



Enterprise Strategy Group | Getting to the bigger truth.™

RESEARCH HIGHLIGHTS

Real-world SLAs and Availability Requirements

Christophe Bertrand, Senior Analyst

JULY 2020

CONTENTS

Research Objectives 3

Research Highlights 4

While IT leads the decision making on data protection technology and sets the SLAs, there is an SLA gap. 5

Downtime is a business—not just IT—problem for which organizations have a low tolerance. 9

Cloud-based applications challenge traditional on-premises deployments in terms of recoverability experience and perception. 12

RPOs are particularly stringent across the board, though disparities exist among SaaS applications. 16

Organizations should engage in recovery testing to ensure data protection success. 20

Data protection and BC/DR processes are becoming increasingly intertwined with the cloud. 23

Research Methodology 27



Research Objectives


Data protection technologies and processes mean nothing unless objectives are not only established and aligned with business and IT objectives, but also measured and improved over time. As organizations increasingly shift to more hybrid and data-intelligent infrastructures, understanding real-world data protection and availability SLAs is becoming more critical for both IT practitioners and the vendors that provide supporting technology.

In order to gain insight into these trends, ESG surveyed 378 IT professionals at organizations in North America (US and Canada) personally responsible for or involved in data protection technology and process decisions for their organizations, specifically those pertaining to the ability to meet SLAs associated with applications/workloads. This research aimed to understand the current state of end-user deployments, identify gaps, and highlight future expectations. Tolerance for downtime, downtime metrics, and real-world SLAs in the context of actual data loss were studied against the backdrop of availability technologies and methods, including hybrid environments.

This study sought to:



Understand the current state of end-users' deployments, identify gaps, and highlight future expectations.




Evaluate the tolerance for downtime as measured by downtime metrics, and real-world SLAs in the context of actual data loss.



Identify organizations' availability technologies and methods, including hybrid cloud environments.

Research Highlights



While IT leads the decision making on data protection technology and sets the SLAs, there is an SLA gap.

With nearly two-thirds identifying IT as the primary drivers setting recovery time objectives (RTOs), it is clear that IT runs the show when it comes to determining the appropriate SLA levels to be met by the infrastructure or services it puts in place. It is also abundantly clear from the data that IT and technology experts not only set SLAs but also make the technical decisions for the determination and deployment of availability mechanisms.



Downtime is a business—not just IT—problem for which organizations have a low tolerance.

Organizations report that 1 in 3 applications are essential to the business, which in turn means more stringent SLAs. The prevalence of public cloud services for hosting vital workloads is evident in the fact that the average organization says more than half of their mission-critical applications run in the cloud.



Cloud-based applications challenge traditional on-premises deployments in terms of recoverability experience and perception.

The “one-hour window” is the crucial time objective in which mission-critical systems must be back up and running for the majority of organizations. While on-premises applications have collectively had the best recovery times, four in ten organizations state they have the highest level of confidence in public cloud infrastructure applications when it comes to meeting recovery SLAs.



RPOs are particularly stringent across the board, though disparities exist among SaaS applications.

Nine in ten respondents reported that their organization could not withstand in excess of an hour's worth of lost data before experiencing significant business impact. When it comes to data restoration in response to an outage, nearly half of respondent organizations indicate that their data is typically at least a week old.




Organizations should engage in recovery testing to ensure data protection success.

When it comes to pressure-testing BC/DR technologies, processes, and policies, only one in four organizations assess these capabilities weekly. This tendency toward infrequent BC/DR testing seemingly manifests itself in the fact that only about one-third of organizations report passing more than 80% of their tests.



Data protection and BC/DR processes are becoming increasingly intertwined with the cloud.

While 42% of respondents still leverage the traditional self-managed secondary site approach to BC/DR, things are changing. Indeed, more than three-quarters of organizations report using a third-party service, including the more than half using a cloud-based solution for BC/DR purposes. When it comes to the current usage trends for IT resiliency solutions, three of the five most commonly used technologies are cloud-based services.

A man with dark hair and a beard, wearing a headset, is looking down in a server room. The background is filled with server racks and computer monitors, creating a professional and technical atmosphere.

While IT leads the decision making on data protection technology and sets the SLAs, there is an SLA gap.

IT Sets the Tone for SLAs and Most Keep Senior Business Leaders Fully Apprised of Data Availability...or Outages

With nearly two-thirds (65%) identifying their IT organizations as the primary drivers setting SLAs, it is clear that IT runs the show when it comes to determining the appropriate SLA levels to be met by the infrastructure or services it puts in place. This makes IT the central player in the determination of the technical solution as well, with the support of business leaders. The data also means that technology vendors and partners have an opportunity to influence the decisions made by IT and provide guidance.


Organizations report a high level of transparency into outages due to the likely robust communication between IT and business stakeholders. This fosters increased coordination between business and IT and subsequently, greater transparency on technology capabilities. This transparency across the business is key as it can help organizations identify issues with and resources for SLA improvements.



65%

of IT organizations are the primary drivers of availability SLAs.

How IT reports system and data availability to business stakeholders.

