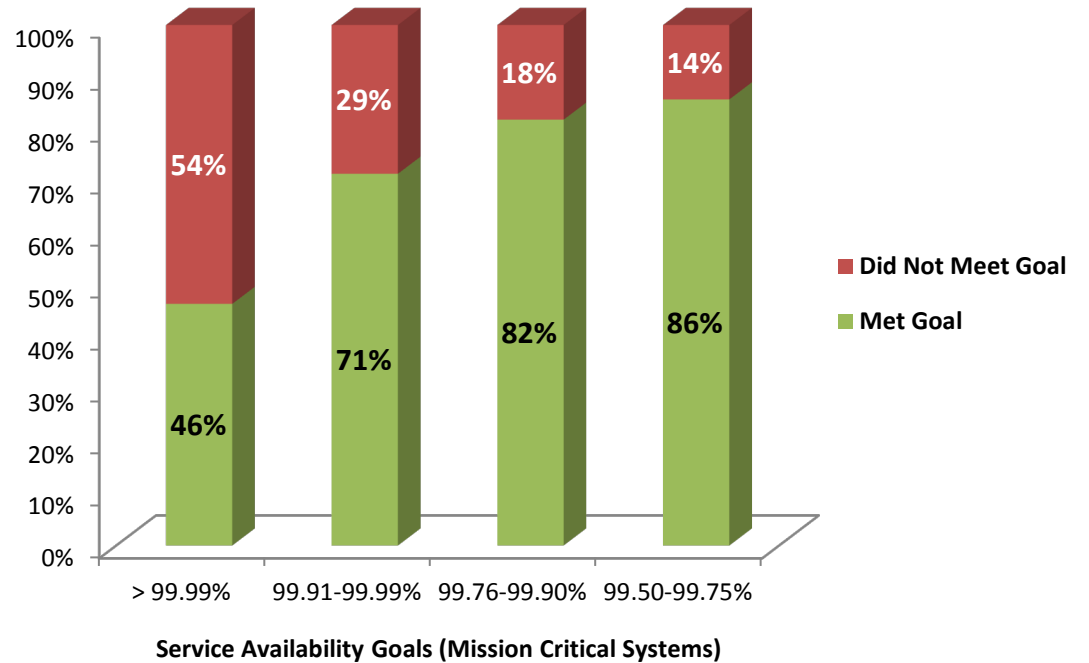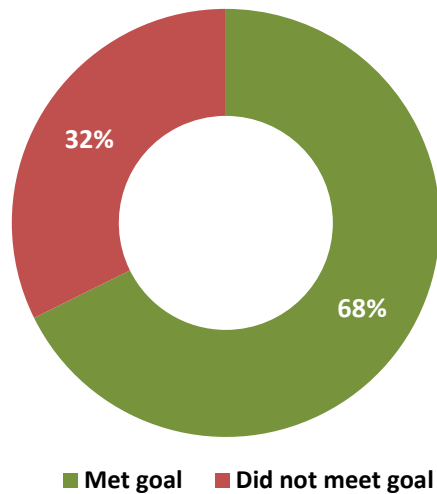
What Were the Causes for Outages?

Hardware failures were the most common reason for system outage, cited by 55% of the respondents.

Other leading causes include upgrades & migrations (47%), human error (46%), power outages (38%), and application error (38%).

2013 results show a significant increase in human error as a cause of outages compared to 2012.

