Uptime Institute | RESEARCH

RISK & RESILIENCY

UI Intelligence report 23

Annual outage analysis

The causes and impacts of publicly recorded IT service and data center outages from 2016-2018

Lead Analysts
Andy Lawrence (alawrence@uptimeinstitute.com)
Rhonda Ascierto (rascierto@uptimeinstitute.com)

UII-23 v1.3 published January 22, 2019 last updated October 3, 2019

This Uptime Institute Intelligence report covers:

Publicly reported IT outages 2016-18

Introduction

Critical systems and data centers are immeasurably more reliable than they were two or three decades ago. In most cases, problems are identified and resolved before users and customers notice. Not only has equipment become far more reliable over time, but management processes are now in place to anticipate failures or limit the consequences. In an era of cloud, distributed architectures, traffic management, and low-cost replication, IT can re-route around many failures, in some cases automatically.

Despite all this, Uptime Institute Intelligence finds that major failures are not only still common, but that the consequences are high, and possibly higher than in the past—a result of our reliance on IT systems in all aspects of life. In 2018, there were major outages of financial systems, day-long outages of 911 emergency service call numbers, aircraft losing services from ground-based IT landing systems, and health-care systems lost during critical hours.

Surprisingly, detailed and reliable research on the causes and impacts of IT service failure is difficult to find. Most studies have small samples and have often concentrated on areas of certain interest to the sponsors, such as security, denial of service, and data center power. But it is the view of the Uptime Institute that many failures have multiple causes or can cascade between data centers and networks, triggering secondary failures, and that a more holistic or service-based approach is needed.

Furthermore, the increasing use of outsourced services and public cloud has led to a loss of visibility: Cloud and service providers can sometimes be disarmingly open in discussing their failures. More commonly they provide little or no commentary, and sometimes they do not admit to outages at all. The lack of transparency has led some people in the data center industry to call for mandatory reporting, as is the case for security breaches in some countries.

However, the likelihood of this happening may be considered remote. The industry will most likely have to accept a certain lack of transparency and openness, making it difficult to learn lessons and, in some cases, to hold operators accountable.

KEY FINDINGS

The findings discussed in this report help to extend the industry's understanding of the type, range, causes, and extent of outages, providing some lessons. However, as our methodology (section below) states clearly, these findings are only part of the total picture and should be treated with caution, as they do not represent a statistically defensible study.

The three main findings in this report are:

- Major and damaging outages continue to trouble the IT industry, despite improvements in technology and management. There is clear evidence that availability does not match marketing claims (service level promises).
- Major publicly recorded outages are now more likely to be caused by IT and network problems than a year or two years ago, when power problems were a bigger cause.
- Public cloud-based services account for a significant number of reported service outages, with causes ranging from power to wide-area synchronization issues. Although the reliability/availability of these services is generally good, their scale and complexity means that outages are likely to have a clear and well-recorded impact.

Methodology

Uptime Institute currently has three sources of data on data center and IT outages or incidents that can potentially lead to outages. These are:

- The Abnormal Incident Report (AIRs) database. This is a confidential system for Uptime Institute Network members to report incidents in detail under a non-disclosure agreement (NDA).
- Uptime Institute's Annual Industry Survey. This global survey, with some 1,300 respondents in 2018, asks detailed questions about outages, and some of the findings are discussed here for context. This represents the most statistically significant data set relating to outages in the data center industry.
- Uptime Institute Intelligence public outages database. Since the beginning of 2016, Uptime Institute has collected data about major IT outages from public media reports and sources on an ongoing basis. This enables us to collect information on major outages that became visible to the public and the media, and, over time, to identify patterns. This data is discussed in this report.