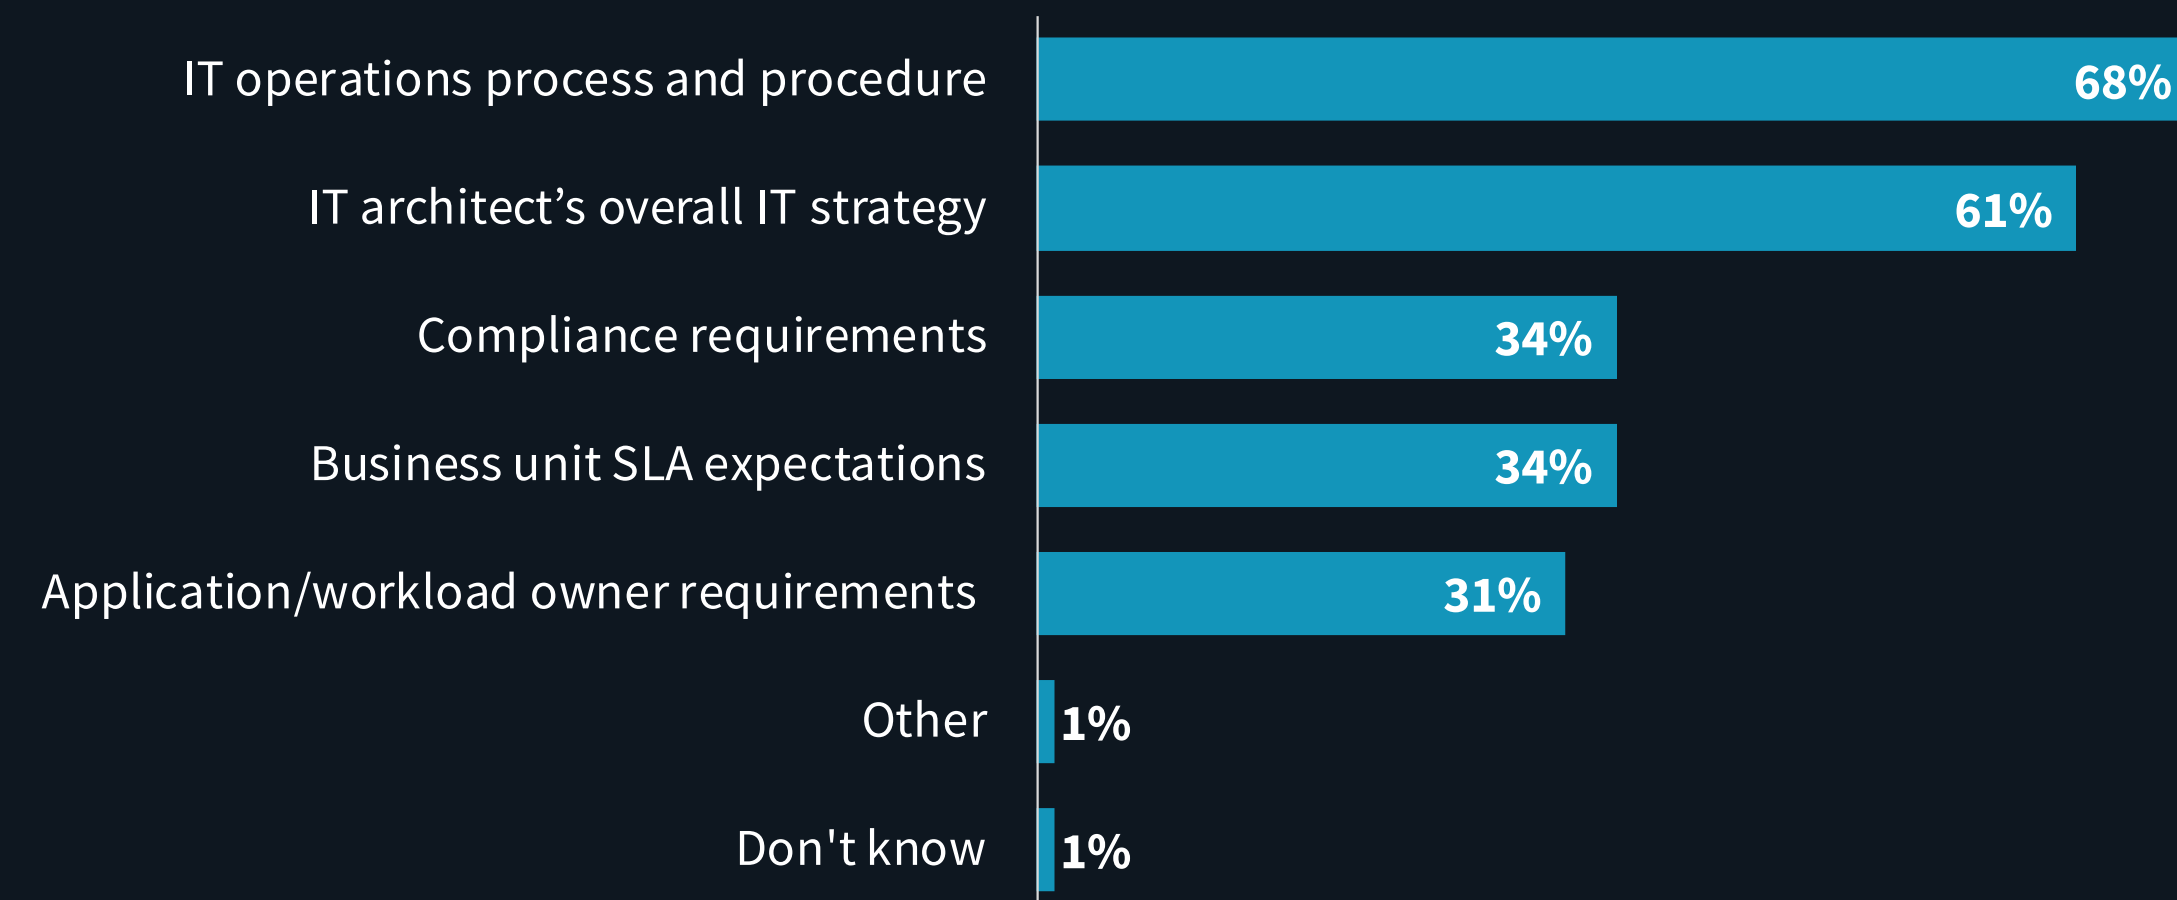


IT Operations and Architects Guide Availability Mechanism Choices, though RTO Tracking Is Not Fully Automated


It is abundantly clear from the data that IT and technology experts not only set SLAs but also make the technical decisions for the determination and deployment of availability mechanisms. In short: there is a clear emphasis on “building it right (architecture) so that it runs right (operations)” with fewer organizations paying attention to meeting line-of-business and/or compliance requirements. This last point is significant, as newer regulations are now taking a closer look at topics associated with availability and data.

Surprisingly, less than half (49%) of organizations have instrumentation in place to report on recovery time objectives (RTOs), which leaves a lot of room for error and misreporting, and saddles operations teams with a significant amount of manual work. To state the obvious, you can’t manage what you can’t measure. In order to get to consistently low RTOs, ESG would expect to see a vast majority of organizations with unified or centralized reporting capabilities.

Factors influencing usage of availability mechanisms.




How organizations track recovery times.



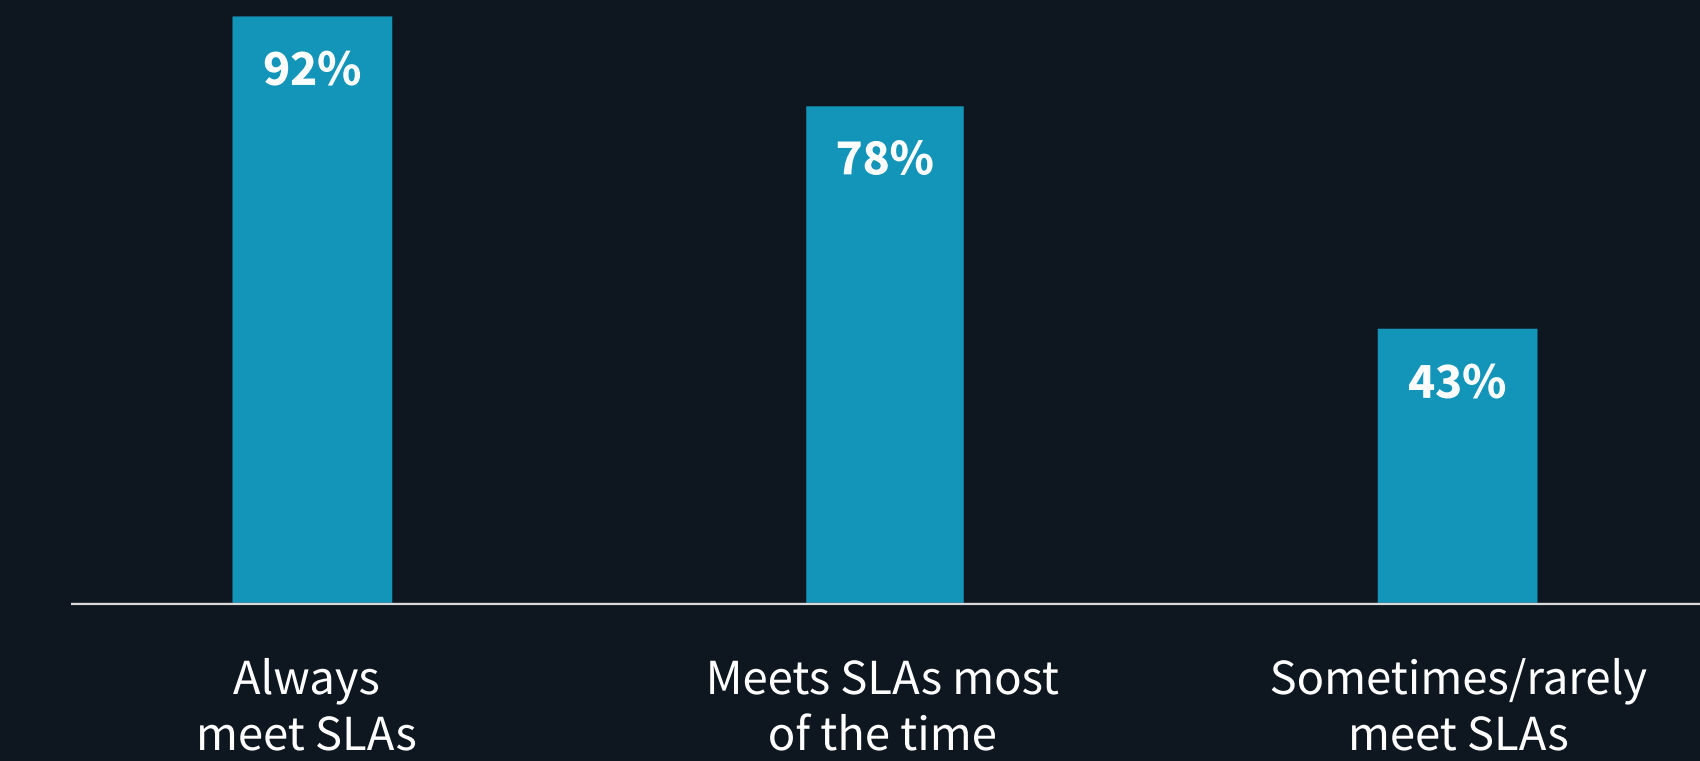
Only 1 in 5 Always Meet Recovery SLAs, but These Organizations Are Much Likelier to Track Outages with Systems Management Tools

Against a backdrop of constant data growth, as well as increased business and economic uncertainty, organizations will need to work to improve their ability to meet SLAs. Indeed, barely 1 in 5 organizations report they always meet their recovery objectives, while one-third say they sometimes or rarely meet them. The need for and benefits of greater reporting and SLA tracking tools are clearly seen in the fact that higher success meeting SLAs comes to those using tools. Specifically, organizations that indicate they always meet their recovery SLAs are more than twice as likely (92% versus 43%) as those saying they sometimes or rarely satisfy these RTOs to use systems management tools to track outages.

Frequency of actual recovery times meeting pre-established SLAs.




Percent of organizations using systems management tools to track outages based on frequency of recovery times meeting SLAs.



Organizations that indicate they always meet their recovery SLAs are more than twice as likely (92% versus 43%) as those saying they sometimes or rarely satisfy these RTOs to use systems management tools to track outages.



