WHY ORGANIZATIONS STILL STRUGGLE TO DIGITALLY TRANSFORM & INNOVATE

Examining the real-world impacts when service Availability goals are not met

2017 Veeam Availability Report

Full Report

Contents

Research Introduction

We've never been as dependent on technology as we are today, nor have we had as many critical business functions and personnel that rely so heavily on their data. For organizations to achieve their business goals, they are looking to digital transformation and the cloud to deliver more efficient, agile, and reliable services to meet user needs. As part of this transformation, IT teams have to do an increasingly better job to ensure their systems' Availability and protection, and are looking at heterogeneous and hybrid environments to drive efficiencies and performance optimization.

Veeam®, which has been helping organizations achieve higher operational Availability for over a decade, commissioned the Enterprise Strategy Group (ESG) to deliver its sixth annual Veeam Availability Report. This reports seeks to (1) quantify whether organizations are meeting their Availability goals, (2) assess the impacts to organizations that are insufficient in their service levels, and (3) understand how these challenges are affecting strategic business such as digital transformation.

Executive Summary

Organizations continue to struggle with Availability assurance within their IT environments:

Four out of five organizations recognize that they have an "Availability Gap." In this year's research, 82% of respondents recognized the inadequacies of their recovery capabilities when compared with SLA expectations of their business units, which is consistent with the last two annual surveys. Even while some organizations are endeavoring to improve, the heightening expectations of the business units, combined with the ever-evolving and diversifying IT landscape and the move to heterogeneous and hybrid environments, continue to create challenges in providing adequate service Availability of IT. This causes wider issues for the business in terms of customer and employee confidence.

82%
of enterprises are facing a gap between user demand and what IT can deliver, or an 'Availability Gap'

On average, enterprises experience direct financial costs of $21.8M (US) due to these Availability and Protection Gaps, with the acknowledgement that these numbers will vary due to industry, company size, and locale. These Availability and Protection Gaps also impact truly strategic business modernization initiatives:

$21.8M
is an average financial cost of Availability and Protection Gaps for the enterprises

- **82% of organizations' virtualization deployments and strategies have been affected** by their data protection solution.
- **66% of organizations report digital transformation initiatives are being hindered** (either significantly or somewhat) by unplanned downtime or insufficient application Availability.

66%
of enterprises admit that digital transformation initiatives are being held back by unplanned downtime

Six out of seven organizations lack a high level of confidence in their ability to reliably protect/recover data within their virtual environments. 85% of respondents rated themselves less than very confident in their organization's current capabilities regarding virtual machine backup and recovery. With virtualization being the underpinning of every modern IT environment, including on-premises and cloud-hosted, *any* response other than "very confident" in 2017 is unacceptable.

Three out of four organizations acknowledge that they have a "Protection Gap." Also consistent with surveys from years past, 72% of respondents this year are unable to protect their data frequently enough to ensure that their business units' expectations against data loss are met.

Moreover, the impacts of downtime and data loss can reach far beyond direct economic loss:

- Externally, half of organizations believe Availability challenges can lead to loss of customer confidence, which can also impact brand integrity, reduce stock price, and cause revocation of licenses/accreditations.
- Internally, many believe Availability challenges can lead to loss of employee confidence, which often results in diverting resources away from long-term or business-critical projects.

The findings of this study are consistent with past ESG research and past Veeam® reports, all clearly illustrating that organizations must reconsider their data Availability, protection, and recovery capabilities. A company's failure to better align these key resiliency capabilities with the expectations of their business constituents will continue to put their organizations at risk and hinder innovation and digital transformation strategies.

Why do Data Availability and Protection Continue to Challenge Organizations?

Too many organizations continue to struggle with data recovery in their efforts to ensure the Availability of their virtualized systems. In fact, only 15% of surveyed decision makers are very confident in their current solution's ability to reliably back up and recover virtual machines within their SLAs.

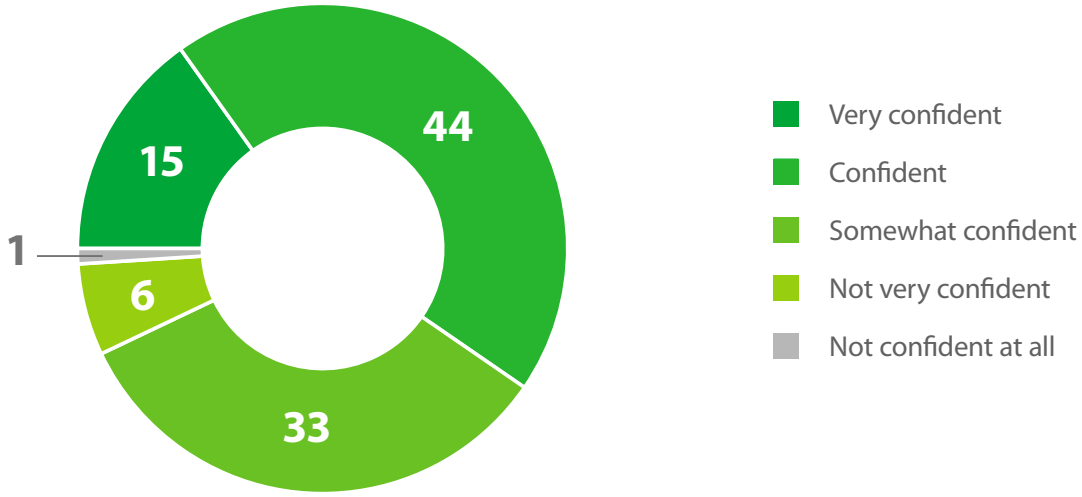

Figure 1. How confident are you in your organization's current primary solution's ability to reliably backup/recover VMs and recover what you need within your SLAs? (Percent of respondents, N=1,060)

That is an appallingly low percentage of real confidence. Any organization that is not "very confident" in its ability to protect the foundational structure of its modern data center, and provide total data and applications availability, should be reexamining its strategy and the technologies that it depends on.

Unfortunately, the pervasive lack of complete confidence is well founded. Consider the fact that the respondents surveyed say they are meeting their recovery time objectives (RTOs) and recovery point objectives (RPOs) only 72% of the time. In more than one out of four attempts, their recovery effort either fails, takes too long, or recovers an inadequate amount of data.

15%

of decision makers are very confident in their current solution's ability to back up and recover virtual machines

Sentiment Regarding the Availability and Protection Gaps

Across many enterprises, respondents almost universally acknowledge that their IT teams cannot recover fast enough, reliably enough, or thoroughly enough. Veeam refers to these challenges as the Availability Gap and the Protection Gap.

- **Availability Gap** refers to the difference between the service levels expected by business units, and an organization's ability to deliver the application and information that users demand.
- **Protection Gap** refers to an organization's tolerance for lost data being exceeded by IT's inability to protect that data frequently enough.

Alarmingly, more than four out of five organizations surveyed recognize they have an Availability Gap, and nearly three out of four organizations recognize they have a Protection Gap.

