Business Resiliency: The Need for 24/7/365 Operations

Prepared for IBM by Tech Republic

CONTENTS:

About this white paper: This document has been prepared by CBS Interactive on behalf of IBM. IBM has specified topic, title and key themes of this guide and may have contributed to and exercised editorial control over the content. This white paper may only be quoted and reproduced by IBM in its entirety.

Executive Summary

Today's customers and business partners expect 24/7/365 uptime. Server and infrastructure outages can cost thousands of dollars every minute, when you calculate the costs of remediation, the lost revenue and productivity, the reputation damage, and the penalties associated with SLA and regulatory non-compliance. Many companies would consider their tolerance for downtime to be zero.

Despite these pressures, power outages, severe weather, hardware issues, and human error continue to cause IT service disruptions regularly, and recovery times average 24 hours. This stagnation is due primarily to flat IT budgets, increasingly complex and interdependent systems, globally distributed business operations, and increasingly mobile workforces.

Technologies like virtualization, data replication, and cloud computing have evolved to help companies overcome these resiliency challenges. Whole servers, and even whole environments, can be ported to standby resources for near-instant failover. Virtualization and replication are nearly ubiquitous in enterprise settings, and adoption of cloud services is growing steadily. Cloud resiliency models include private, shared, on-site, and offsite options. Many companies employ hybrid models to ensure high availability for their most critical business systems and cost-effective recovery for non-critical components.

Because cloud services are relatively new and confusing in their service and security parameters, it's important to work closely with an expert resiliency service provider, such as IBM. IBM offers a comprehensive portfolio of continuity and resiliency services – plus expert consultation – to suit the needs of any business. Through its combination of data centers, hardware, software, technical experts, and cloud services, IBM can generate a well-crafted resiliency plan and implement a tailored solution that ensures seamless failover and fail-back, data security, event remediation, and regulatory compliance. Even better, IBM can manage the entire solution without burdening the customer's in-house IT resources.

Introduction

“The server is down.” These are four dreaded words that may mean different things to different people, but every connotation is negative. For the IT staff, it means scrambling to ensure the affected resources are brought back up and running and to restore data from backup, if necessary. For the rank-and-file, it means minutes or hours of lost productivity, frustration, and stress. For executive management, it means thousands or more in costs, damage to the company’s reputation, and even penalties if mandated service levels are not met.

In this white paper, we’ll discuss the various factors that drive the need for resiliency and the potential costs of downtime. We’ll look at the factors that are making high availability and continuity more difficult for today’s businesses, and we’ll explore new models that enable companies to failover and fail-back their IT resources quickly and reliably.

WHY RESILIENCY?

In the last few years, companies have seen a marked decrease in their tolerance for downtime. Gone are the days when a weekend service outage for maintenance could be considered routine. In order to remain competitive, today’s businesses must be engaged – transacting, processing, servicing customers, and reacting to changing conditions – 24/7/365. At the simplest level, business continuity is critical because every minute of downtime equates to lost revenue. And in the longer term, if your operations are disrupted, customers will likely take their patronage elsewhere.

SLAS

In addition, many companies are subject to service level agreements that include penalties for non-compliance. In his “Guide to SLAs,” IT management consultant Barclay Rae explains that SLAs may be set up between providers and customers, partners, and even internal departments.[1] They may also be formalized as service level contracts, with enforceable liabilities and penalties.

SLAs force organizations to quantify and consider each element of their IT service operation, including:

- Descriptions of each service
- Escalation paths for remediation of problems
- Hours of support
- Speed of resolution for incidents
- Speed of non-emergency service delivery
- Tiers of incident severity
- Outlines of customer responsibilities
- Special considerations

Every minute of downtime equates to lost revenue.

[1] Barclay Rae, “A Guide to SLAs,” July 2012

Drilling down into the tiers of severity, Rae outlines a sample SLA that puts top-priority incidents at 1 hour recovery, second-tier incidents at 4 hours, and third-tier incidents at one business day. But even this degree of responsiveness is outdated, and bound to cost thousands of dollars or more, as we'll see.