**Downtime is a business
—not just IT—problem
for which organizations
have a low tolerance.**




Mission-critical Applications Are on the Rise, especially in the Cloud

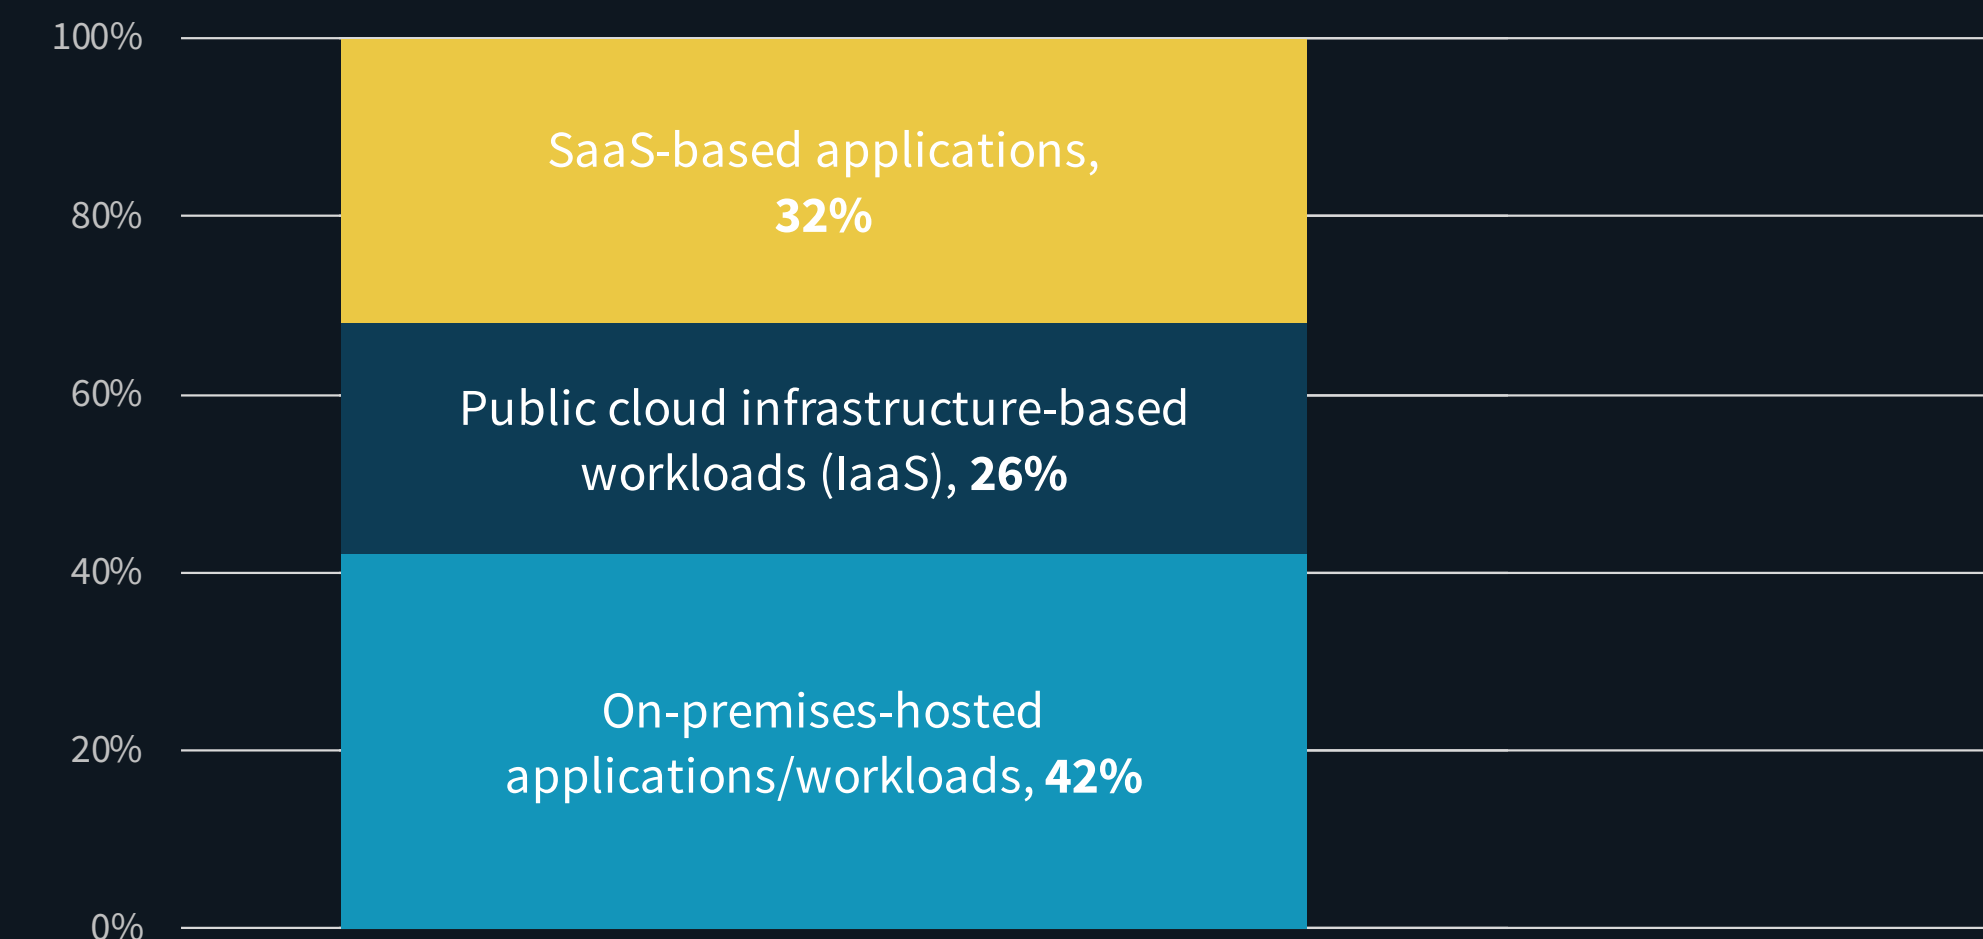
Data and applications run the business, but not every application is equal! Organizations report that 1 in 3 applications are essential to the business, which in turn means more stringent SLAs. This proportion is likely to increase over time as organizations continue on their path to digital transformation, making more data central to their business and processes. Given the existing and growing amount of data in most organizations' infrastructure, the ability of data protection solutions—including BC/DR—to handle scale is becoming more and more critical.

The prevalence of public cloud services for hosting vital workloads is evident in the fact that the average organization says 58% of their mission-critical applications run in the cloud. This leads to a number of challenges for the data protection infrastructure: can the solutions in place credibly deliver mission-critical SLAs for applications running on public cloud infrastructure or in SaaS environments? While organizations have full control over their processes and IT assets on-premises, the degree of control changes drastically as workloads migrate to cloud, making the protection and availability of data and applications fundamentally different.

1 in 3 production applications are deemed mission-critical.



Percent of “mission-critical” applications/workloads currently operated/run in each environment.



Most Concerning Impacts of Application Downtime

Organizations report that downtime has significant impact, ranging from economic to operational to legal in nature. Among the most visible are the direct operational efficiency consequences of an outage on IT, with one in five citing the diversion of resources from other business-critical projects as the most concerning impact. Effects more obvious to the bottom line are evident and top of mind for more than one-quarter of IT professionals who report that loss of customer confidence (15%) or revenue (11%) are their most concerning downtime impacts. While IT is directly affected by and responsible for downtime, it's really the whole business that suffers.

Most concerning impact of application downtime.



