Veeam Data Center Availability Report 2014

The Challenge of The Always-On Business

VEEAM
IT JUST WORKS!™

AVAILABILITY™
for the Modern Data Center

Contents

Executive Summary

The Veeam Data Center Availability Report 2014 investigates the increasing demands organizations face to provide an 'Always-On Business', what actions they are taking to meet those demands, and how successful their actions are. Following on from previous Veeam Data Protection Reports, this report investigates whether existing solutions can provide the always-on availability that businesses demand in the 21st Century.

In particular, the report shows that there are clear demands for 24/7 access to IT services and applications, with over 90 percent of enterprises increasing their requirements for minimizing downtime and guaranteeing access to data. It also demonstrates how organizations are modernizing their data center infrastructure, in order to meet these requirements.

Despite investments in the modern data center, there is an "availability gap" between the Always-On Business requirements and what legacy backup solutions can deliver in terms of Recovery Time Objective (RTO) and Recovery Point Objectives (RPO). Indeed, to meet the demands of the Always-On Business, organizations would need to recover mission-critical data in 60 percent of the time it takes them now and perform backup 1.5 times more often.

This report shows the scale of this gap and the financial impact of failing to meet business demands. Currently, organizations suffer application downtime 13 times a year, with the total cost of downtime and data loss reaching up to $10,163,114. In addition, organizations experience backup recovery failure twice a year. Because of this, organizations will lose at least $2 million a year through data loss and factors such as lost productivity and missed opportunities.

The concluding sections of the report break down the reasons why legacy backup solutions do not yet meet RTO and RPO requirements. It was highlighted that these solutions lack capabilities such as high-speed recovery (wanted by 60 percent of organizations), data loss avoidance (53 percent); verified protection (47 percent); using backup data as a production-like test environment for new patches or updates (38 percent); and complete visibility, including proactive monitoring and alerting of issues before any operational impact (36 percent).

It then outlines the actions enterprises are taking to deliver Availability for their Modern Data Centers. Considering the lack of capabilities above, it comes as little surprise that 78 percent of organizations plan to change their backup solution in the next 2 years.

This report is based on an online survey conducted in August and September 2014 by Vanson Bourne, an independent market research organization, of 760 senior IT decision makers from organizations across the United States, United Kingdom, Germany, France, Italy, the Netherlands, Switzerland, Brazil, Australia and Singapore that employ more than 1,000 people.

Survey Background

Respondents were selected from a cross-section of industries. These were: Retail, Distribution and Transport, comprising 23 percent of respondents; Manufacturing (22 percent); Financial Services (22 percent); Business & Professional Services (17 percent); and other commercial sectors (16 percent). As a result, the survey findings were evenly distributed across a variety of enterprises and sectors, with less risk of one sector skewing the statistics (Figure 1).

Figure 1: Respondent industries (%)

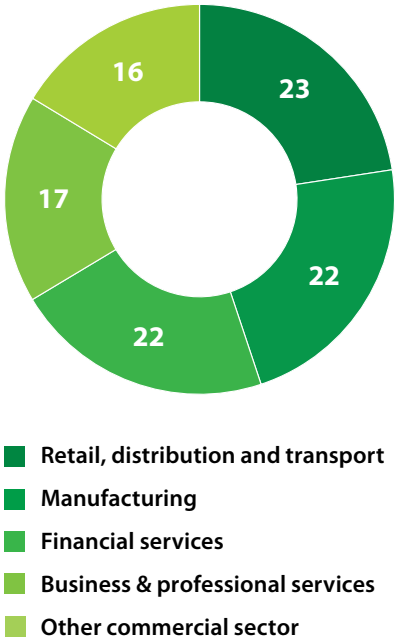

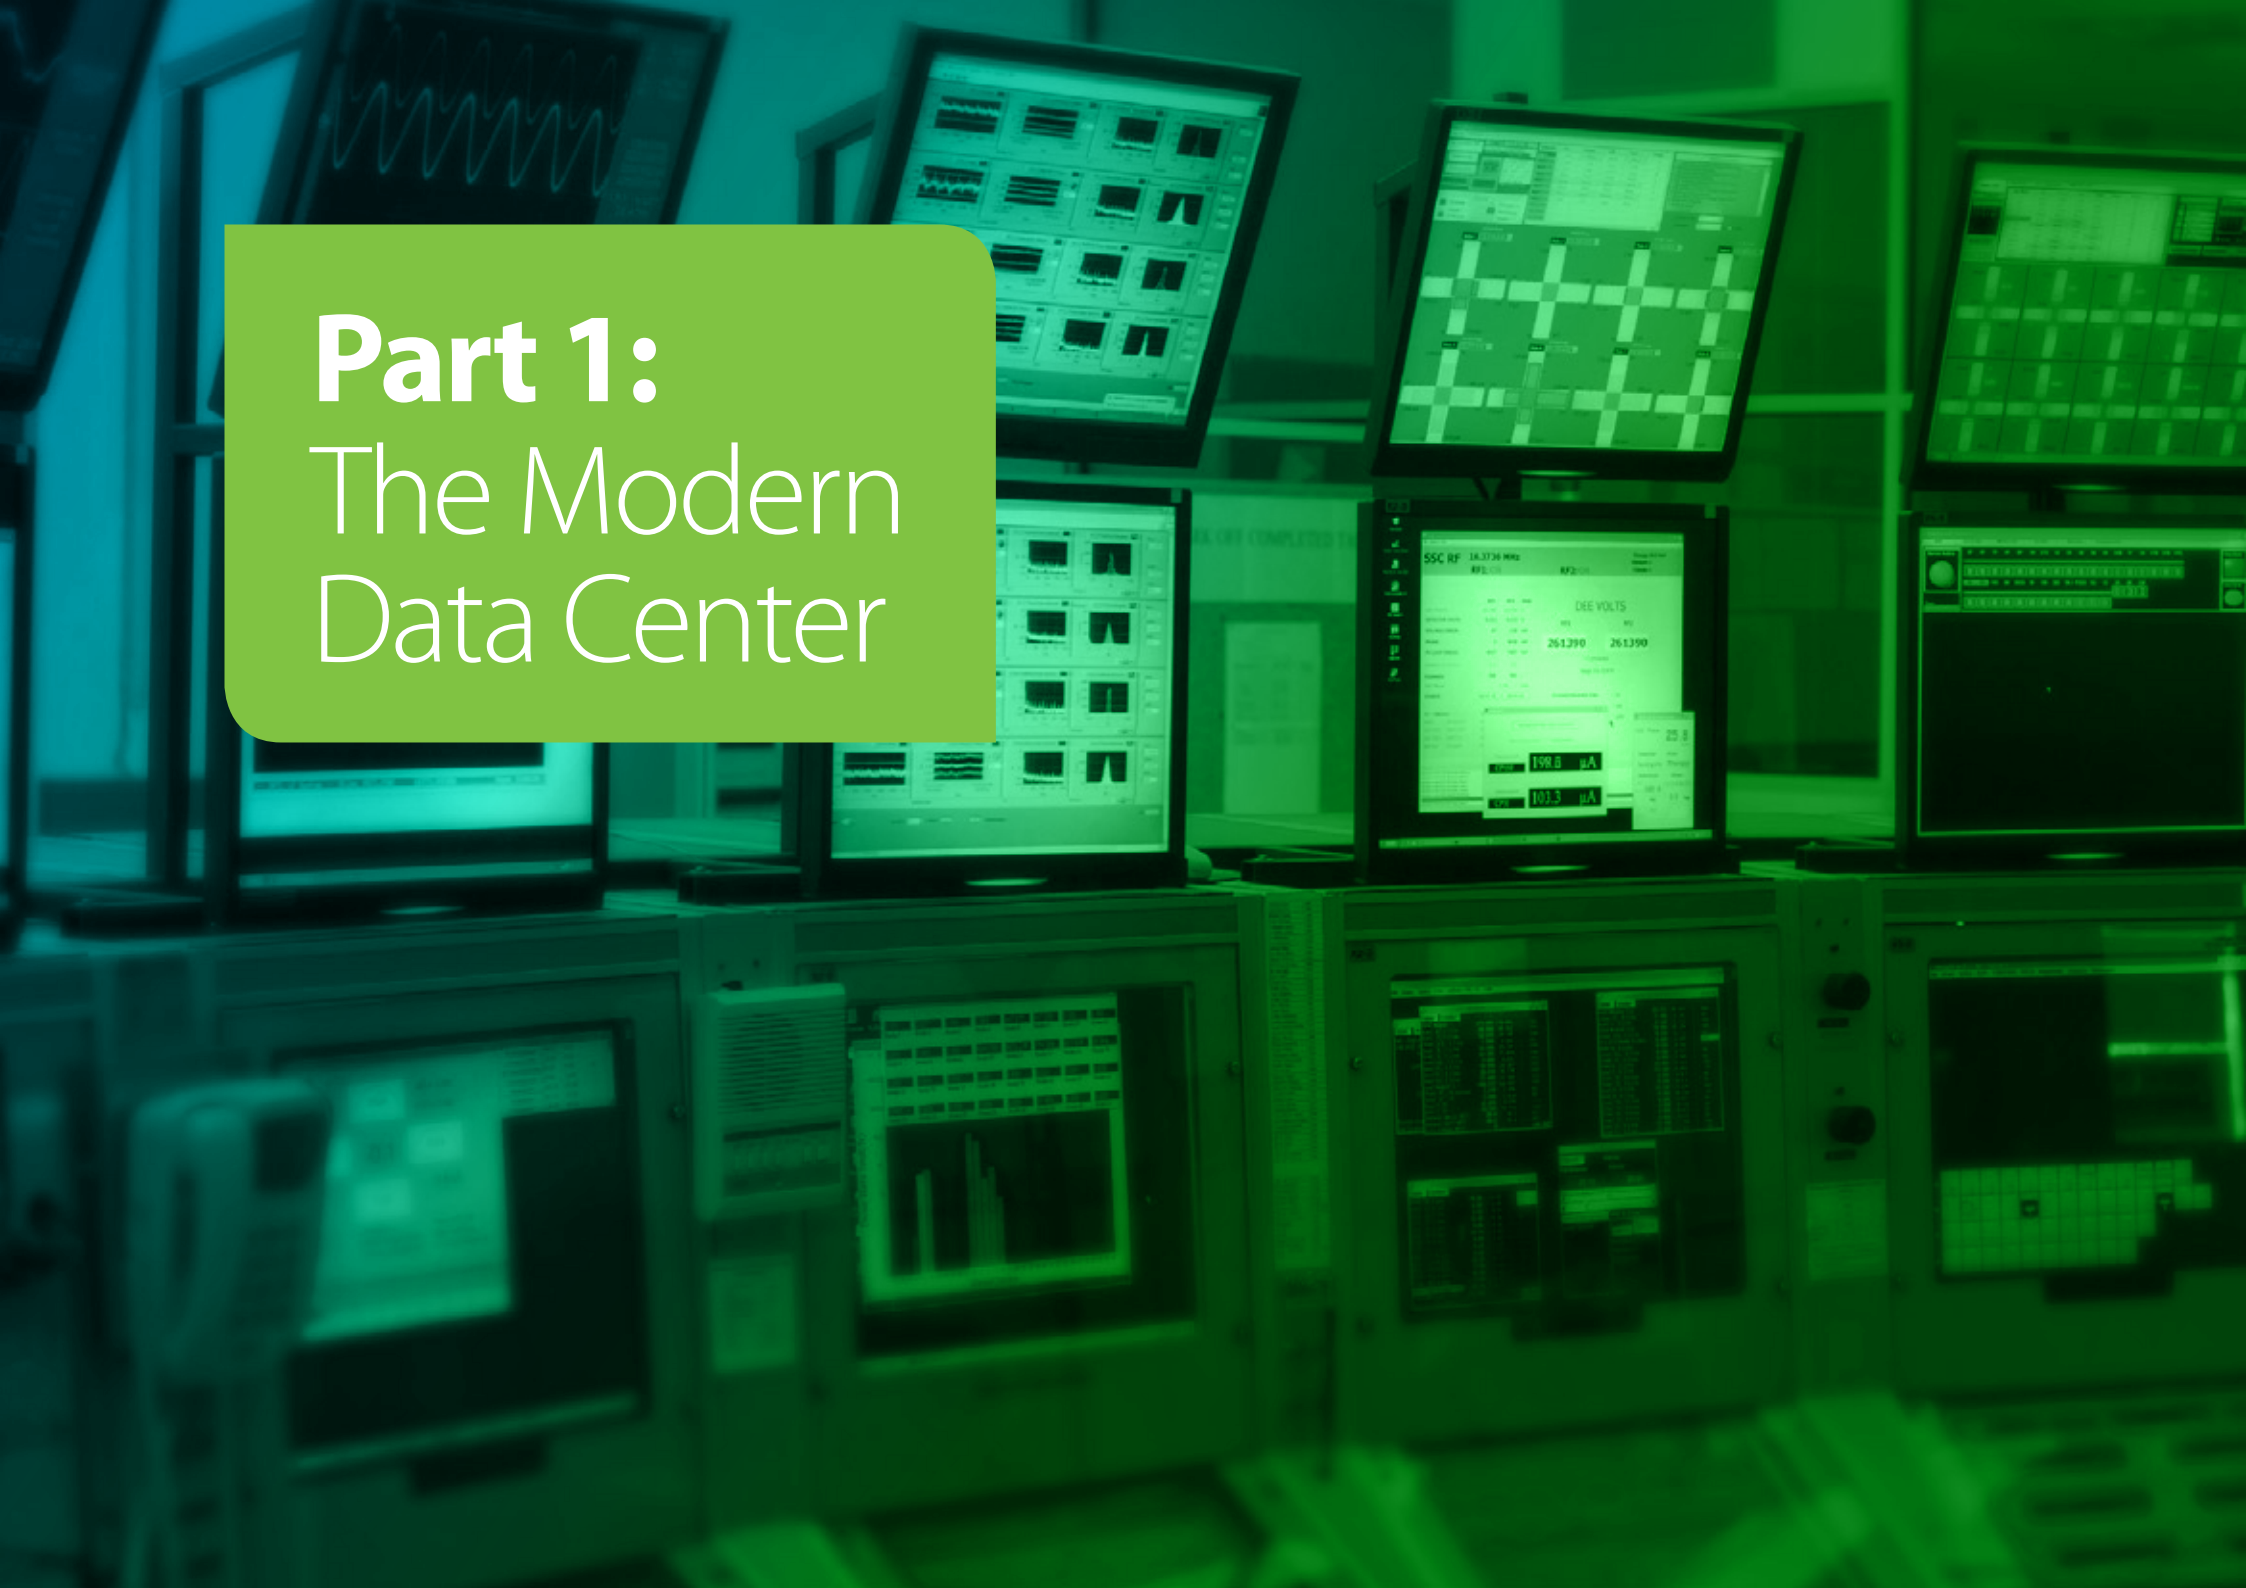

Part 1:
The Modern Data Center

Part 1: The Modern Data Center

Business requirements have changed dramatically in a relatively short time. IT has become strategic for every organization, and business requirements for IT have dramatically changed. To stay competitive today businesses need to provision IT services faster, strengthen security and control, lower operational costs and increase business agility. To meet these requirements businesses are building a modern data center by investing into modern technologies such as virtualization, modern storage and cloud.

