WHITE PAPER

Damage Control: The Impact of Critical IT Incidents

Written by Bob Tarzey, analyst and director, Quocirca
November 2017

splunk> turn data into doing™

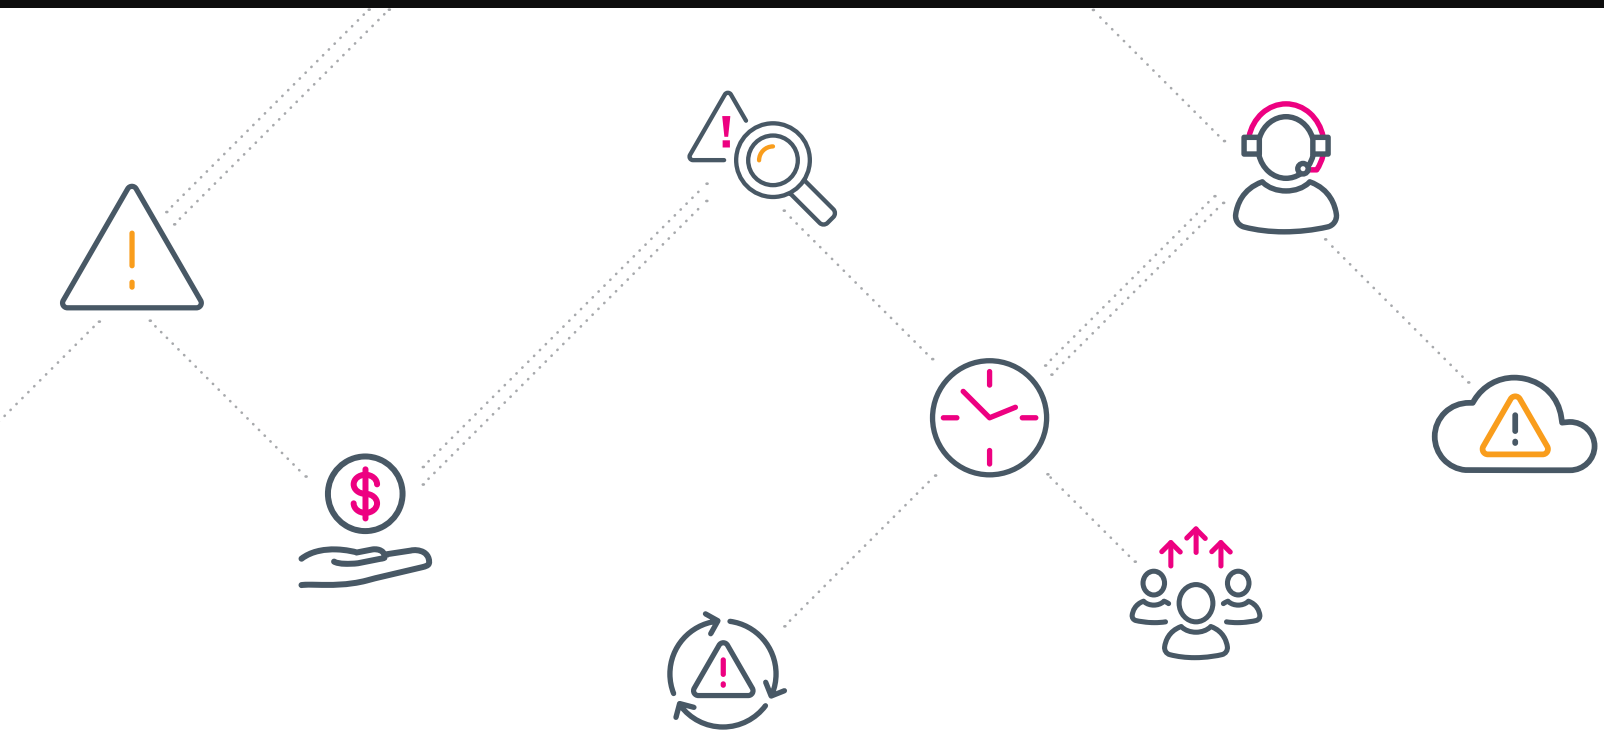

Damage Control

The impact of critical IT incidents

Written by Bob Tarzey, analyst and director, Quocirca
November 2017

IT incidents result in annual costs running into millions of dollars for organizations across all industry sectors. Dealing with incidents distracts IT staff from other activities; beyond the IT department, incidents impact business productivity and the customer experience.

The average organization is dealing with hundreds of incidents every week. The IT systems and infrastructure that underlie these incidents generate vast amounts of data. Collecting and analysing these data from the various IT systems involved, leads to the faster, more efficient detection of incidents, improved troubleshooting, and shorter mean time to repair. Event data can also assist with root cause analysis, helping to prevent problems recurring. Time saved from dealing with IT incidents leaves IT staff with more time to focus on digital innovation, rather than constantly fighting to keep the lights on.

Executive Summary

Incidents and critical incidents

The average organization logs about 1,200 IT incidents per month, of which 5 will be critical. It is a challenge to wade through all the data generated by the events that lead to these incidents and prioritize dealing with them. Seventy percent say a past critical incident has caused reputational damage to their organization, underlining the importance of timely detection and to minimize impact.

The cost to IT and to the business of critical incidents

The mean cost to IT of a critical incident is US $36,326, the mean downstream cost to business is an additional US $105,302. These two costs rise together, suggesting high cost to IT is a proxy for poor event and incident management, which has a knock-on effect for business operations.

Mean time to detect, repair and perform root cause analysis

Eighty percent say they could improve their mean time to detect incidents, which would lead to faster resolution times and decrease the impact on businesses. The mean time to repair for critical incidents is 5.81 hours — this reduces if there are fewer incidents to manage in the first place. On average, a further 7.23 hours are spent on root cause analysis, which is successful 65% of the time.

Unnecessary incidents

Duplicate and repeat incidents are a persistent problem. Ninety-seven percent say their event management process leads to duplicates, where multiple incidents are created for the same IT problem; 17.2% of all incidents are duplicates. Ninety-six percent say failure to learn from previous incidents through effective root cause analysis leads to repeat incidents; 13.3% of all incidents are repeats.

Visibility into IT infrastructure

The monitoring of IT infrastructure to log events and identify incidents could be improved; 80% admit they have blind spots, leading to delayed detection and investigation of incidents. The complexity of IT systems and the tools that monitor them leaves many organizations without an adequate, holistic end-to-end view of their IT infrastructure.

Coping with the volume of events and event management

Dealing with the volume of events generated by IT monitoring tools is a challenge. Fifty-two percent say they just about manage, 13% struggle, and 1% are overwhelmed. Those with event management processes which enable them to easily manage the volume of events have a faster mean time to detect incidents and fewer duplicate and repeat incidents.