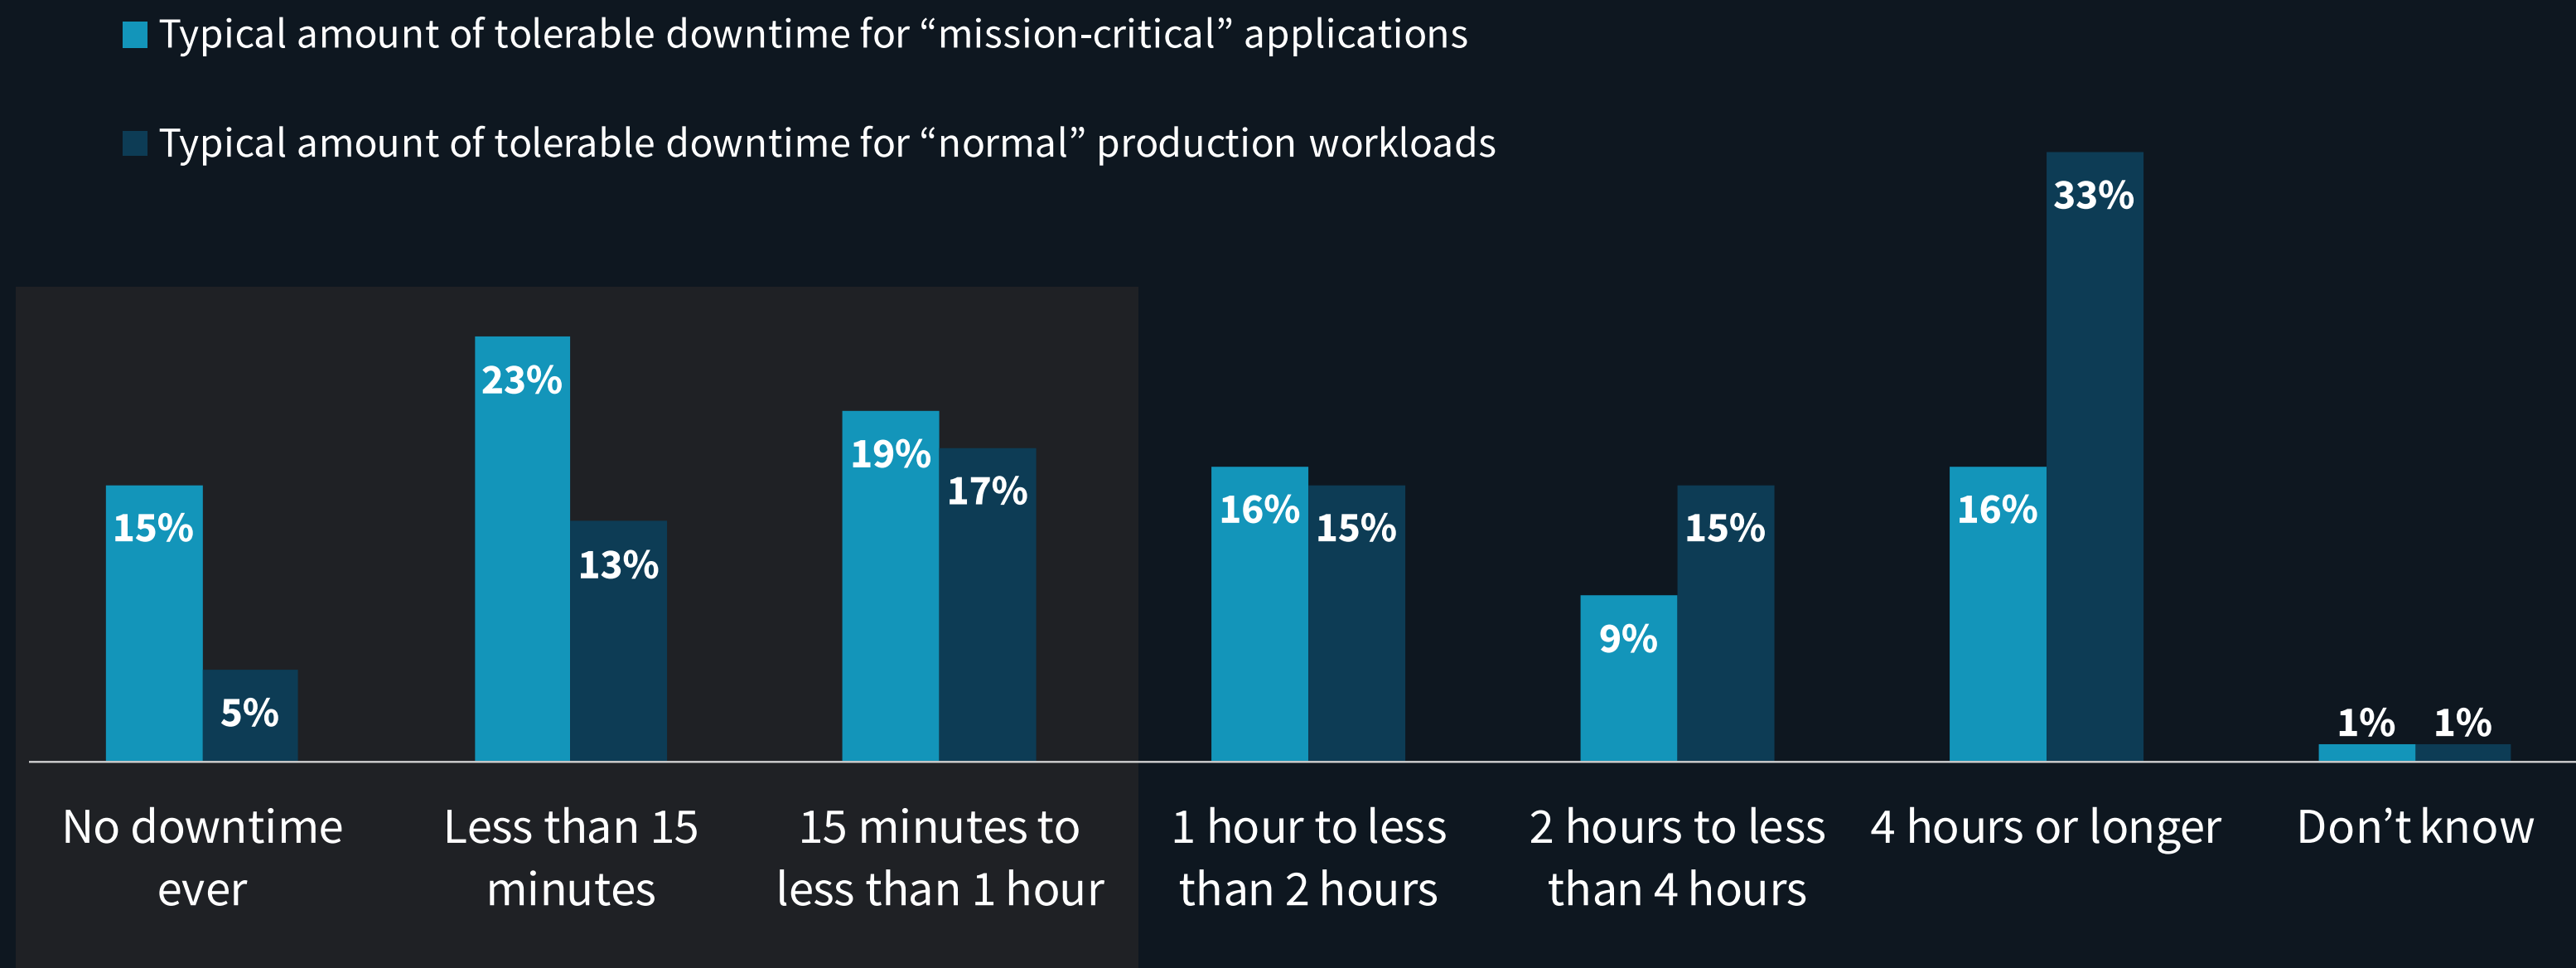


**Cloud-based applications
challenge traditional
on-premises deployments
in terms of recoverability
experience and perception.**

Most Can't Handle More than an Hour of Downtime for Mission-critical Applications

Time flies when systems are down and/or data is unavailable. The “one-hour window” is the crucial time objective in which mission-critical systems must be back up and running for the majority (57%) of organizations. It is also worth noting that 15% can tolerate no downtime for their mission-critical applications. Looking at all organizations collectively, the estimated mean for acceptable downtime for mission-critical applications is 2 hours, meaning that solutions that are deployed for recovery or failover must meet stringent requirements in a timely manner.

Amount of downtime applications/workloads tolerate before making the decision to “failover/recover” to a BC/DR secondary site.



ONE-HOUR WINDOW

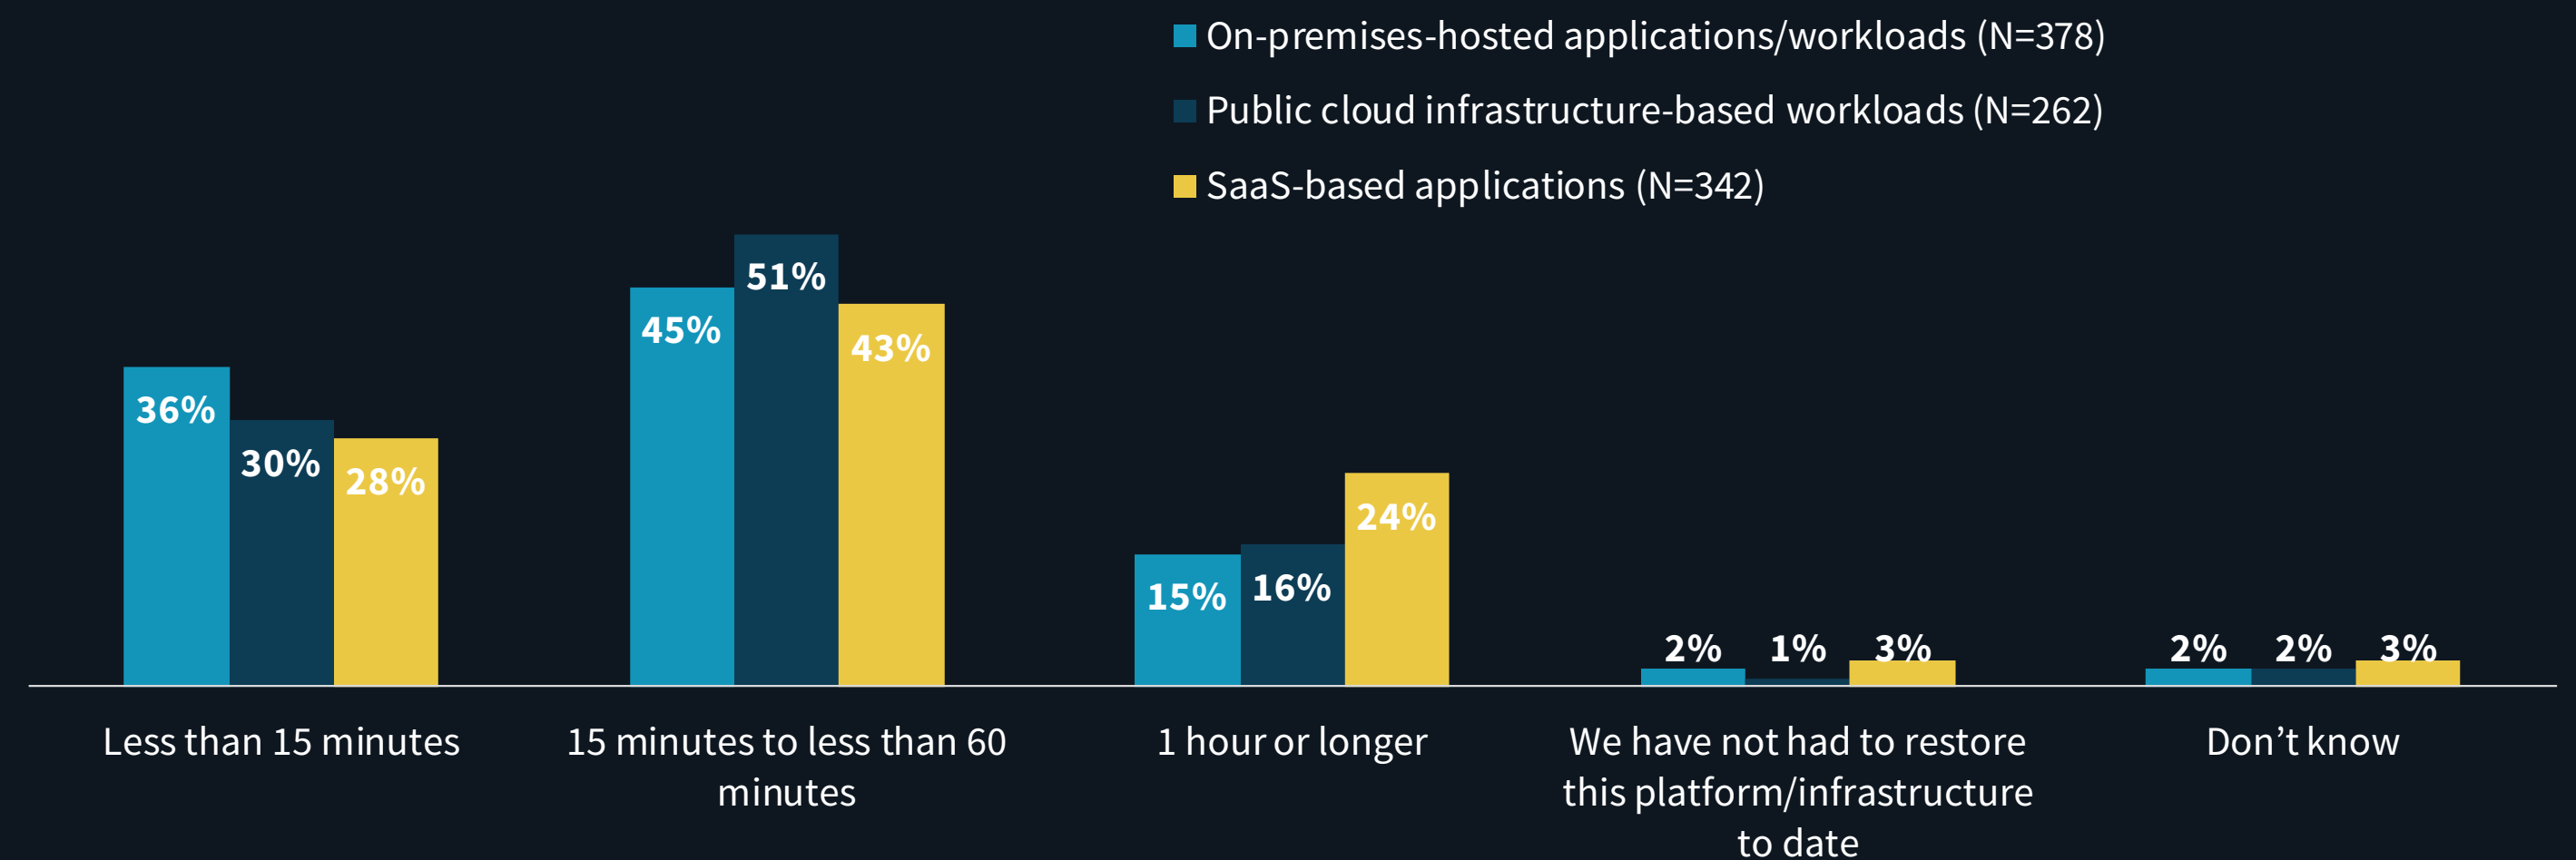
The “one-hour window” is the crucial time objective in which mission-critical systems must be back up and running for the majority (57%) of organizations.