Currently, 81 percent of organizations are either modernizing or have already modernized their data centers, with a further 16 percent planning to do so within the next 2 years (Figure 2).

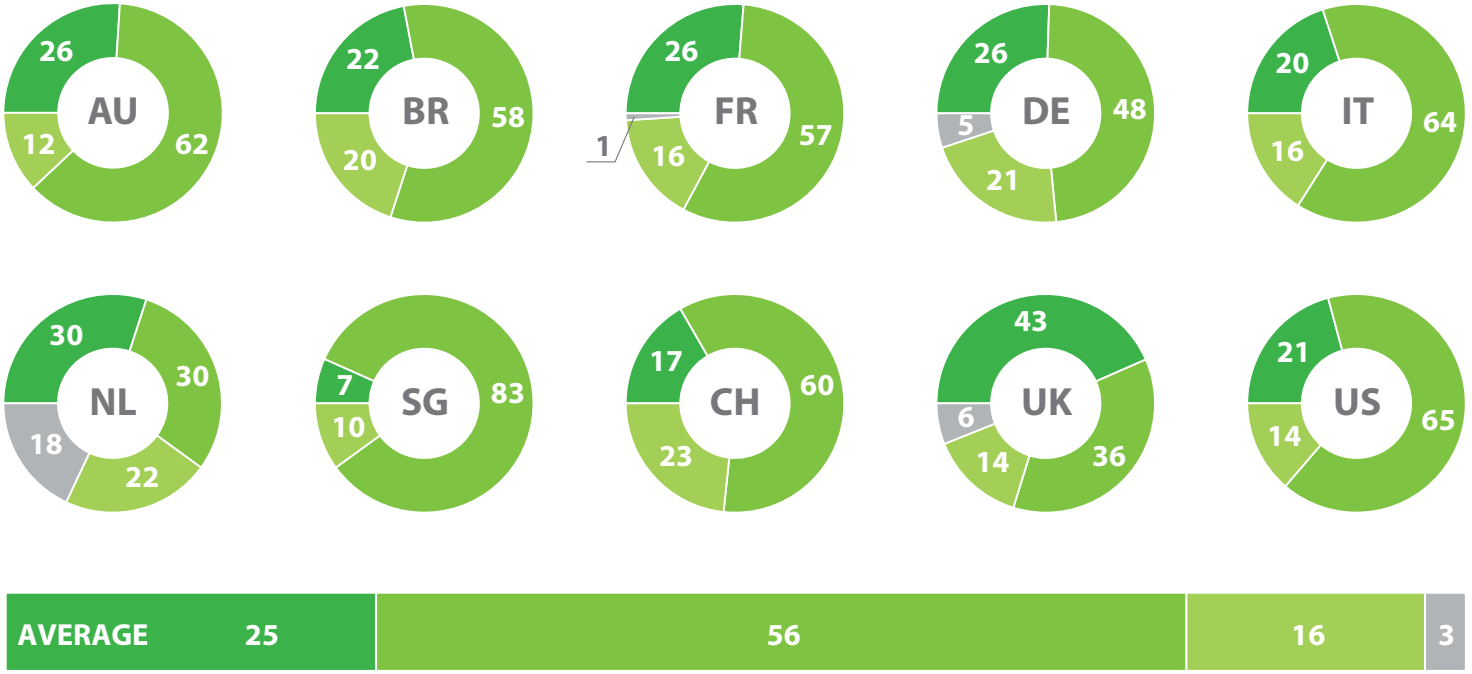

Figure 2: Organizations that have modernized, are modernizing or are planning to modernize their data centers (%)

- Data center is already modernized
- Modernizing the data center now
- Planning to modernize within 2 years
- No plans to modernize data center

68 percent of those organizations modernizing their data centers are doing so in order to enable 24/7, always-on business operations (Figure 3). There appears to be a consensus on the technology needed to fully modernize a data center. 97 percent of organizations engaged in data center modernization are either investing or planning to invest in server virtualization. Other technologies focused on are storage upgrades (95 percent), OS upgrades (94 percent), and data protection & disaster recovery (93 percent) (Figure 4).

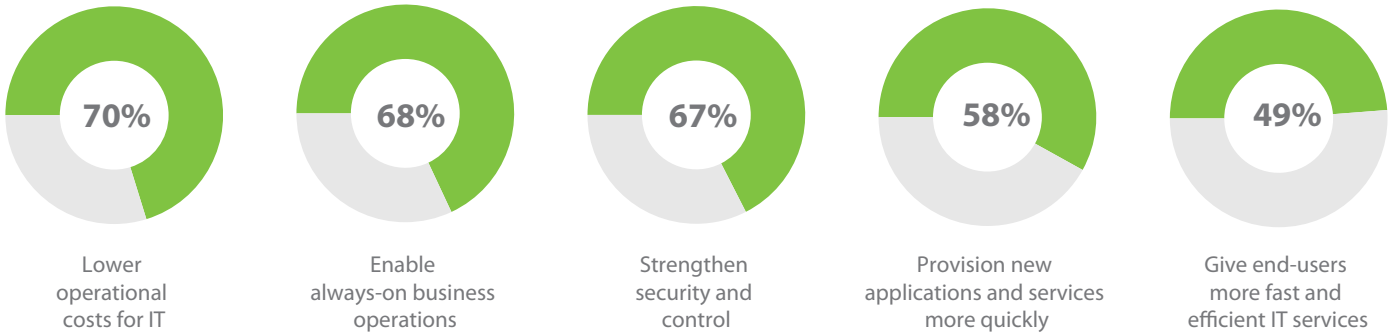

Figure 3: Business drivers for data center modernization (%)

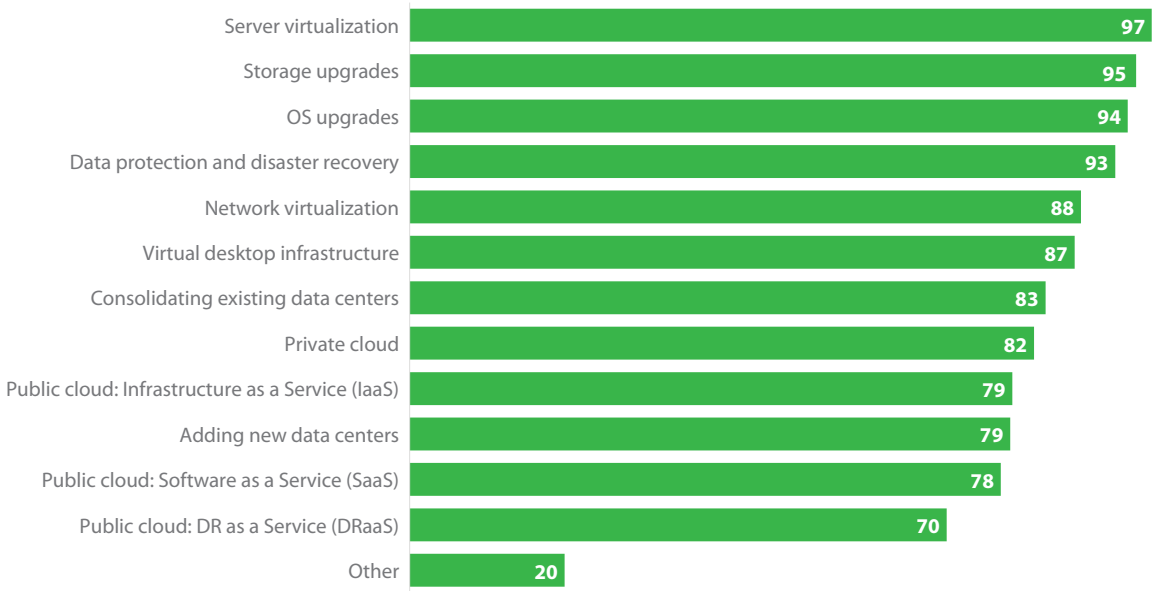

Figure 4: Technologies organizations are investing or planning to invest in (%)

Part 2: The Always-On Business

Part 2: The Always-On Business

Over the last 10 years, businesses have seen end-users demand increased access to data and applications for many reasons, including:

- Workers departing from the traditional nine-to-five working day
- Globalization allowing businesses to branch out across multiple time zones
- Customers conducting business online at any time
- Supply chain and logistics integration and automation requiring constant access to operational systems and data
- The rise of the Internet of Things (IoT) meaning devices are permanently connected and monitored

Taken in concert, these demands mean that businesses must be "always-on." More data and applications are considered to be mission-critical, and businesses have less patience for downtime, driving the need for increased Availability for their Modern Data Centers.

Over 90 percent of organizations surveyed were increasing their availability requirements to meet the always-on needs of their business. Specifically, 93 percent of respondents are increasing their requirements for minimizing downtime, while 92 percent are increasing their requirements for guaranteeing access to data (Figure 5).

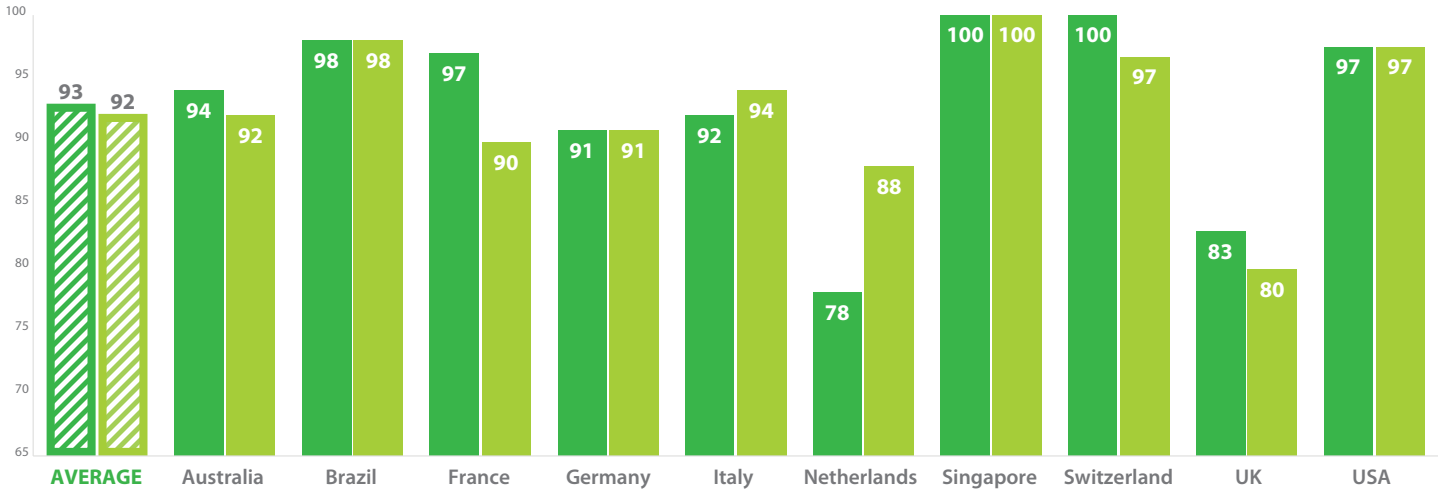

Figure 5: Organizations increasing their availability requirements over the past 2 years (%)

Organizations that have increased requirements for minimizing application downtime

Organizations that have increased requirements for guaranteeing access to data

This is being driven by the demands of end-users, the most common of which is more frequent, real-time interactions between customers, partners, suppliers and employees (65 percent). The need to access applications across time zones (56 percent), increased adoption of mobile devices (56 percent), employees working outside regular hours (54 percent) and an increasing level of automation for decision making and transactions (53 percent) (Figure 6), were also flagged as key requirements.

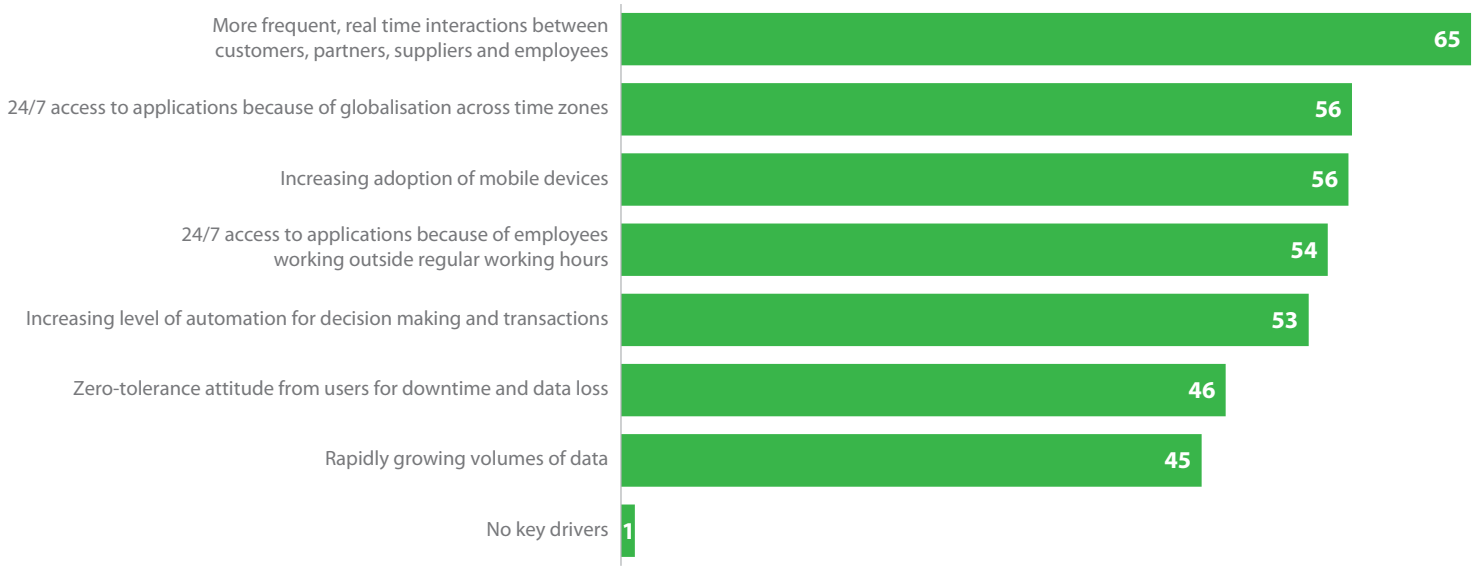

Figure 6: Key drivers for minimizing downtime and guaranteeing access to data (%)

Part 3:
The Availability Gap

Part 3: The Availability Gap

Data center modernization does not automatically result in increased availability of all data and applications. In fact, despite investing in virtualization, advanced storage and the cloud, many businesses are not able to meet their RPO and RTO service level agreements (SLAs). This opens an "availability gap" between the availability requirements of the always-on business and what the company's backup solution can actually deliver.

To help measure availability businesses look at RPO and RTO for critical and non-critical applications. Business must be confident that its most mission-critical applications will be available 24/7. Currently, 46 percent of workloads are mission-critical, although this is expected to be 53 percent by 2016 (Figure 7).