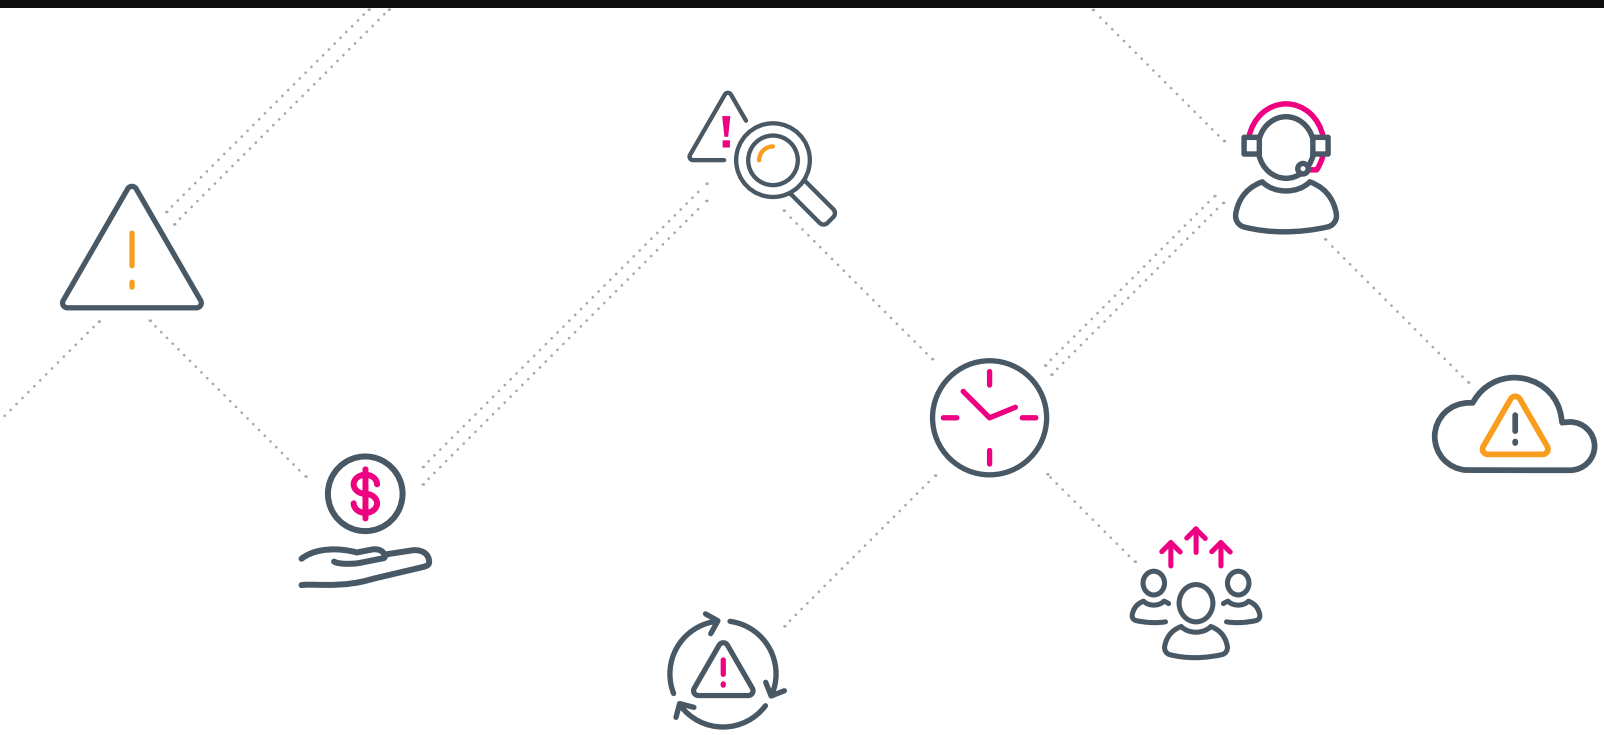

Introduction

This report presents new research into the impact that the volume of IT incidents has on organizations, especially the critical incidents that can halt business processes and impact users and customers. It looks at the capabilities in place to process the vast volume of event data generated by the tools that monitor their IT, including their ability to avoid duplicate incidents, and to reduce mean time to detect (MTTD) and mean time to repair (MTTR). It also looks at how root cause analysis can avoid the future repetition of incidents. Early detection of incidents reduces the impact of the downstream factors that follow, and the cost of incidents to both IT departments and the businesses they serve.

The research was carried out across 9 countries: USA, Japan, Singapore, Australia, Sweden, Netherlands, Germany, France and UK. The research covered a range of business sizes and sectors (see appendix).

IT Incidents and Critical Incidents

As organizations become more reliant on IT, their IT infrastructure has become increasingly complex. The pressure to support new digital initiatives and remain competitive, leads to the addition of new layers of IT infrastructure — including virtualization, containerisation and cloud services. This leaves users frustrated, incurs unwanted business costs, and can lead to reputational damage.

The research shows that incidents are of greatest concern when they impact the user experience or slow down project delivery (Figure 1). For IT management

Organizations suffer 5.1 critical incidents per month

teams, all the event data generated by IT infrastructure monitoring tools needs to be filtered to discover what is relevant in order to troubleshoot problems and perform timely maintenance. IT teams need help to see through noise and recognize the problems that need attention. If not detected and addressed early enough, small problems can become critical incidents.

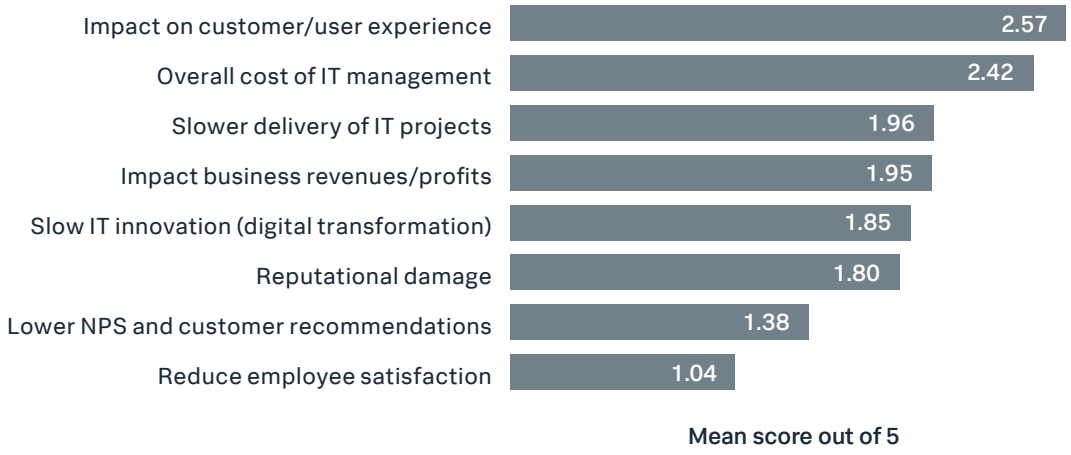

Figure 1: Potential consequences of IT incidents which are of most concern

The highest priority must be given to critical incidents (variously labelled by service desk systems as severity 1, priority 1, or P1). The average organization logs about 1,200 incidents per month, 5.1 of which are critical (Figure 2). Seventy percent say a past critical incident has caused reputational damage to their organization. It can be difficult for IT teams to know what is critical and what is not; they need help to prioritize and triage as effectively as possible.