Time to recover

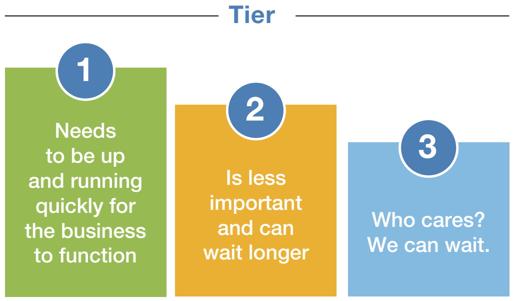
Tier
1
Needs to be up and running quickly for the business to function
2
Is less important and can wait longer
3
Who cares? We can wait.

5 years ago:	< 4 hours	24 - 72 hours	> 72 hours
now:	0	< 4 hours	24 - 72 hours

Source: Forbes Insights, "How the cloud is changing resilience in the expanding universe of digital data," 2014

REGULATORY COMPLIANCE

Beyond service level agreements with customers and internal stakeholders, many IT organizations are bound by government and industry regulations to maintain redundancy and availability for critical systems and sensitive data. Global companies must comply with relevant data protection laws in every country in which they operate. Just a sampling of these regulations includes:

- Basel II, requiring systems availability for international financial institutions
- European Union Data Protection Directive, covering data backup and availability
- Sarbanes-Oxley, covering financial data in publicly traded companies in the US
- HIPAA, requiring data availability for US health information

These guidelines require that companies maintain redundant resources for failover, plus secure and comprehensive data backup for information integrity. They include significant financial penalties for companies that are found in violation. These fines must be considered when companies evaluate the cost of potential downtime.

CALCULATING THE COST

According to the Ponemon Institute's 2014 Cost of Data Breach Study, the dollar amount that companies ascribe to data incidents must include direct, indirect, and opportunity costs in order to reflect the monetary damage accurately.[2] These factors include:

- Detection of the incident
- Containment of the incident
- Recovery of networks, data, and/or core systems
- Forensics associated with post-event investigation
- Third parties engaged to help remediate the problem and audit systems
- Legal costs associated with customer and compliance breaches
- Training/increased workload for helpdesk and support staff
- Lost revenue from business disruption
- Lost opportunities from customer turnover

60 percent of these incidents were caused by human error and system glitches.

Taking these into consideration and applying them across a survey sample of 314 companies in 10 countries, Ponemon found that the average cost per leaked record was $145 globally, up 9 percent from 2013. The average total incident costs range from $3 million to nearly $6 million, depending on country. It's important to note that these breaches aren't limited to cyber attacks. On average, 60 percent of these incidents were caused by human error and system glitches – in other words – IT outages.

MILEAGE MAY VARY

Naturally, not every company will see costs this high. Some may see them even higher. Depending on industry, breaches and outages can cost far more than the average. Healthcare breaches, for example, cost $359 per leaked record, Ponemon found.

For credit card applications, which are in place at many companies, an hour of downtime could cost as much as $2.6 million, and a 49-minute Amazon.com homepage outage in January of 2013 cost the company nearly $5 million, according to estimates cited by the Neverfail Group.[3]

[2] Ponemon Institute, "2014 Cost of Data Breach Study," May 2014

[3] Neverfail Group, "Downtime Report: Top Ten Outages in 2013," December 2013

Disaster recovery plan updates need to change

"As a result of your most significant disruption, which of the following turned out to be the greatest impacts to your organization?"
(Rank 1-3)