Hardware failure
- 2012: 52%
- 2013: 55%

System upgrades & migrations
- 2012: 44%
- 2013: 47%

Human error
- 2012: 37%
- 2013: 46%

Power outages
- 2012: 43%
- 2013: 38%

Application error
- 2012: 39%
- 2013: 38%

Unknown
- 2012: 8%
- 2013: 17%

Cyber attacks
- 2012: 3%
- 2013: 7%
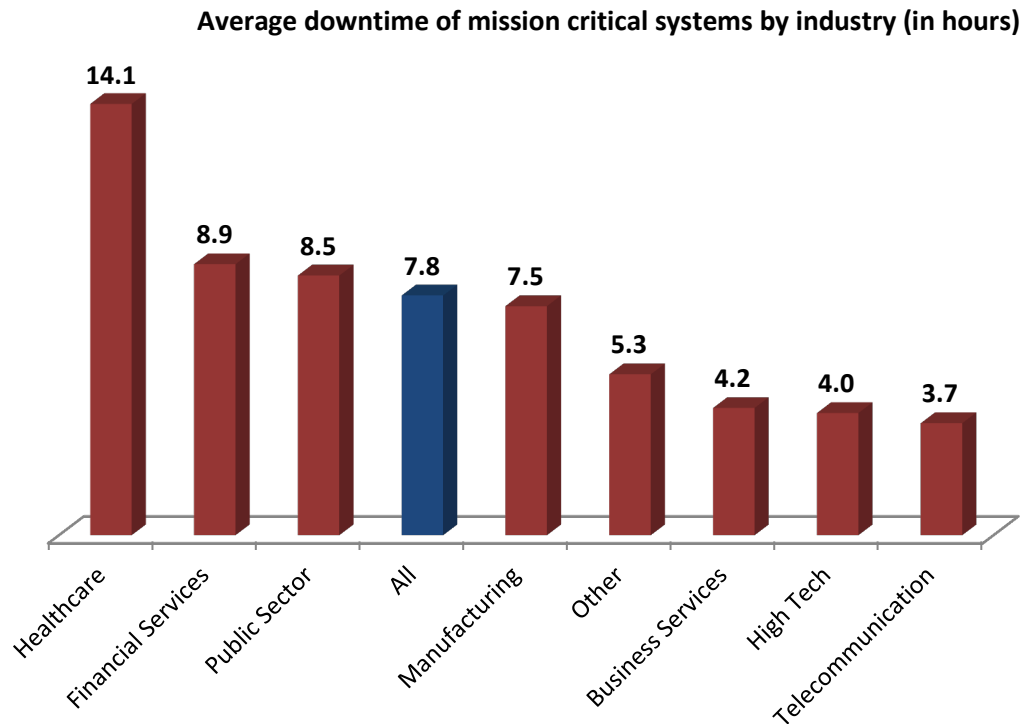
Other
- 2013: 9%

2012
2013

Multiple answers per respondent, results add up to more than 100%

Which Industries Experience the Most Downtime?

The average downtime reported for mission-critical systems across all industries was 7.8 hours.

The lowest average downtime was reported by respondents from the telecommunication industry (3.7 hours for the year), followed by high-tech (4.0) and business services organizations (4.2).

Highest downtime figures were reported by respondents from healthcare (14.1), financial services organizations (8.9) and public sector organizations (8.5).

Average downtime of mission critical systems by industry (in hours)

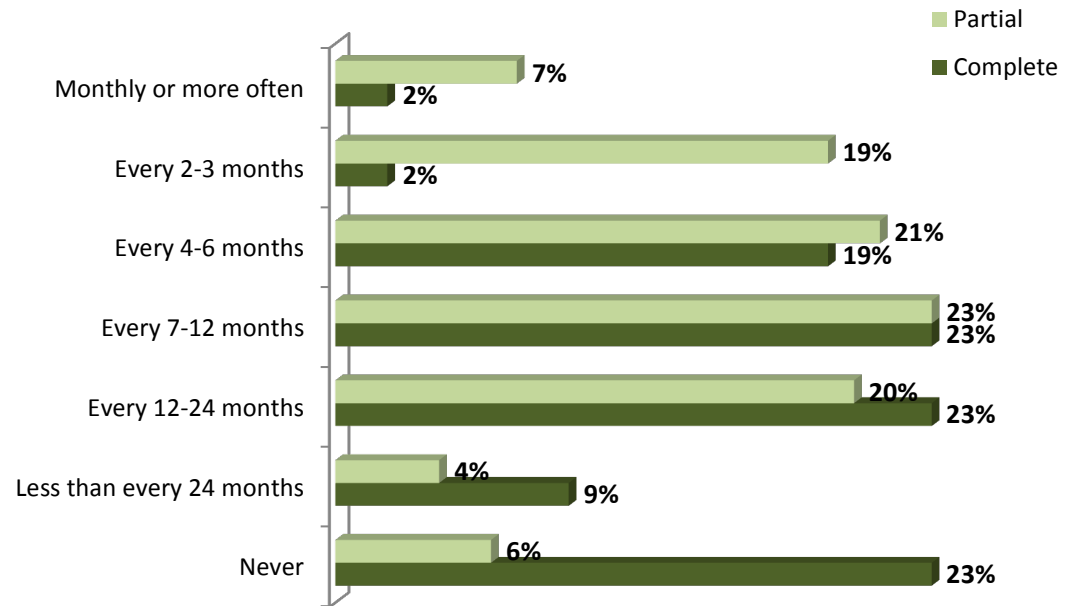| Industry	Hours
Healthcare	14.1
Financial Services	8.9
Public Sector	8.5
All	7.8
Manufacturing	7.5
Other	5.3
Business Services	4.2
High Tech	4.0
Telecommunication	3.7

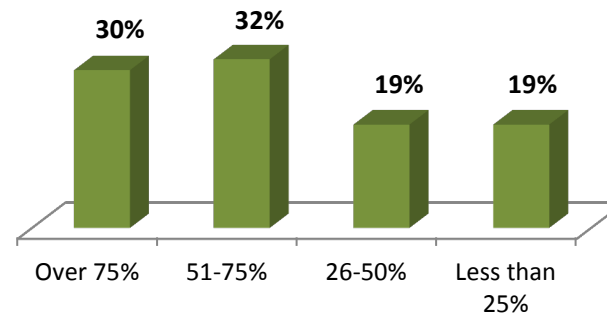# How Often Does Disaster Recovery Testing Take Place?