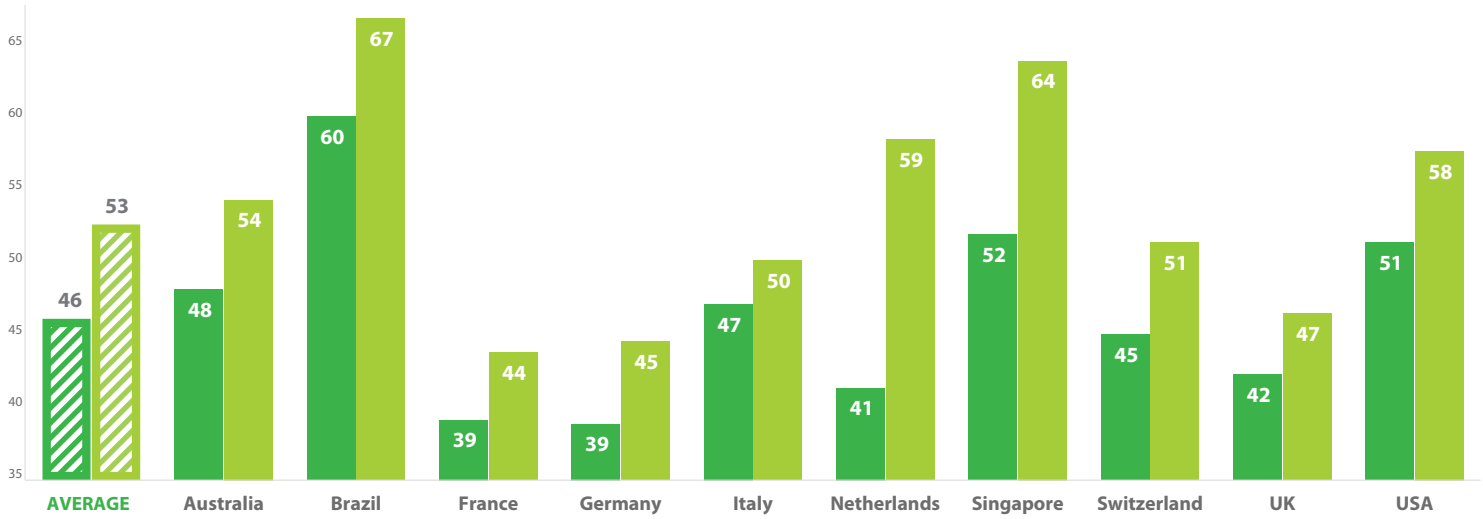

Figure 7: Proportion of mission-critical workloads (%)

- Current mission-critical percentage of workload
- Expected mission-critical percentage of workload in two years' time

The first SLA critical to the always-on business is the RTO; i.e. how quickly applications can be recovered. Faster recovery of applications means less downtime and less impact to the business in terms of lost sales and productivity. Currently, mission-critical applications take an average of 2.86 hours to recover, against an RTO of 2.69 hours (Figure 8), while non-mission-critical applications take an average of 8.45 hours to recover against an RTO of 10.02 hours (Figure 9).
As we can see, the average enterprise is either meeting or close to meeting its SLAs for recovering data.

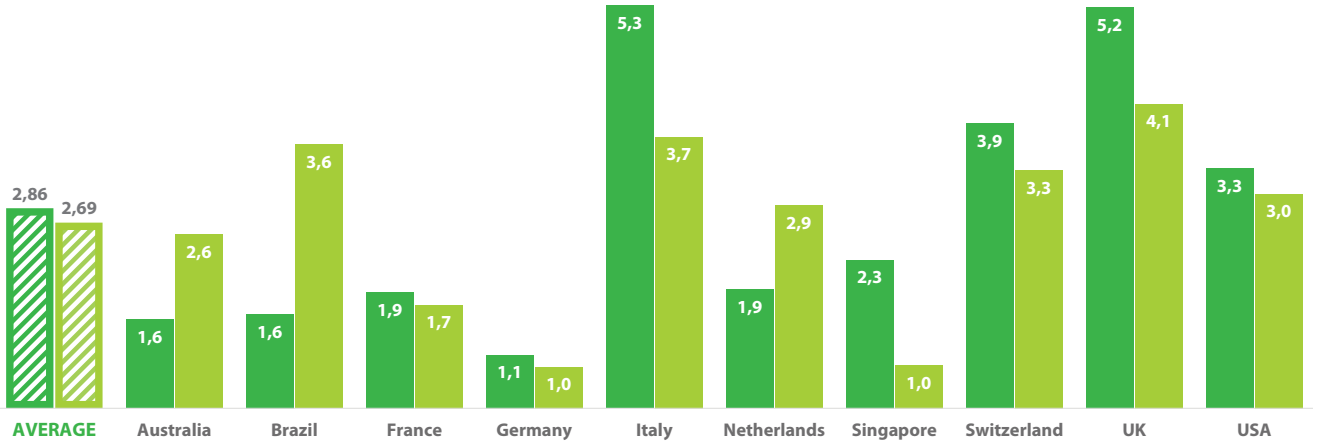

Figure 8: Average recovery time against current RTO: mission-critical applications (hours)

Average recovery time
Current RTO

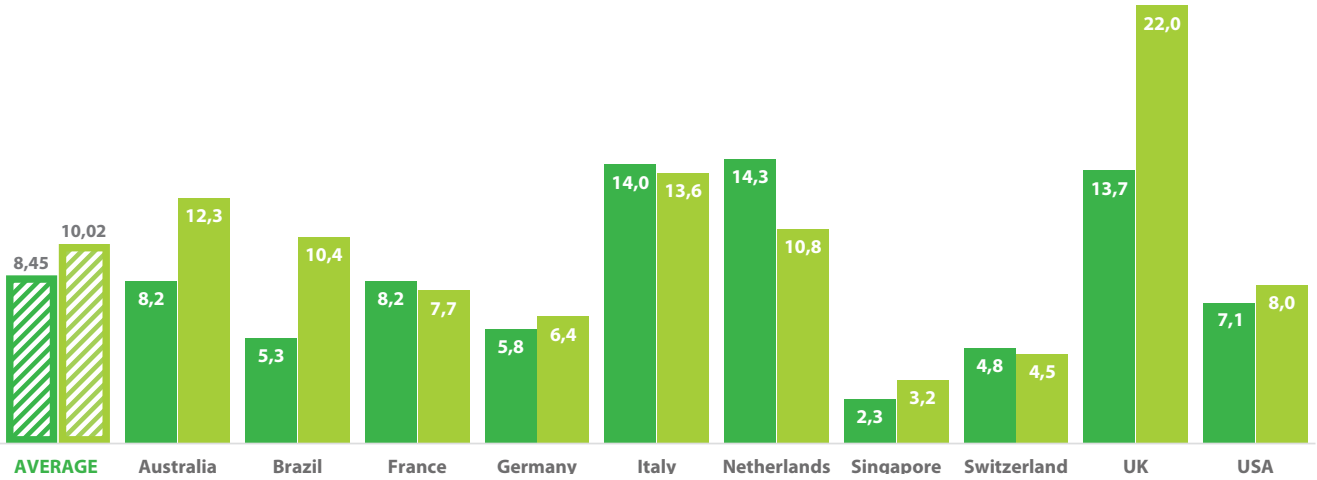

Figure 9: Average recovery time against current RTO: non-mission-critical applications (hours)

Average recovery time
Current RTO

The second critical SLA is the RPO; i.e. in the event of an IT failure, what is the most recent data that can be recovered and so how much will be irrecoverably lost? The more often an organization backs up its data, the smaller its RPO, and therefore the lower its risk exposure to data loss. Currently, organizations have less success meeting their RPO SLAs, putting them at risk of excessive data loss.

Mission-critical applications are backed up every 4.81 hours against an RPO of 3.53 hours (Figure 10). Non-mission critical applications are backed up every 14.46 hours against an RPO of 11.53 hours (Figure11). Looking at the RPO and RTO figures together it is clear that organizations can recover their applications within the agreed time, but they are at risk of excessive data loss.

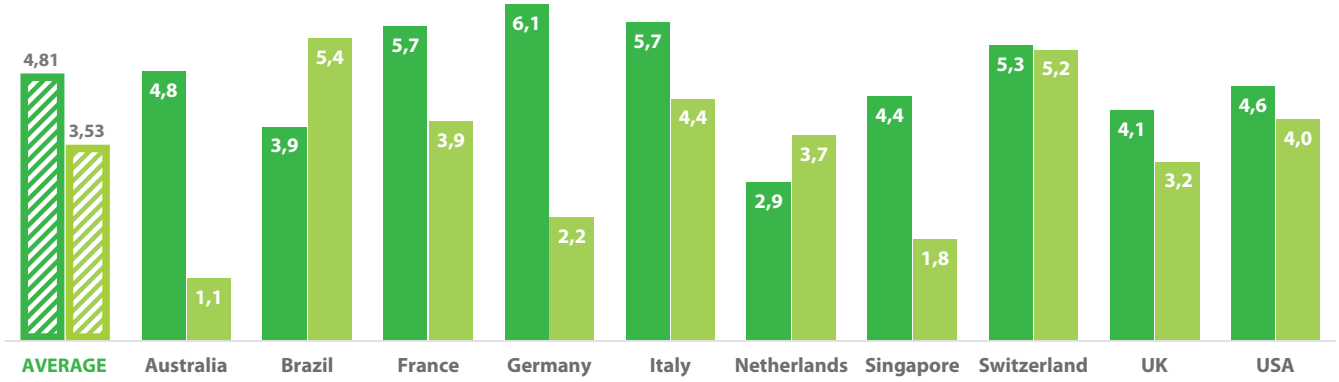

Figure 10: Average backup period against current RPO: mission-critical applications (hours)

Average backup period

Current RPO

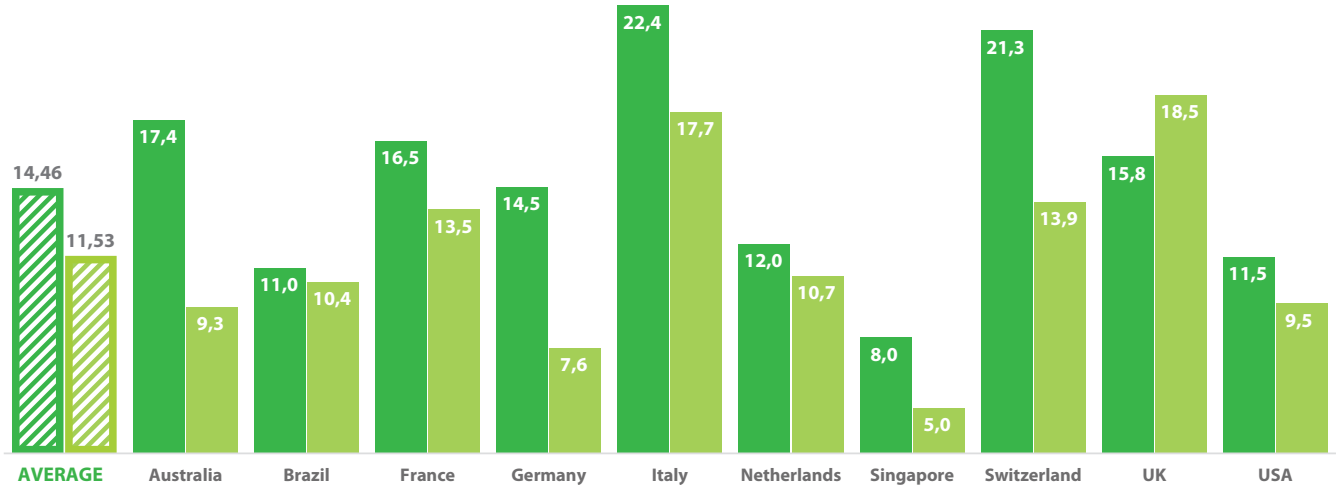

Figure 11: Average backup period against current RPO: non-mission-critical applications (hours)

Average backup period

Current RPO

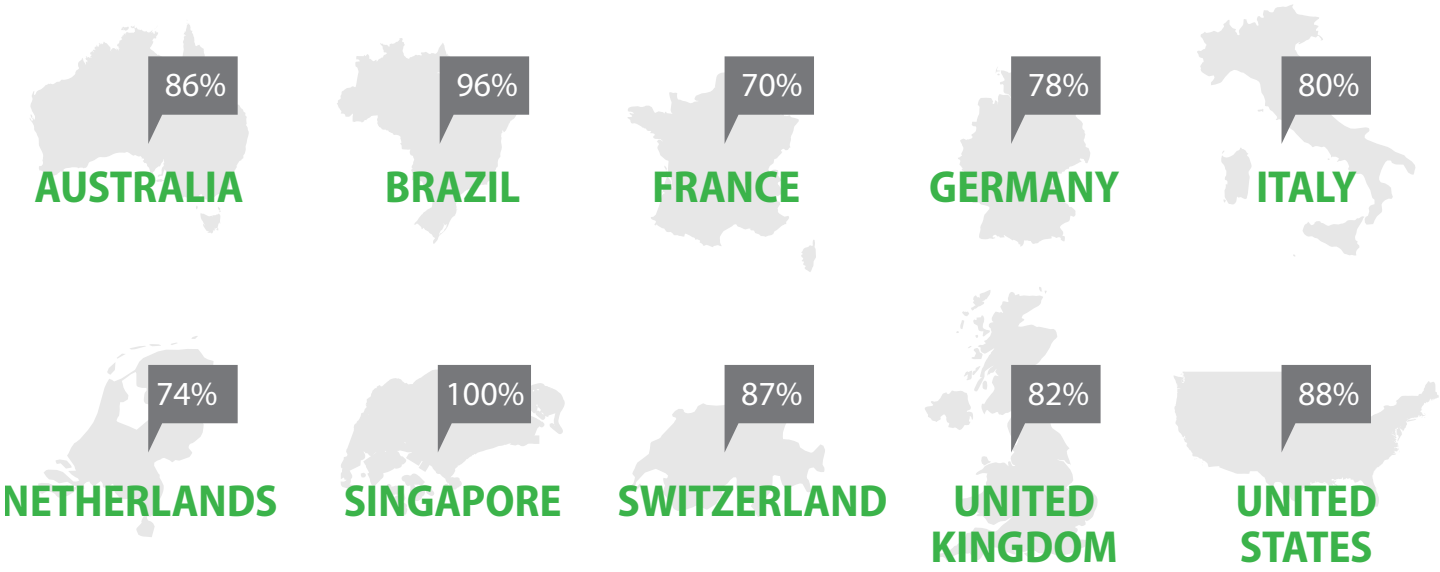

Figure 12: Percentage of organizations identifying an "availability gap" (%)

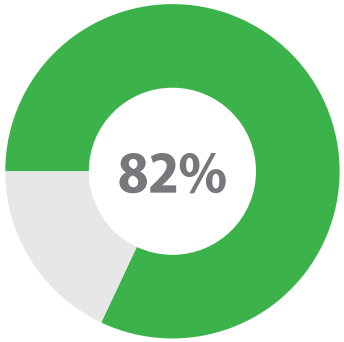

Yet even this performance is currently not enough to meet a company's availability demands. 82 percent of respondents say that there is an "availability gap" between the level of availability they can provide now and what end-users demand in order to provide an always-on business (Figure 12).