70% say a critical incident has caused reputational damage to their organization

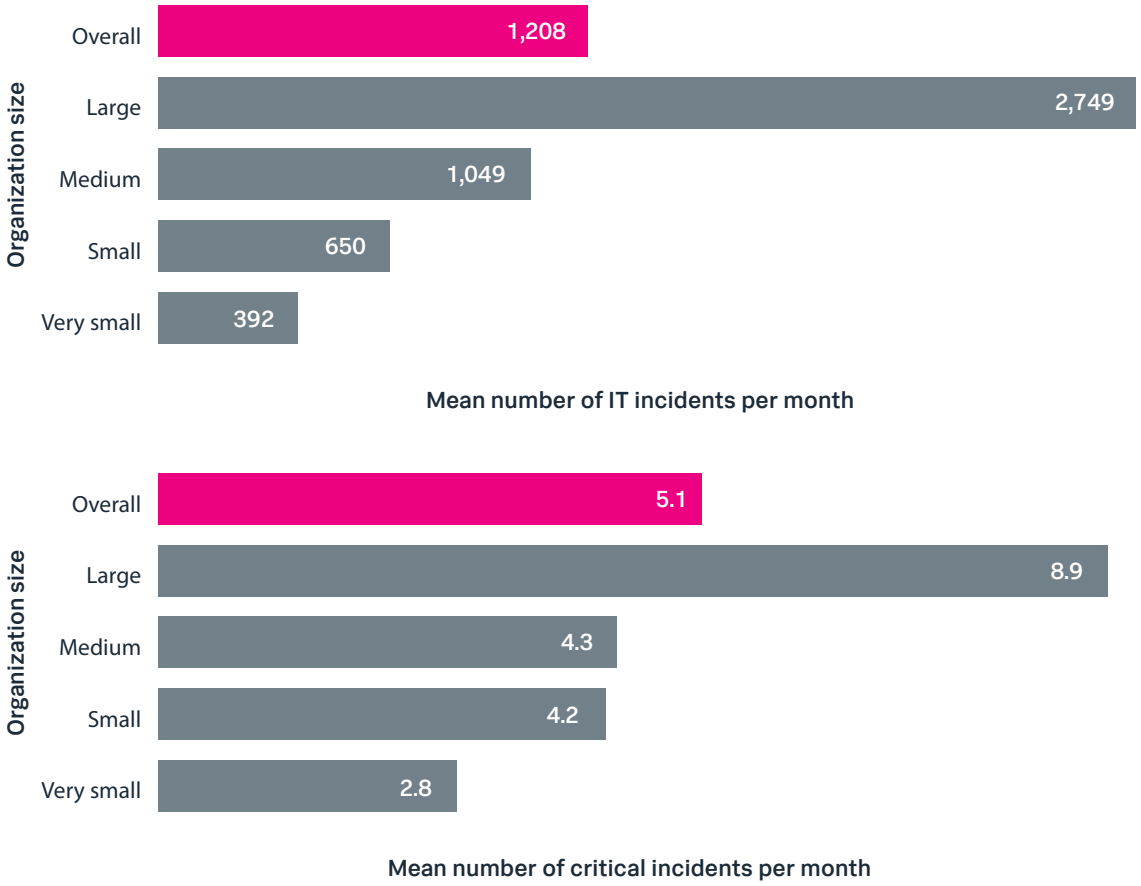

Figure 2: Mean number of IT incidents and critical incidents logged per month

The Costs of Critical Incidents

The mean cost to IT of a critical incident is US $36,326. On top of this, the mean cost to the business is US $105,302 (Figure 3). These two costs rise together, suggesting high cost to IT is a proxy for poor incident management which has a knock-on effect for the business operations.

Those with the slowest MTTD, the time it takes to identify an incident from the point a failure first occurs, have the highest cost to IT of critical

Mean cost to IT of a critical incident is US $36,326

Mean cost to business of a critical incident is US $105,302

Total mean cost of a critical incident is $141,628

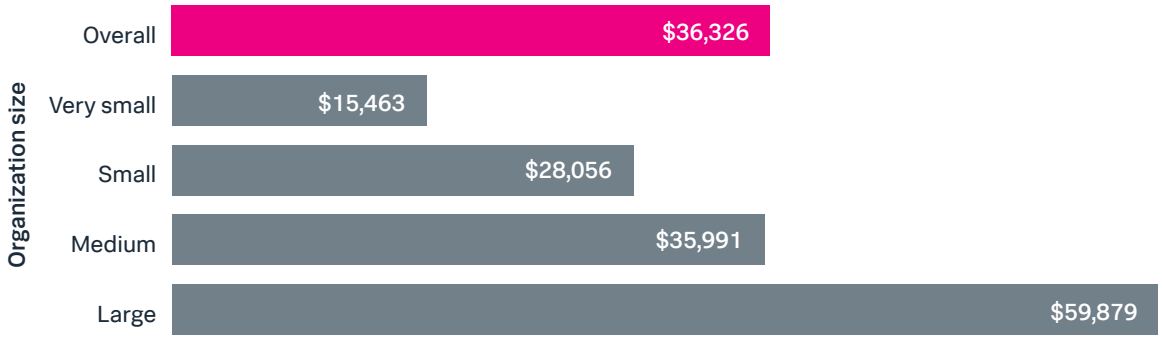

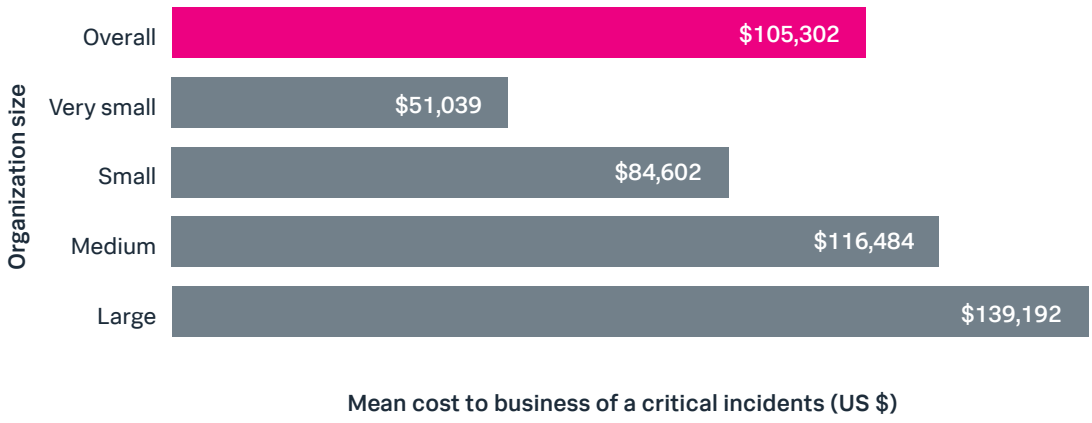

Figure 3: Mean cost to IT and business of a critical incident

incidents (Figure 4). This shows how a problem early in the event management process can escalate through to incident management.

Incident volumes drives up IT costs: those dealing with the most incidents incur a cost to IT more than four times that incurred by those dealing with the least. This is reflected by the fact that more duplicate and repeat incidents also correlate with higher costs to IT (Figure 5). This again reflects poor event and incident management, and a failure to learn from effective root cause analysis.

Anything that can be done to reduce these costs should be welcome. The research presented in this report shows that earlier detection of incidents and improvements to the efficiency of the incident investigation process reduce the impact of IT incidents. It will also show that effective root cause analysis avoids the recurrence of incidents. All these factors reduce the costs of incidents to IT and the even higher consequent costs and damage to business operations.