On-premises Has Better Average Restore Times, though IaaS Has Perception Advantage

Looking at actual restore times, it is important to note that most organizations report meeting the desired one-hour window, but there are variations based on the type of infrastructure being leveraged. Specifically, applications running in on-premises data centers appear to be more easily recoverable than those hosted on public cloud infrastructure or SaaS-based applications. This is not necessarily surprising as having direct access over the infrastructure on-premises is an advantage for IT professionals who likely have more hands-on control over any potential issues in the data center. SaaS lags in recoverability which means that as organizations continue their journey to the cloud and leverage the associated infrastructure and services, in addition to ensuring they are employing the proper data protection mechanisms, they must consider the potential impact of recoverability of their systems and data, and whether the tradeoff is acceptable.

Confidence—or perceived confidence—is a major influencer in buying decisions, and this is especially true when it comes to data and application recoverability. And while on-premises applications have collectively had the best recovery times, 40% of organizations state they have the highest level of confidence in public cloud infrastructure applications when it comes to meeting recovery SLAs. This raises the question of what will happen with on-premises applications if the general availability and recoverability sentiment keeps favoring cloud platforms.

Time it typically take to restore the functionality of applications/workloads and associated data after an outage.



Application model in which organizations have the highest level of confidence in meeting recovery SLAs.



40%

Public cloud infrastructure-based workloads (IaaS)



36%

On-premises-hosted applications/workloads



24%

SaaS-based applications

Duration Of Longest Outage Reveals a Noticeable SLA Gap


For most organizations, applications must be back up and running in less than one or two hours depending on their criticality, yet some outages can last much longer, resulting in organizations missing their SLAs. It is almost inevitable that organizations will be affected by an outage, as evidenced by the fact that 80% report having experienced one within the last year. So the focus should be less on if and more on when an outage will occur, especially when a bad outage can extend for hours beyond an RTO. And while nearly one-third of organizations report their longest outage was less than one hour, the overall average was six hours, which is significantly longer than the stated objectives for mission-critical or normal workloads.



79%

of organizations had at least one production application outage in the past year.

Duration of longest application outage in past year.



ESTIMATED MEAN:
6 hours






RPOs are particularly stringent across the board, though disparities exist among SaaS applications.

Mission-critical Data Loss Tolerance Is (Understandably) Low

Mission-critical data loss tolerance is—not surprisingly—low. Indeed, 90% of respondents reported their organization could not withstand in excess of an hour’s worth of lost data before experiencing significant business impact, equating to an estimated mean RPO of 22 minutes. It is worth noting that 15% of organizations actually report objectives of no data loss at all, which means putting in place availability technologies and highly redundant infrastructure and solutions that support this objective.

Amount of mission-critical data that can be lost without significant impact to the business.



SaaS RPO Tolerance

SaaS workloads have become critical in many organizations and for this reason, organizations have high RPO expectations for these cloud-based applications. It is important to distinguish that availability of the service itself, and what the provider does for their own backup purposes, should not be conflated with the need and responsibility for organizations to perform their own backups to ensure the recoverability of their data. Microsoft Office 365 and Salesforce.com top the list of most stringent SLAs, with more than one-third of these users reporting RPOs of less than five minutes' worth of lost data.

■ Microsoft O365 (N=301)

🕒 ESTIMATED MEAN:
27 minutes

■ Salesforce.com (N=171)

🕒 ESTIMATED MEAN:
30 minutes

■ Netsuite (N=84)

🕒 ESTIMATED MEAN:
44 minutes


■ Dropbox (N=148)

🕒 ESTIMATED MEAN:
47 minutes

■ Slack (N=73)

🕒 ESTIMATED MEAN:
52 minutes

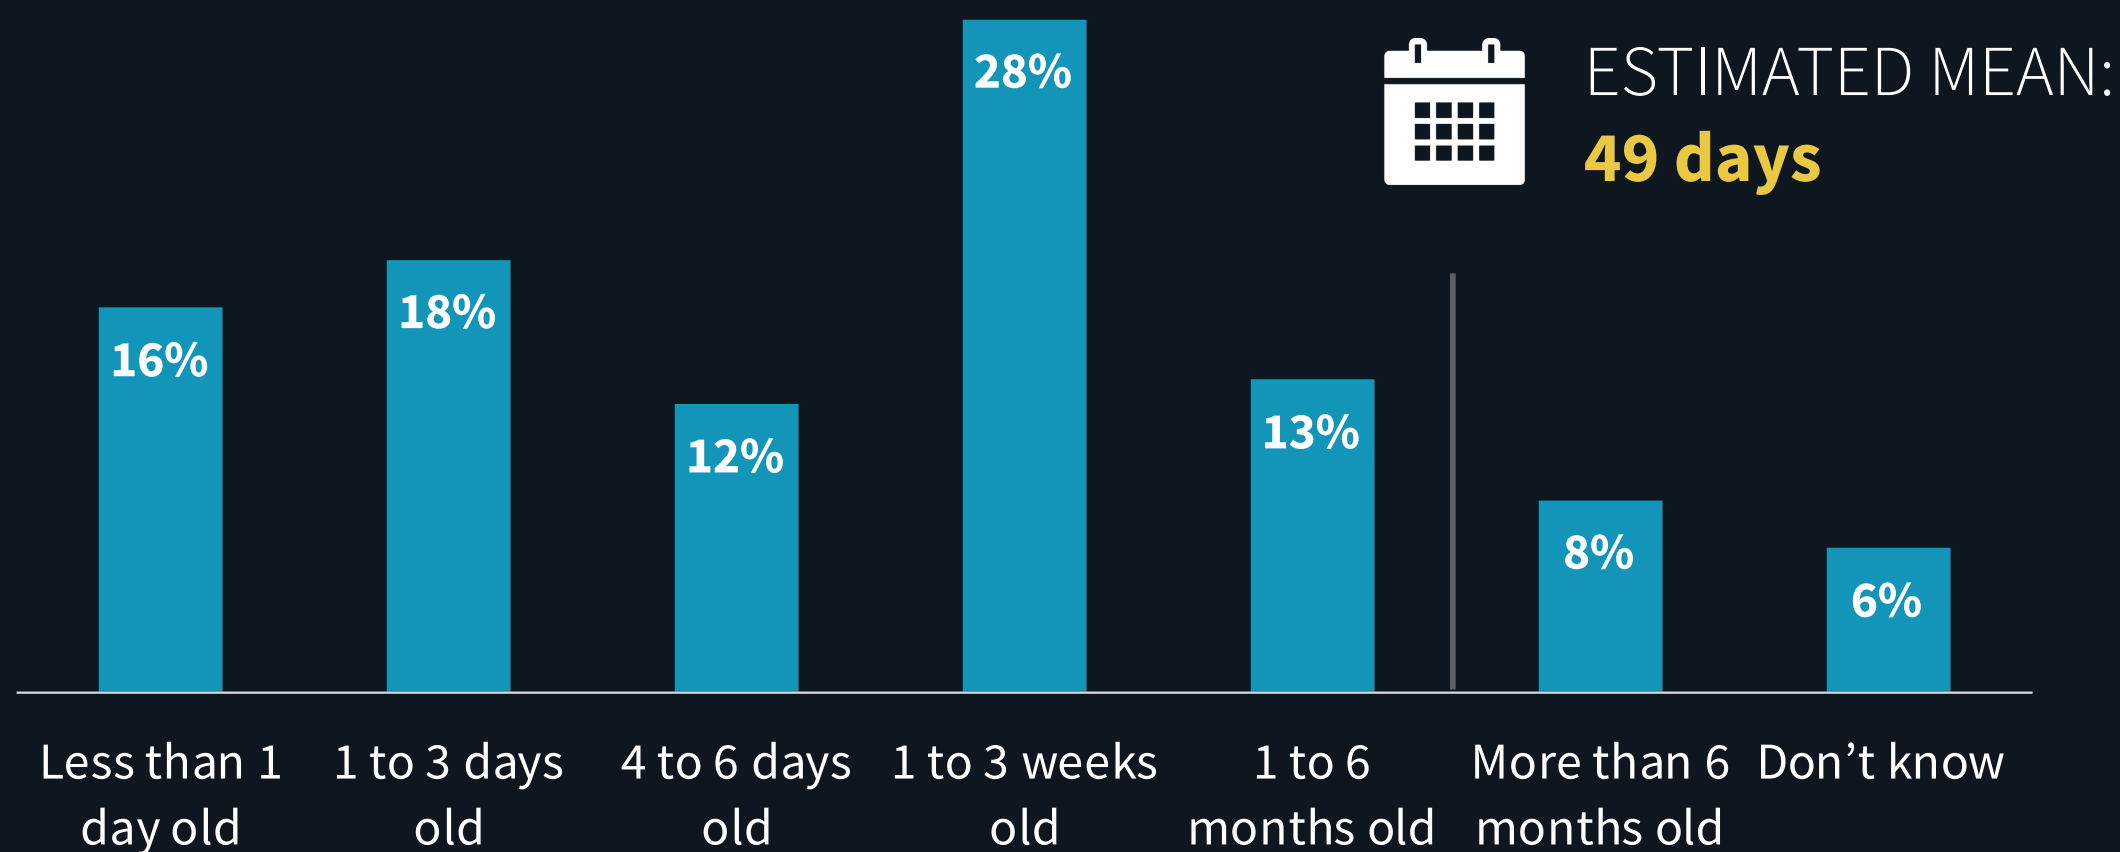
Amount of acceptable data loss for different SaaS applications.