4 of 5

organizations surveyed recognize they have an Availability Gap

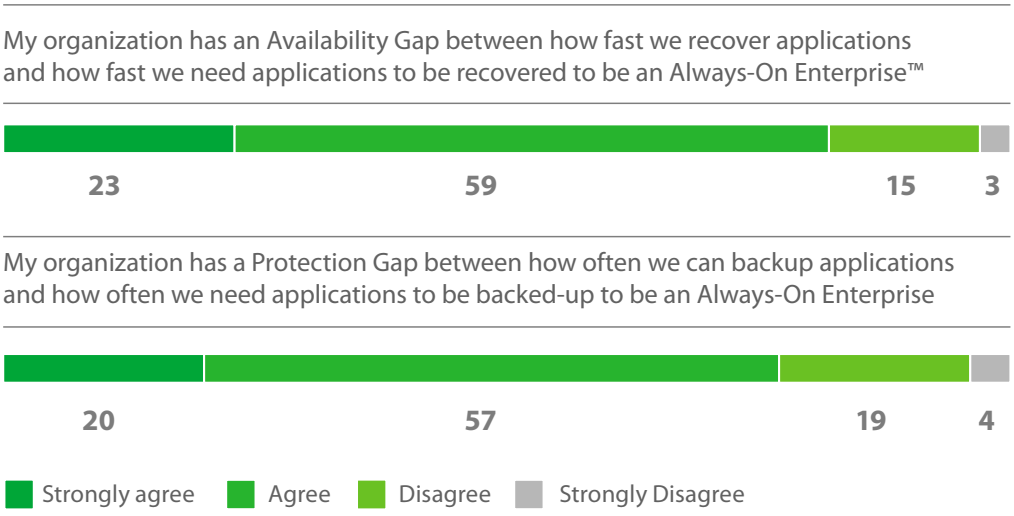

Figure 2. Please rate your agreement with the following statements. (Percent of respondents, N=1,060)

It is notable that so many decision makers are acknowledging, for the third year in a row, that they continue to suffer an Availability Gap, with similar trending apparent in relation to the Protection Gap—each with nearly equal year-over-year occurrence.

This is more than just a pervasive problem; it is also an ongoing one.

3 of 4

organizations surveyed recognize they have a Protection Gap

Recoverability Realities and Ramifications

It is vitally important for an organization to recognize the precariousness of its IT systems and avoid trivializing downtime when it happens.

On average, more than one in four (27%) servers suffer at least one unplanned outage each year. To understand how ESG calculated averages, means, and medians in this report, please see appendix one[1].

1 in 4

servers suffer at least one unplanned outage each year

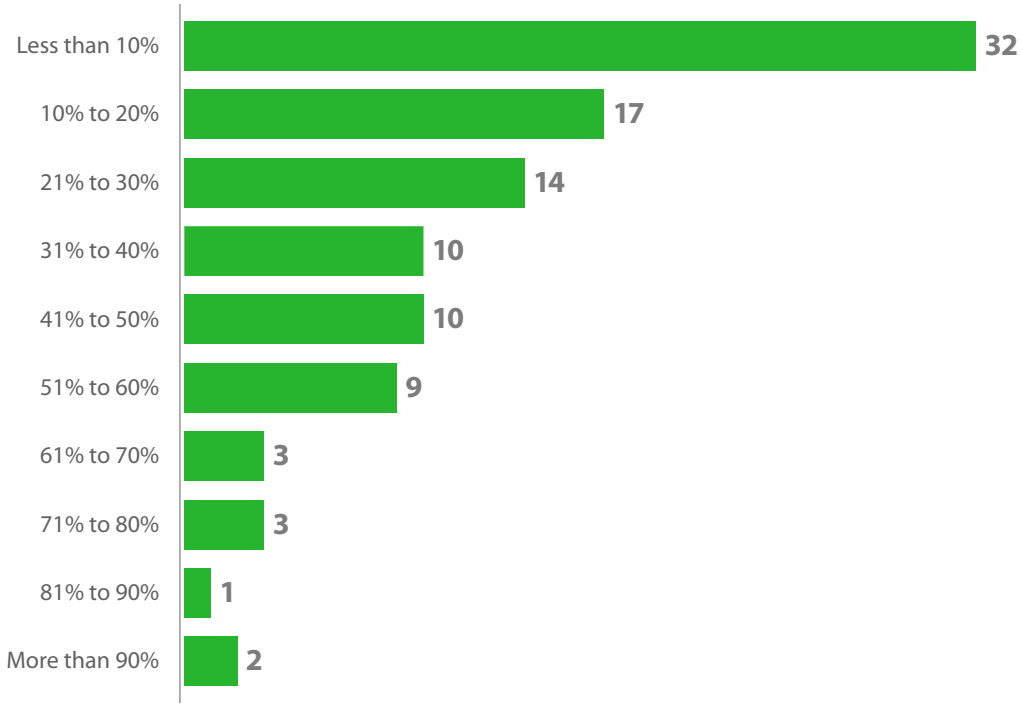

Figure 3. What percentage of your organization's production servers suffer at least one unplanned outage per year? (Percent of respondents, N=1,005)

And, some of those outages last for quite some time.

[1] See Notes Regarding the Calculations and Data Displayed within This Report in Appendix: Research Methodology and Respondent Demographics

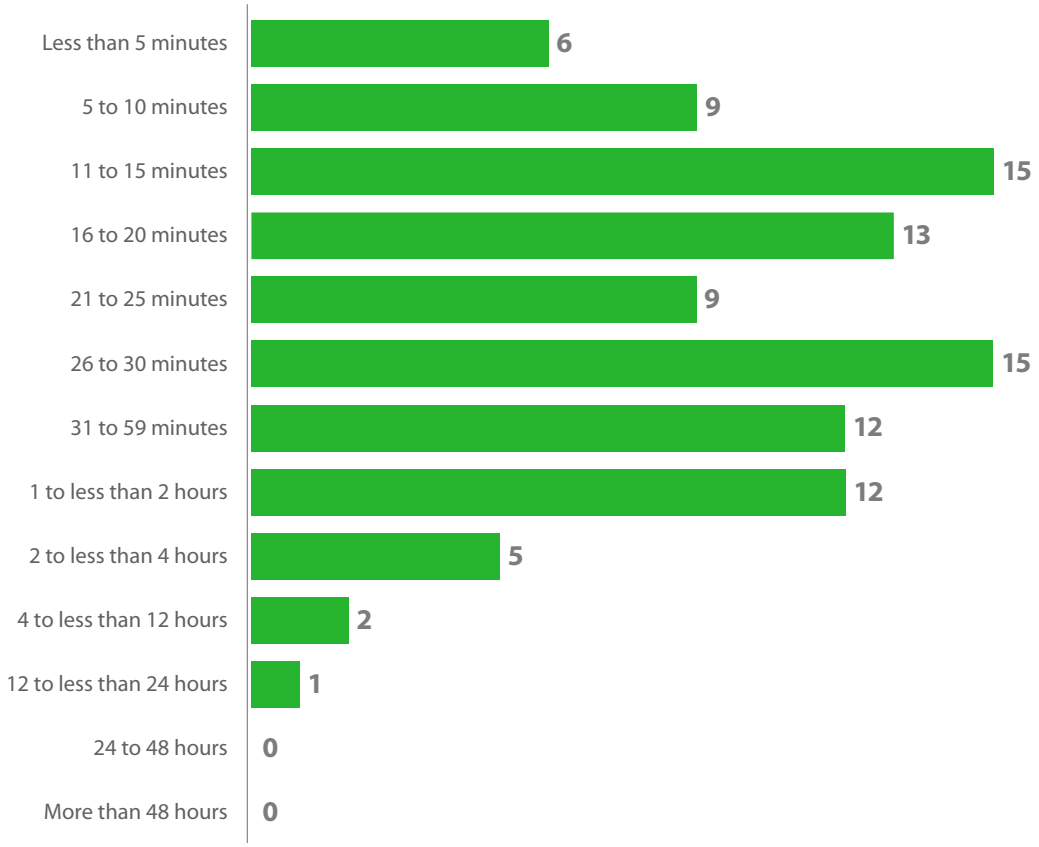

Figure 4. On average, how long do unplanned outages last? (Percent of respondents, N=989)

The median[2] length of an outage is 23 minutes. While that may not sound like much, consider:

- What is the impact to thousands of passengers of an airline whose planes are grounded for just 23 minutes?
- What is the impact to customers and an online retailer whose website is offline for even 23 minutes?
- What is the impact to a patient in a hospital whose data is unavailable for only 23 minutes?