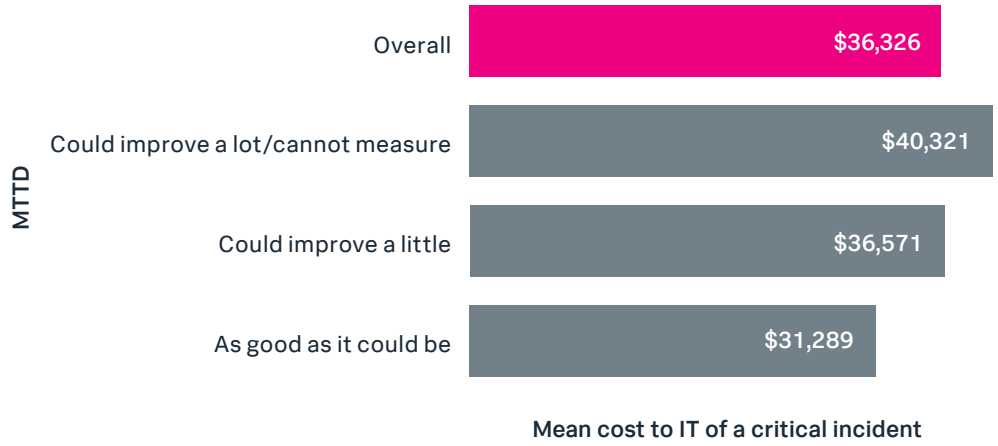

Figure 4: MTTD and cost to IT of critical incidents

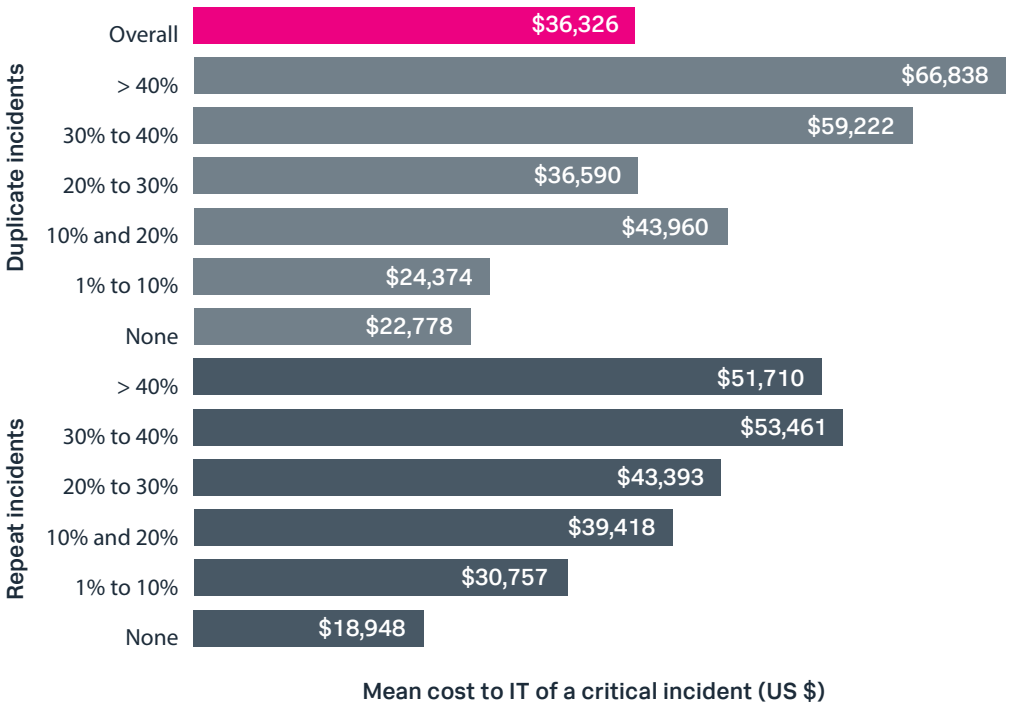

Figure 5: Duplicate and repeat incidents and cost to IT of critical incidents

Dealing With Incidents

Incidents can only be dealt with once they have been detected, be it via automated event management or human reporting, by IT practitioners, end users or customers. Eighty percent say they could improve their MTTD, leaving 20% that believe MTTD is as good as it could be. However, that number more than doubles for those that say they can easily manage the volume of events generated by their IT monitoring tools (Figure 6).

IT teams must ensure incidents are prioritized according to the impact they are likely to have on business processes. This requires equipping IT operators with the necessary visibility and analytics across all the components of IT infrastructure.

After detection, the process of fixing problems and incident closure can be addressed. MTTR is the mean time from the moment an incident is detected

80% say they could improve their MTTD, leaving 20% that believe MTTD is as good as it could be

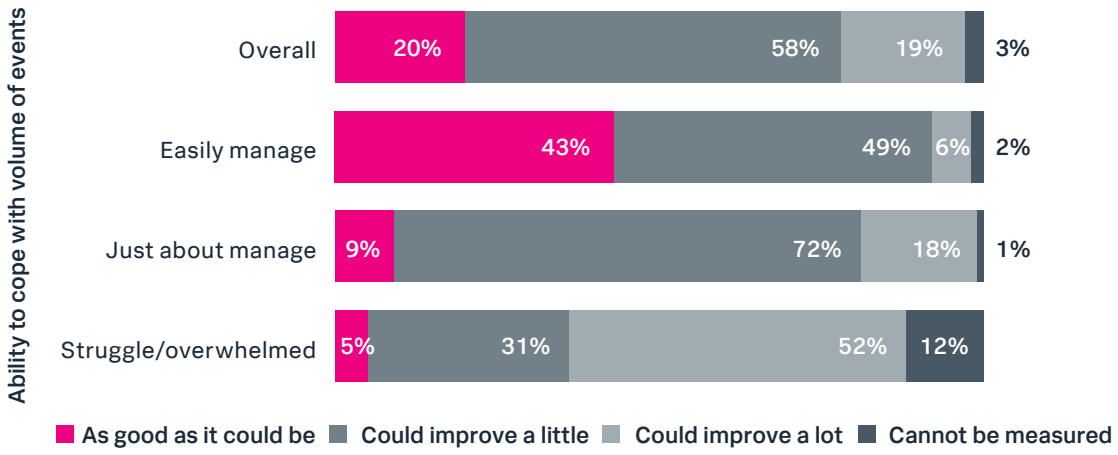

Figure 6: Coping with volume of events and mean time to detect

to the point that the problems are repaired, and the case is closed. MTTR for critical incidents is a key metric for many organizations. The overall MTTR for critical incidents is 5.81 hours. However, the noise created by the volume of incidents generated, including unnecessary duplicates and repeats, slows MTTR (Figure 7).