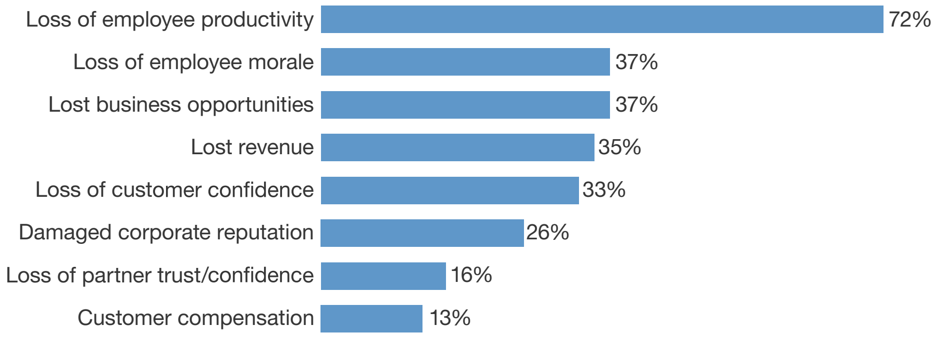

Base: 66 global disaster recovery decision-makers and influencers who have declared a disaster or had a major business disruption (multiple responses accepted)

Source: Forrester Research, "The State of Business Technology Resiliency, Q2 2014," May 2014

In a study conducted last year by Continuity Software, 43 percent of respondents said every hour of downtime costs $100,000 or more, and 12 percent said each hour costs more than $1 million. The vast majority of respondents – 90 percent – said service availability is of critical importance to their customers. In addition, 73 of the respondents to Continuity's survey said their availability goals are higher than 99.91 percent, or less than eight hours of unplanned downtime each year. That number is trending upward compared with 68 percent in 2013.[4]

The cost of downtime will differ from one company to the next, but it's been rising steadily across industries over the last four years, and calculating it is a good exercise in understanding and communicating the need for resiliency measures.

CURRENT TRENDS

It's interesting to note that despite all the pressures we've discussed, resiliency efforts at many companies have remained relatively stagnant over the last few years.

Forrester Research's "State of Business Technology Resiliency, Q2 2014" report shows that, while demand for rapid recovery time objectives (RTO) may be rising on the business side, actual recovery times averaged around 24 hours in 2013, up from 18.5 hours in 2010. In fact, only 2 percent of respondents in a 2013 Forrester survey said they could recover in less than 1 hour.[5]

Meanwhile, the causes of downtime have remained consistent. Even though large-scale disasters such as hurricane Sandy dominate the headlines, power failures, IT failures, and human error are the most common causes of business downtime, according to Forrester.

[4] Channel Insider, "Helping Combat Downtime, On-Premise and in the Cloud," June 2014

[5] Forrester Research, "The State of Business Technology Resiliency, Q2 2014," May 2014

Top causes of downtime are mundane events, not disasters

"What was the cause(s) of your most significant disaster declarations(s) or major business disruption?"