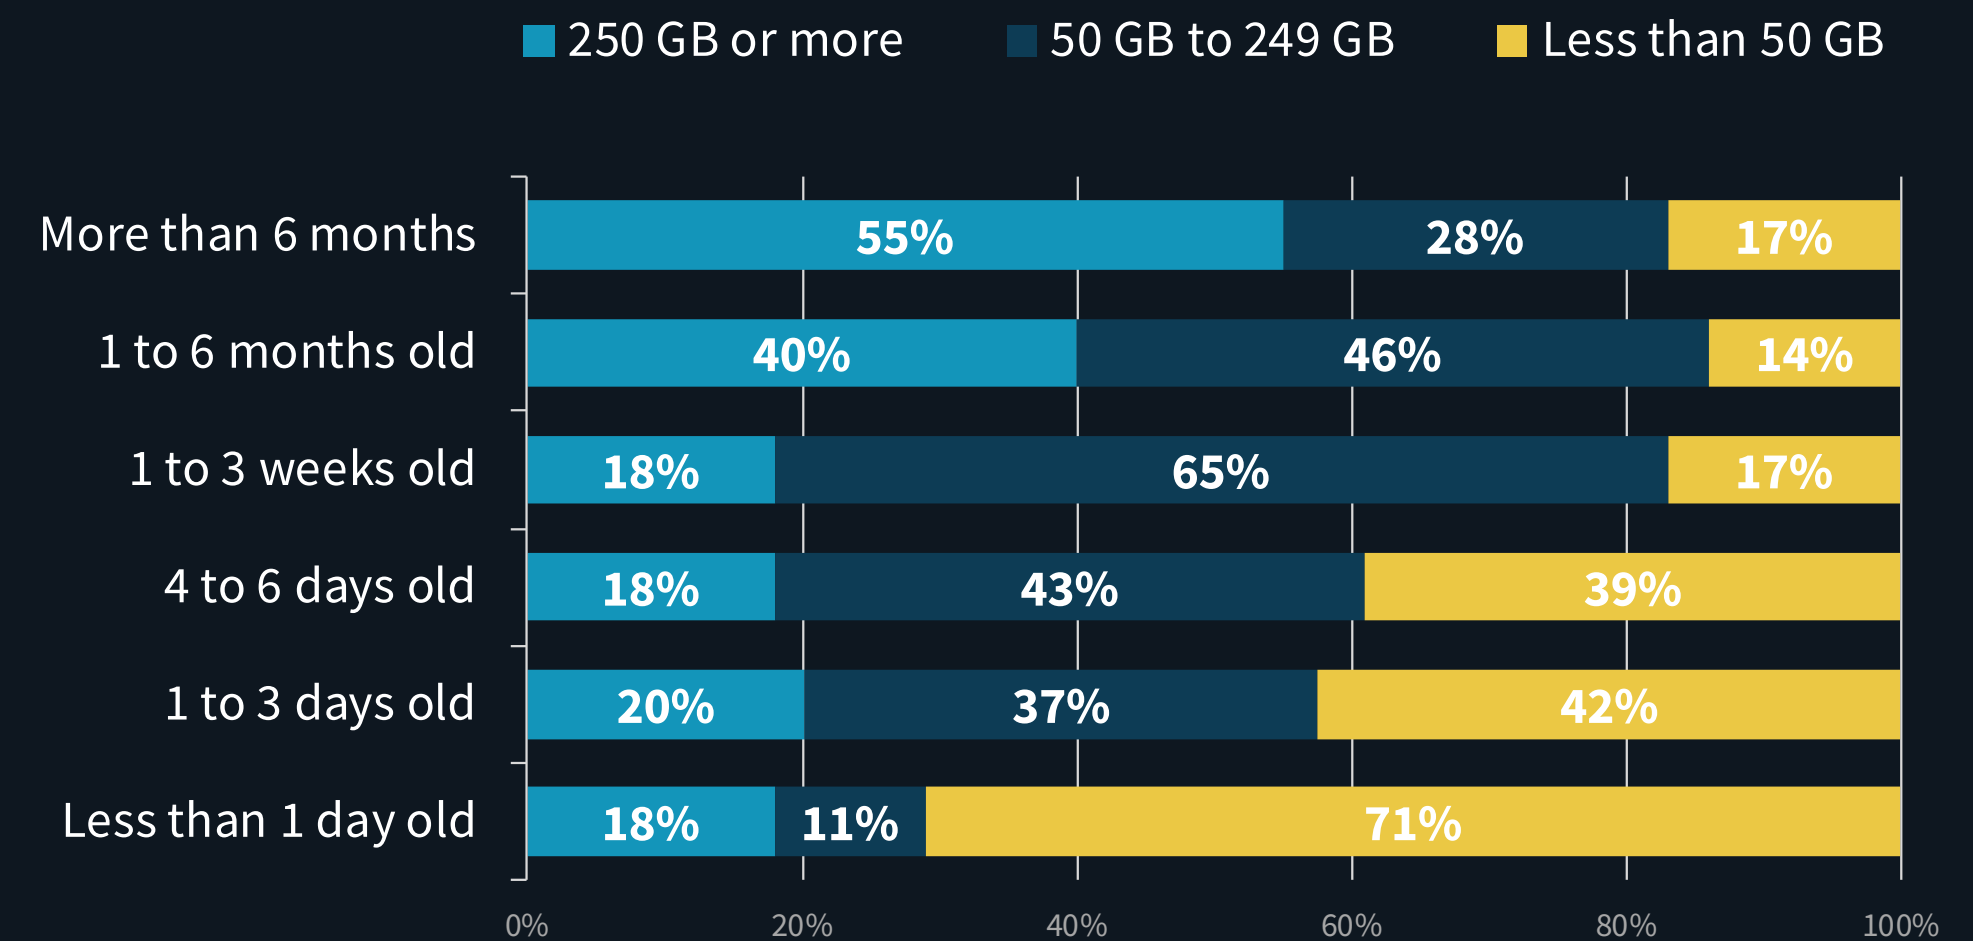
Recovered Data Tends to Skew Older, Especially Large Recoveries

When it comes to data restoration in response to an outage, nearly half (49%) of respondent organizations indicate that their data is typically at least a week old, with the overall average age being 49 days. The amount of data restored as part of a typical recovery effort tends to skew older, especially large recoveries. The vast majority (71%) of one-day old recoveries are less than 50 GB, which could simply be explained by the nature of these types of recoveries, such as corrupted tables, deleted files, etc., and make it much easier to deliver on operational recovery SLAs. Past the one-day window, a significant and progressive jump occurs towards larger recoveries—time means more data, and likely more recovery time and resources.

Average age of data restored after an outage or incident.



Amount of data restored as part of a typical recovery based on the average age of recovered data.



A woman with dark hair and glasses is shown in profile, looking towards the left. She is wearing a dark green jacket. In the background, there is a large screen displaying various data visualizations, including line graphs and bar charts in shades of blue, green, and yellow. The overall scene is dimly lit, suggesting a professional or office environment.


**Organizations should engage
in recovery testing to ensure
data protection success.**

BC/DR Recovery Testing Is Being Done, but There Is a Big Opportunity for Improvement


When it comes to pressure-testing BC/DR technologies, processes, and policies, only one in four organizations assess these capabilities weekly. This means the majority of organizations are testing on a monthly cadence or even less frequently, though few report never conducting testing. Given the frequency of cyber-attacks and other infrastructure crippling issues, it would be more encouraging to see a larger number of organizations testing on a weekly basis, since in this area, the faster organizations can uncover issues, the faster they can fix them.