In order to address this gap, respondents state that their data protection solutions would need to offer an RTO of 1.73 hours and an RPO of 3.19 hours. As a result, organizations would need to recover mission-critical data in 60 percent of the time it takes them now (Figure 13 on page 14) and perform backup 1.5 times as often as they do now (Figure 14 on page 14). However, it is highly likely that RTO and RPO requirements will continue to shrink as businesses attempt to guarantee 24/7 IT services. This means that, despite their data center modernization efforts, organizations will fall further behind the SLAs that the always-on business demands.

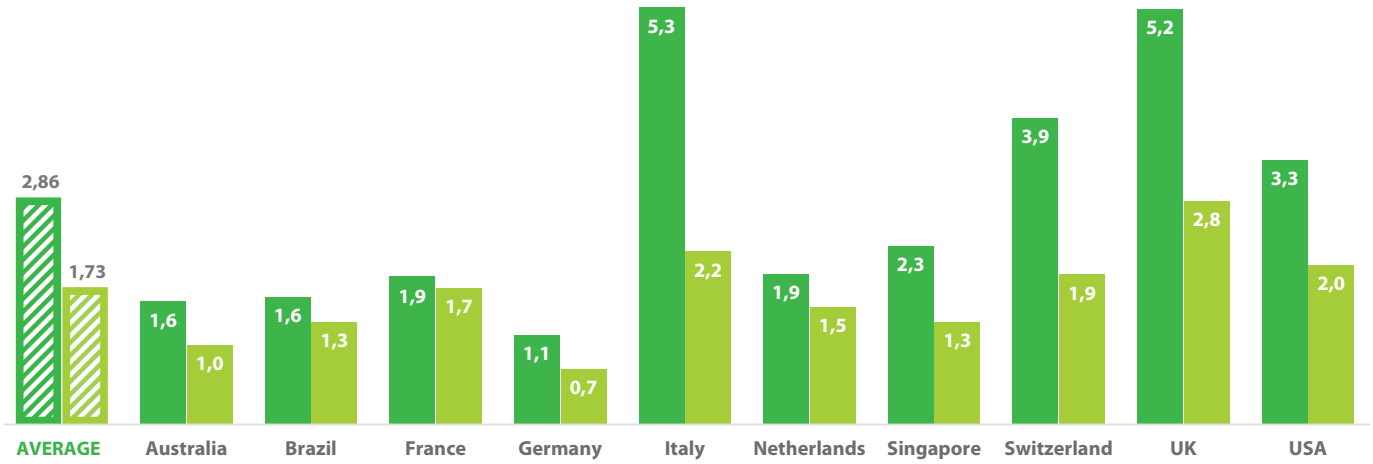

Figure 13: Difference between desired RTO and average recovery time: mission-critical applications (hours)

- Average recovery time
- Desired RTO

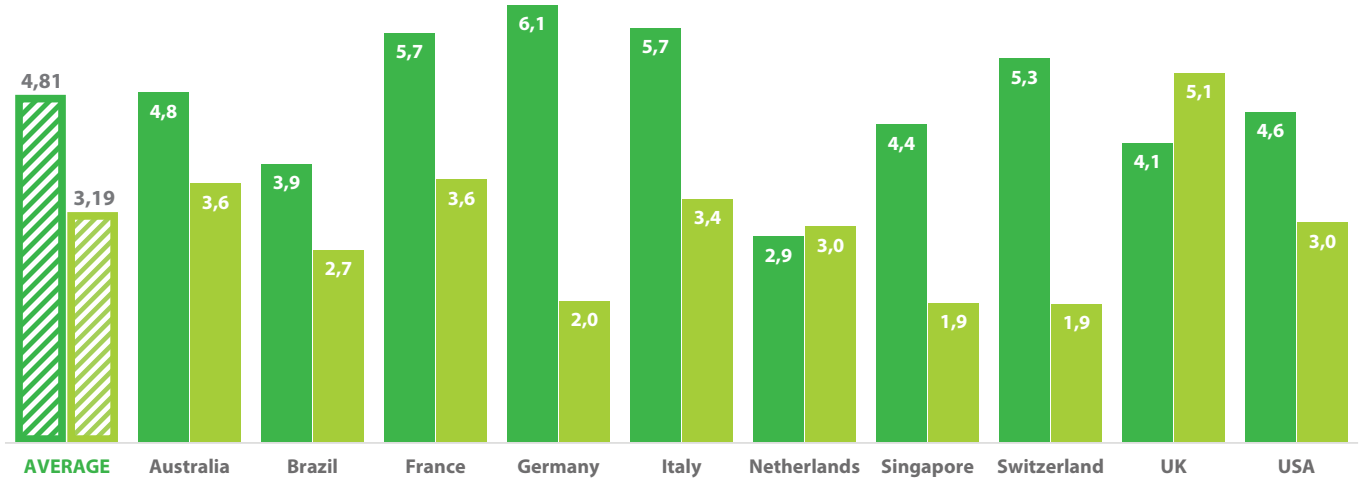

Figure 14: Difference between desired RPO and average backup period: mission-critical applications (hours)

- Average backup period
- Desired RPO

Part 4:
The Financial Cost of Downtime

Part 4: The Financial Cost of Downtime

One significant effect of failing to meet data and application availability demands is that organizations are exposed to unnecessary costs. For instance, businesses may miss sales opportunities or have to restructure their operations at some expense when critical applications fail; not to mention the costs of lost productivity during any downtime. Since there is no longer a "safe" period of downtime outside normal nine-to-five working hours, these costs are magnified. On average, organizations encounter unplanned downtime 13 times per year (Figure 15).

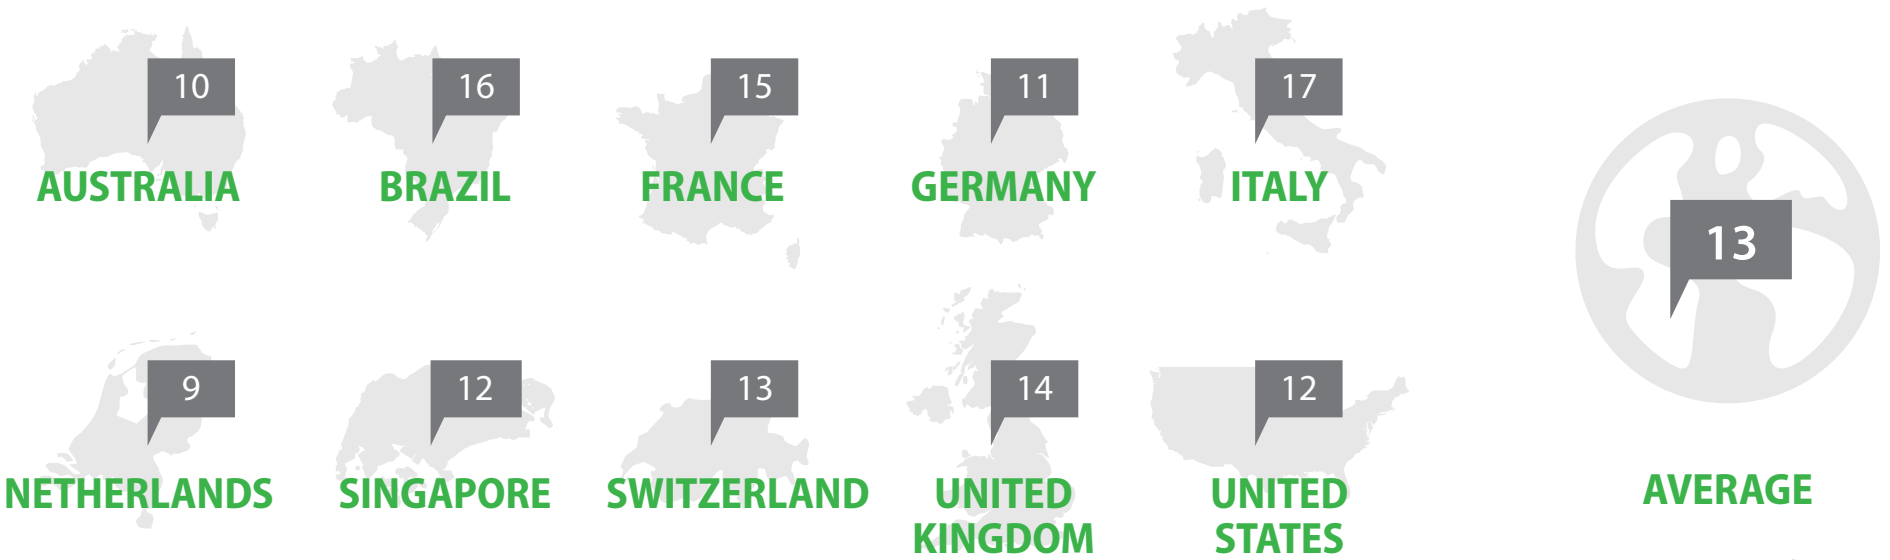

Figure 15: Number of incidents of unplanned application downtime per year

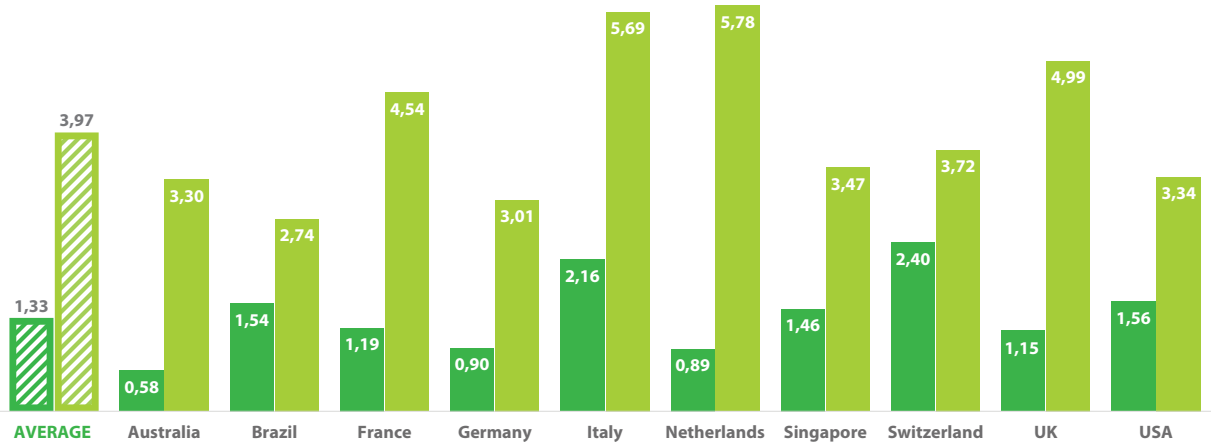

Figure 16: Length of unplanned application downtime (hours)

Mission-critical applications

Non-mission-critical applications

This downtime lasts 1.33 hours for mission-critical applications and 3.97 hours for non-mission-critical applications (Figure 16). The average cost of one hour of downtime for a mission-critical application is $82,864, and for a non-mission-critical application is $43,886 (Figure 17). This means that an incident of mission-critical application downtime costs, on average, $110,209; while non-mission-critical downtime costs, on average, $174,227 (Figure 18).

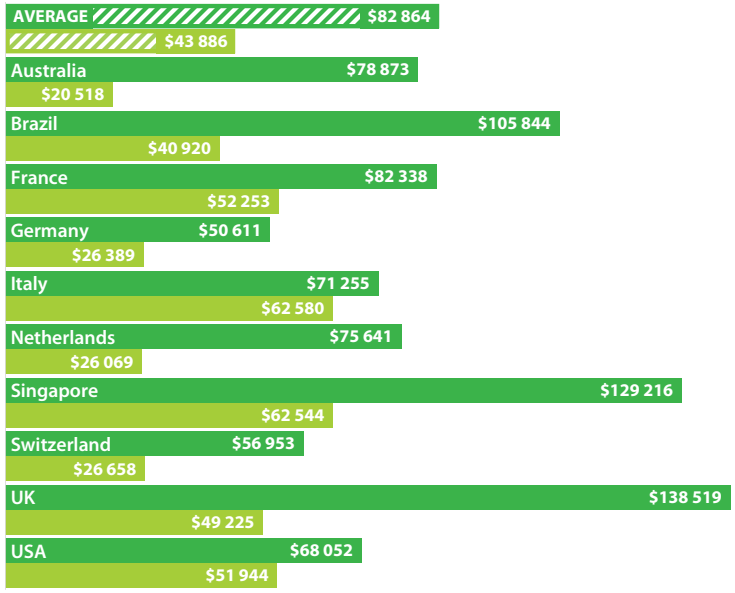

Figure 17: Cost per hour of application downtime (US$)

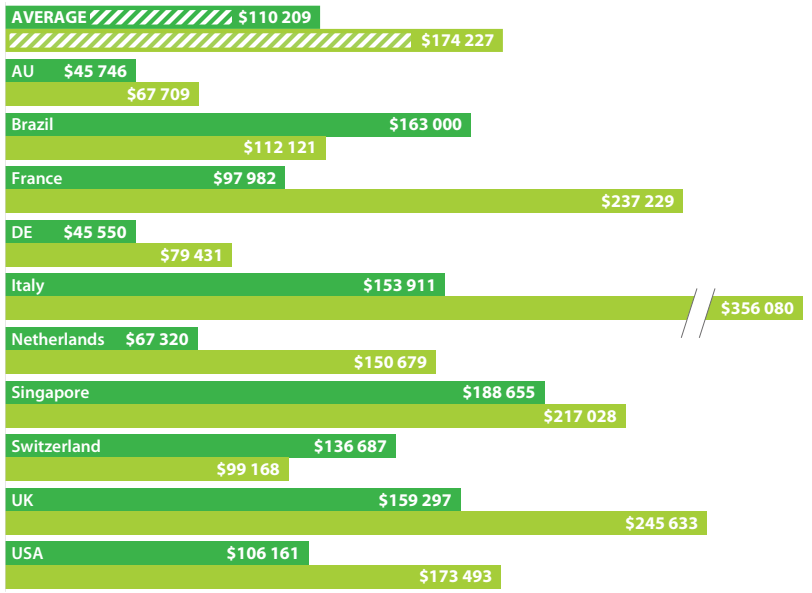

Figure 18: Cost per incident of application downtime (US$)

Mission-critical applications

Non-mission-critical applications

As well as the costs of downtime itself, there is also a cost associated with data loss – i.e. data that has not been backed up and so cannot be recovered in the event of downtime. Depending on the importance of the data itself, this can be a huge cost for an organization in terms of missed sales opportunities and lost productivity. Data loss for mission-critical applications costs on average $70,913 per hour of data lost. For non-mission critical applications, data loss costs an average of $42,016 per hour of lost data (Figure 19).

