

2016 Infrastructure Resiliency Survey

Published by



Executive Summary

This survey presents infrastructure resiliency insights that can help IT infrastructure, business continuity, and disaster recovery executives to benchmark their organization's performance and practices against their peers. The results presented here are based on responses from 230 IT professionals from a wide range of industries and geographies collected through an online survey.

Some of the key findings of the survey include:

- **Service availability goals are becoming more ambitious.** As many as 81% of the survey respondents have a service availability goal of less than 8 hours of unplanned downtime a year (compared to 73% in 2014), and 37% have a goal of less than one hour a year.
- At the same time, **as many as 39% of the respondents fell short of meeting their goal.** 34% of the organizations surveyed had an unplanned outage in the past month, and 13% had one in the past week.
- **While cyber attacks make the headlines, they only cause a small fraction of system downtime.** The most common causes are application error and system upgrades, each responsible for over four hours a year on average.
- Although the majority of the survey respondents have moved some of their mission-critical systems to the cloud, **those that have mission-critical system in the cloud were less successful in meeting their service availability goals** compared to organizations that have not made the move.
- **The top challenge in meeting infrastructure resiliency goals is the knowledge gap and inability to keep up with vendor best practices.** This challenge is significantly more prominent in the cloud environment and one of the reasons companies with larger cloud footprint are struggling to meet their goals.

Table of Contents

Avoiding Productivity Loss is the Top Resiliency Driver	4
Service Availability is Highly Critical to Customers	5
Service Availability Goals	6
Service Availability Commitment to Customers	7
Reality is Short of the Goal	8
Larger IT Footprint Means More Challenges for IT Resiliency.....	9
What is Downtime?	10
Downtime Events in the Past Week have Doubled	11
Leading Strategies for Ensuring Infrastructure Resiliency.....	12
Knowledge Gap is the #1 Challenge.....	13
Large Companies Have a Higher Price Tag on Downtime	14
Larger Companies Spend More on Business Continuity	15
Cyber Attacks are Not the Main Reason for Unplanned Outages	16
Mission-Critical in the Cloud.....	17
Which Industries Lead the Way to the Cloud?.....	18
Primary Cloud Provider	19
Cloud Systems Are Less Resilient	20
Industries with Higher Cloud Adoption Lag in Meeting Their Goals.....	21
Cloud Testing	22
Frequent Testing Leads to Better Resiliency.....	23
Respondent Demographics	24
Respondent Demographics	25
Table of Figures.....	26
See how IT Operations Analytics can help you prevent infrastructure outages.....	28
About Continuity Software.....	28

Avoiding Productivity Loss is the Top Resiliency Driver

Avoiding productivity loss is the top driver for infrastructure resiliency initiatives, cited by 44% of the survey respondents.

Additional drivers include ensuring customer satisfaction (22%), protecting company reputation (17%) and regulatory compliance (13%).



Figure 1: Top drivers for infrastructure resiliency initiatives

Service Availability is Highly Critical to Customers

At the same time, the vast majority (82%) of the survey respondents indicated service availability is highly critical to their customers (at least 8 on a scale of 1-10).




Figure 2: How critical is the availability of your service to your customers?

Service Availability Goals

Service availability goals are becoming more ambitious. As many as 81% of the organizations have a service availability goal of less than 8 hours of unplanned downtime a year, compared to 73% that had this goal in 2014.

37% have a goal of less than one hour a year.




Figure 3: Service availability goals for mission-critical systems 2016 vs. 2014

Service Availability Commitment to Customers

Although organizations are setting the bar for Service Availability goals higher than in previous years, a significant portion of the organizations are still hesitant to present these goals as a commitment to their customers.

While there is an increase in the number of organizations that don't have a formal service availability commitment to their customers compared to 2014, there is also an increase in the number of organizations that have a commitment of one hour or less of unplanned downtime a year, reaching 21% of the respondents in 2016.




Figure 4: Service availability commitment to customers 2016 vs. 2014

Reality is Short of the Goal

While the goals are getting more ambitious, a large portion of organizations fall short of their targets. As many as 39% of the survey respondents missed their service availability goal, a similar figure to what respondents reported in 2014.