23
minutes is the median length of an outage

Any seasoned IT professional can share war stories related to system outages and how those outages resulted in disruptions. Moreover, the media publishes new examples almost every week (hint: don't be one of them). From life-changing events to a simple inability to communicate with a colleague, customer, or partner, all business processes are at risk when IT fails its users.

Considering the potentially high impact of outages, coupled with the insufficiency of legacy IT-based efforts to back up and recover data, the desire among executives to seek out better Availability tools is warranted.

[2] Ibid

The Availability Gap and Downtime

If you take a closer look at what surveyed organizations deem "high-priority" workloads compared with "normal" workloads, a startling difference becomes apparent:

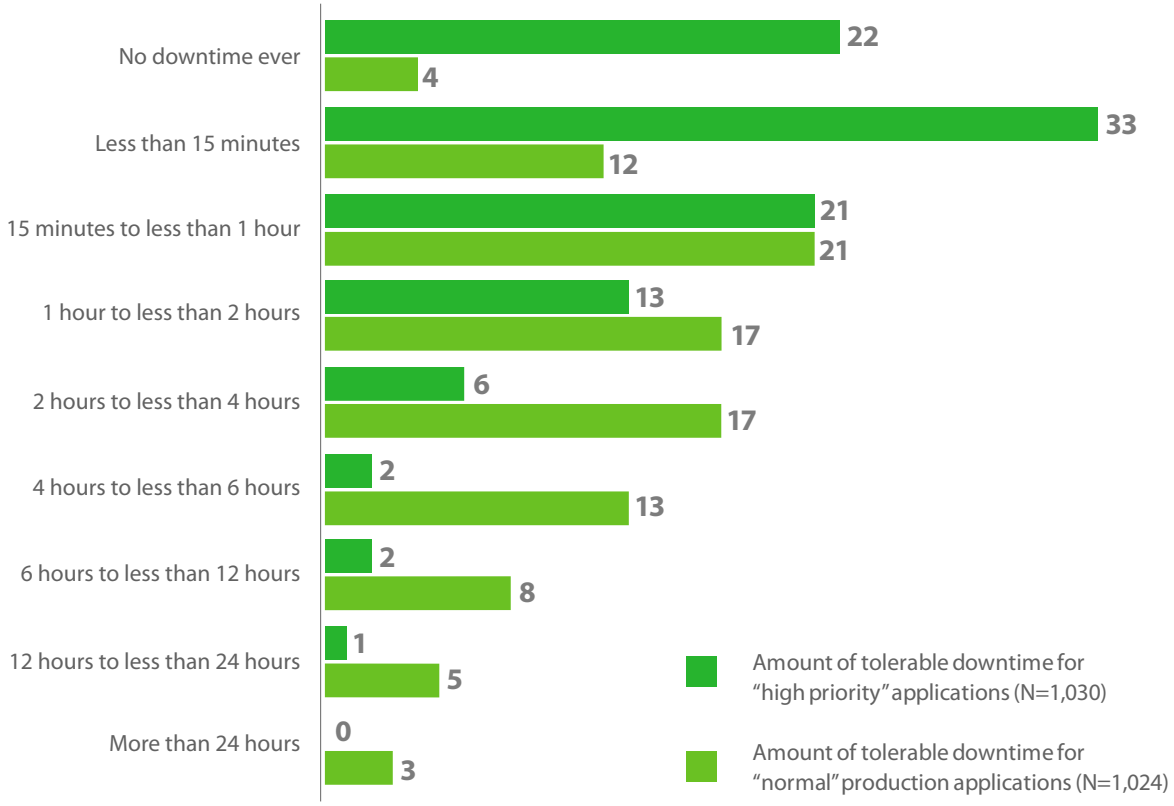

Figure 5. What is the amount of downtime your organization can tolerate from its "high-priority" production applications compared to "normal" production applications? (Percent of respondents)

- The median tolerable downtime among *high-priority* applications is 7.5 minutes, whereby an "only 23-minute" outage would exceed the limit for the majority of high-priority applications.
- The median tolerable downtime among *normal* applications is 90 minutes. While 23 minutes seems somewhat more tolerable, many normal applications will also have breached their SLAs with such an outage.

The Protection Gap and Data-loss SLAs

The disparity between the speed at which IT can recover platforms/workloads and the Availability expectations of business units and other end-users is not the only concern. Organizations are also inadequately protecting data:

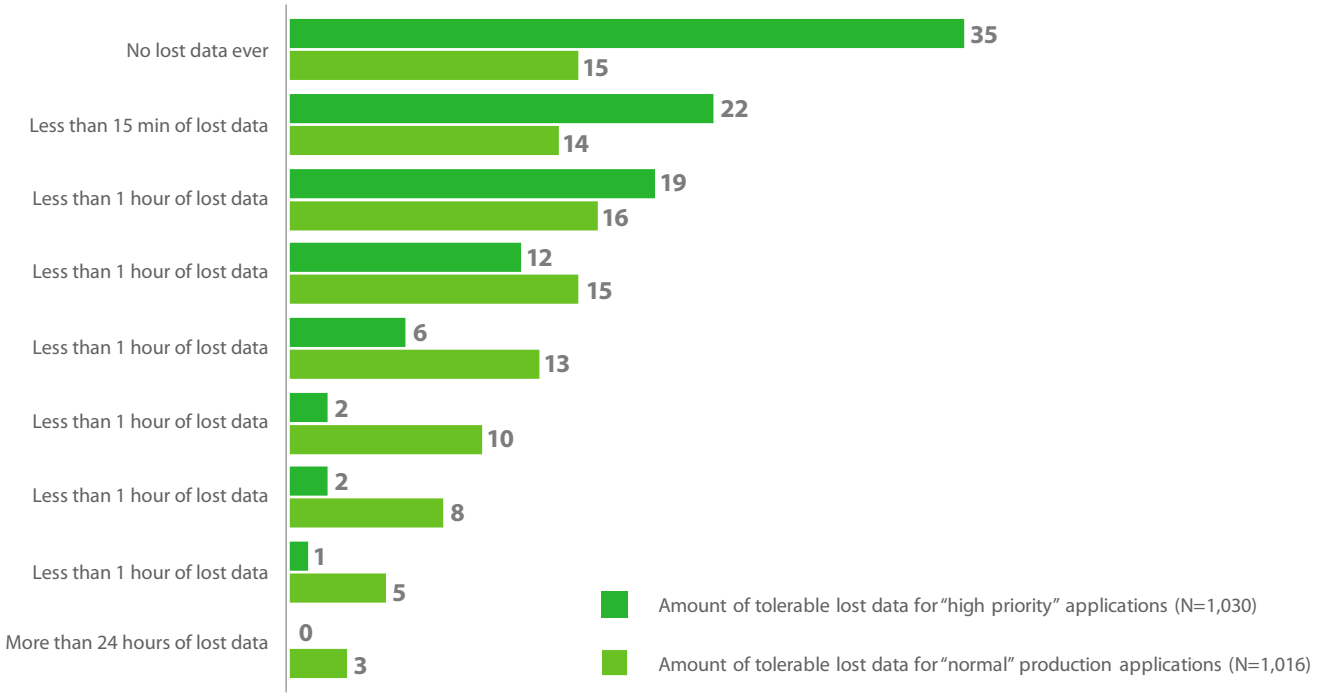

Figure 6. What is the amount of lost data that your organization can tolerate from its "high-priority" production applications compared with "normal" production applications? (Percent of respondents)

- The average acceptable data loss among *high-priority* applications is 72 minutes, as shown in Figure 6. But, the surveyed organizations only protect their high-priority data approximately every 127 minutes, on average.
- Similarly, while the average acceptable data loss among *normal* applications is 240 minutes, surveyed organizations only protect their normal data approximately every 352 minutes.