In fact, this tendency toward infrequent BC/DR testing seemingly manifests itself in the very lackluster results of these tests. Specifically, only about one-third of organizations report passing more than 80% of their tests, and overall, 63% of BC/DR tests are successful, which is low relative to the expectations established by typical enterprise-class SLAs.

Frequency with which BC/DR recovery tests are executed.




Percentage of routine tests for BC/DR preparedness that succeed or “pass.”

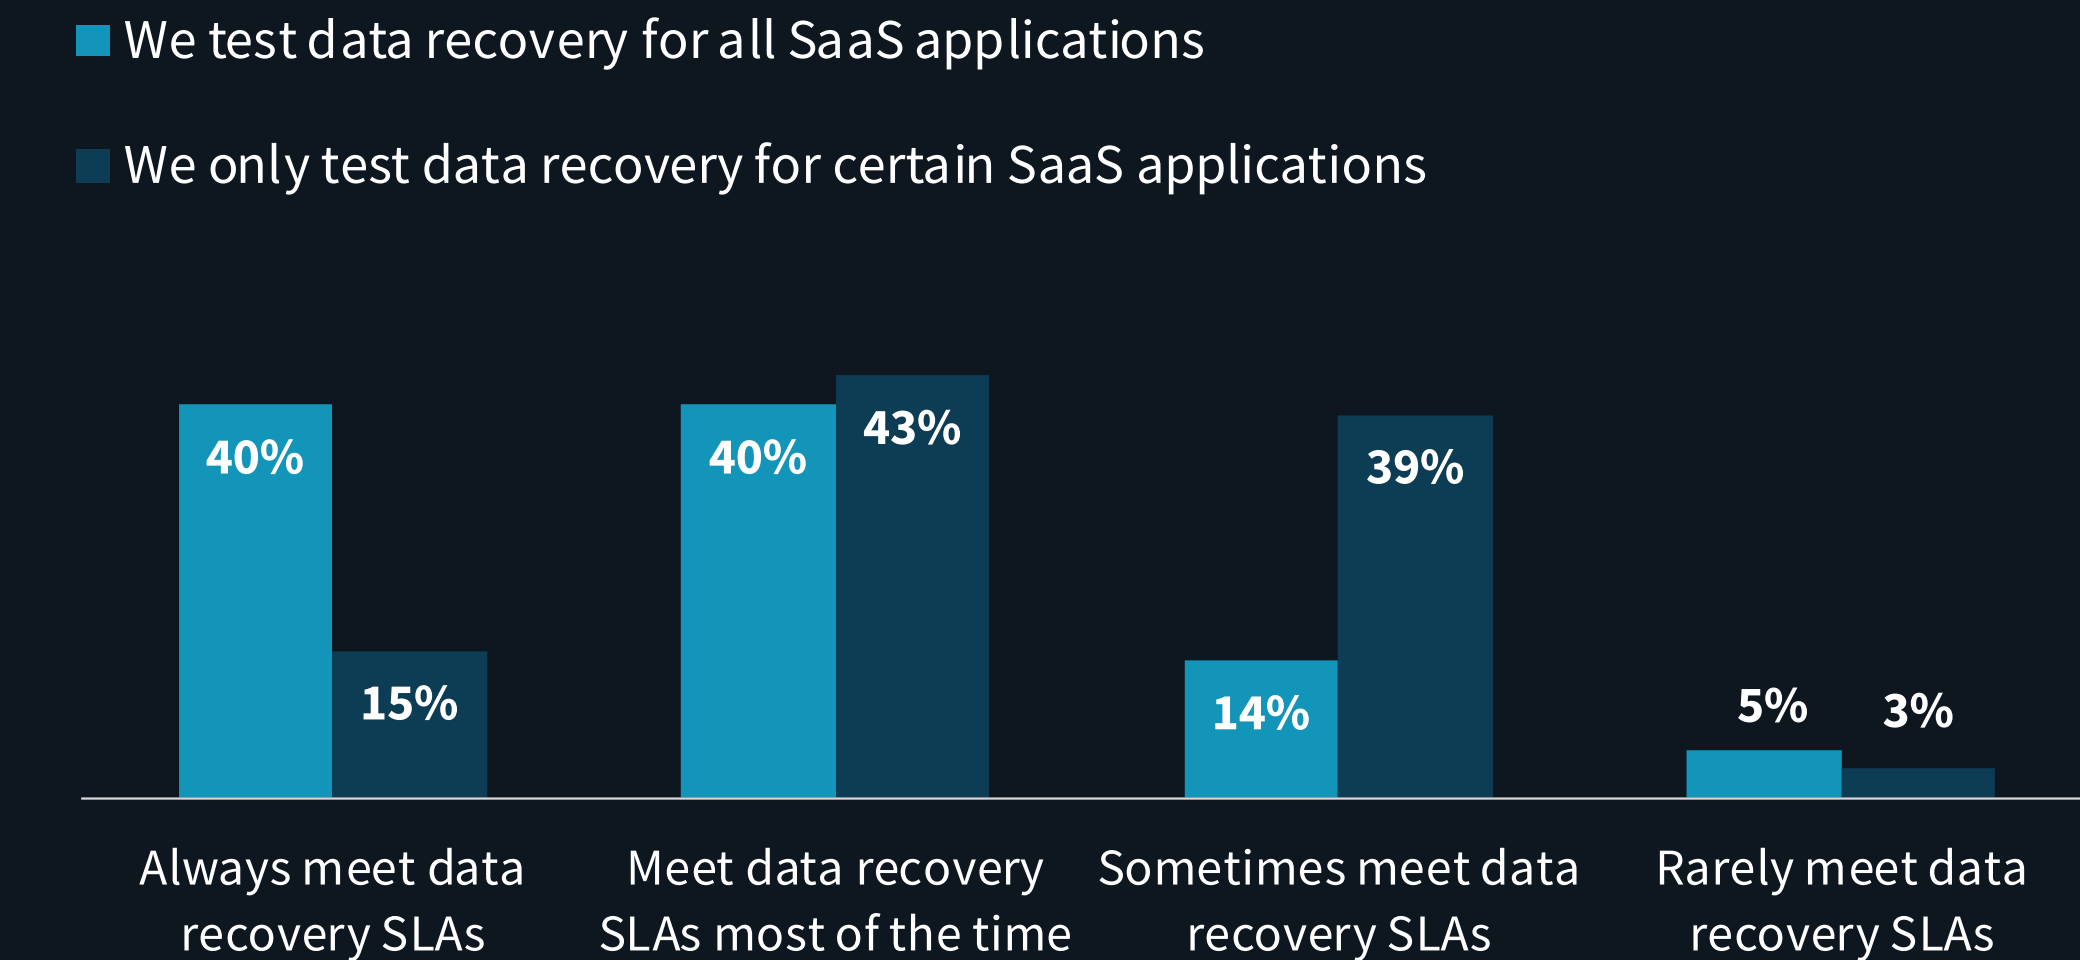


More Extensive Data Recovery Testing Yields Higher SLA Success Rates for SaaS Applications

While public cloud services are relative newcomers in the world of IT, a majority of organizations opt to test their applications or workloads in these environments for data recovery. In an ideal situation, the proportion of “all” testers should be closer to 90%, since many applications and workloads—whether mission-critical or not—are exposed to stringent recovery and availability SLAs. Further investigation into those organizations using SaaS shows that testing is a good thing for ensuring recovery success. Specifically, organizations that test all their SaaS applications are nearly three times likelier to always meet data recovery SLAs (40% versus 15%) than organizations that cherry pick the applications they test. The lesson is obvious: test everything to achieve higher levels of success.



Success rate of SaaS data recoveries based on extent of recovery testing.



A complex network diagram with blue nodes and white lines on a dark background. The nodes are small blue dots, and the lines are thin white lines connecting them. The network is dense and multi-layered, with many nodes and connections. The overall appearance is that of a large, interconnected system, possibly representing a cloud network or a data center.