As can be expected, organizations with higher service availability goals had lower rates of achieving their goals.




Figure 5: Service availability goals: met vs. missed

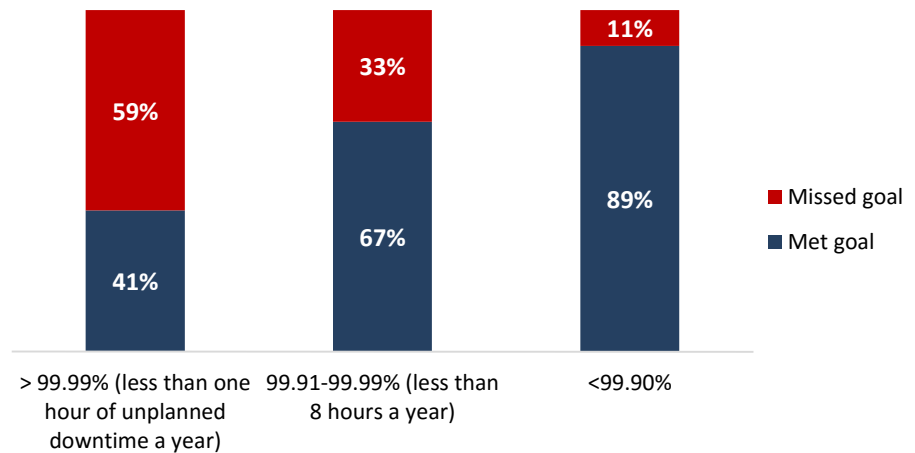


Figure 6: Service availability goals for mission critical systems

Larger IT Footprint Means More Challenges for IT Resiliency

Companies with a larger IT footprint have a harder time meeting their service availability goals.

The majority of companies with over 500 servers (55%) fell short of their service availability goals for mission-critical systems.


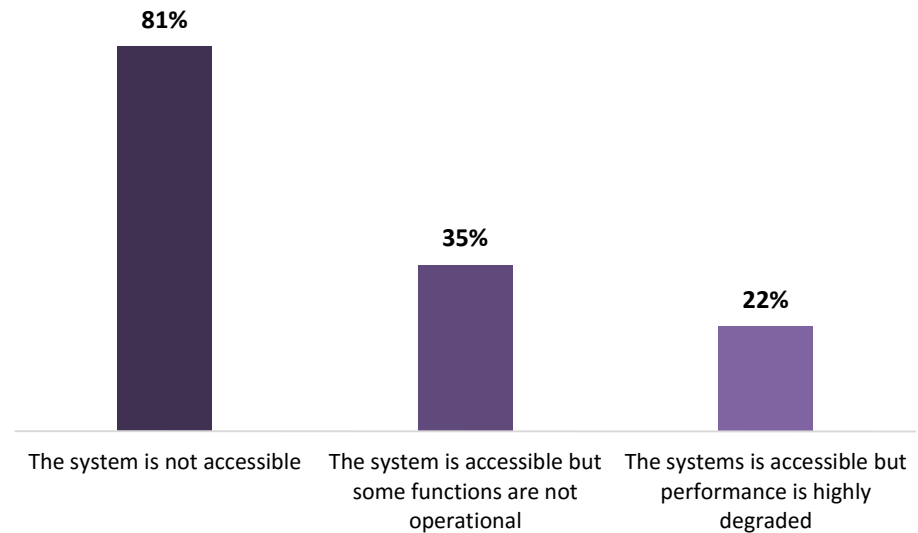


Figure 7: Meeting service availability goals by IT footprint

What is Downtime?

The majority of the respondents (81%) define downtime as an inaccessible system.

Over a third (35%) of the organizations also include in the definition of downtime instances when a system is accessible but some functions are not operational or when performance is highly degraded (22%).



*Figure 8: Definition of downtime
(Respondents could select multiple options)*

Downtime Events in the Past Week have Doubled

34% of the survey respondents experienced an unplanned outage in the past month, and 13% had an outage in the past week, which is twice as many as reported in 2014.