This is a quantifiable example of a Protection Gap. To be clear, most organizations believe that they have an Availability Gap, a Protection Gap, or both. To overcome these gaps, they must start with increasing the frequency of protection and boosting the agility and reliability of recovery.

The Costs of Availability and Protection Gaps

In an age when mission-critical VMs and nonessential VMs might live on the same host today but not tomorrow, and the number of users per VM varies widely, calculating downtime can be daunting for organizations of any size. For the purposes of this report, downtime costs included the following inputs (mostly from elsewhere in this report):

- The average total number of production servers (1,200) deployed at organizations
- The percentage of servers experiencing at least one outage per year (27%)
- The average length of unplanned outages (85 minutes)
- The average hourly costs for business-critcal ($108,000) and non-business-critical applications ($48,000)—adjusting for the average ratio of business-critical to non-business-critical applications
- ESG's industry average application to server ratio (.81)

$21.8M

Organizations participating in this research suffer direct costs of $21.8M annually, on average

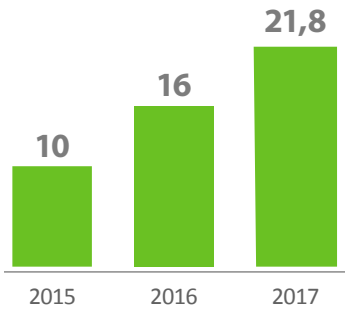

Figure 7. Estimated annual downtime costs per respondent organization (in $Millions USD)

Using these inputs, ESG calculates that organizations participating in this research suffer direct financial costs of $21.8M annually, on average. This continues the trend toward rising costs from downtime as seen in 2016 ($16M) and 2015 ($10M).

But wait, there's more …

How Organizations are Dealing with these Gaps

When you examine the discrepancies between archaic IT approaches and the expectations of business units, three "levels of realization" arise:

- **In theory**, the acknowledgement of the existence of the Availability Gap and the Protection Gap also represents a conceptual or intellectual acknowledgement that one's data protection and recovery mechanisms and strategy must evolve.
- **In practice**, irrefutable, quantifiable gaps do exist between IT protection and recovery capabilities and business units' expectations, and they are pervasive.
- **In reality,** the costs associated with downtime and data loss also lead to a diverse range of other negative impacts. Yes, the economic impacts are the easiest to envision, but other ramifications are equally, if not more, harmful.

41%

acknowledge that Availability issues around loss of customer and employee confidence is a most concerning potential impact

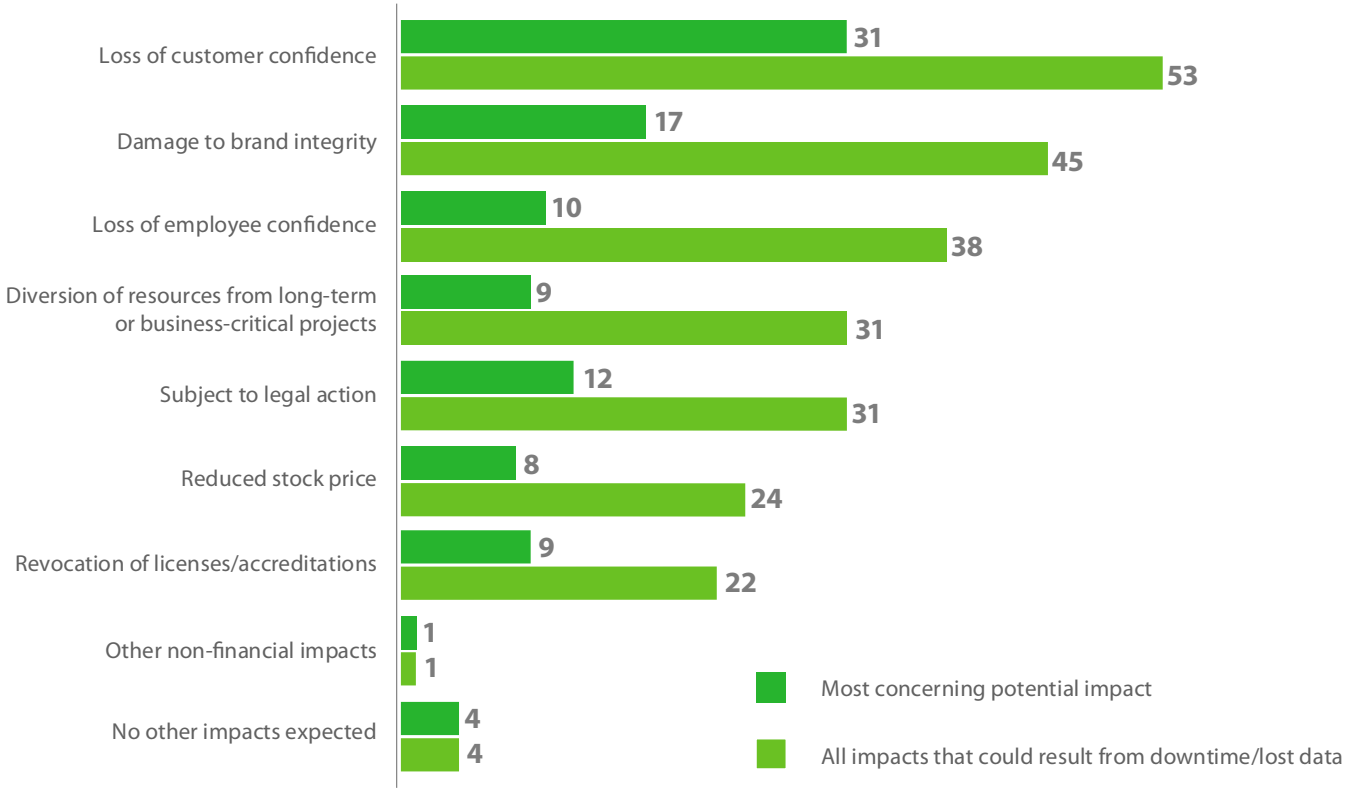

Figure 8. What other impacts – if any – could result in your organization from application downtime or lost data? Which impact is most concerning for you? (Percent of respondents, N=943)

Consider, for example, that only 4% of executives believe their organizations experience only monetary impacts as a result of downtime or data loss.

Most acknowledge that Availability issues could cause their organizations to suffer from problems such as a reduction in customer and employee confidence or damage to brand integrity.

Coincidently, the stacking of these impacts is nearly identical to previous years' research, from customer confidence and brand integrity being the most concerning to the miniscule number of people in denial of these non-financial impacts.

The problem will only get worse over time. Only 13% of respondents surveyed expect their organization's cost of downtime or data loss to decrease in the future. For everyone else, as business unit Availability expectations continue to increase and IT continues to struggle, the ancillary impacts of downtime and data loss will likely increase as well.

Hindrances to Organizational Virtualization Strategies

It is extremely important to recognize that inadequate protection and recovery mechanisms don't just hinder today's systems and business processes. They will also hinder an organization's ability to continue to modernize its IT environment as part of evolving for the sake of its business.

Virtualized servers are the foundation upon which most modern IT infrastructures are built. Most surveyed respondents (82%) acknowledge some relationship between the viability of their backup solution and the relative success of their virtualization deployment strategy:

- A non-trivial amount (33%) recognize that inadequacies in their VM backup solution have ***slowed*** their organization's virtualization deployment efforts.
- On a brighter note, many (49%) recognize that an effective VM backup solution has enabled them to ***significantly accelerate*** their virtualization deployment strategy.