Most organizations surveyed (70%) conduct a partial disaster recovery test at least once a year.

At the same time, less than half the organizations (46%) conduct a complete test annually or more often, and **as many as 23% don't ever conduct a complete test.**

Partial tests cover over three-quarters of the business services footprint in 30% of the organizations and over half of the business services in 62% of the organizations surveyed.

Partial ■
Complete ■

	Partial	Complete
Monthly or more often	7%	2%
Every 2-3 months	19%	2%
Every 4-6 months	21%	19%
Every 7-12 months	23%	23%
Every 12-24 months	20%	23%
Less than every 24 months	4%	9%
Never	6%	23%

Multiple answers per respondent, results add up to more than 100%

Over 75%	51-75%	26-50%	Less than 25%
30%	32%	19%	19%

What Are the Top Challenges in Disaster Recovery Testing?

Lack of resources and risk to service availability are the top two challenges to disaster recovery testing, cited by 51% and 50% of survey respondents respectively.

Other challenges include inadequate documentation of disaster recovery components and topology (29%) and inadequate system capacity for testing (21%).

Challenge	Percentage
Lack of resources for testing	51%
Risk to service availability	50%
Inadequate documentation of disaster recovery components and topology	29%
Inadequate system capacity	21%
Other	15%
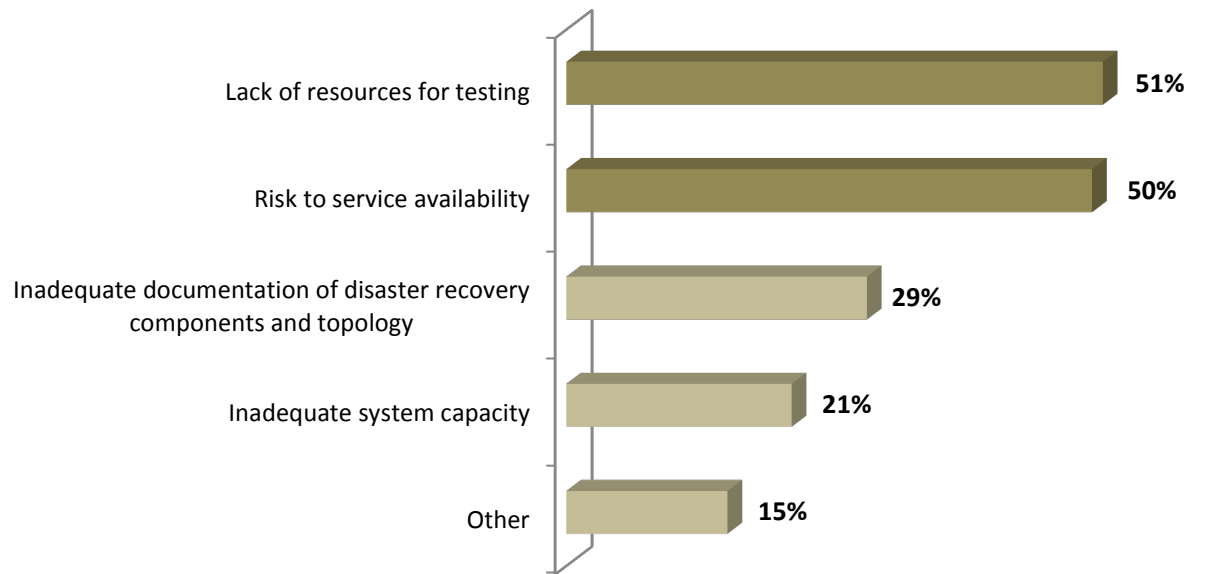
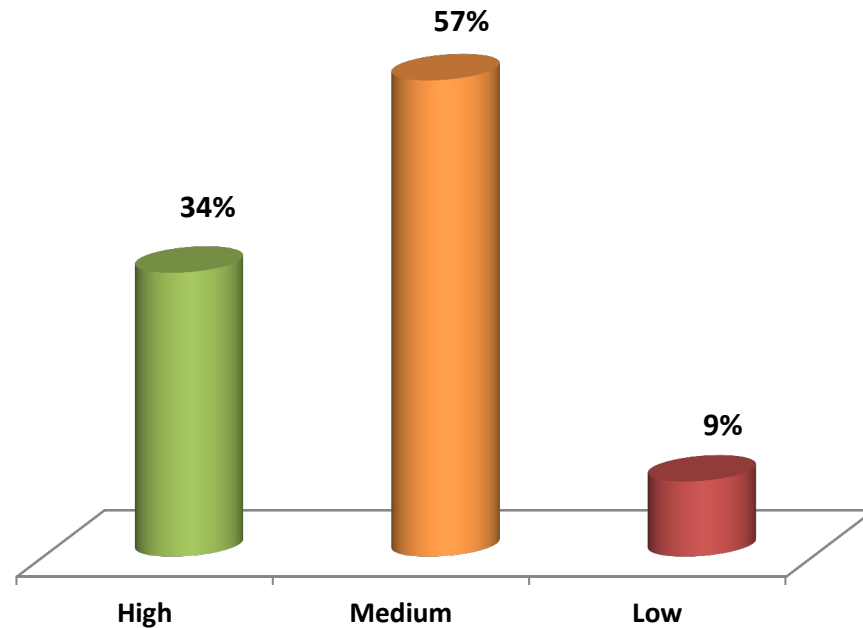*Multiple answers per respondent, results add up to more than 100%*