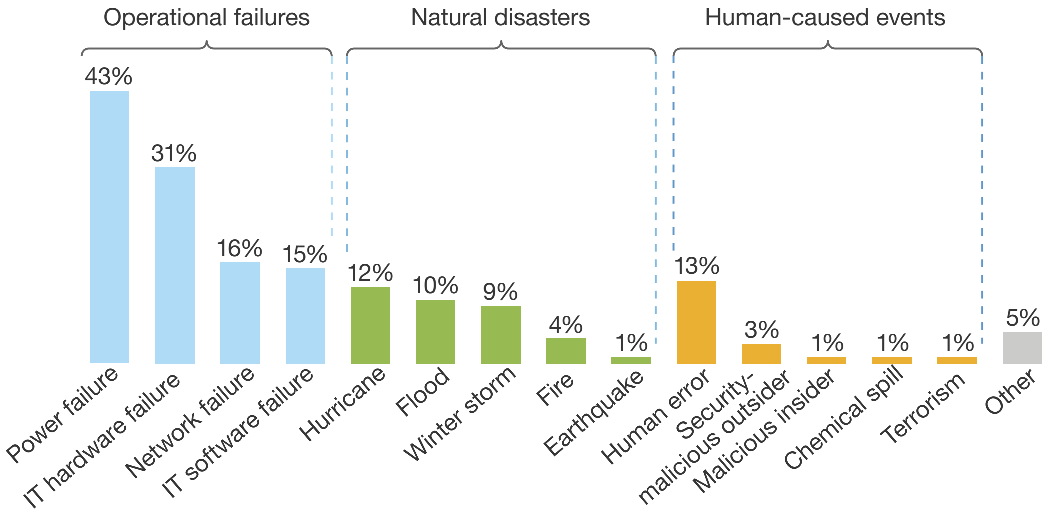

Base: 94 global disaster recovery decision-makers and influencers (does not include "don't know" responses; multiple responses accepted)

Source: Forrester Research, "The State of Business Technology Resiliency, Q2 2014," May 2014

Other analyst firms concur. In the above-mentioned Continuity Software study, 87 percent of respondents had experienced a downtime event within 3 months of the survey, and the most common causes were hardware failure, equipment upgrades, human error, and power disruptions.

Likewise, KPMG's "2013-2014 Continuity Insights" report cited severe weather and power outages as the top causes of downtime, with IT-related errors in third place.[6]

So if the causes of downtime are well understood, why are recovery times and costs worsening from one year to the next?

Power failures, IT failures, and human error are the most common causes of business downtime.

CHANGING LANDSCAPE

The answer is that IT environments have grown far more complex and distributed as companies keep pace with global competition. Increasingly mobile and remote workers around the world need access to critical business systems at all times, on multiple devices, and as we've discussed, any interruption creates a snowball of financial damage.

Forrester refers to this trend as "the era of now." In the above-referenced report, analyst Stephanie Balaouras explains that communication applications have become mission-critical elements of every business's operations, customer expectations are at an all-time high, and interdependencies between applications (a sales tool that relies on a particular database via a business process management framework, for example) create an environment where more components than ever are considered critical.

[6] KPMG, "The 2013-2014 Continuity Insights and KPMG LLP Global Business Continuity Management (BCM) Program Benchmarking Study," April 2014

Mission-critical and business-critical tiers are increasing

"What percentage of your applications and data fall into the following tiers?"

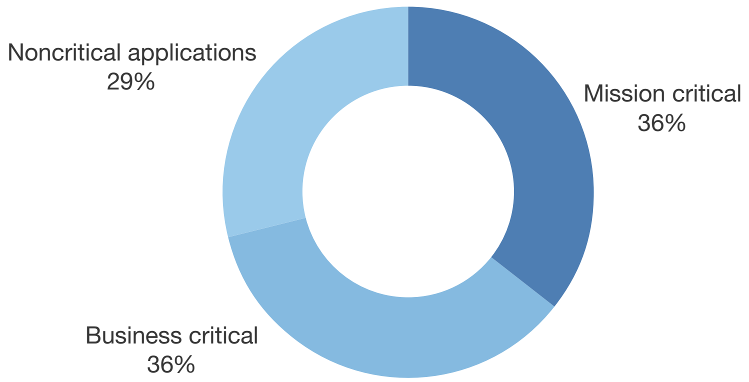

Base: 94 global disaster recovery decision-makers and influencers
(does not include "don't know" responses; percentages do not total 100 because of rounding)

Source: Forrester Research, "The State of Business Technology Resiliency, Q2 2014," May 2014

At the same time, IT budgets have remained virtually flat since the economic downturn of 2008, with recovery only starting to appear in 2015, according to research firm Corporate Executive Board.[7] That means companies haven't had the resources to invest in upgrading or streamlining their resiliency strategies.