The overall MTTR for critical incidents is **5.81** hours

The IT teams most likely to be involved in fixing critical incidents tend to be from traditional areas of IT management (Figure 8). IT security is high on the

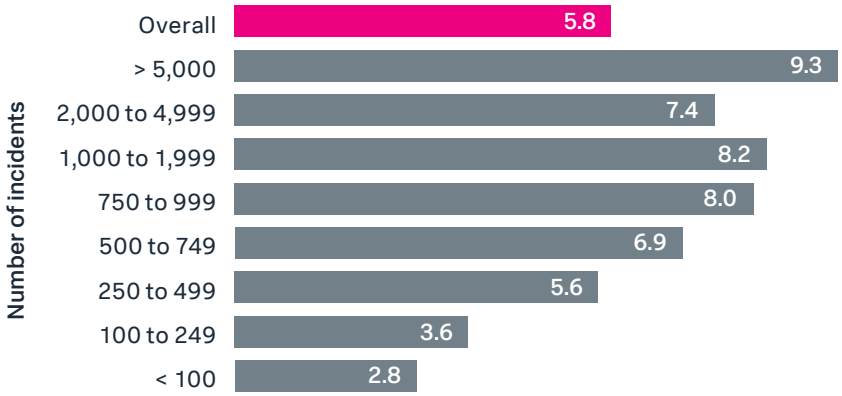

Figure 7: Impact of volume of incidents on mean time to repair (MTTR) for critical incidents

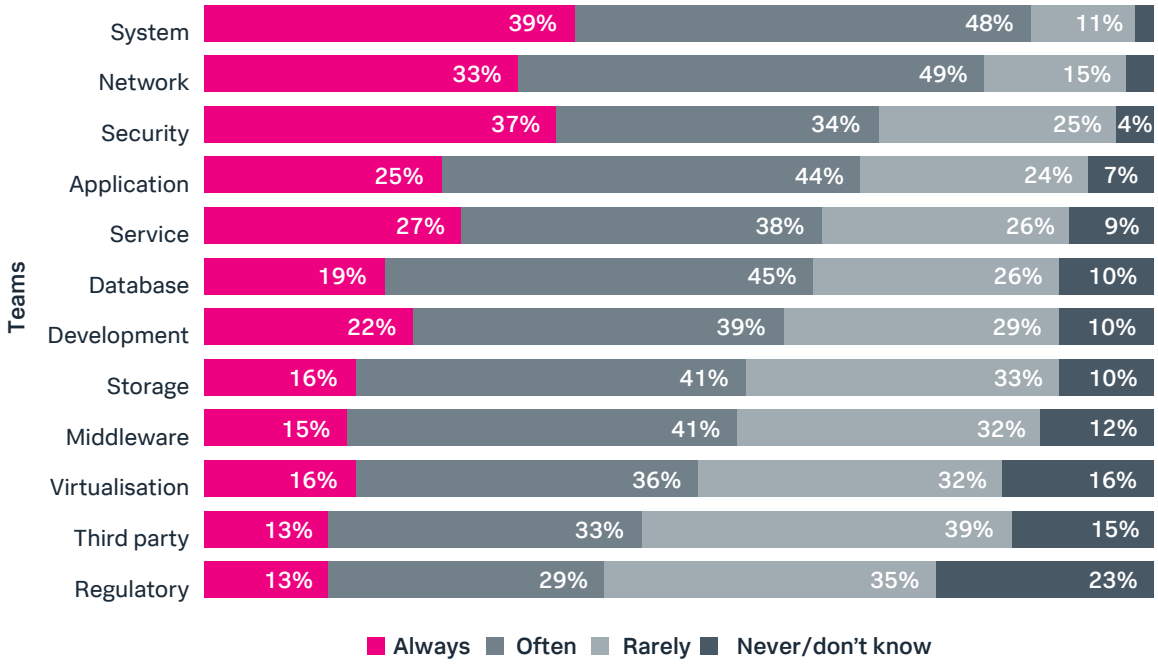

Figure 8: Involvement of engineers from IT teams in resolving critical incidents

list, which is not surprising, since security is the top IT management concern identified in the current survey (Figure 9). An incident is more likely to be critical if there is a security problem (such as a denial-of-service attack or ransomware infection) or a possible data breach.

The typical team size that comes together to fix a critical incident is reported as six (Figure 10). This varies little by business size as the range of skills required for any given incident will be similar. It is hard to reduce the number of subject matter experts needed in the teams that come together to investigate incidents. Improving IT infrastructure visibility can help as IT staff can be given access to data from areas of infrastructure beyond their own remit.

An incident is more likely to be critical if there is a security problem

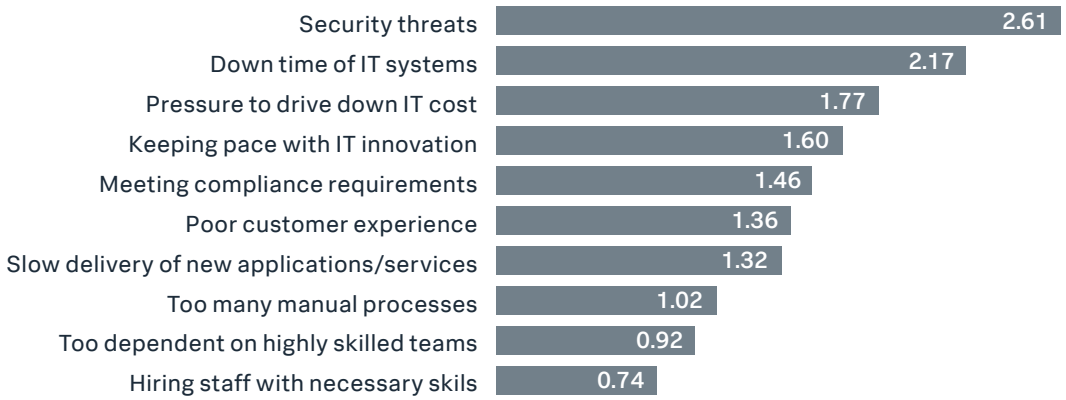

Figure 9: General IT management concerns

Effective Root Cause Analysis Reduces the Cost of Critical Incidents

The quicker an incident is recognized and fixed, the sooner investigations can start into what caused it in the first place. This is root cause analysis (RCA), which is essential to improve the ongoing management of IT infrastructure and avoid repeat incidents. The mean time spent on RCA is 7.23 hours and RCA is determined 65% of the time. The average team size reported for RCA is also about 6 (Figure 10).

Those that say they nearly always determine RCA have about one third the number of repeat incidents, compared to those which only determine RCA 50% of the time. Avoiding repeats reduces the overall cost of IT incidents.