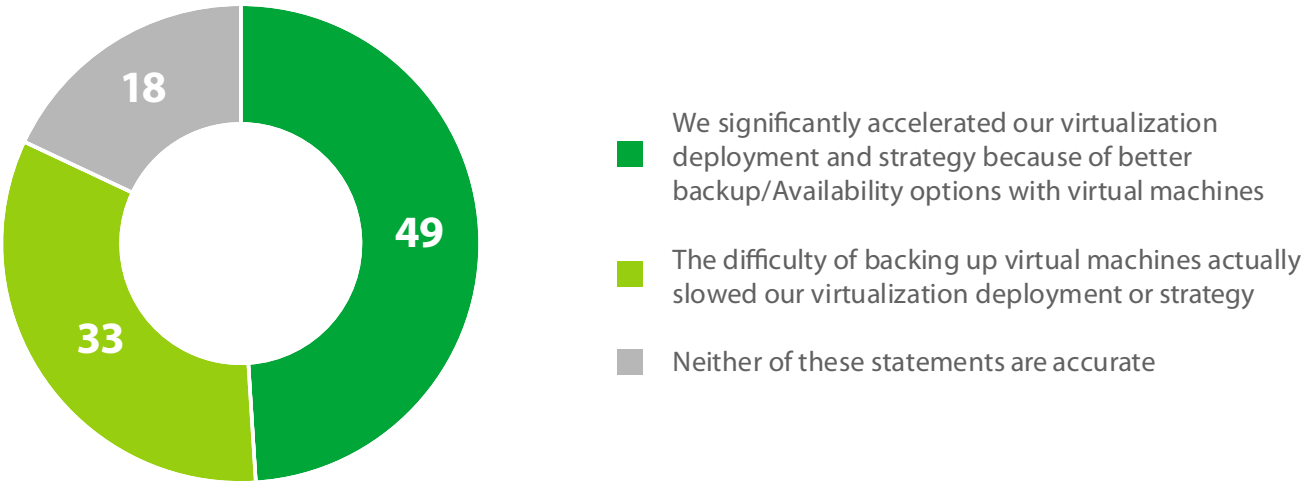

Figure 9. Which of the following statements about the relationship between server virtualization and data protection is most accurate? (Percent of respondents, N=964)

Hindrances to Organizational Cloud Strategies

Just as server virtualization forced new approaches to data protection and recovery, so too does "the cloud"—in each of its many consumption models:

- Those moving production workloads to hosted IaaS or PaaS services, or embracing SaaS, will need to rethink their data protection and recovery scenarios; and many will find changing vendors necessary.
- Meanwhile, cloud storage enables new options for data retention, especially when combined with turnkey backup services (BaaS) or failover mechanisms (DRaaS).

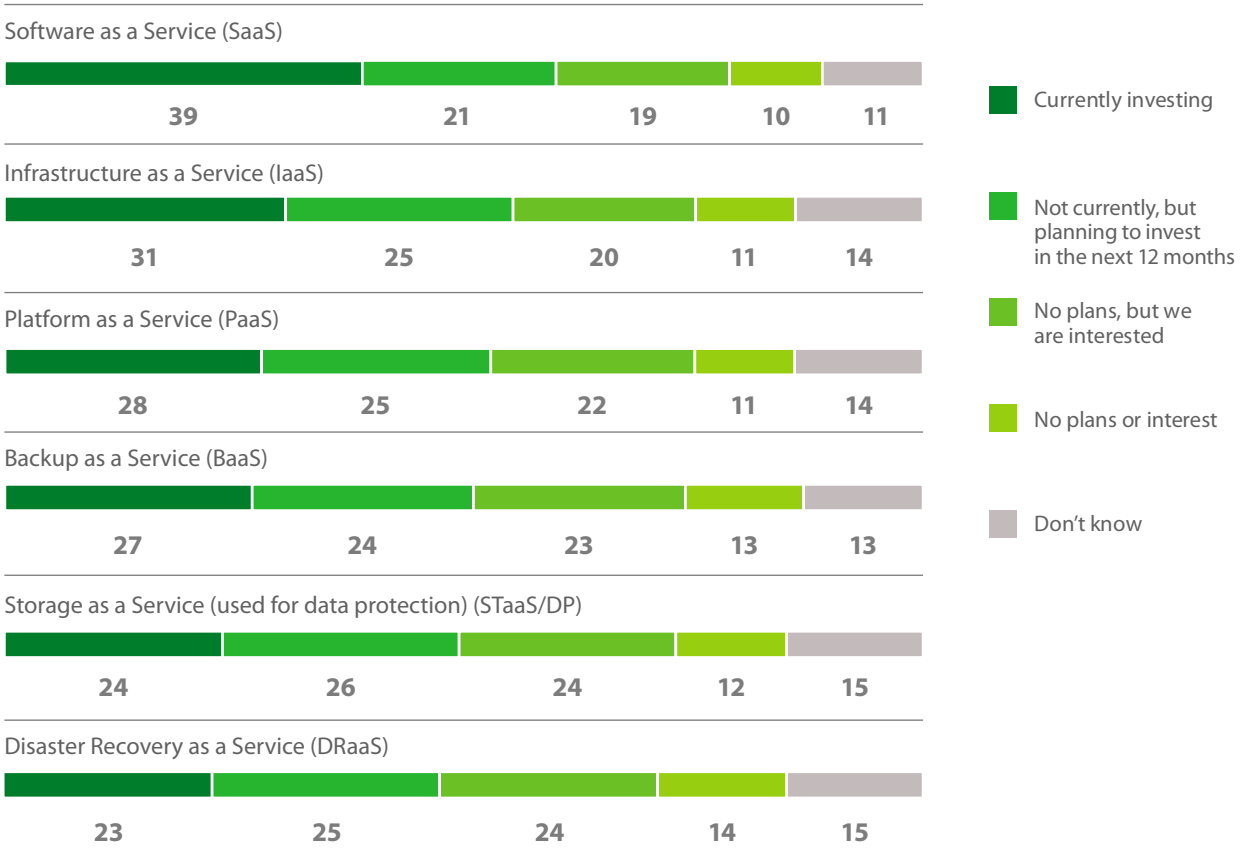

Figure 10. What types of cloud-based services is your organization either currently using or planning to use over the next 12 months (if any)? (Percent of respondents, N=1,060)

Given the material level of cloud investment observed, it's clear that cloud services will invariably change how IT accomplishes both production and protection goals, but not all data protection vendors are cloud-ready.

Hindrances to Organizational Digital Transformation Initiatives

Although some organizations are still modernizing their foundational infrastructures for virtualization, many others recognize that a *digital transformation* strategy will take them much further than simple infrastructure modernization.

- More than two-thirds of the surveyed respondents (69%) recognize that digital transformation is critical or very important to their organizations moving forward.
- That being said, about half of them (45%) say they are still in the planning or just beginning phases of digital transformation initiatives.

It is alarming that more than half of respondents whose organizations have digital transformation initiatives on their roadmap (66%) report that those initiatives are being inhibited because of unplanned downtime or insufficient application Availability.