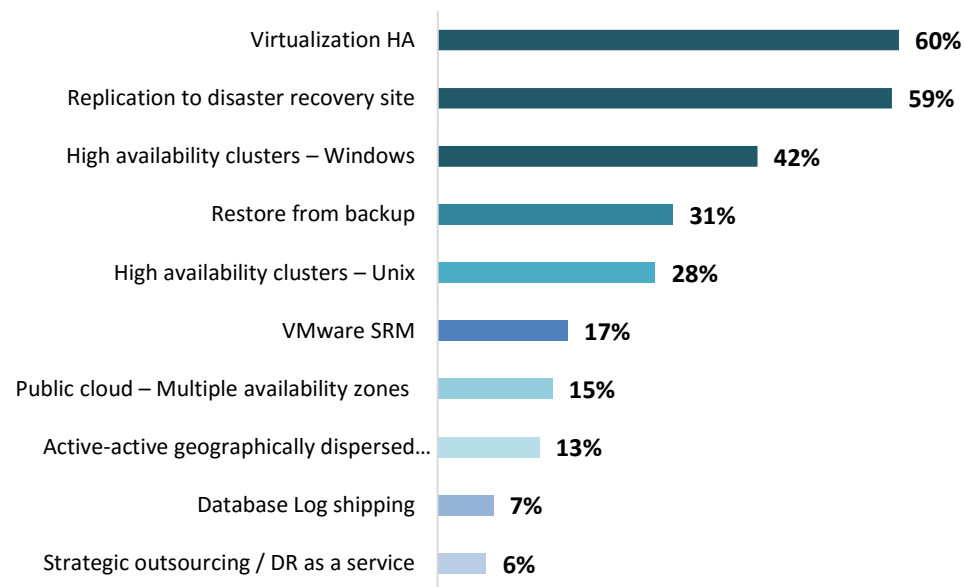


Figure 9: Last downtime event 2016 vs. 2014

Leading Strategies for Ensuring Infrastructure Resiliency

Virtualization HA (60%) and replication to disaster recovery site (59%) are the leading strategies for ensuring infrastructure resiliency.

Other common strategies include: high availability clusters – Windows (42%), restore from backup (31%), high availability clusters – Unix (28%) and VMware SRM (17%).



*Figure 10: Strategies and tools to ensure infrastructure resiliency
(Respondents could select multiple options)*

Knowledge Gap is the #1 Challenge

When asked to rate the top challenges they are facing in ensuring infrastructure resiliency, respondents pointed out to following:

1. Knowledge gap / keeping up with best practices
2. Lack of resources for testing
3. Cross-domain / cross-team coordination
4. Measuring service availability risk and readiness
5. Inadequate documentation of infrastructure topology