Confidence in Disaster Recovery Testing

Just over a third of the respondents (34%) have high level of confidence in the completeness and accuracy of disaster recovery testing.
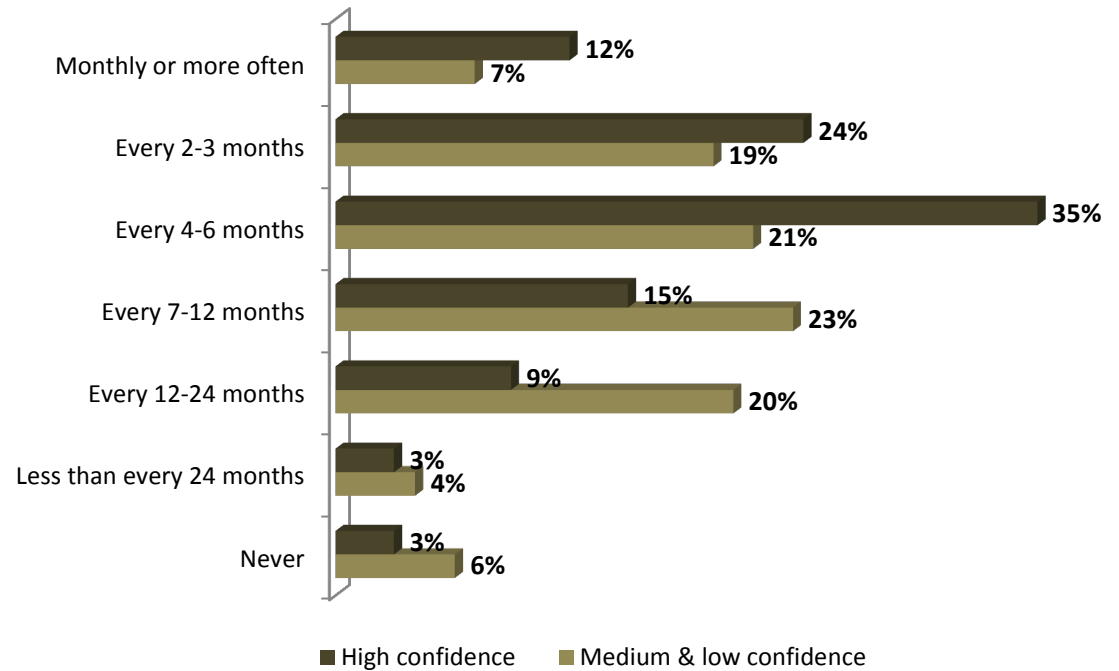
Over a half of the survey respondents (57%) describe their confidence as medium, while 9% have low confidence in their testing.