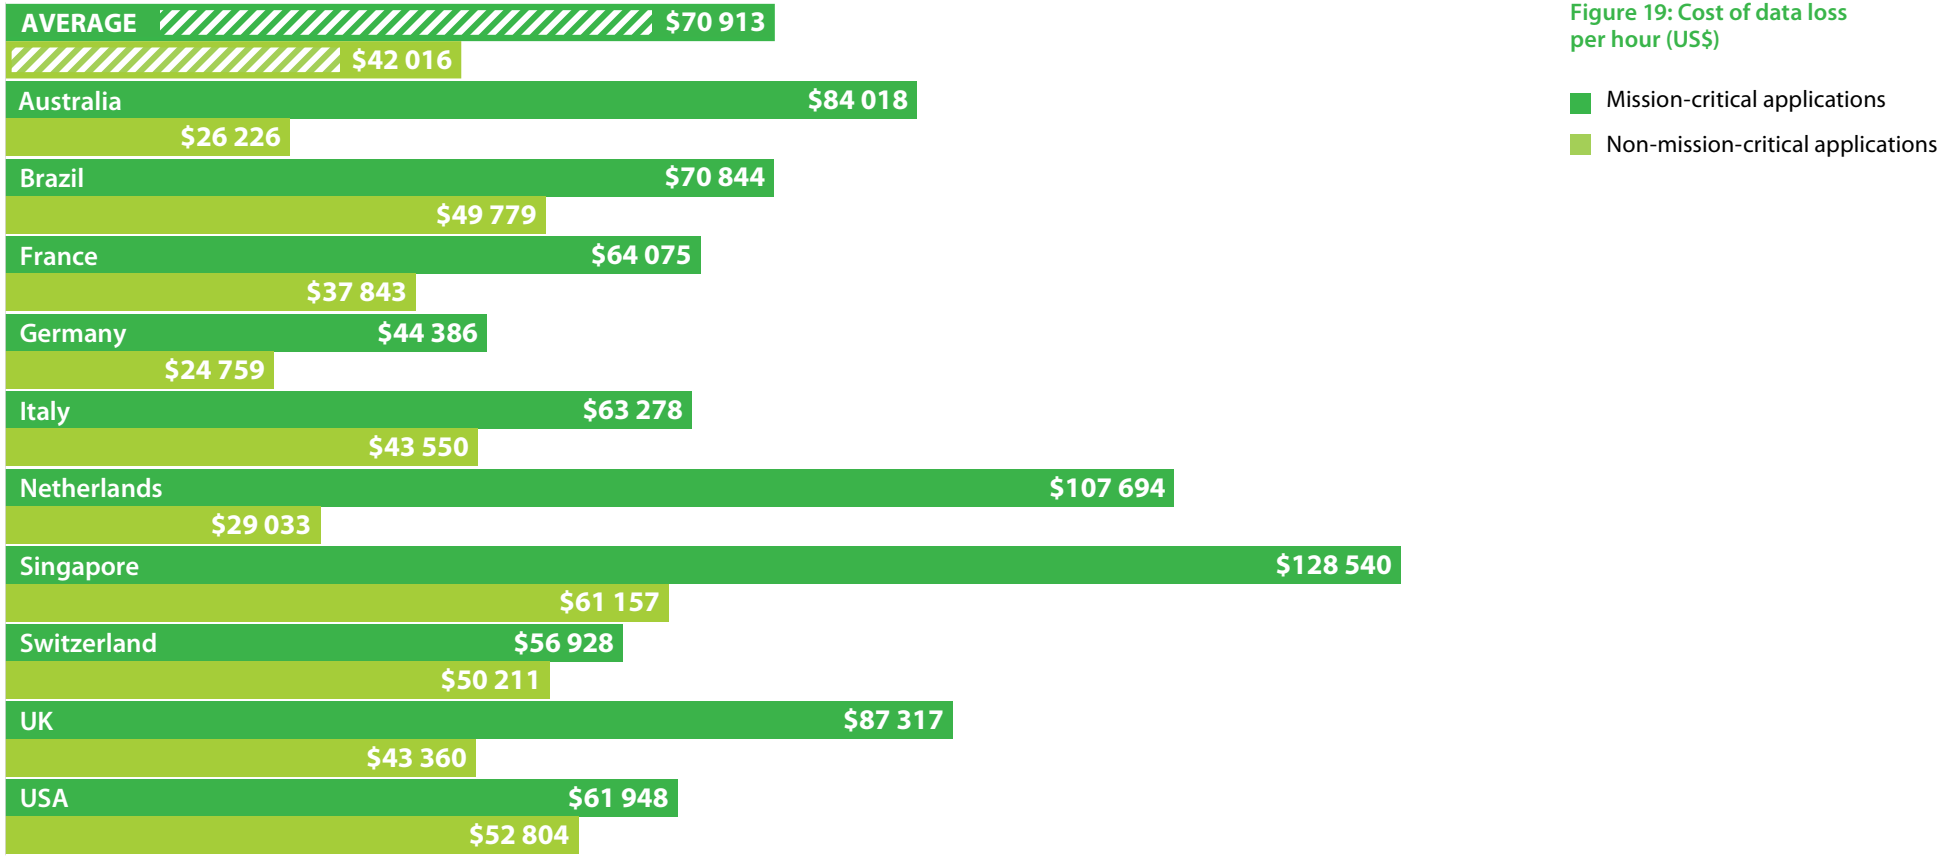

Figure 19: Cost of data loss per hour (US$)

Comparing the frequency of back-up for applications, as shown in Figures 10 and 11, this shows that a single incident can cost organizations up to $341,091 in lost data for mission-critical applications, and up to $607,551 for non-mission critical (Figure 20).

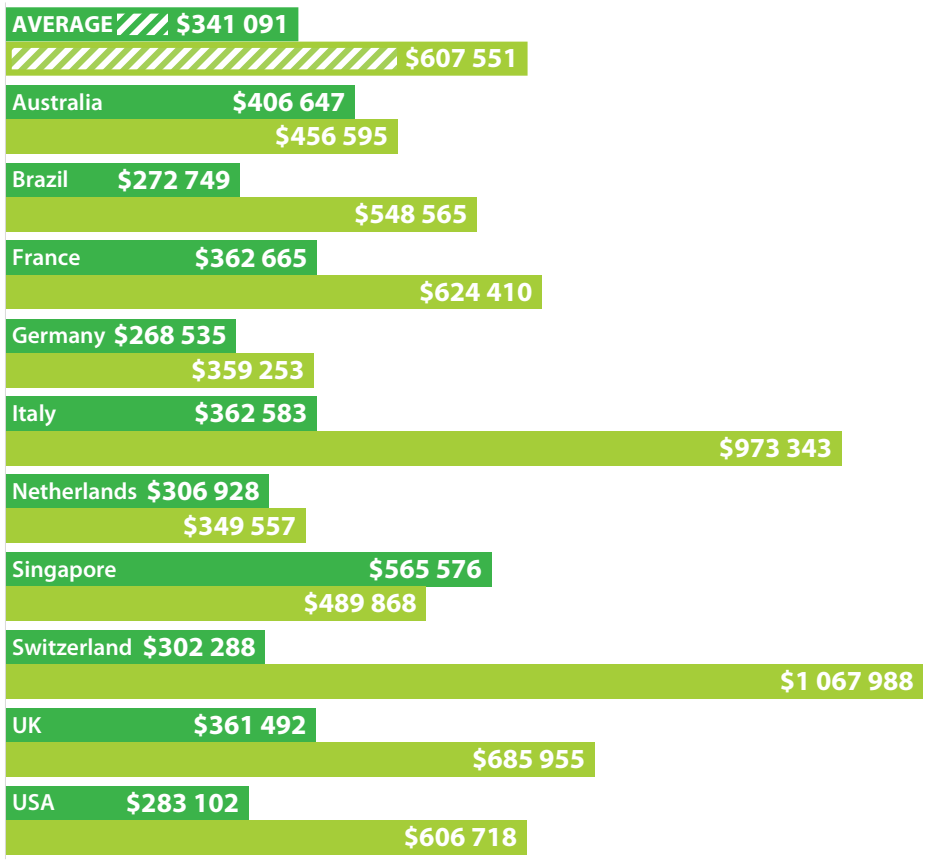

Figure 20: Maximum cost of data loss per incident (US$)

- Mission-critical applications
- Non-mission-critical applications

In total, a single incident of downtime can cost organizations $451,300 for mission-critical applications and $781,778 for non-mission critical, adding the costs of downtime and data loss (Figure 21).

With an average of 13 incidents per year, enterprises face an average annual cost of up to $10,163,114, depending on the nature of the application and how much data is lost in each case (Figure 22).

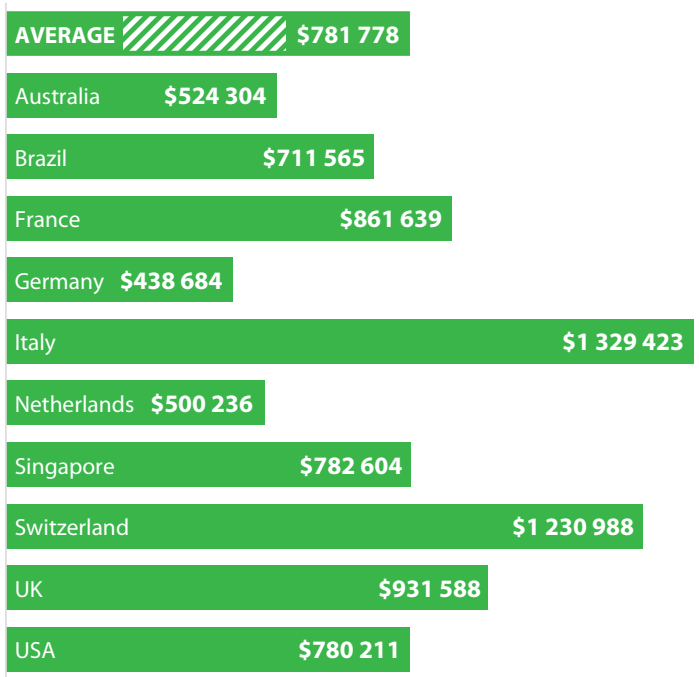

Figure 21: Maximum cost per incident of downtime (US$)

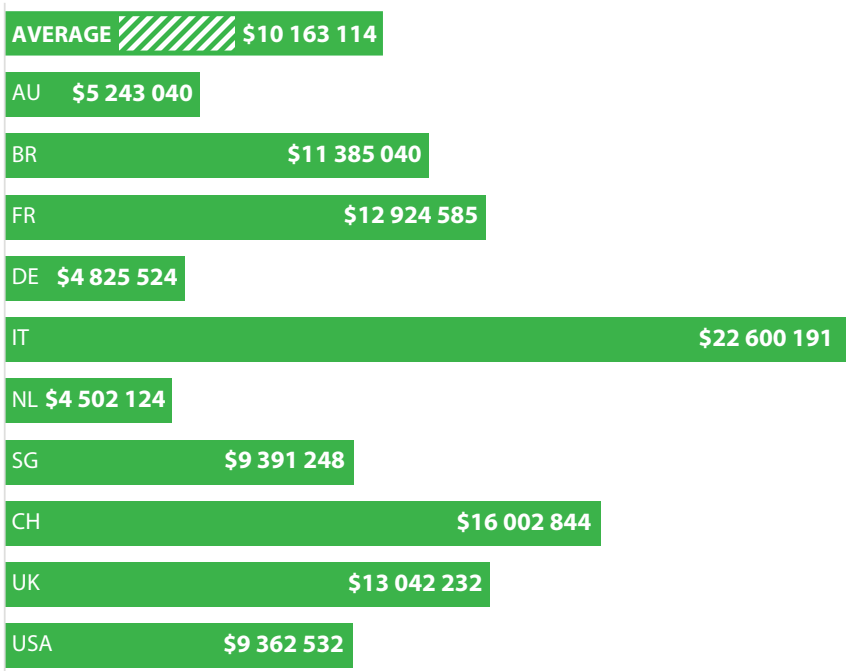

Figure 22: Maximum annual cost of downtime (US$)

These costs are likely to be a significant reason why businesses are demanding constant availability of IT services. For instance, if businesses could meet the data loss SLAs currently demanded by the Always-On Business, as illustrated in Figure 14, for both mission-critical and non-mission-critical applications, then assuming the cost per hour of data loss, as shown in Figure 19, remains constant, the maximum risk of data loss would be $226,212 for mission-critical applications and $134,031 for non-mission critical. This represents reduced risk of at least $100,000 per incident, or $1.3 million a year, from improved RPOs alone (Figure 23). Since improved RTOs are also likely to affect downtime, the cost savings can be significant.

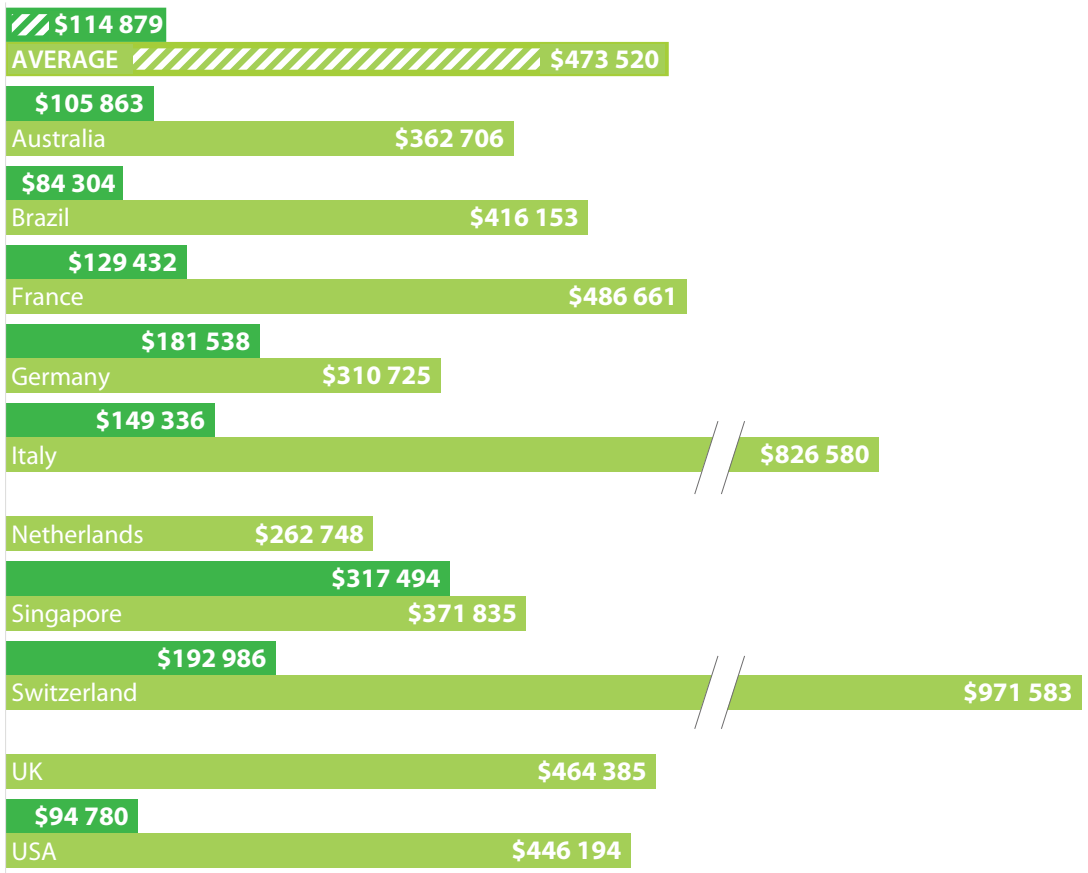

Figure 23: Minimum reduced data loss risk per incident from improved RPOs (US$)

- Per incident of mission-critical application downtime
- Per incident of non-mission-critical downtime

Indeed, if businesses could meet a target RTO and RPO of 15 minutes or less, which modern data protection tools can provide, then the savings would be significant. Again assuming that the cost-per-hour of both data loss and application downtime, as in Figure 17, remains constant, then the maximum cost of an application failure would be $38,444 (Figure 24). If unexpected downtime happens 13 times a year, this represents a maximum annual cost of $499,772, meaning a bare minimum saving of $932,945.