66%

report that their digital transformation initiatives are being inhibited because of unplanned downtime or insufficient application Availability

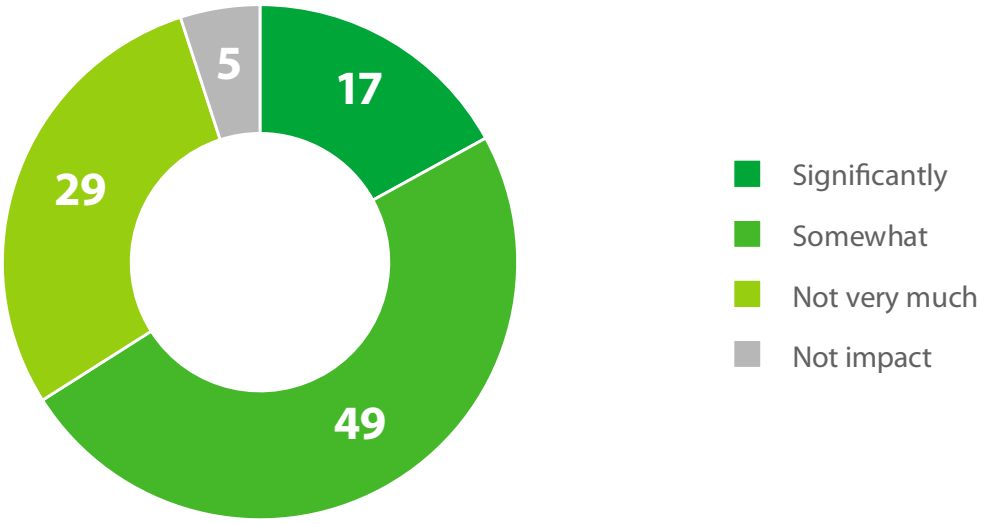

Figure 11. To what extent are your organization's digital transformation initiatives being inhibited by unplanned application downtime or insufficient application Availability? (Percent of respondents, N=970)

In order to advance these critical digital transformation initiatives past their predominantly nascent stages, many organizations must resolve their uptime and Availability shortfalls.

Conclusion

Most organizations do (and others don't, but should) recognize that they have gaps in their Availability and protection capabilities, resulting in a failure to meet the expectations of their business units, their executives, their employee colleagues, and their customers. Reasons for this situation include:

- Most IT infrastructures are in a perpetual state of modernization, which includes digital transformation initiatives, aggressive virtualization strategies, increased though sometimes experimental adoption of Hybrid Cloud services, diversification of production platforms, and heightening SLAs—all without commensurate increases in budget.
- Many organizations align neither their protection frequency nor their recovery mechanisms with the established SLAs of their business units, resulting in inadequate Availability.
- Many organizations are not able to effectively quantify the myriad costs and impacts of downtime or data loss, thereby hindering their ability to garner economic and operational support for better mechanisms and results.

These challenges are not trivial, but they are surmountable. Organizations must address the Availability and Protection Gaps that they have, or they put their employees and their institutions at risk of a wide range of productivity, economic, sentiment, and operational failings.

- **Any organization that cannot recover granular data or whole VMs faster than the established SLAs related to acceptable downtime has an *Availability Gap*.** That Availability Gap will result in lost user productivity, non-compliance with assured access mandates (both between business partners and in applicable regulated industries), and lost confidence by employees, customers, and the markets that they serve.
- **Any organization that does not protect its data at a frequency greater than the mean of its SLAs related to data loss has a *Protection Gap*.** That Protection Gap will result in lost data, which will cause employee productivity challenges both in the recreation of the data and in serving customers without complete information.

Gaps in either Availability or protection invariably hinder today's operating environments, the virtualization strategies and deployments that are modernizing today's data centers, and ultimately the digital transformation initiatives that so many institutions are relying on to ensure their market relevance moving forward.

Next Steps

The first and most crucial step in ensuring the viability of your IT systems in service to your business units and customers is to accept that you have an Availability Gap (until you can prove otherwise). Too many organizations that lack accurate metrics or monitoring processes presume that their systems are sufficient and are therefore hindering their organizations through their naiveté. Instead, presume that you have the problem, and then quantify it. Only a minority of organizations (less than one in five) will determine otherwise, and many of those are likely resilient because of their heightened Availability efforts in years past.

Next, quantify your business unit's SLAs and assess your own protection mechanisms and recovery capabilities. Only by comparing your Availability and protection expectations with your real-world capabilities will you be able to determine the size of the gaps in your strategy.

Convert your gaps into impact analyses. In the Business Continuity/Disaster Recovery (BC/DR) world, this is referred to as a business impact analysis (BIA), which is accomplished by simply asking, *"If [system] were to fail, what would that cost us [in economics, process, perception, etc.]?"* By looking at past systems logs, most will discover that those systems have had interruptions in the past, which can now be quantified as business impact.

With an accurate understanding of the frequency and duration of outages within your environment, compared with the SLA expectations of your constituents, and an assessment of the economic and perception impacts specifically to your organization, **you are ready to reimagine what it would take to become an Always-On Enterprise.**

1. *Recognize that virtualization will almost certainly be the underpinning of your infrastructure*, and therefore you must ensure that your protection and recovery capabilities for highly virtualized systems exceed your business SLAs. This alone can solve a significant portion of your Availability and Protection Gaps.

2. *Understand that cloud services will undoubtedly play a bigger role in your strategy moving forward,* although which types of cloud services will vary greatly between cloud storage, cloud-based protection services, cloud-based infrastructure in production and BC/DR scenarios, and cloud-based applications (such as Office365). Each of these platform choices will affect your protection and recovery options, which again must be measured first and foremost against the SLAs to ensure reduced Availability and Protection Gaps.
3. And lastly, but perhaps most importantly, a*cknowledge that downtime and data loss are not just theoretical concepts, and that RPO/RTO are not just metrics for an IT scorecard.* The lack of agile and reliable recovery/Availability mechanisms will impact your virtualization underpinnings today and will hinder the digital transformation initiatives that are supposed to bring you into tomorrow. *All of that starts with a commitment to being Always-On.*

Appendix: Research Methodology and Respondent Demographics

Research Methodology

Veeam commissioned The Enterprise Strategy Group, a leading IT analyst, research, and strategy company, to develop and execute the survey upon which this report is based.

To gather data for this report, ESG conducted a comprehensive online survey of 1,060 ITDMs from private and public sector organizations with at least 1,000 employees in 24 different countries between November 18, 2016 and December 31, 2016.

The geographic representation of the respondent base is shown in Figure 12.

United States	N=158
United Kingdom	N=103
France, Germany	N=78
Benelux (Belgium, Netherlands), Hong Kong	N=75
Australia, Japan, China, Brazil, Singapore	N=50
Canada	N=49
Italy, Nordics (Sweden, Denmark, Finland), Russia, Thailand, India, Middle East (UAE, Saudi Arabia, Israel), Mexico	N=16-30

Figure 12. Number of Qualified Respondents per Country/Region

To qualify for this survey, respondents were required to be employed in an IT role with day-to-day knowledge of and/or familiarity with their organization's data/file backup and recovery environment and strategy. All respondents were provided an incentive to complete the survey in the form of cash awards and/or cash equivalents.

All respondents were subject to a stringent quality assurance process, which included filtering out unqualified respondents, removing duplicate responses, and screening the remaining completed responses (on a number of criteria) for data integrity.

Please see the Respondent Demographics section of this report for more information on these respondents.

Notes Regarding the Calculations and Data Displayed within This Report

In this report, calculated means and medians are estimated for applicable questions in which response options were presented as numeric ranges. This is done by using the midpoint of each data range selected by each respondent as the assumed respondent value and calculating the average (be it median or mean) based on the aggregate distribution of respondents' responses to the relevant question. References to any averages cited in this report refer to the mathematical mean, unless median is explicitly cited.

In addition, totals in figures and tables throughout this report may not add up to 100% due to rounding.

Respondent Demographics

The data presented in this report is based on a survey of 1,060 qualified respondents. Figures 13 through 17 detail the demographics of the respondent base, including respondents' current role, as well as respondent organizations' total number of employees, primary industry, and server footprint.

Respondents by Role

Respondents' current role within their organization is shown in Figure 13.