Organizations with High Confidence in Disaster Recovery Testing Conduct Testing More Often
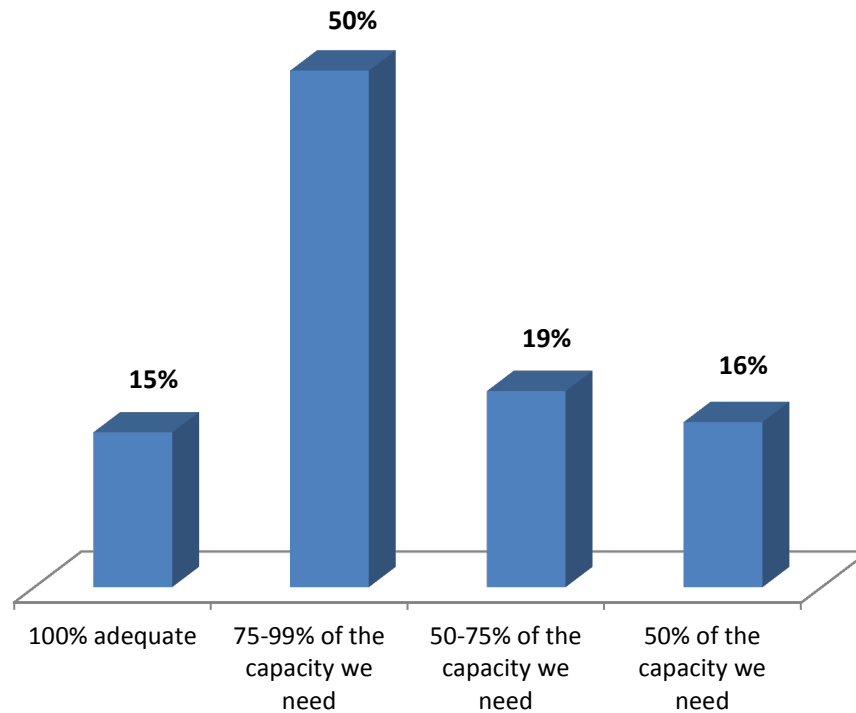
Close to three-quarters (71%) of organizations with high confidence in disaster recovery testing conduct a partial test at least once every six months.

Only 47% of organizations with medium and low confidence in disaster recovery testing conduct a partial test once in six months.

Monthly or more often — 12% / 7%
Every 2-3 months — 24% / 19%
Every 4-6 months — 35% / 21%
Every 7-12 months — 15% / 23%
Every 12-24 months — 9% / 20%
Less than every 24 months — 3% / 4%
Never — 3% / 6%

■ High confidence   ■ Medium & low confidence
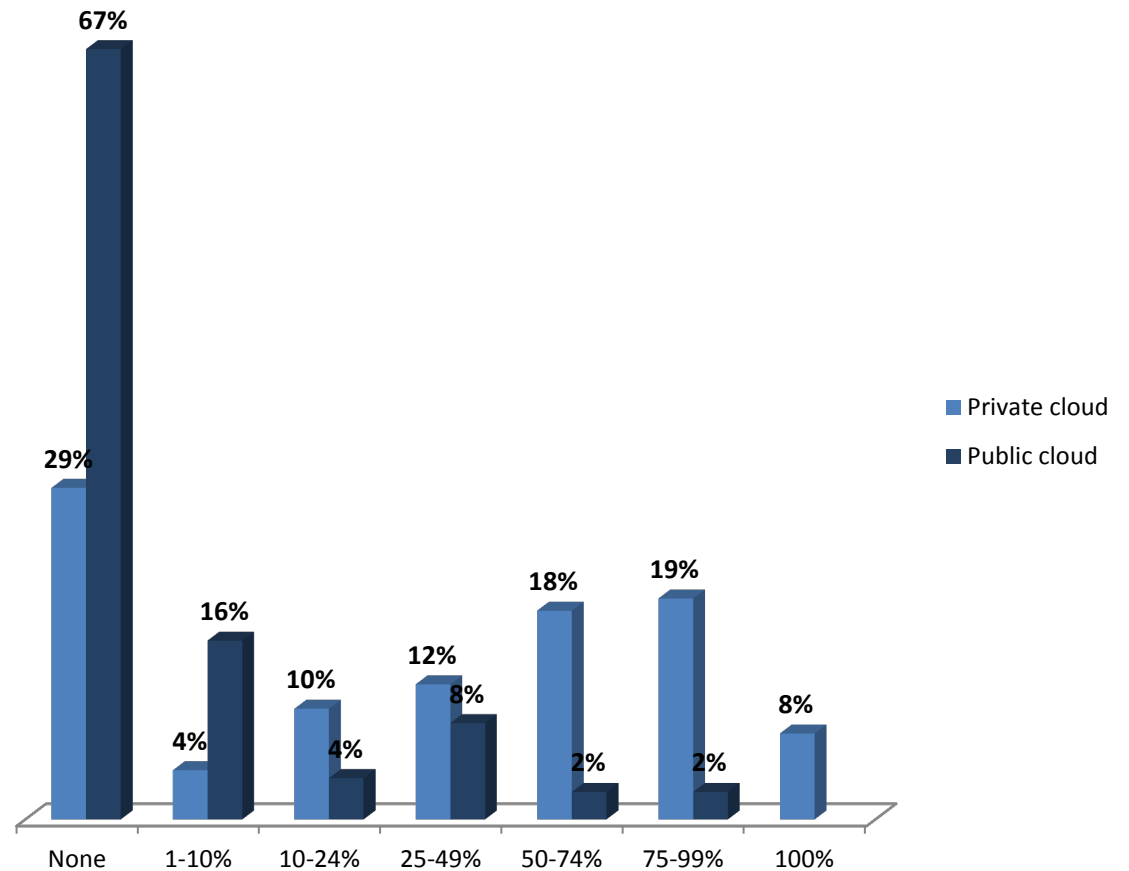
Disaster Recovery / Failover Capacity

Only 15% of the organizations surveyed have adequate disaster recovery capacity, while just over a third of the organizations (35%) have less than 75% of the required capacity.