These statistics show the true costs of the "availability gap", and these costs will only increase as demand for the Always-On Business grows. Businesses need to act immediately in order to ensure that the costs of the availability gap don't grow from tens to hundreds of thousands.

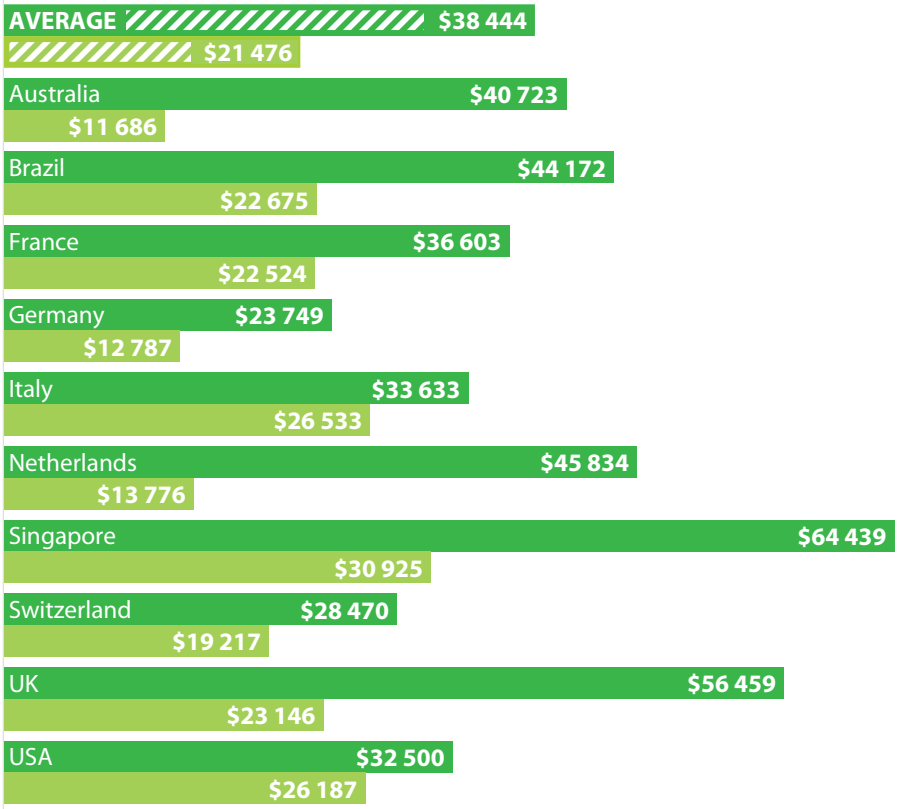

Figure 24: Maximum cost per downtime incident from meeting 15-minute RTPO (US$)

Maximum mission-critical cost

Maximum non-mission-critical cost

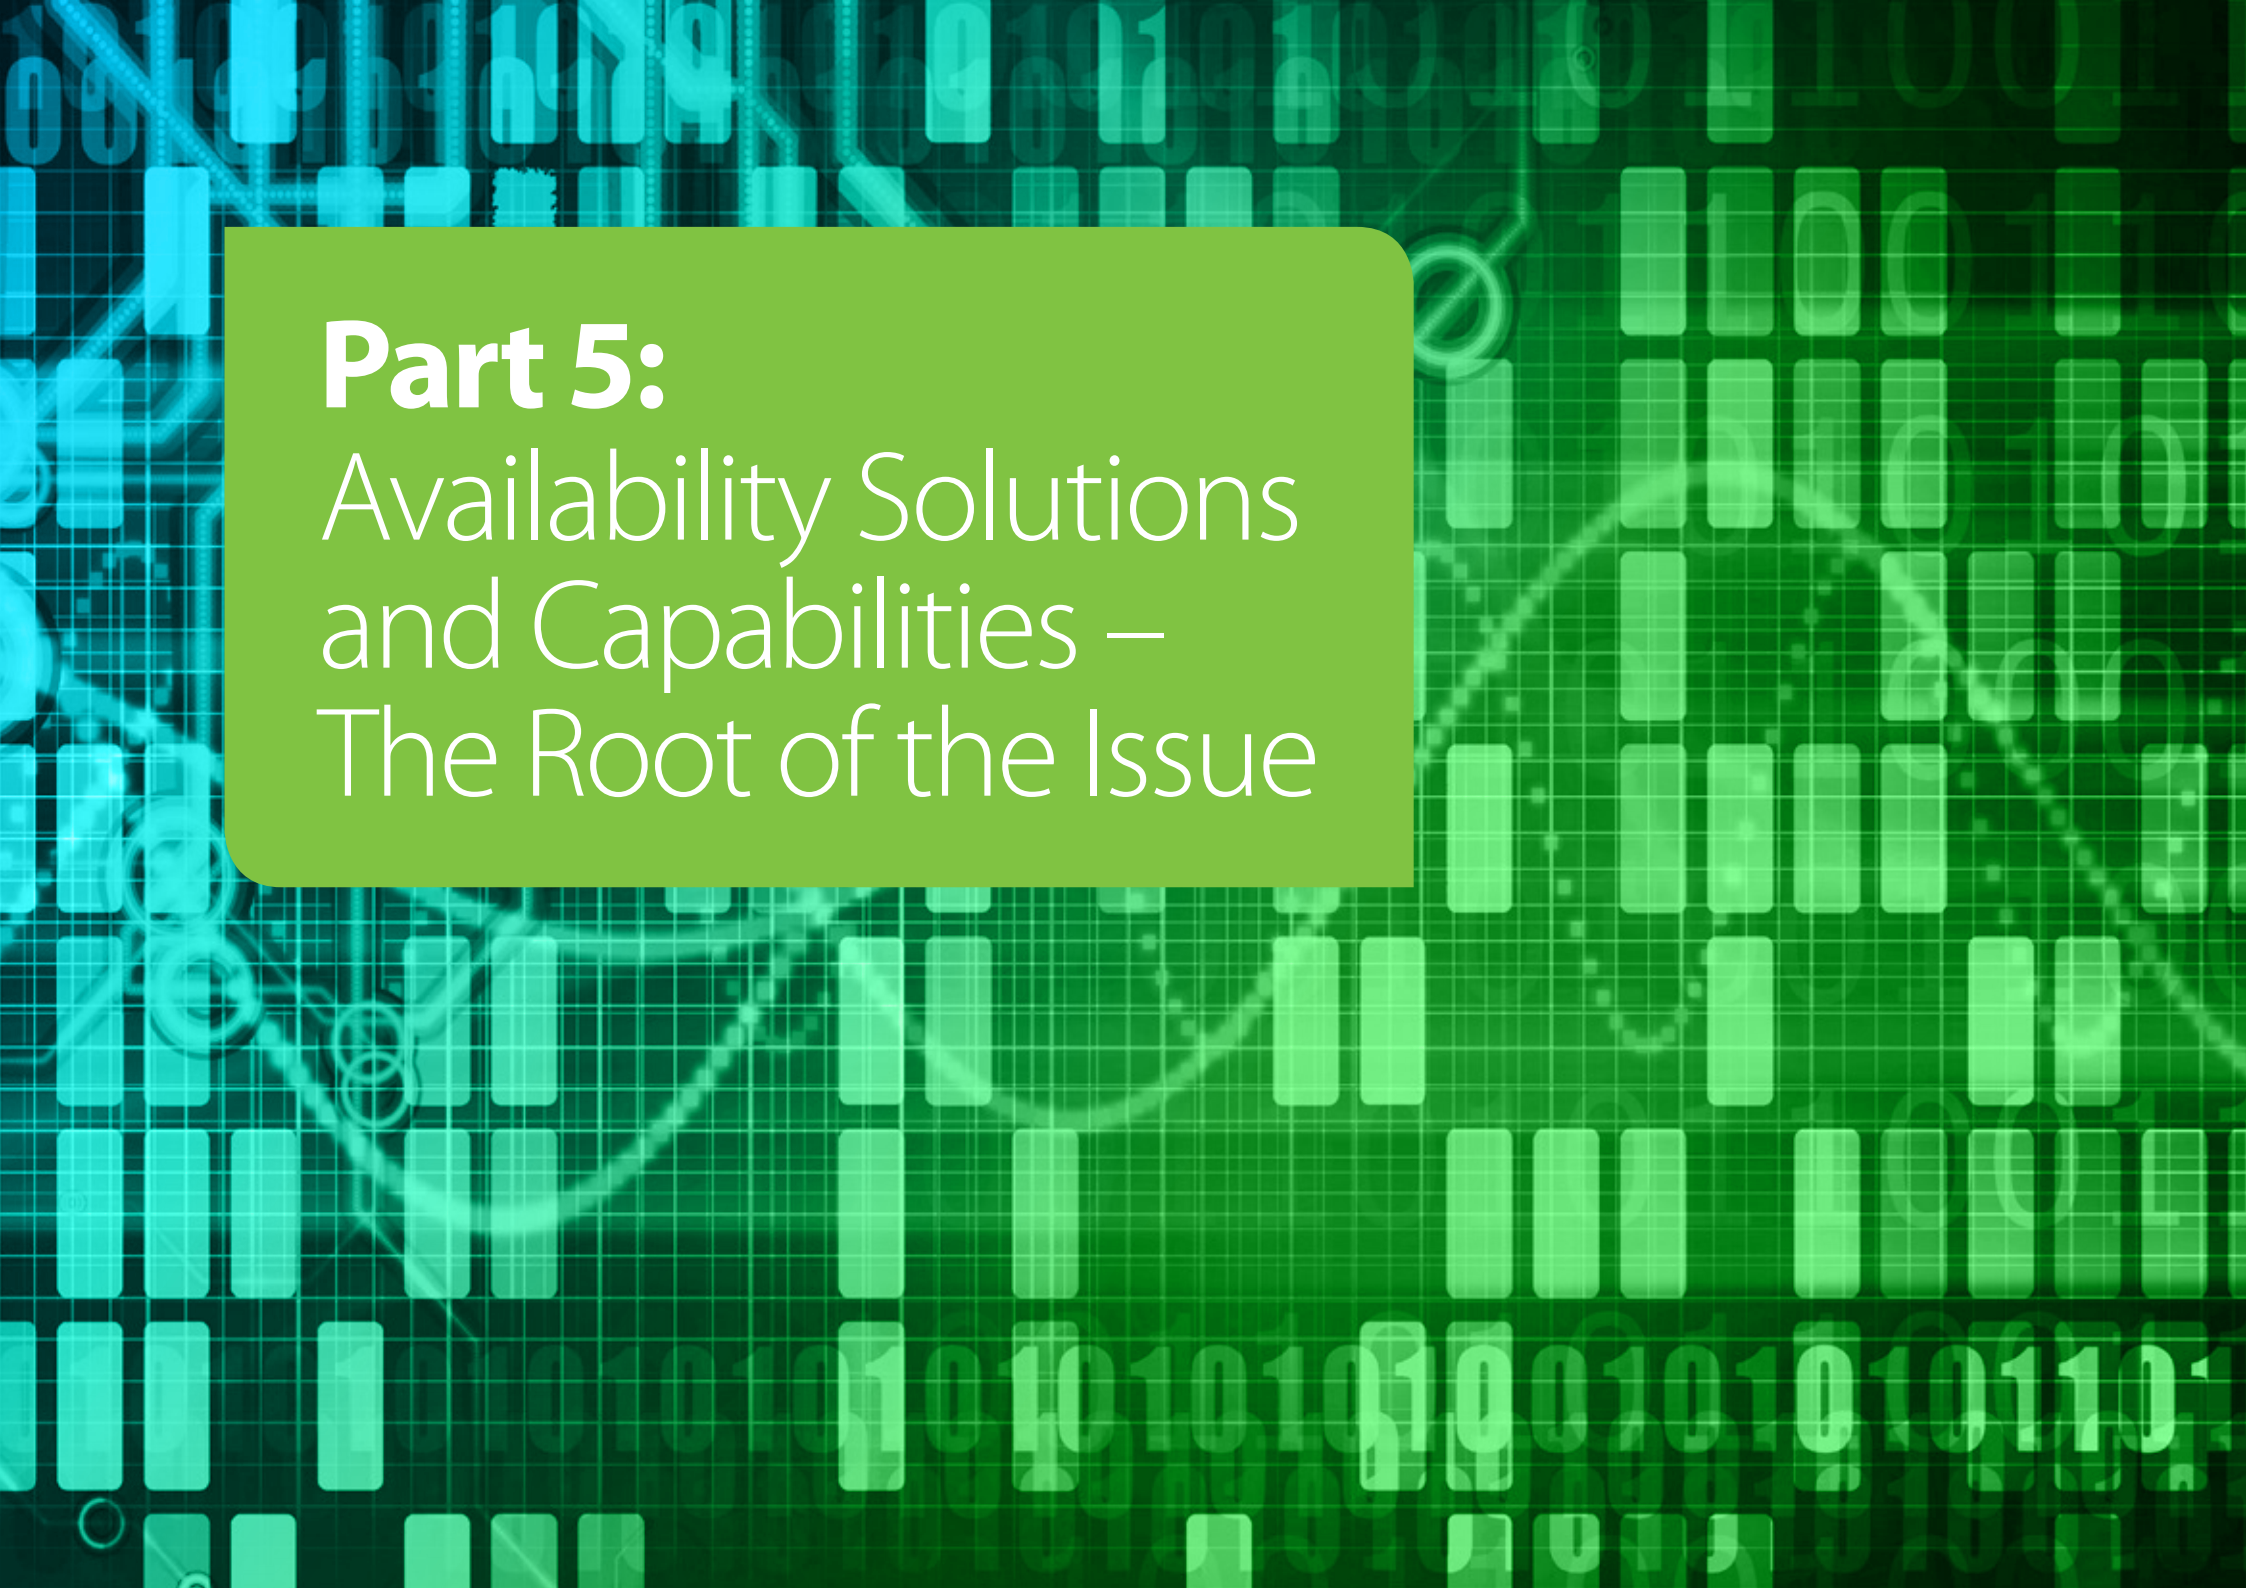
Part 5:
Availability Solutions and Capabilities – The Root of the Issue

Part 5: Availability Solutions and Capabilities – The Root of the Issue

An inability to support the Always-On Business ultimately comes down to a business's legacy backup solution: without sufficient capabilities, IT departments cannot guarantee the RTOs and RPOs that the business demands.