In addition, mobility is changing the way workers do their jobs. Back in 2012, IDC predicted that the number of mobile workers would reach 1.3 billion by 2015, representing over 37 percent of the global workforce.[8] More recently, Gartner released a forecast saying that by 2016, 38 percent of employers would require their workers to use their own mobile devices for work, and this number would trend close to 50 percent by 2017.[9]

All of these factors put stress on traditional resiliency setups, which involve redundant hardware (which is expensive to maintain) and tape backups (which are time-consuming in recovery scenarios). They also create multiple points of potential failure, as device dependencies and network paths must be mapped carefully and accurately; any one mistake during a failover can bring down the whole infrastructure.

In this rapidly evolving and increasingly interconnected business landscape, companies must look for smarter resiliency solutions to protect their infrastructure, their data, their operations, and their workers.

[7] Corporate Executive Board, "IT Budget Benchmark Key Findings 2014-2015," 2014

[8] Reuters, "Mobile Worker Population to Reach 1.3 Billion by 2015, According to IDC," January 2012

[9] Gartner, "Bring Your Own Device: The Facts and the Future," May 2013

VIRTUALIZED AND REPLICATED

Thanks to virtualization – the abstraction of server, networking, and application workloads from the hardware on which they're running – resiliency is a more attainable goal for companies across the spectrum of size and industry. Whole servers, and even whole environments, can be ported to standby hardware with minimal downtime, and many companies replicate all business activities in real time, so the standby environment will be ready for near-instant failover.

Data Protection Technologies: 2010 to 2013

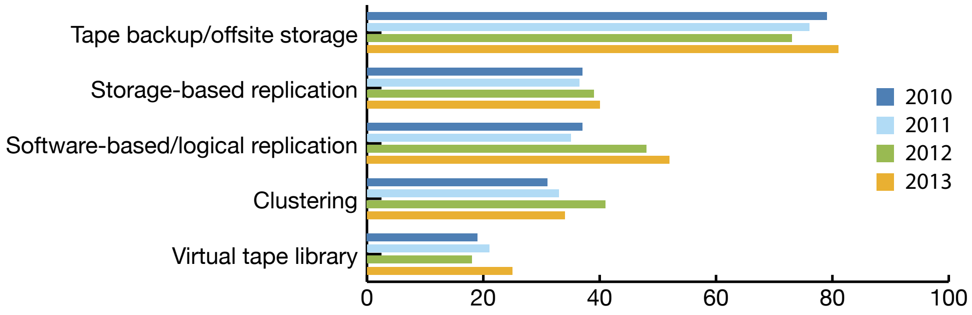

Source: Vision Solutions, "State of Resilience 2013," January 2014

Continuity Software's 2014 survey found that 72 percent of respondents use virtualization for high availability in their environments, up from 63 percent in 2013.

Vision Solutions, in its "The State of Resilience 2013" survey of over 3,500 IT professionals, shows a marked increase in software-based replication deployments from 2010 to 2013.[10] Forrester found a similar trend when comparing 2010 to 2013, with replication usage growing from 35 percent to over 50 percent. Tape backup, the report adds, is still the most popular method for protecting non-critical systems.

Data between primary recovery sites

"How do you copy data between your primary recovery sites?"

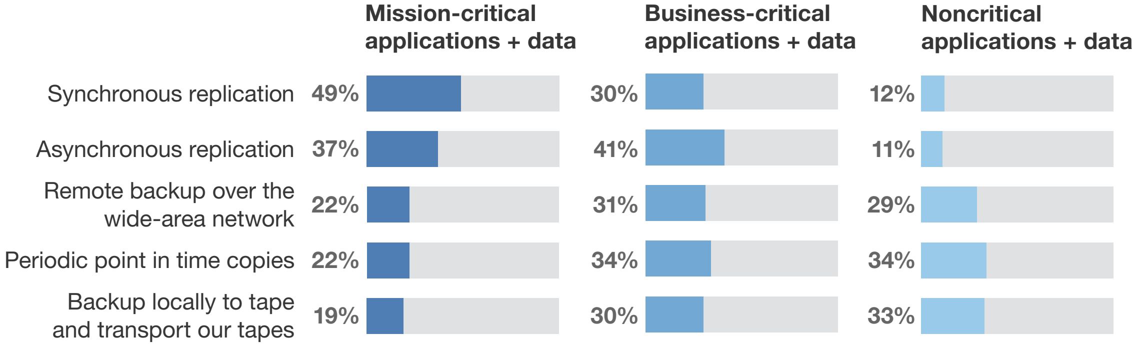

Base: 94 global disaster recovery decision-makers and influencers
(does not include "don't know" responses; multiple responses accepted