Mean time spent on RCA is 7.23 hours

RCA is determined 65% of the time

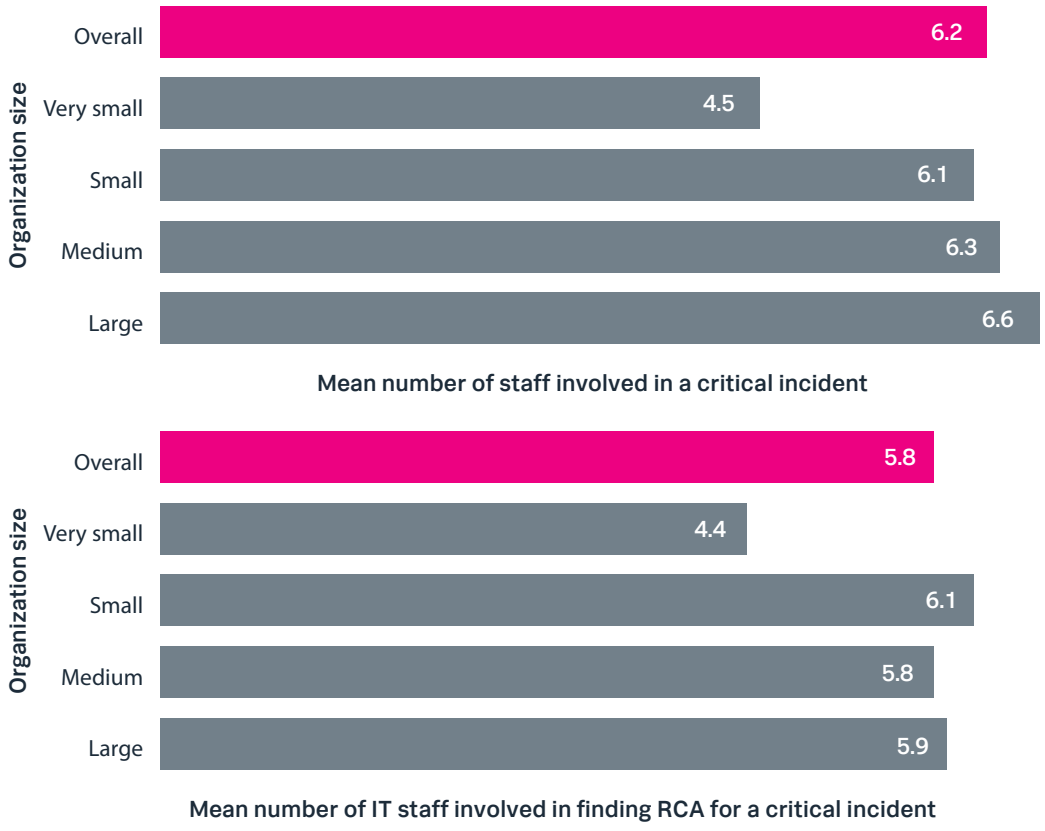

Figure 10: IT staff involvement in resolving critical incidents

Coping With the Volume of Events and Event Management

Two thirds admit that dealing with the volume of events generated by their IT monitoring tools as part of their event management process is a problem; 20% have no event management process at all. Sixty-six percent find dealing with the volume of events a challenge: 52% just about manage, 13% struggle and 1% are overwhelmed.

66% find dealing with the volume of events a challenge: 52% just about manage, 13% struggle and 1% are overwhelmed

Those that keep on top of the volume of events have fewer duplicate and repeat incidents (Figure 11). Much of this is due to the capability of event management processes and the tools that support them; for example, the ability to triage an event effectively through good visibility of IT infrastructure. Poor event management leads to false alarms, putting pressure on the downstream incident management process, delaying problems being detected and fixed.

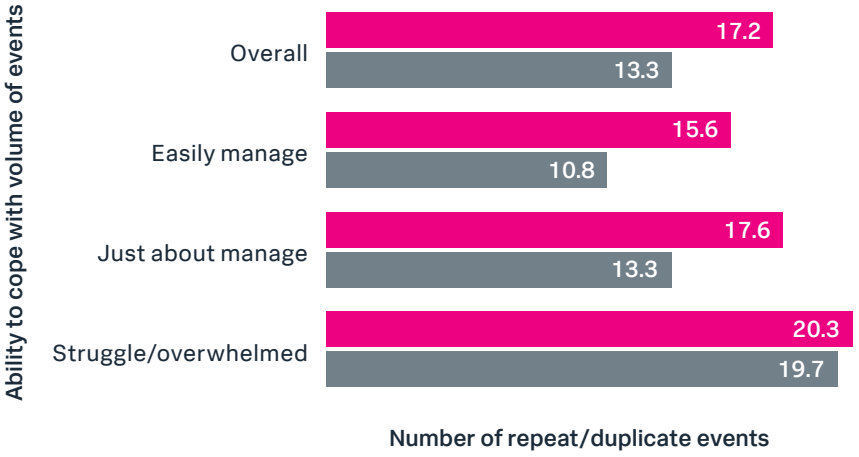

Figure 11: Ability to cope with volume of events and reducing duplicate and repeat incidents

Unnecessary Incidents

Poor end-to-end visibility and IT infrastructure complexity leads to duplicated incidents: the same problem being logged multiple times by a service desk and in the worst case separate IT teams addressing the same incident in parallel. Ninety-seven percent experience duplicate incidents; nearly a fifth (17.2%) of all incidents are duplicates.

97% experience duplicate incidents; 17.2% of all incidents are duplicates

Once the short-term crisis that an incident may herald has been dealt with, ineffective root cause analysis can mean the recurrence of the same incidents later due to failure to learn from past problems. Ninety-six percent say failure to learn from previous incidents leads to repeat incidents; 13.3% of all incidents are repeats.

96% say failure to learn from previous incidents leads to repeat incidents; 13.3% of all incidents are repeats

The scale of duplicate and repeat incidents should not be underestimated. The average organization can expect 374 unnecessary (duplicate) or avoidable (repeat) incidents a month; for a large organization this rises to 852. The causes of this wasted IT effort, and the subsequent costs it incurs, lie in problems with event and incident management processes.

This is recognized by IT management teams: of the impact various issues may have on the efficiency of IT operations, duplicate and repeat incidents top the list (Figure 12). Fix the problem and around one third of logged IT incidents (and the time spent investigating them) could be eliminated.

Figure 12: The impact issues have on the efficiency of IT operations

Visibility into IT Infrastructure

Having tools that can provide good visibility of IT infrastructure leads to faster detection and investigation of incidents and improves root cause analysis. However, coverage varies: holistic end-to-end visibility of IT processes was rated lowest of all, with just 45% saying they had excellent or good visibility (Figure 13).

Visibility tends to be better in traditional areas such as servers, storage and networking, and worse with next-generation technologies such as cloud computing, containers and mobility. This suggests that monitoring capabilities have not kept up with the pace of change required to meet business demands. Poor visibility into the holistic delivery of IT and the end-user experience makes it harder to address the top concern when it comes to the potential consequences of IT incidents – the customer and user experience.

Only 2.5% have full visibility across all relevant infrastructure, whilst 0.7% said they have no visibility of anything

The ratings given in Figure 13 can be used to calculate an overall visibility score (see appendix for calculation). The average visibility out of a maximum of 4 is 2.56, so there is plenty of room for improvement. Only 2.5%