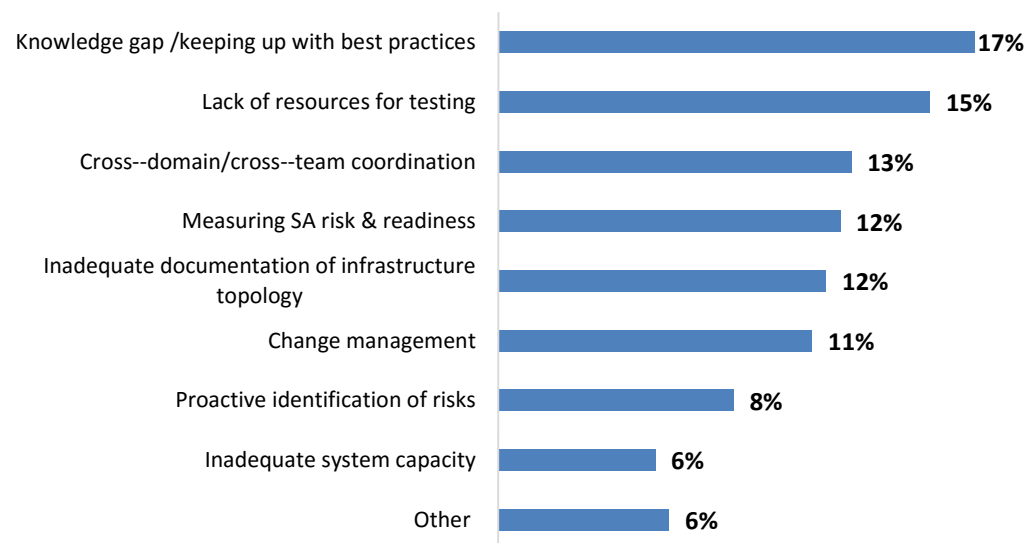


Figure 11: Top challenges in ensuring infrastructure resiliency

Large Companies Have a Higher Price Tag on Downtime

Organizations with over 10,000 employees have a higher price tag on each hour of downtime compared to organizations with fewer than 10,000 employees.

For 36% of the larger organizations, every hour of downtime costs over \$100,000.

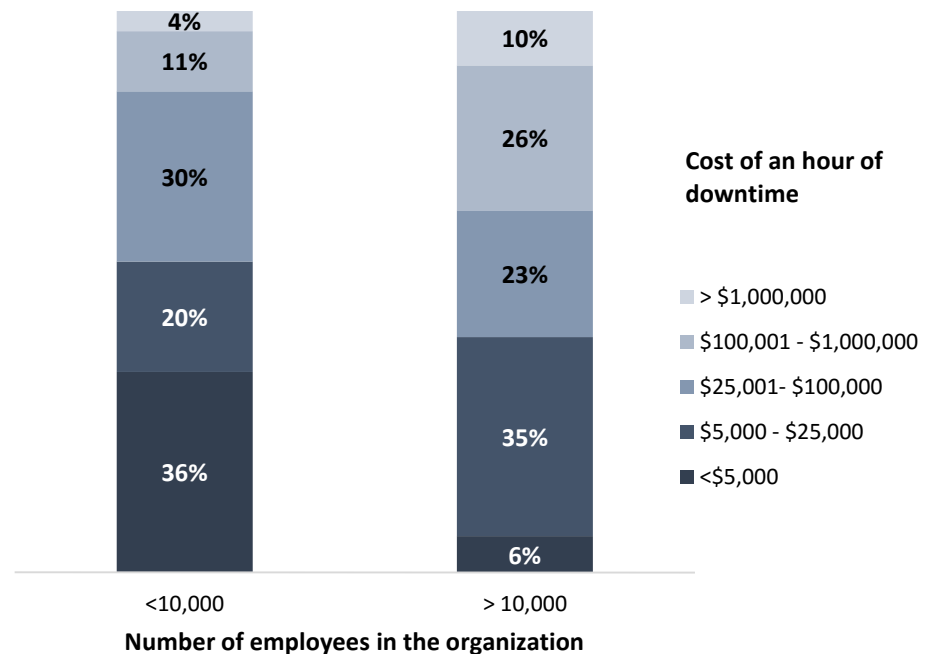


Figure 12: Cost of an hour of downtime by company size

Larger Companies Spend More on Business Continuity

62% of the respondents have initiatives for improving infrastructure resiliency in 2016.

Business continuity / disaster recovery budget reflects the size of the company.

37% of the organizations with over 10,000 employees have an annual budget of over \$50M.




Figure 13: Have initiatives for infrastructure resiliency improvement




Figure 14: Business continuity / disaster recovery budget by company size

Cyber Attacks are Not the Main Reason for Unplanned Outages

While cyber attacks make the headlines, they only cause a small fraction of system downtime. The most common causes are application error and system upgrades, each responsible for over 4 hours a year on average.

It is noteworthy that for many outages the reason is actually unknown (over 5 hours on average for the year).

Other reasons include human error (3.4), hardware failure (3.2) and power outage (1.6).




Figure 15: Causes of unplanned outages (in hours)

Mission-Critical in the Cloud

64% of the organizations run some mission-critical applications in their private cloud and 42% run such applications in the public cloud.