Source: Forrester Research, "The State of Business Technology Resiliency, Q2 2014," May 2014

[10] Vision Solutions, "State of Resilience 2013," January 2014

ENTER THE CLOUD

The final ingredients in next-generation resiliency plans are cloud services. The ability to host resources on redundant, offsite infrastructure – and to transfer data and virtualized components at WAN speeds or better – creates a variety of failover options for companies that want to ensure rapid recovery in case of an outage.

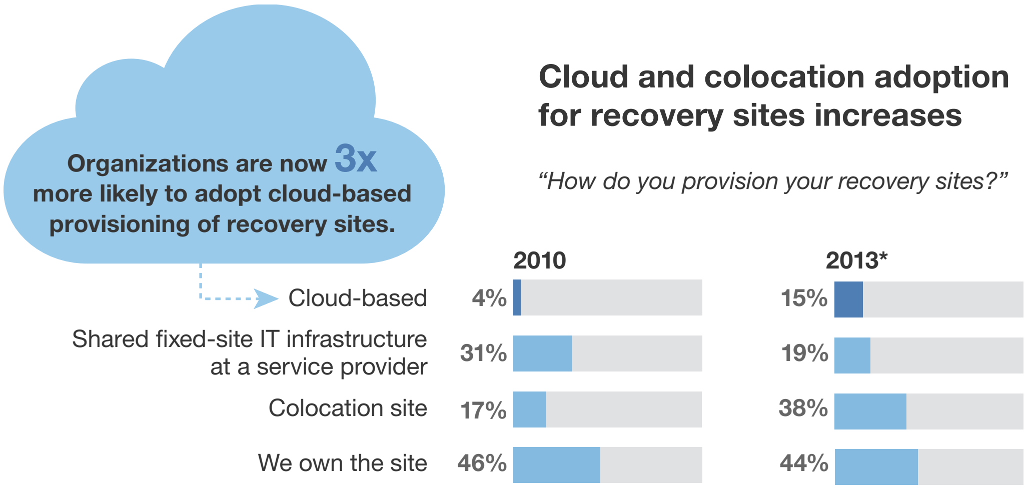

Cloud and colocation adoption for recovery sites increases

“How do you provision your recovery sites?”

Base: 180 global disaster recovery decision-makers and influencers
*Base: 85 global disaster recovery decision-makers and influencers
(does not include “don’t know” responses; multiple responses accepted)

Source: Forrester Research, “The State of Business Technology Resiliency, Q2 2014,” May 2014

In addition to traditional internal resources, these options include:

- Dedicated offsite resources in a collocation facility
- Shared infrastructure in an offsite facility
- On-premises or hosted private cloud
- Public cloud

72 percent of respondents use virtualization for high availability.

Many companies are pursuing strategies that combine several of these options into a hybrid resiliency solution. According to Forrester, cloud-based recovery adoption rose from 4 percent in 2010 to 15 percent in 2013, and collocation is rising in popularity, as well. One in five told Forrester that they use a mix of models to provide fast failover for critical systems and a more phased recovery scheme for non-critical elements.

KPMG found slightly higher adoption for hybrid and cloud disaster-recovery solutions, with nearly 20 percent citing private cloud and just over 8 percent reporting that they use public cloud services for some capabilities.