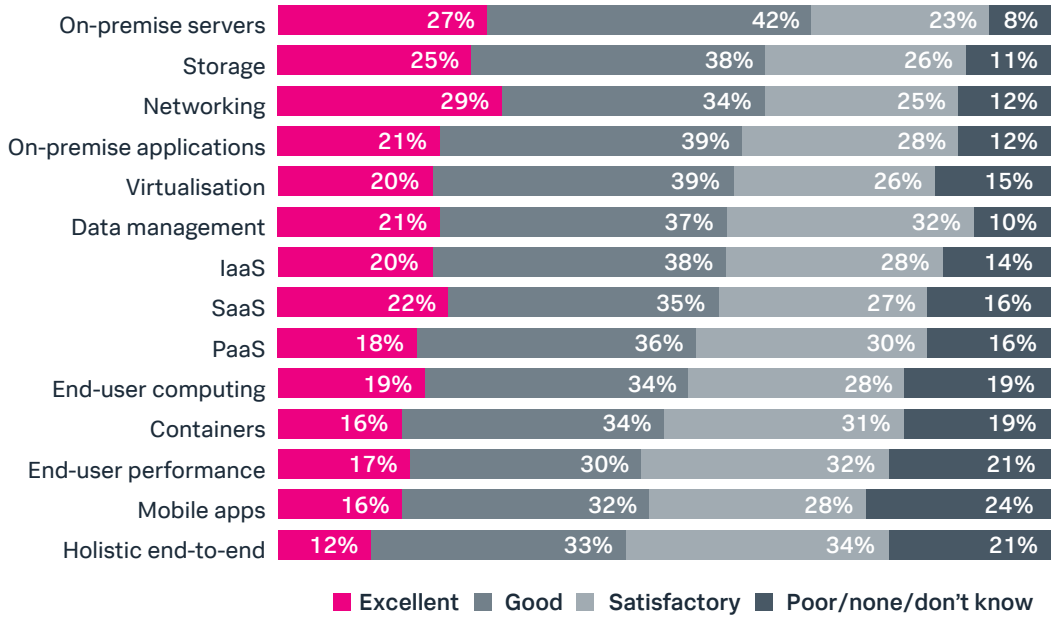

Figure 13: Visibility into IT infrastructure

have full visibility across all relevant infrastructure, while 0.7% said they have no visibility of anything.

The average organization has 19.8 monitoring tools, a large organization may have 100s. Better visibility through effective monitoring should mean fewer blind spots, however, 80% admit they have blind spots.

Better visibility correlates with reduced duplicate and repeat incidents. This is mostly likely through improvement of the overall incident detection process (Figure 14). The number of repeats can also be expected to decrease as IT staff will have more time to look at the root cause of incidents, learn from mistakes, and prevent recurrence.

The average organization has 19.8 monitoring tools, a large organization may have 100s

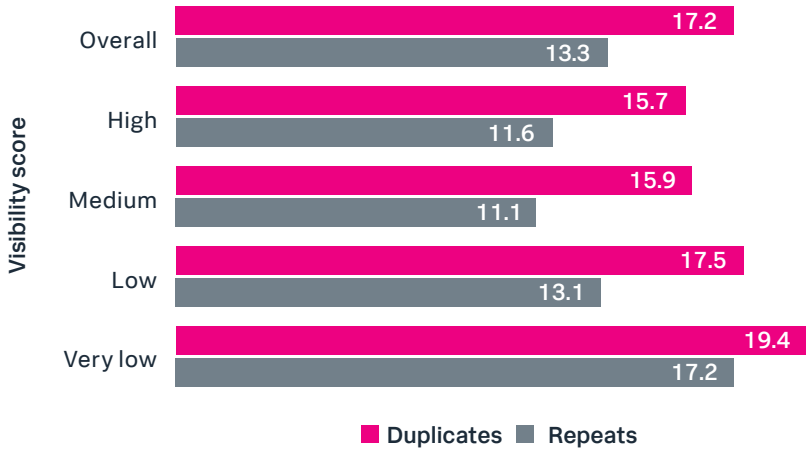

Figure 14: Visibility and repeat duplicate incidents

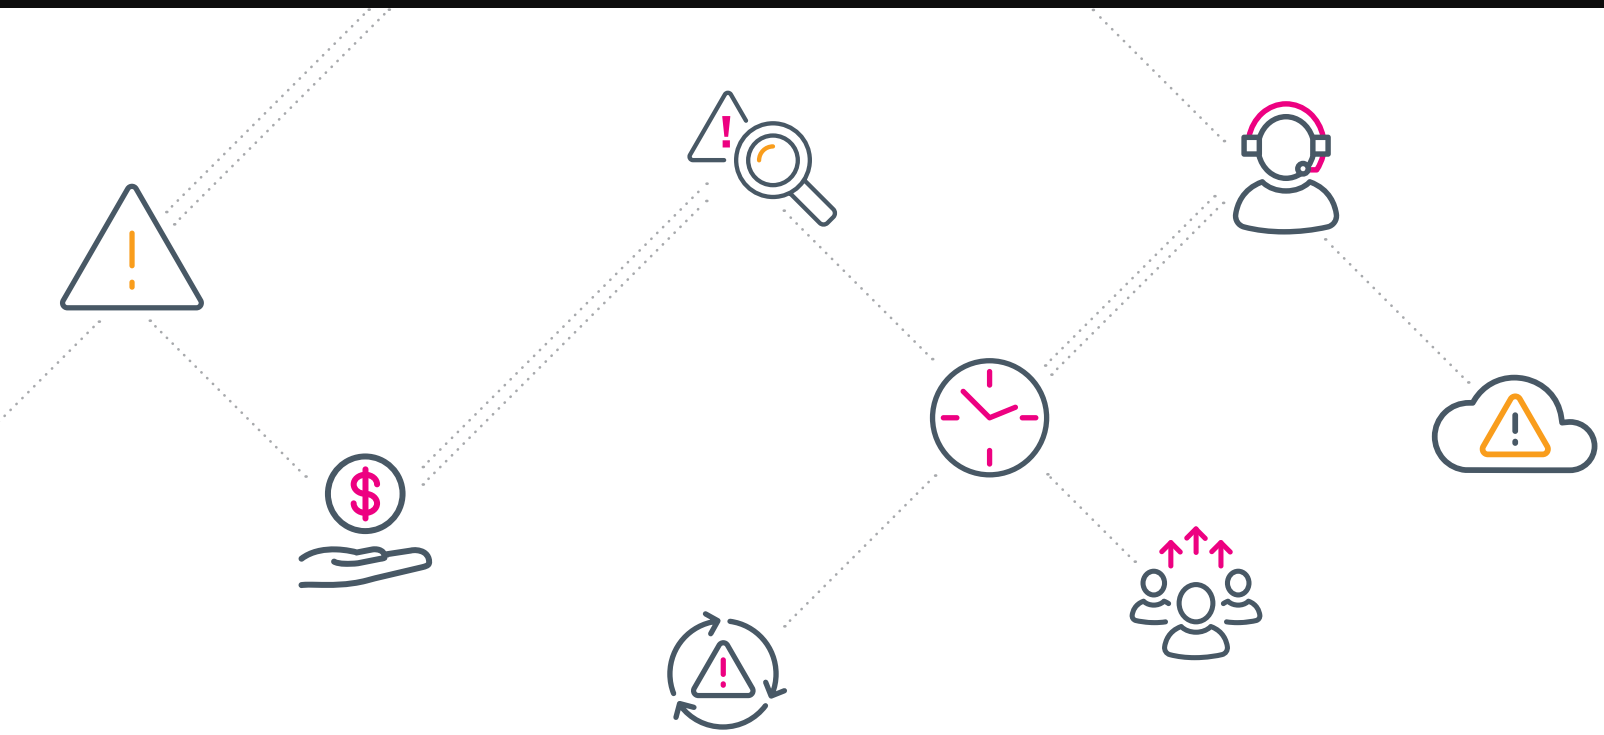

Conclusion

IT departments are under pressure to accelerate the development of high-quality digital services for users and customers to remain competitive. This must be achieved while keeping IT costs under control.

Increased IT complexity leaves IT staff struggling to effectively manage incidents, resulting in higher than necessary costs to both IT departments and to the businesses they serve. In a world increasingly reliant on digitally driven processes, if the process for managing IT incidents is not improved, more users and customers will be impacted, which will lead to higher risk of reputational damage.