Organizations recognize this: 92 percent of respondents identified availability capabilities they would like to have in their data center, but were currently unable to implement. These capabilities included high-speed recovery, i.e. the ability to recover any application or server in under 15 minutes, which 60 percent of organizations wanted. Other capabilities demanded were data loss avoidance, i.e. reducing data loss to 15 minutes or less (53 percent); verified protection, i.e. guaranteed recovery of every file and application every time (47 percent); using backup data as a production-like test environment for new patches or updates (38 percent); and complete visibility, including proactive monitoring and alerting of issues before any operational impact (36 percent) (Figure 25).

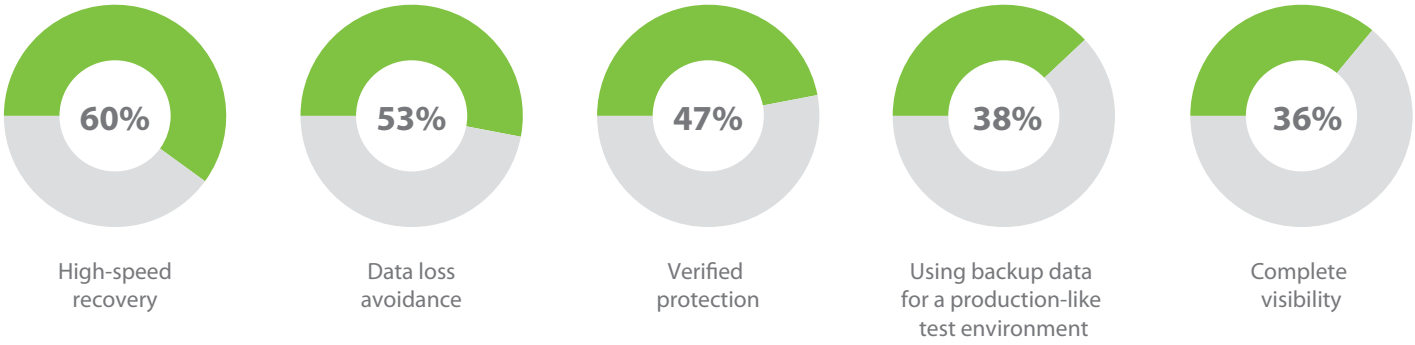

Figure 25: Capabilities organizations would like in their data center, but cannot implement (%)

Organizations also recognized the barriers preventing them from implementing these capabilities. The primary barrier, in every case, is the cost of new technology, followed by the complexity of development or a lack of expertise; the current product not providing the capabilities needed; and human resources constraints (Figure 26).

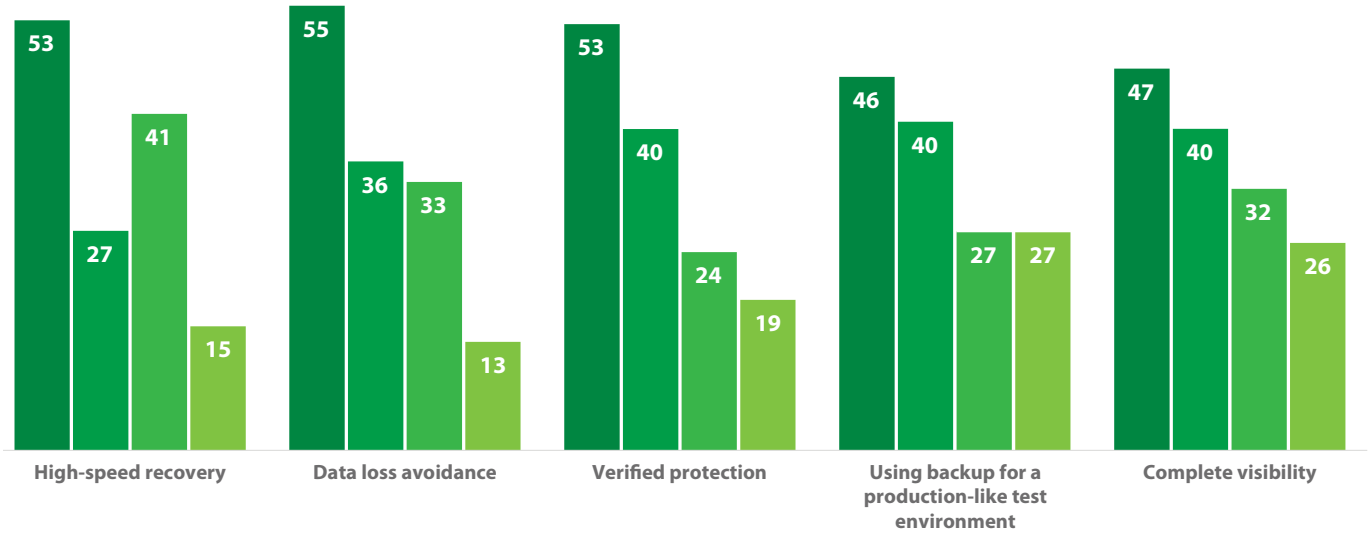

Figure 26: Factors preventing organizations from implementing capabilities (percentage of organizations reporting factor) (%)

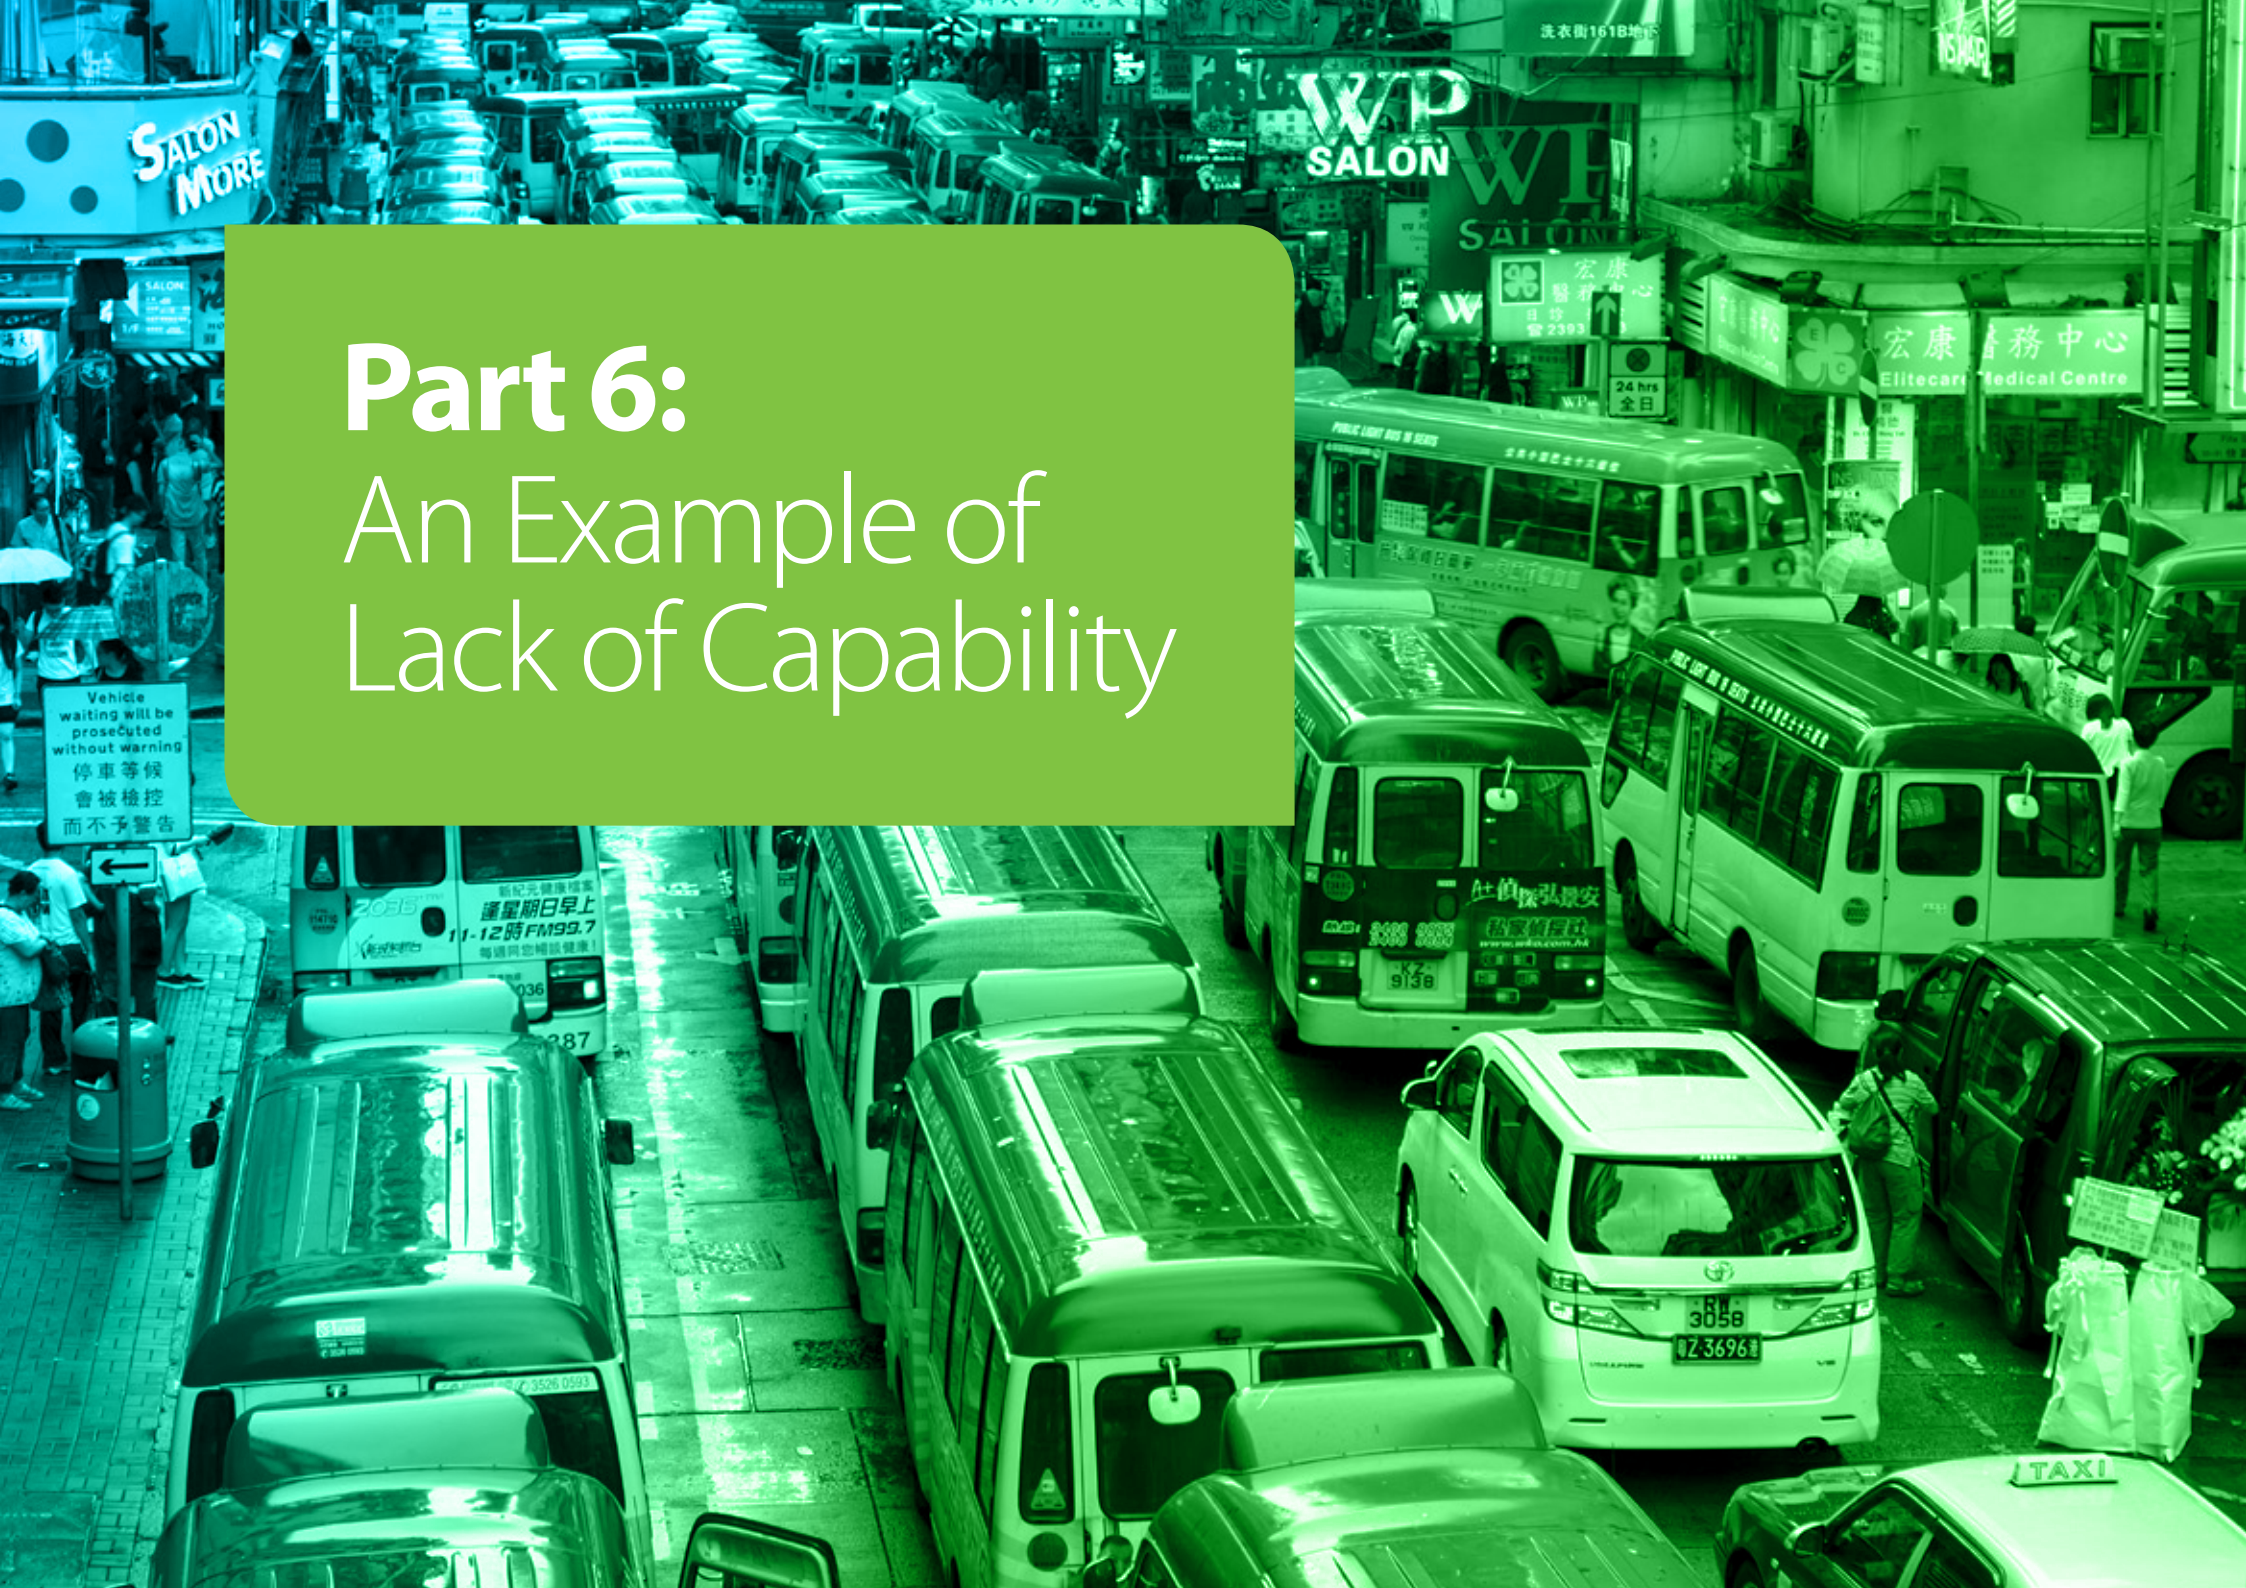

Part 6:
An Example of Lack of Capability

Part 6: An Example of Lack of Capability

An example of how this lack of capability affects data and application availability in businesses can be found in testing and verification. When a backup is made, there is always a chance that it is damaged and will not recover when needed. By testing backups, organizations can verify that they will recover correctly and that nothing will be lost. However, without the correct capabilities, verification is a time-consuming task meaning that only a fraction of backups will be verified.