Organizations' current IT DR strategies

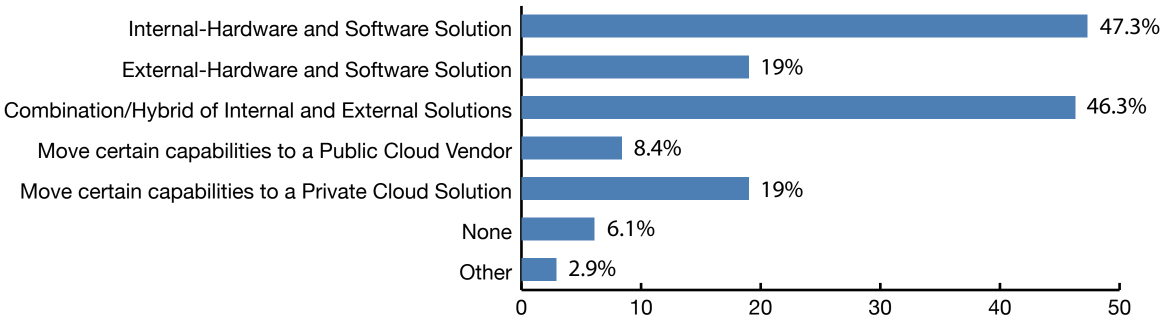

Source: KPMG, "The 2013-2014 Continuity Insights and KPMG LLP Global Business Continuity Management (BCM) Program Benchmarking Study," April 2014

CLOUD CONFUSION LINGERS

Because cloud services are relatively new and their adoption – especially in large organizations – is often sporadic and decentralized, there is a fair amount of confusion around the parameters of cloud SLAs; the security, ownership, and location of data; and who is responsible for outages that affect cloud providers.

For example, in the Continuity Software study, 44 percent of respondents said their cloud service availability is on a par with their internal systems. Some 30 percent said the availability of their cloud services is lower than their other systems, and 26 percent said cloud service availability is superior. This kind of spread seems to indicate that companies don't necessarily have a solid grasp on the availability metrics for their internal and cloud-based services.

This conclusion is borne out by KPMG's findings, which indicate that a whopping 68 percent of respondents don't understand the nature of their cloud-based IT disaster recovery measures. While recovery services often include on-demand infrastructure (IaaS) and hosted software solutions (SaaS) – and hosted compute platforms (PaaS) and networking solutions (NaaS) to a lesser extent – most of the professionals surveyed had no idea what their cloud resiliency plans included.

Organizations' currently implemented IT DR plans in the cloud

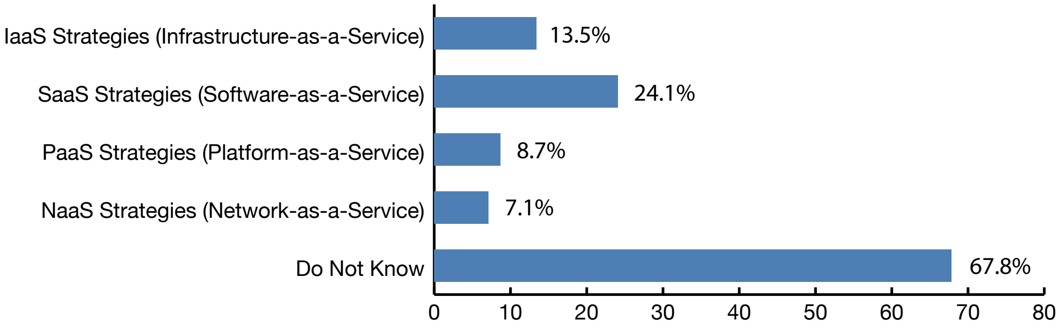

Source: KPMG, "The 2013-2014 Continuity Insights and KPMG LLP Global Business Continuity Management (BCM) Program Benchmarking Study," April 2014

PARTNER UP

We’ve established that resiliency is more important and in many ways more difficult that ever before, and that cloud services can take virtualization and replication to the next level by providing multiple types of offsite resource hosting.