Dealing with critical IT incidents needs to be a top IT priority. IT teams should be equipped with tools to provide end-to-end visibility of processes and the IT infrastructure that supports them. The tools are needed to enable rapid detection and investigation of IT incidents, streamline root cause analysis, and reduce the size of teams involved in fixing problems. Achieving all of this reduces IT costs and the much higher consequential cost and impact of IT incidents on businesses.

Click HERE to calculate the cost of critical IT incidents to your organization.

Appendix

Calculating the scores in Figures 1, 9 and 12

With certain questions respondents were asked to rank the top 3 or 5 out of a range of issues, putting one first, another second, and so on until 3 (or 5) had been selected. In the analysis a value was assigned to each choice as follows:

Five issues (Figures 1 and 9)

- Score = 5 for top issue
- Score = 4 for second issue
- Score = 3 for third issue
- Score = 2 for fourth issue
- Score = 1 for fifth issue
- Score = 0 for each issue not selected

Three issues (Figure 12)

- Score = 3 for top issue
- Score = 2 for second issue
- Score = 1 for third issue
- Score = 0 for each issue not selected

These scores were used to calculate the weighted average concerns about the various issues used throughout the report. As each respondent is forced to select five (or three) issues in order, if one issue rises, another must fall.

Visibility score calculated from data in Figure 13

The overall infrastructure visibility score derived from the data in Figure 13 was calculated by awarding 4 for excellent, 3 for good, 2 for satisfactory, 1 for poor, and 0 for none/don't know and then calculating an overall average score.

Exchange rates

The relevant financial data is reported in US $. The questions were all asked using local currency. The following exchange rates were used: US $1.00 =

- Singapore $1.36
- Japanese Yen 112
- Australia $1.28
- Swedish Krona 8.05
- Euro €0.85 Euro
- UK £0.76

Demographics

All respondents to the survey were senior IT managers. The countries and business sectors covered are shown below, with the breakdown by organization size. The fieldwork was carried out by Quocirca's research partner, Vanson Bourne.

Demographics

All respondents to the survey were senior IT managers. The countries and business sectors covered are shown below, with the breakdown by organization size. The fieldwork was carried out by Quocirca's research partner, Vanson Bourne.

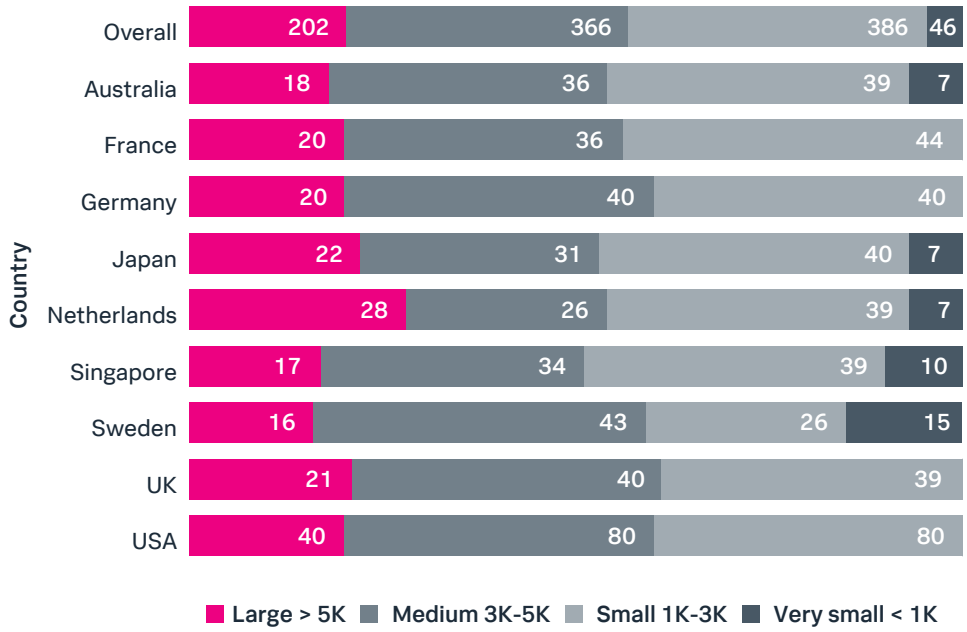

Figure 15: Countries by business size

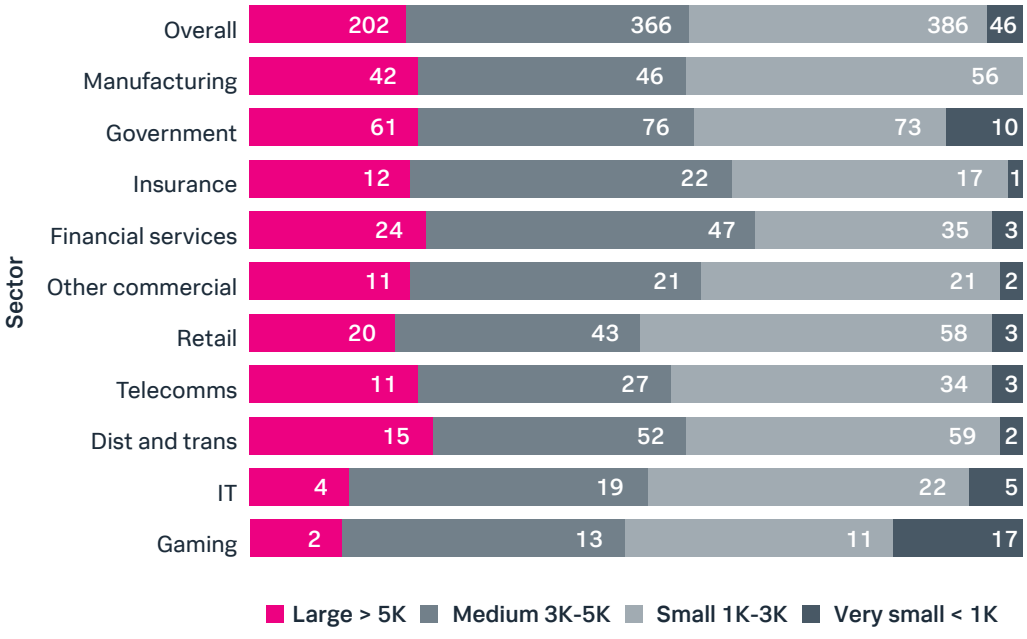

Figure 16: Business sectors by business size

splunk>

QUOCIRCA

About Splunk

Splunk Inc. (NASDAQ: SPLK) turns data into action. Organizations use market-leading Splunk solutions with machine learning to solve their toughest IT, Internet of Things, and security challenges. Join millions of passionate users and discover your “aha” moment with Splunk today: **http://www.splunk.com**.

About Quocirca

Quocirca is a UK-based research and analysis company. Quocirca produces free-to-market content aimed at IT decision makers and those that influence. Much of the content Quocirca produces is based on primary research across Europe, the Americas and Asia, sponsored by a broad spectrum of IT industry organizations. Quocirca content is written from an independent standpoint and addresses the use of IT within the context of an organization, rather than specific products. Through its close relationships with the media, Quocirca articles and reports reach millions of influencers and decision makers: **www.quocirca.com**.

2020-Splunk-IT-Quocirca-Damage-Control-109-WP