50%

15%

19%

16%

100% adequate | 75-99% of the capacity we need | 50-75% of the capacity we need | 50% of the capacity we need

Mission-critical in the Cloud

71% of the survey respondents run at least some mission-critical applications in a private cloud environment, while a third (33%) run such applications in the public cloud. These numbers are very similar to the results from last year's survey.
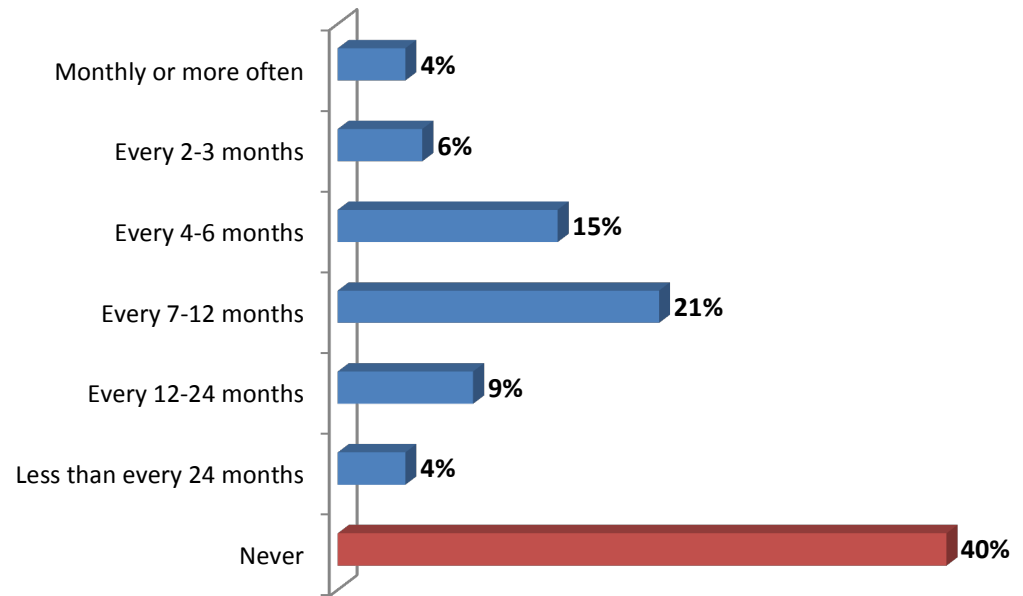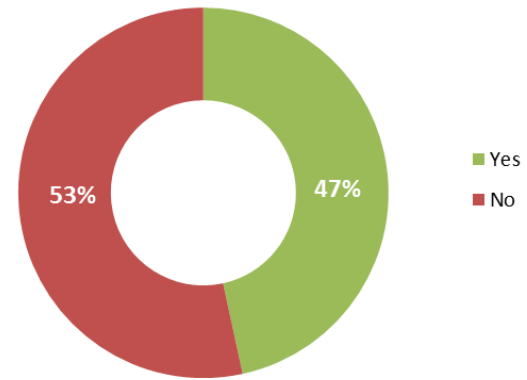
As many as 45% of the organizations run more than half of their mission-critical applications in the private cloud, while only 4% do so in the public cloud.

Legend: ■ Private cloud ■ Public cloud

Chart data:
- None: Private cloud 29%, Public cloud 67%
- 1-10%: Private cloud 4%, Public cloud 16%
- 10-24%: Private cloud 10%, Public cloud 4%
- 25-49%: Private cloud 12%, Public cloud 8%
- 50-74%: Private cloud 18%, Public cloud 2%
- 75-99%: Private cloud 19%, Public cloud 2%
- 100%: Private cloud 8%