Figure 16: Mission-critical applications running in the private cloud




Figure 17: Mission-critical applications running in the public cloud

Which Industries Lead the Way to the Cloud?

High tech, telecommunication and retail are the industries that lead the way to cloud adoption. The education sector joins them in public cloud adoption.

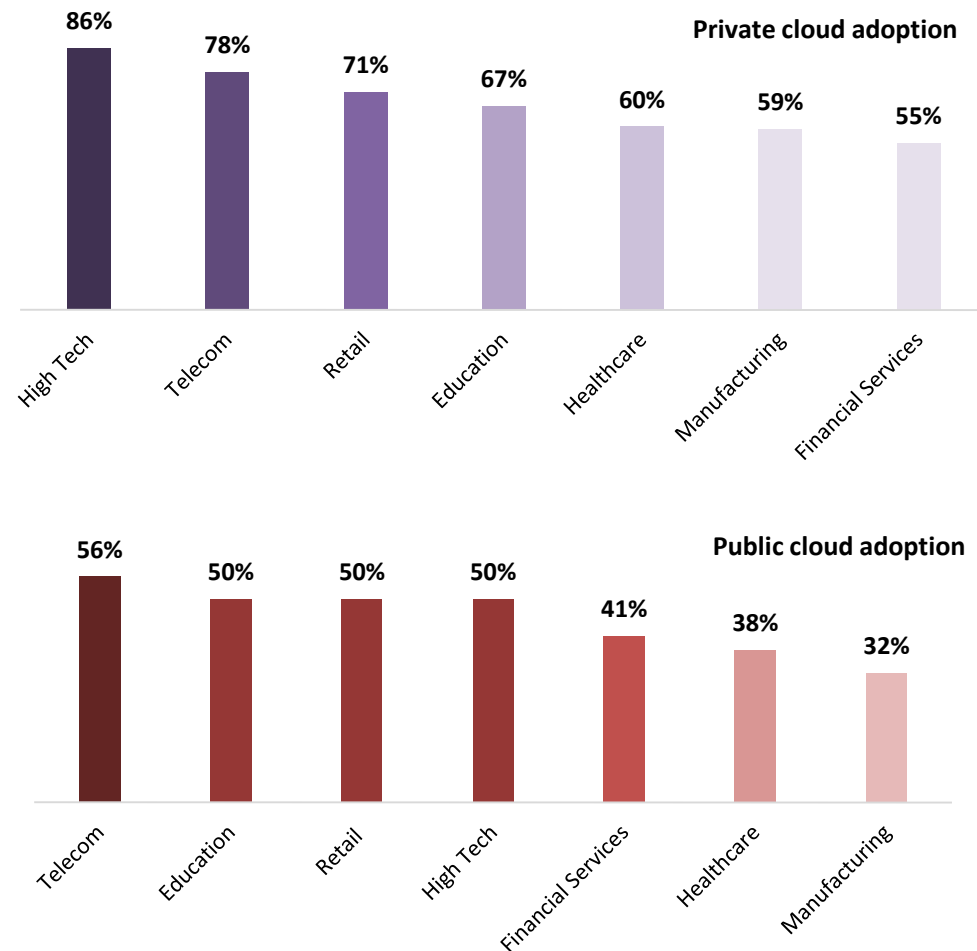


Figure 18: Which industries lead the way to the cloud?

Primary Cloud Provider

Amazon (39%) and Azure (35%) are the primary cloud providers of the organizations surveyed.

26% of the respondents use other cloud providers.

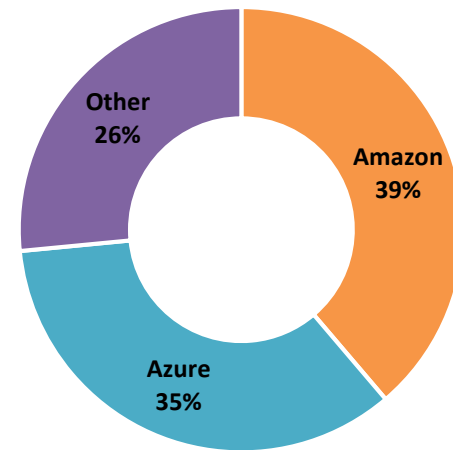


Figure 19: Primary cloud provider

Cloud Systems Are Less Resilient

Companies that have mission-critical system in the cloud were less successful in meeting their service availability goals compared to organizations that don't store any of their mission-critical data in the cloud.