Organizations test their backups for recoverability on average every eight days (Figure 27).

However, each quarter, organizations only test an average of 5.26 percent of their backups (Figure 28); meaning that the vast majority of backups are not verified and so could fail.

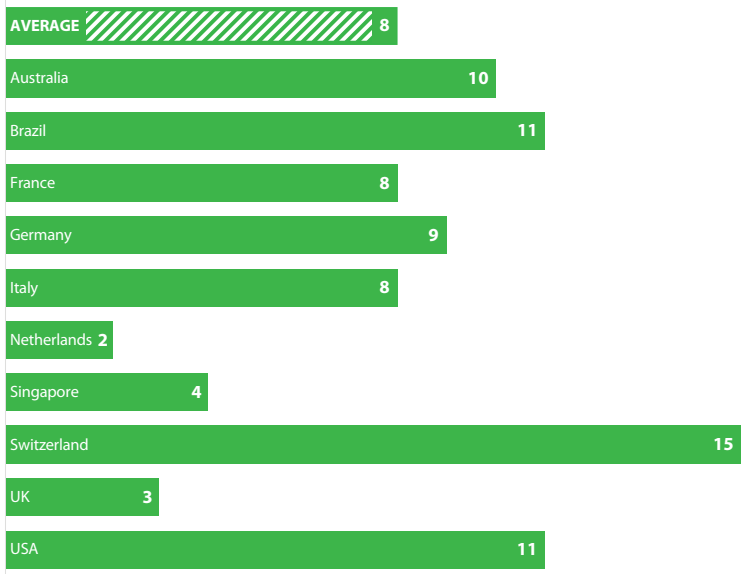

Figure 27: Frequency of backup testing (days)

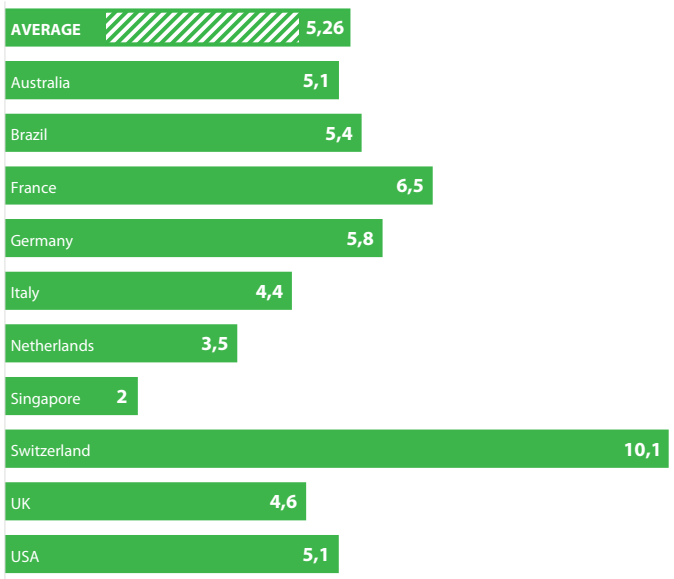

Figure 28: Percentage of backups tested each quarter (%)

This is borne out by the fact that 16.74 percent of backups fail to recover (Figure 29). With unplanned downtime occurring 13 times a year, this means that every year organizations' recovery will fail twice – greatly increasing the duration, data loss and cost of downtime.

Indeed, because of these failures, data loss will cost organizations a minimum of $682,182 a year; since the best case situation is that they will have to roll back to the last valid backup (Figure 30).

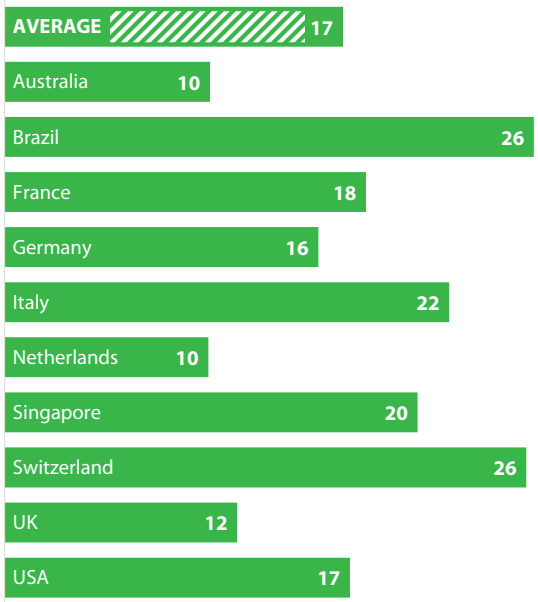

Figure 29: Percentage of backups that fail to recover (%)

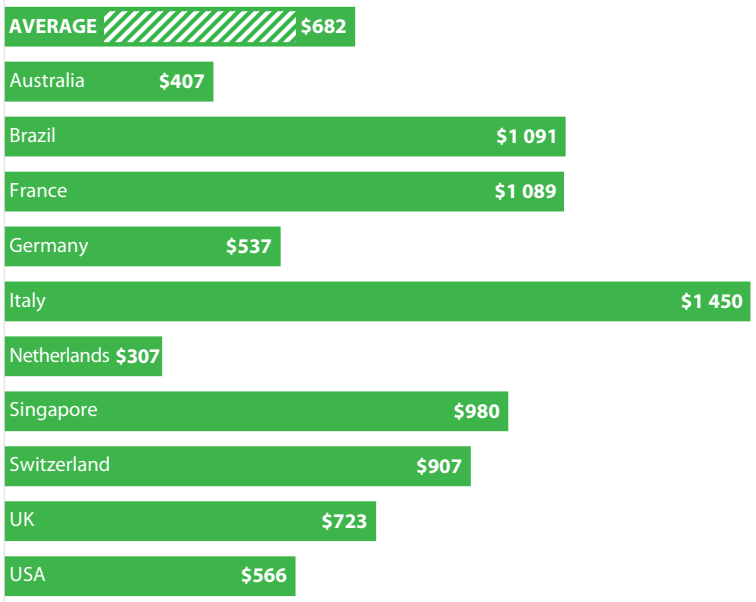

Figure 30: Minimum annual cost of unavoidable data loss (thousands US$)

Adding the cost of this data loss to the minimum average cost of downtime, enterprises will lose at least $2 million from application failure every year (Figure 31).

Testing is not only necessary for validating the recoverability of backups. Testing patches or application updates in a production-like sandbox environment before rolling them out into production can also ensure that these application patches or upgrades will perform as expected, and that the business will not suffer more downtime than expected. However, this is currently not the case. 87 percent report more downtime than expected when they perform patches or upgrades to applications (Figure32).

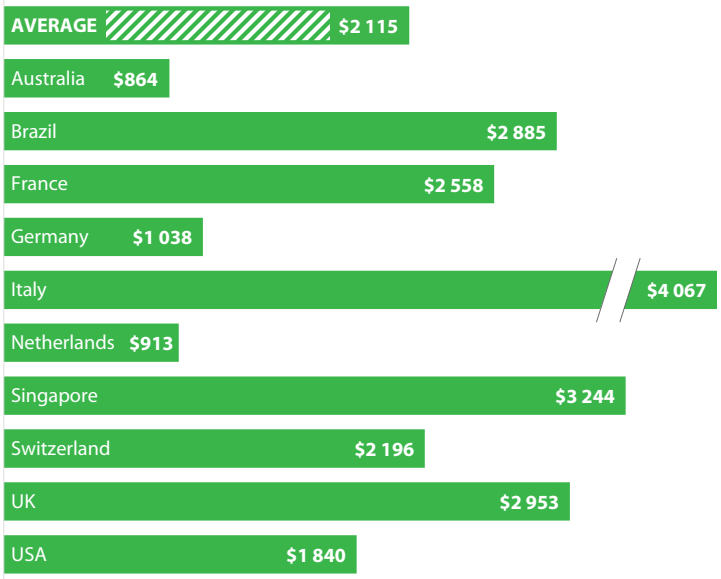

Figure 31: Minimum annual cost of application failure (thousands US$)

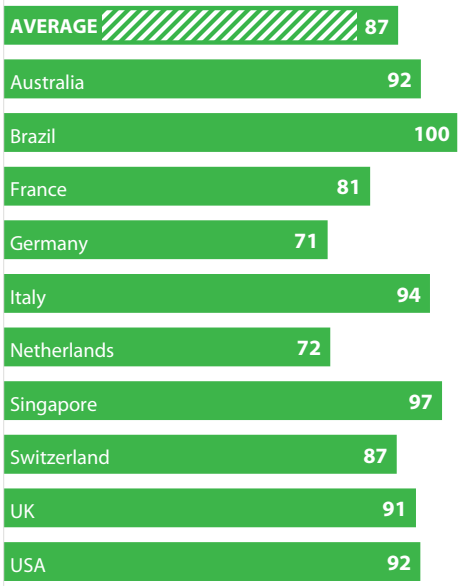

Figure 32: Percentage of organizations reporting more downtime than expected when performing patches or upgrades to applications (%)

Part 7:
Looking Ahead

Part 7: Looking Ahead

As we have seen, organizations are well aware of the need to deliver Availability for the Modern Data Center for the Always-On Business, and that they are not yet ready to do so. Indeed, in two years' time the landscape should look very different.

78 percent of organizations plan to change their data protection product in the next two years, with the average timeframe being six months (Figure 33).

This will be critical if business demands are to be met. It is clear that the current trend of data center modernization is not giving organizations the capabilities they need to make the always-on business a reality. Instead, IT departments need to be certain that recovery time is as short as possible; that data loss is minimized; and that backups will recover as expected, when they are needed. Without these capabilities, businesses will have no option but to bear the increasing costs of the availability gap.

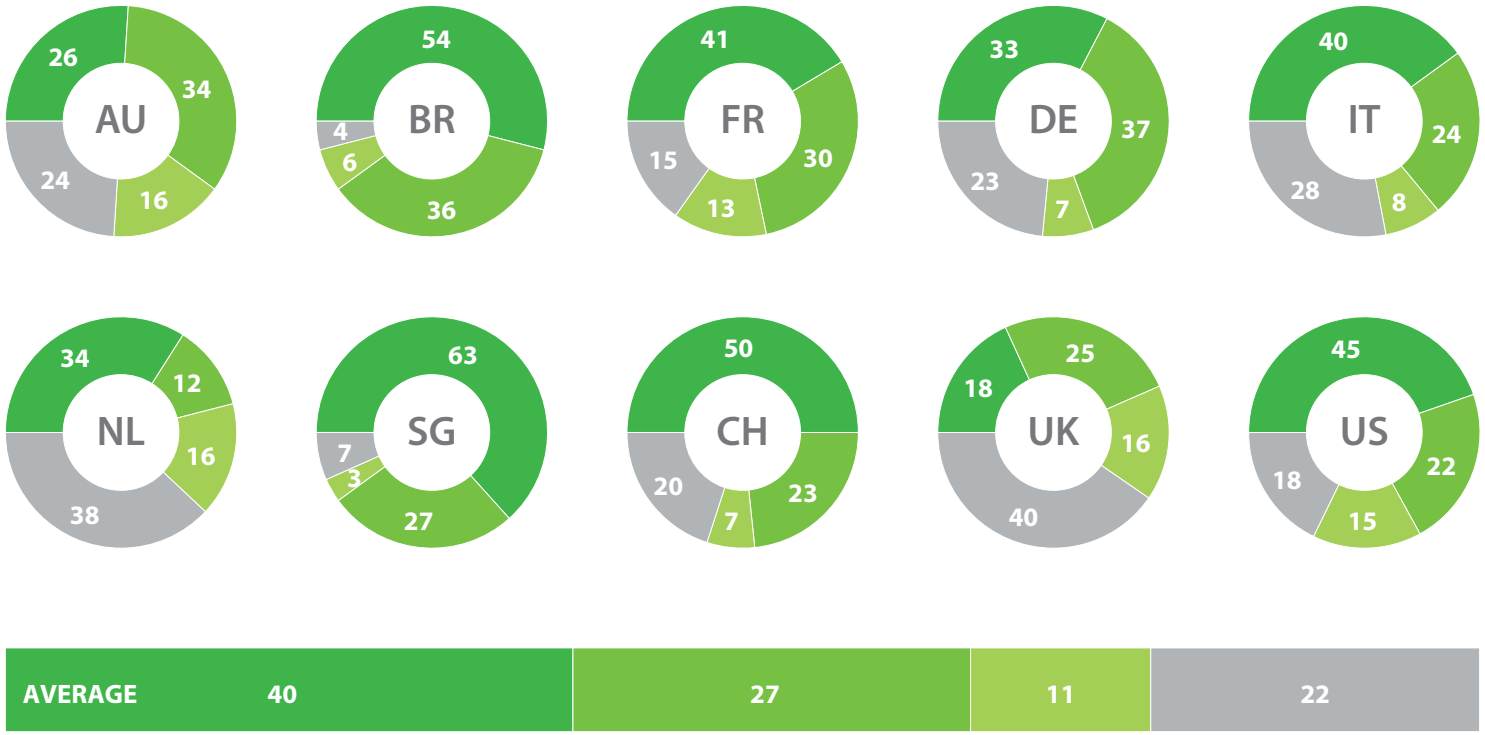

Figure 33: Percentage of organizations planning to change their data protection product in the next 2 years (%)