Private Cloud Disaster Recovery and Availability Testing

Less than half (47%) of respondents have a disaster recovery solution in place for their private cloud systems.
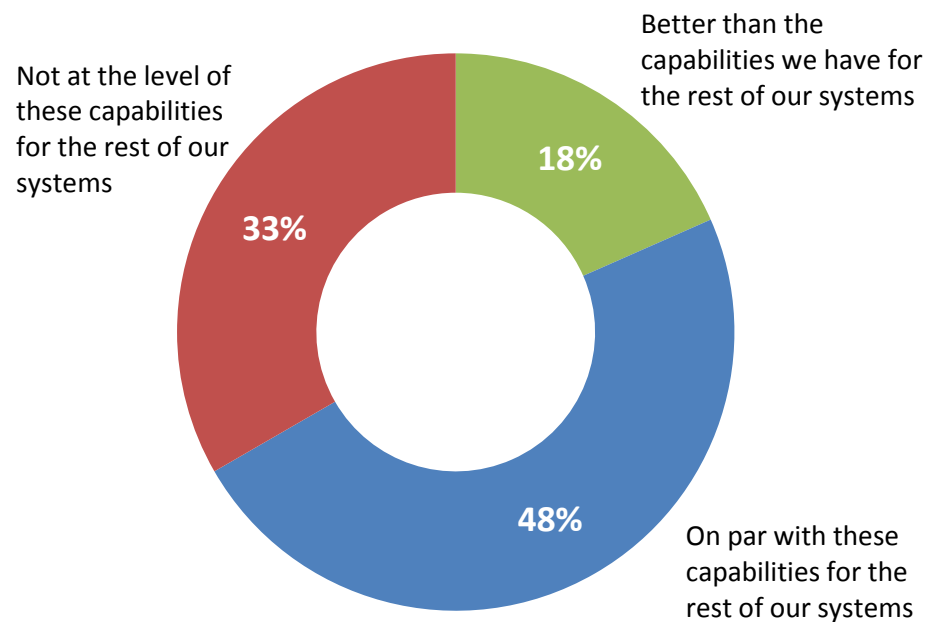
Less than a half of the organizations surveyed (47%) test their private cloud availability at least once a year, and 40% never conduct such a test.

- 53%
- 47%
- Yes
- No

Monthly or more often	4%
Every 2-3 months	6%
Every 4-6 months	15%
Every 7-12 months	21%
Every 12-24 months	9%
Less than every 24 months	4%
Never	40%

Disaster Recovery and high Availability for Cloud Applications

A third of the respondents (33%) describe their disaster recovery and high availability capabilities for the cloud below the level of those capabilities for the rest of their systems.
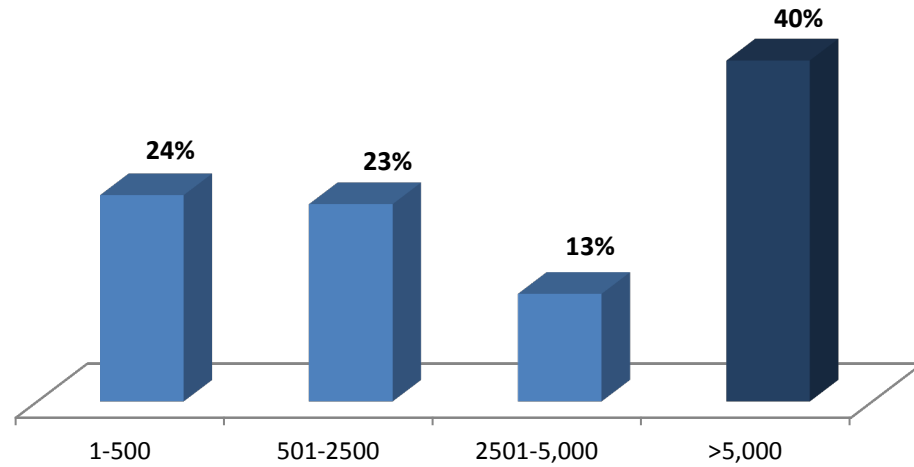
18% describe their cloud capabilities as better than those for the rest of their systems.

Not at the level of these capabilities for the rest of our systems

Better than the capabilities we have for the rest of our systems

18%

33%

48%

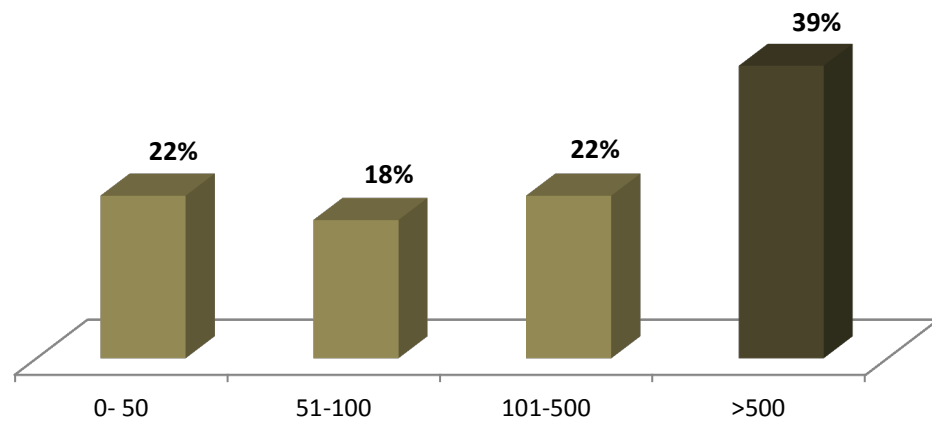On par with these capabilities for the rest of our systems

Respondent Demographics

40% of the survey respondents come from organizations of 5,000 or more employees and 39% have more than 500 servers in their data center.