The figures are similar for private and public cloud. About half of the companies with cloud systems did not meet their goals, compared to about a third of the companies that don't have any cloud-based mission-critical systems.

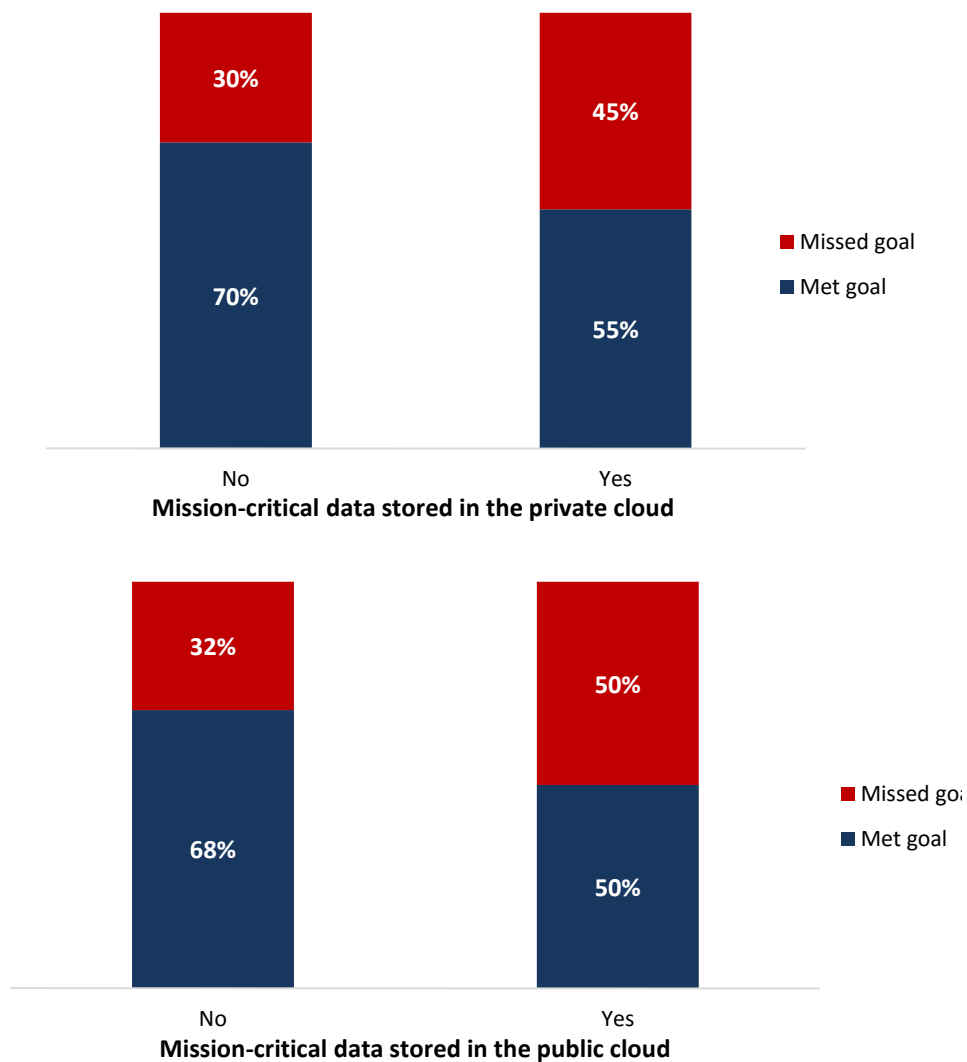


Figure 20: Mission-critical data stored in the cloud vs. met goal

Industries with Higher Cloud Adoption Lag in Meeting Their Goals

Industries with higher levels of cloud adoption – high tech, telecommunication, and retail – generally have a harder time meeting service availability goals.