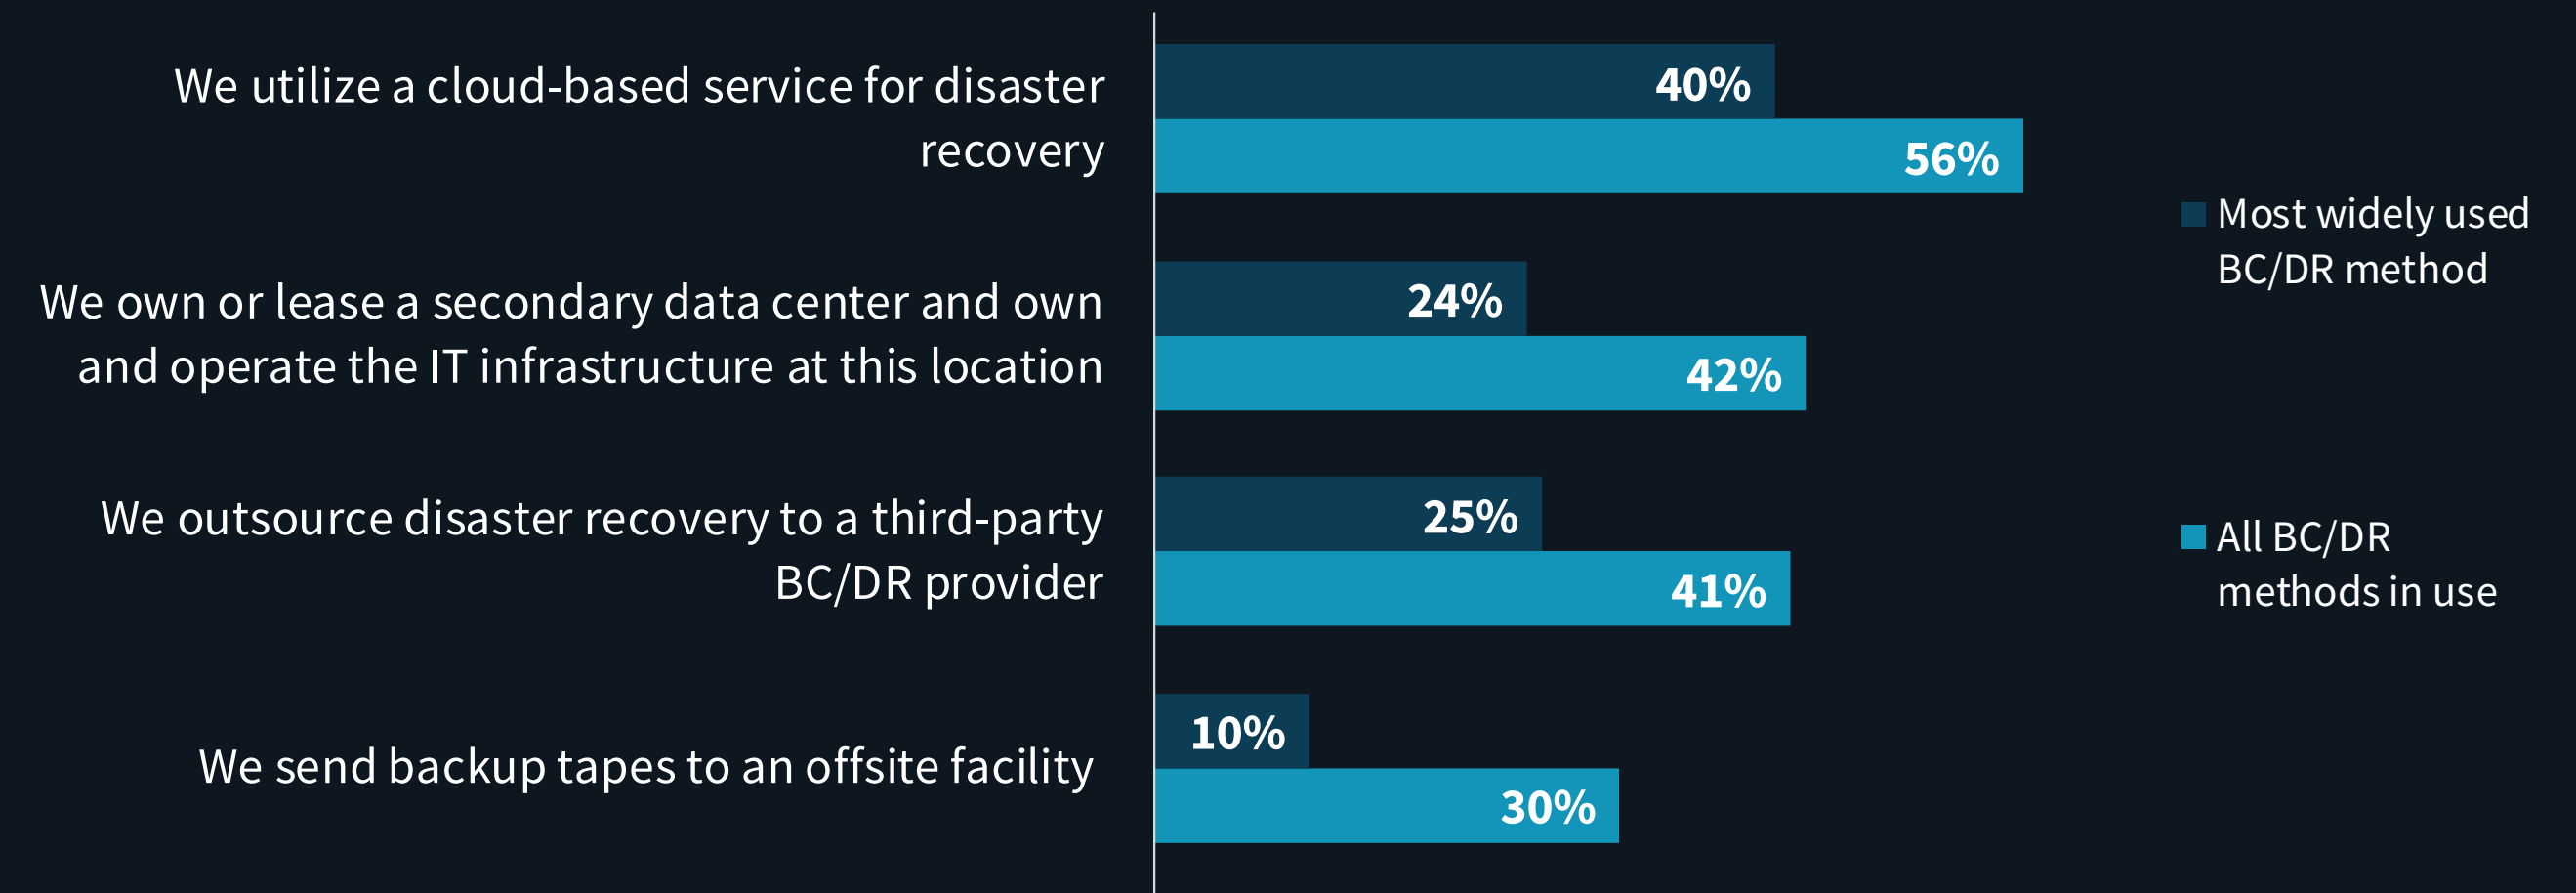
Data protection and BC/DR processes are becoming increasingly intertwined with the cloud.

Cloud Is Becoming the Go-to Approach for Secondary DR Sites


Secondary sites have always been the foundation of BC/DR strategies and have traditionally relied on a secondary data center directly owned and managed by organizations themselves. And while 42% of respondents still leverage this approach to some extent, things are changing. Indeed, 79% of organizations report using a third-party service, including the more than half (56%) using a cloud-based solution for BC/DR purposes. It should also be noted that a trusted legacy approach like tape offsite services (30%) is still in the picture.

Focusing on cloud services for BC/DR, organizations use a variety of consumption models, often concurrently, though in time, more standardization or focus on one primary solution is likely. Moving forward, data protection services like BaaS and DRaaS will be by far the most favored options over leveraging IaaS, which is typically more DIY in nature. Also, among data protection services, organizations anticipate favoring DRaaS over BaaS, which could be attributed to the growing need for more stringent SLAs and business interruption mitigation mandates. In time, expect to see backup and DR merge in some form of SLA continuum that can be dialed up or down, and delivered as a service. It is also important to note that these services can yield ecosystem opportunities for channel or vendor partners as exemplified by the provider-managed DRaaS option, which is challenging IaaS for the top spot.

Current approaches to disaster recovery.



Cloud services used for BC/DR.



Resiliency Technologies Are Becoming Cloudy

Cloud-based solutions are clearly gaining favor with IT professionals when it comes to data protection mechanisms, and this extends to IT resiliency technologies. In fact, when it comes to the current use of these solutions, three of the five most commonly used technologies are cloud-based services in the form of cloud storage, SaaS with native resiliency capabilities, and BaaS.

Top five currently used IT resiliency technologies.

**42%**

IT hardware with “fault tolerant” architecture.

**38%**

Cloud-based storage for data survivability.

**38%**

Cloud-based applications (SaaS) with native resiliency between sites.

**30%**

Software with “failover” or “orchestration” recovery capabilities.

**28%**

Backup-as-a-service (BaaS) with cloud-based restoration capabilities.



Pure Storage (NYSE: PSTG) gives technologists their time back. Pure delivers a modern data experience that empowers organizations to run their operations as a true, automated, storage as-a-service model seamlessly across multiple clouds. One of the fastest-growing enterprise IT companies in history, Pure helps customers put data to use while reducing the complexity and expense of managing the infrastructure behind it. And with a certified customer satisfaction score in the top one percent of B2B companies, Pure's ever-expanding list of customers are among the happiest in the world.

Modern data protection isn't just an insurance policy. It's a vital component of the contemporary data center that encompasses multiple platforms and technologies, delivers efficient protection of critical data and applications with blazing fast restores, and lets you derive real business value from your protection data. It can also be a critical component in risk mitigation, protecting organizational reputation and customer satisfaction.

If your data protection experience isn't all it could be, let Pure Storage show you the easiest route to the modern data protection experience.

LEARN MORE

About ESG

Enterprise Strategy Group is an IT analyst, research, validation, and strategy firm that provides market intelligence and actionable insight to the global IT community.




Research Methodology


To gather data for this report, ESG conducted a comprehensive online survey of IT professionals from private- and public-sector organizations in North America (United States and Canada) between March 20, 2020 and March 28, 2020. To qualify for this survey, respondents were required to be IT professionals personally responsible for or involved in data protection technology and process decisions for their organizations, specifically those pertaining to the ability to meet SLAs associated with applications/workloads. All respondents were provided an incentive to complete the survey in the form of cash awards and/or cash equivalents.

After filtering out unqualified respondents, removing duplicate responses, and screening the remaining completed responses (on a number of criteria) for data integrity, we were left with a final total sample of 378 IT professionals.


RESPONDENTS BY NUMBER OF EMPLOYEES



RESPONDENTS BY AGE OF COMPANY



RESPONDENTS BY INDUSTRY



All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.



Enterprise Strategy Group is an IT analyst, research, validation, and strategy firm that provides actionable insight and intelligence to the global IT community.

© 2020 by The Enterprise Strategy Group, Inc. All Rights Reserved.