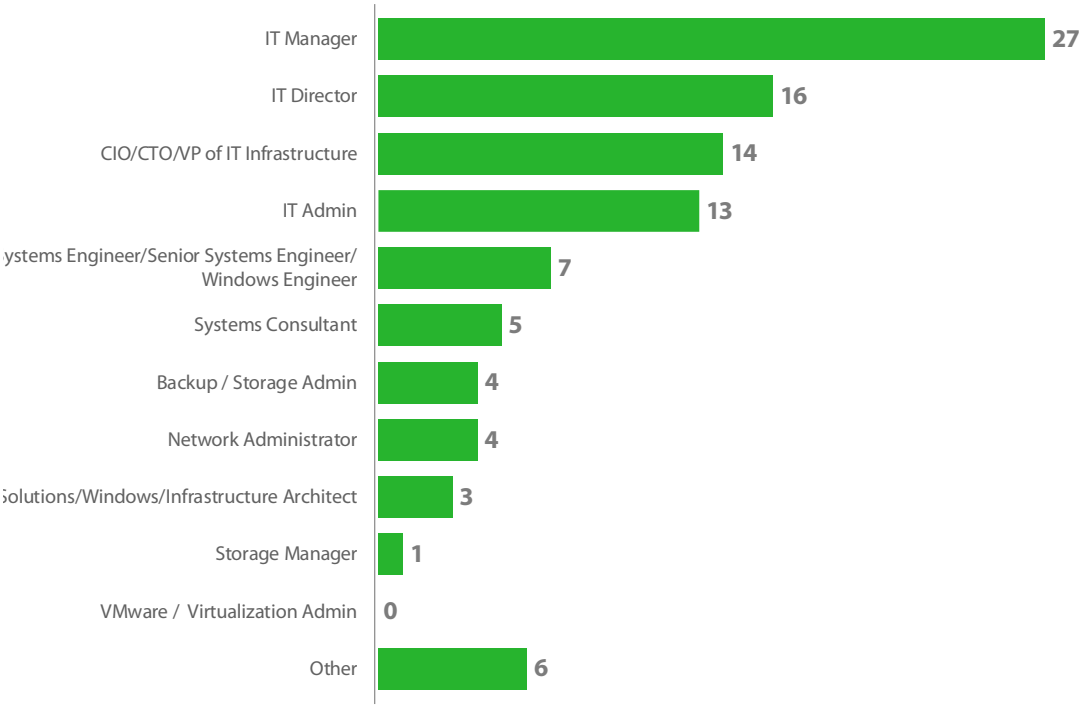

Figure 13. Which of the following best describes your role within your organization? (Percent of respondents, N=1,060)

Respondents by Number of Employees

The number of employees in respondents' organizations is shown in Figure 14.

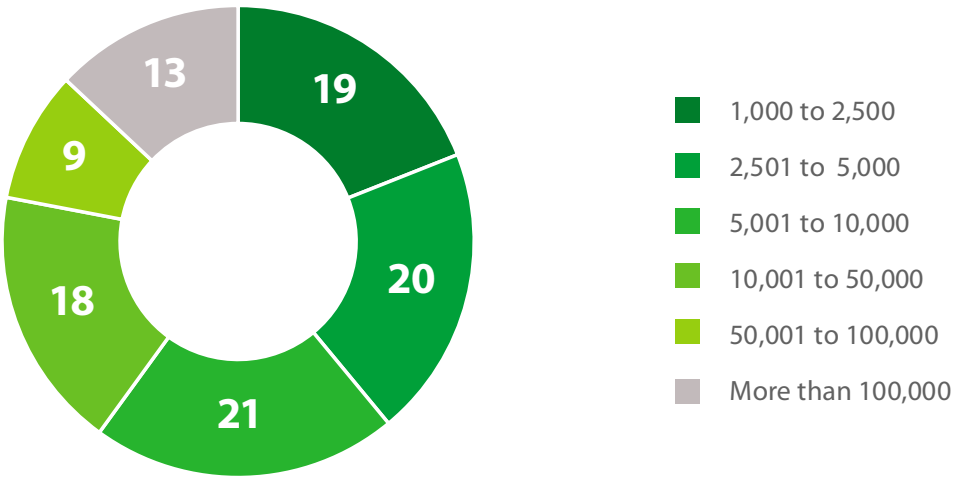

Figure 14. How many total employees does your organization have worldwide? (Percent of respondents, N=1,060)

Respondents by Industry

Respondents were asked to identify their organization's primary industry. In total, ESG received completed, qualified responses from individuals in 12 distinct vertical industries, plus an "Other" category, shown in Figure 15.

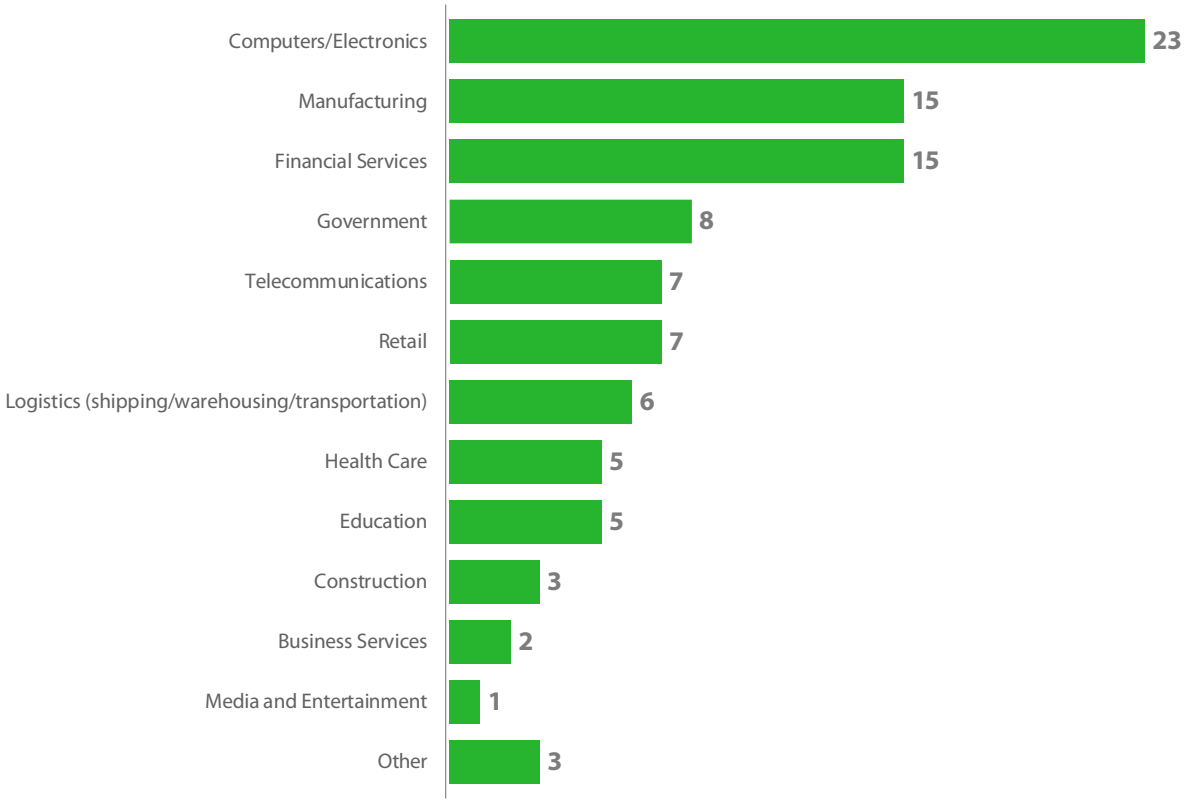

Figure 15. What is your organization's primary industry? (Percent of respondents, N=1,060)

Respondents by Number of Production Servers

Respondent organizations' number of production physical and virtual servers is shown in Figure 16.