One of the exceptions is the education sector, which is one of the leaders in public cloud adoption yet has the best record for meeting service availability goals.




Figure 21: Meeting service availability goals by industry

Cloud Testing

Only 35% of the organizations surveyed have a disaster recovery solution in place for their cloud systems.

Almost half of the respondents (47%) never test their private cloud availability. Only 18% test it on a quarterly basis or more often.




Figure 22: Have a DR solution in place for cloud systems




Figure 23: Frequency of testing private cloud availability

Frequent Testing Leads to Better Resiliency

There is a positive correlation between the frequency of testing cloud availability and meeting service availability goals.

73% of the organizations that test their cloud availability every 3 months or more often were successful in meeting their service availability goals, compared to just 50% of the organizations that conduct such testing less often.




Figure 24: Resiliency of cloud infrastructure

Respondent Demographics

Most survey respondents come from mid-size and large companies, with 40% of the survey respondents coming from organizations of over 10,000 employees.

Over half of the respondents (54%) have more than 500 servers in their datacenter.




Figure 25: Number of employees




Figure 26: Number of servers

Respondent Demographics

A third (33%) of the survey respondents are in IT infrastructure management and 31% in other IT roles.

Survey respondents represents a broad cross-section of industries. The highest representations are from the manufacturing (19%), financial services (17%) and education (9%) sectors.




Figure 27: Job responsibility




Figure 28: Industry

Table of Figures

Figure 1: Top drivers for infrastructure resiliency initiatives.....	4
Figure 2: How critical is the availability of your service to your customers?	5
Figure 3: Service availability commitment to customers	7
Figure 4: Service availability goals for mission-critical systems	6
Figure 5: Service availability goals: met vs. missed.....	8
Figure 6: Service availability goals for mission critical systems	8
Figure 7: Meeting service availability goals by IT footprint	9
Figure 8: Definition of downtime	10
Figure 9: Last downtime event	11
Figure 10: Strategies and tools to ensure infrastructure resiliency	12
Figure 11: Top challenges in ensuring infrastructure resiliency	13
Figure 12: Cost of an hour of downtime by company size	14
Figure 13: Have initiatives for infrastructure resiliency improvement	15
Figure 14: Business continuity / disaster recovery budget by company size	15
Figure 15: Causes of unplanned outages (in hours)	16
Figure 16: Mission-critical applications running in the private cloud	17
Figure 17: Mission-critical applications running in the public cloud.....	17
Figure 18: Which industries lead the way to the cloud?	18
Figure 19: Primary cloud provider.....	19
Figure 20: Mission-critical data stored in the cloud vs. met goal	20
Figure 21: Meeting service availability goals by industry	21
Figure 22: Have a DR solution in place for cloud systems	22
Figure 23: Frequency of testing private cloud availability	22
Figure 24: Resiliency of cloud infrastructure	23
Figure 25: Number of employees.....	24
Figure 26: Number of servers.....	24

Figure 27: Job responsibility 25
Figure 28: Industry..... 25

See how IT Operations Analytics can help you prevent infrastructure outages

Sign up for a 30-minute demo and find out how AvailabilityGuard can improve your ability to meet critical KPI goals:

- Identify single-points-of-failure and other configuration risks
- Comply with vendor best practices across all layers of the infrastructure
- Verify configuration changes before they impact the business
- Track KPI's that support continuous improvement
- Establish safer and more agile best practices

[Sign up today](#)

About Continuity Software

Continuity Software helps the world's leading organizations prevent unplanned IT outages. Our award-winning AvailabilityGuard software enables IT teams to proactively identify and eliminate single-points-of-failure and other misconfigurations across the entire infrastructure.

Using AvailabilityGuard's advanced IT Operations Analytics, organizations are able to deliver the highest levels of IT service availability while improving operational efficiency.

Website: www.continuitysoftware.com

Email: info@continuitysoftware.com

Tel: 1-888-782-8170 or +1-646.216.8628