Number of Employees

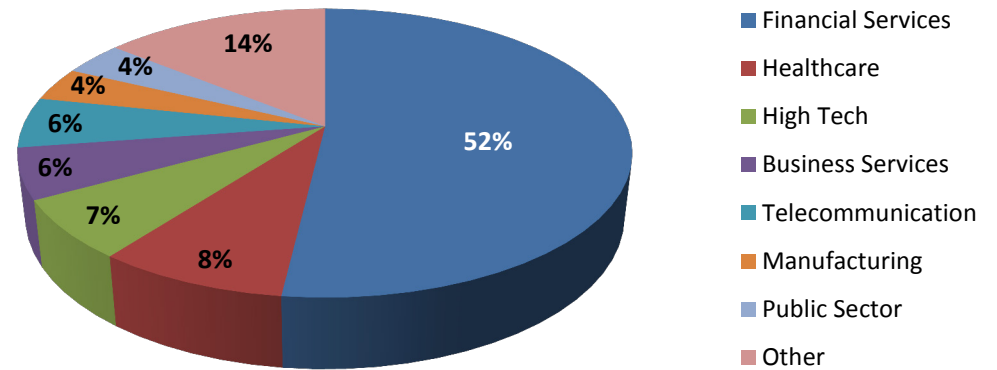| 1-500	501-2500	2501-5,000	>5,000
24%	23%	13%	40%

Number of Servers

0- 50	51-100	101-500	>500
22%	18%	22%	39%

Respondent Demographics (cont.)

Over half of the respondents (52%) represent financial services organizations.

Industry

- Financial Services — 52%
- Healthcare — 8%
- High Tech — 7%
- Business Services — 6%
- Telecommunication — 6%
- Manufacturing — 4%
- Public Sector — 4%
- Other — 14%

Over a third of the respondents (36%) are in business continuity roles, 18% are in disaster recovery and another 17% in IT infrastructure management roles.

Role in the Organization

- Business Continuity — 36%
- Disaster Recovery — 18%
- IT Infrastructure Management — 17%
- IT - Other — 7%
- Risk Management — 5%
- Other — 16%
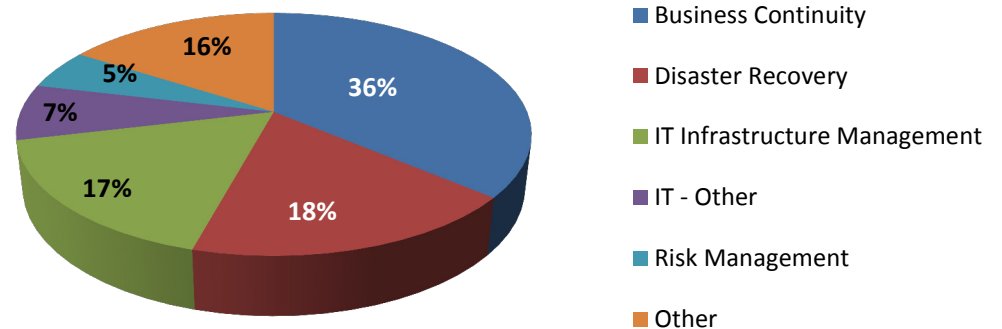
Sign up for a
Service Availability Assessment

- Find hidden risks that can jeopardize your service availability

- Test your environment against a database of 5,000+ documented availability risks

- Get actionable recommendations that will help you eliminate availability risks before they impact your business

100% of the companies that have performed the assessment uncovered vulnerabilities that were previously undetected!

Sign up today

About Continuity Software

Continuity Software is a leading provider of Service Availability Management solutions.

Our software helps many of the world's largest organizations mitigate downtime and data loss risks by monitoring production and remote replication environments to detect hidden vulnerabilities and gaps.

Using our solutions, you will be confident your service availability and data protection goals can be met on a consistent basis.

For more information

Website: www.continuitysoftware.com
Email: info@continuitysoftware.com
Tel: 1-888-782-8170 or +1-646.216.8628

CONTINUITY
SOFTWARE