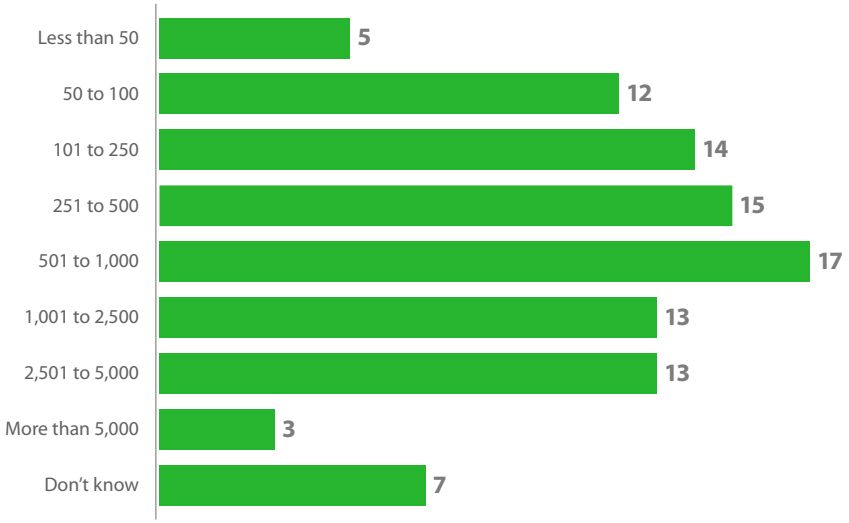

Figure 16. Approximately how many total production servers (i.e., both physical and virtual, but not including test/development) are currently deployed in your organization? (Percent of respondents, N=1,035)

Respondents by Percentage of Virtualized x86 Servers

Respondent organizations' percentage of x86 servers that have been virtualized to date, and how this percentage is expected to change in two years, is shown in Figure 17.

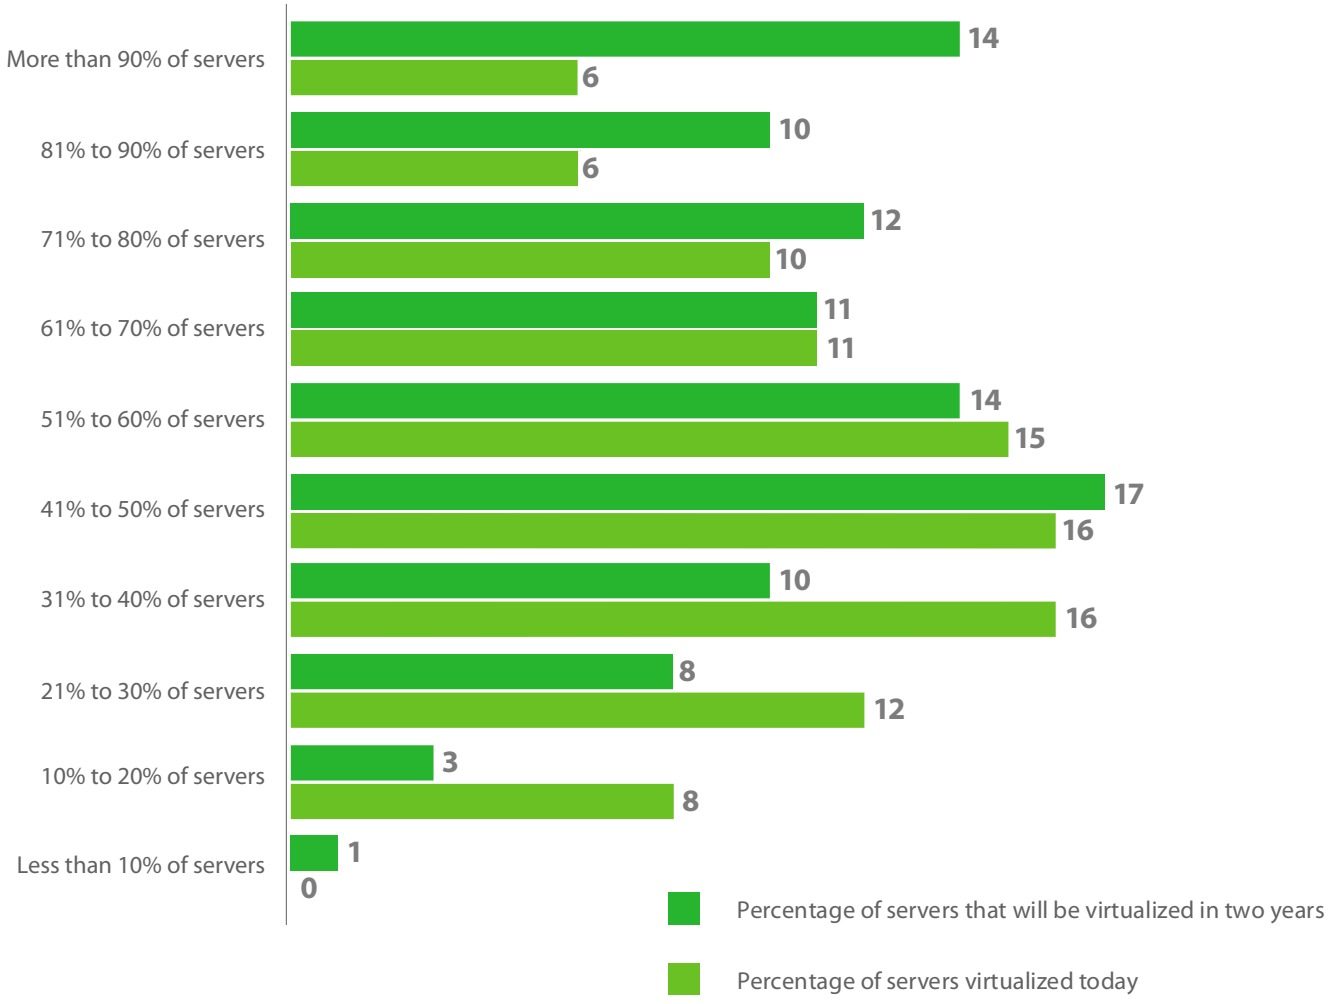

Figure 17. Of all the x86 servers in your organization that can be virtualized, what percentage has been virtualized? Looking ahead two years, what percentage of servers do you believe will be virtualized? (Percent of respondents, N=1,060)

About Veeam Software:
Veeam® recognizes the new challenges companies across the globe face in enabling the Always-On Enterprise™, a business that must operate 24.7.365. To address this, Veeam has pioneered a new market of *Availability for the Always-On Enterprise™* by helping organizations meet recovery time and point objectives (RTPO™) of less than 15 minutes for all applications and data, through a fundamentally new kind of solution that delivers high-speed recovery, data loss avoidance, verified recoverability, leveraged data and complete visibility. Veeam Availability Suite™, which includes Veeam Backup & Replication™, leverages virtualization, storage, and cloud technologies that enable the modern data center to help organizations save time, mitigate risks, and dramatically reduce capital and operational costs, while always supporting the current and future business goals of Veeam customers.

Founded in 2006, Veeam currently has 45,000 ProPartners and more than 230,000 customers worldwide. Veeam's global headquarters are located in Baar, Switzerland, and the company has offices throughout the world.
To learn more, visit https://www.veeam.com.

About ESG
ESG is an IT analyst, research, and strategy company, founded in 1999, with headquarters in Milford, Massachusetts. It conducts research with and for IT vendors, IT professionals, business professionals, and channel partners. ESG maintains ongoing analyst coverage in cloud computing, networking, storage, data protection, cybersecurity, data management and analytics, enterprise mobility, systems management, and channels.

About the Principal Analyst for This Study
Jason Buffington is the Principal Analyst at ESG focusing on all forms of data protection, preservation, and availability. He has actively deployed or consulted on data protection and storage solutions for 28 years, working at channel partners, multiple data protection software vendors, and Microsoft. Jason has been a featured speaker at many server-infrastructure, business continuity, and storage events around the world, and his articles have appeared in numerous IT industry journals.