Because of the tremendous confusion that lingers around cloud services, however, it’s important to work closely with an expert resiliency service provider, rather than attempting to hash out the multiple connections, SLAs, and interdependencies required for a successful hybrid recovery deployment.

Fortunately, businesses can partner with IBM, the recognized and undisputed leader in business technology solutions, for resilience services that run the gamut from on-premises to collocated to cloud-based resiliency.

IBM RESILIENCY SERVICES

IBM offers a comprehensive portfolio of continuity and resiliency services – plus expert consultation – to suit the needs of any business. Through its combination of data centers, hardware, software, technical experts, and cloud services, IBM can generate a well-crafted resiliency plan and implement a tailored solution that ensures seamless failover and fail-back, data security, event remediation, and regulatory compliance. Even better, IBM can manage the entire solution without any need for the customer to expend internal IT resources.

IBM’s Resiliency Services include:

Resiliency Consulting: Leverage IBM’s decades of experience and track record of successful deployments in planning, designing, integrating, and testing your continuity and resilience strategy across the entire company, taking applications and infrastructure at headquarters and all office locations into consideration.

Infrastructure Recovery: IBM conducts a thorough assessment of your IT systems and sets up a data protection and recovery strategy, complete with offsite backups and on-site recovery assistance in case of an incident. This service is ideal for companies with regulatory compliance requirements and complex environments.

Availability Management: Expert IBM consultants analyze previous incidents and create a strategy that aligns with your business processes and helps avoid future downtime. A dedicated program manager supervises every phase of the project, leaving your in-house IT resources free to support business operations and innovation.

Managed Resiliency: Leave the driving to IBM. After you’ve designated the information and systems you need operational in case of a disruption, IBM provides the resiliency management for you, with proactive and event-response services that can be tailored based on criticality.

Cloud Resiliency: IBM’s cloud services include backup, recovery, and virtualized data management for flexible access and comprehensive control. Recover servers, applications, and data in a matter of minutes to reduce risk, improve compliance, and preserve productivity.

Cloud Virtualized Server Recovery: Insuring your virtual and non-virtual resources from disruption, IBM’s cloud recovery services automate failover procedures and improve reliability for business operations. Optional monitoring of the recovery environment is a boon to customer assurance and compliance. This service supports a wide range of server operating systems and platforms.

Cloud Managed Backup: If you're looking for a comprehensive and secure backup solution that automatically includes on-site and offsite replication, look to IBM's managed backup services. These support private, public, and hybrid cloud models; deliver dynamically scalable solutions that reduce total cost of ownership; and include encryption and deduplication for enterprise-class security and efficiency.

THE IBM DIFFERENCE

IBM Resiliency Services is the only truly global provider to enable resiliency across all layers of the enterprise. It supports client facilities in numerous ways:

- IBM maintains more than 312 cloud resiliency centers around the world that provide compute and networking to support the technology layer, replication and data protection at the data and application layers, and world-class resiliency consulting methodology.
- These facilities comprise more than 10 million square feet of data center space for resiliency and disaster recovery operations, with more than 41,000 work area seats for recovery of workplace operations.
- IBM Resiliency Services are currently serving some 6,000 clients in 68 countries.
- IBM has a 100 percent success rate in meeting commitments to clients that have declared a disaster.

Partner with IBM and take full advantage of its five-decades-plus track record of data protection and business resiliency services.

For more information on IBM Resiliency Services, please visit IBM Data Center Services.

About IBM

IBM is a globally integrated technology and consulting company headquartered in Armonk, New York. Operating in more than 170 countries, IBM helps solve problems and provide an edge for businesses, governments and non-profits. The company develops and sells software and systems hardware and a broad range of infrastructure, cloud and consulting services. Today, IBM is focused on three strategic imperatives – to transform industries and professions with data, to remake enterprise IT infrastructure for the era of cloud, and to enable “systems of engagement” for enterprises.