The methodology used in this report is limited: if a failure is not reported or not picked up by media or Uptime Institute, it will not be recorded. This immediately means there is a bias toward coverage of large, public-facing services, in geographies with a well-developed and open media. We also eliminate many small, short failures where the business or reputational impact is clearly negligible. Secondly, the amount of information available varies widely from outage to outage, and sometimes there is very little information available at all. Third, we limit failures to those that had a noticeable impact—a major fire during data center commissioning, for example, may never be registered. Finally, while we include IT system failures, we do not generally include security breaches; this is partly because the impact does not necessarily affect availability. We do, however, include DDOS (distributed denial of service) attacks, since these are related to data center capacity resiliency and do not involve a security breach.

Note: Uptime Institute is also a supporter of the nascent Data Center Incident Reporting Network. At present, this network has not compiled sufficient data to report.

Outages by number

How common are outages and are they increasing? During 2018, Uptime Institute published the results of its Annual Industry Survey on the prevalence of outages (global). We found that almost one third (30.8%) of the IT service and data center operators surveyed had experienced an IT downtime incident or “severe degradation of service” in the past year (see Figure 1).

Figure 1. **Almost one third of companies surveyed suffered a recent outage**

Has your organization experienced an IT service outage or severe service degradation in the last year?

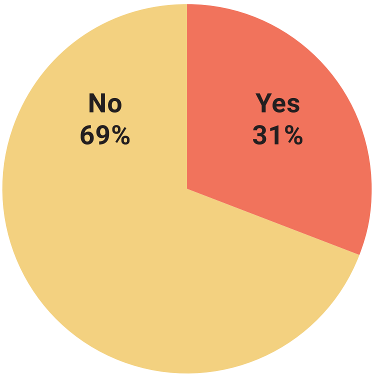

Source: Uptime Institute Global Survey of Data Center Operators and Managers, 2018, n=664

We also asked: “Have you experienced an outage in the past three years?” The level of outages/incidents was higher than we might have expected—48.1% said yes. Managers who had responsibility for end-to-end IT service delivery, rather than just data centers, reported a higher number—as much as 58% over three years – which suggests they have greater visibility into outages than those only responsible for data centers.

These findings do not definitively support a conclusion that outages are rising—although we did see a year-on-year increase—but it does show that outages are common, that most data center or IT managers can expect to suffer from a fairly major outage every few years at least, and, critically, that the 99.99% or even 99.999% availability claims made by the industry are out of line with the reality.

When we look only at outages that were publicly recorded in the media, there was a clear rising trend, with nearly three times as many recorded in 2018 as in 2016 (see Figure 2): However, we are not arguing that this represent more outages—necessarily—but it does represent increased visibility, increased reporting by the media, and improved data collection. Even so, it is safe to say that outages are continuing to be a major challenge for IT services operators and their customers, in spite of technology/management improvements.

Figure 2. **Public outages: 2017 - 2018**

Number of Outages: **162 Total**

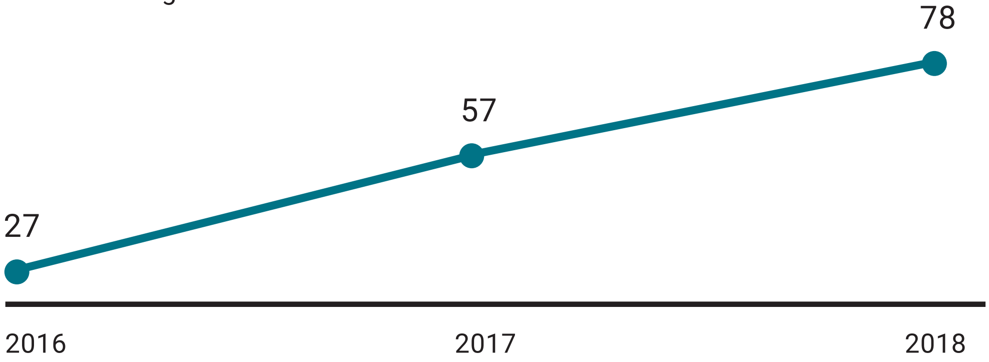

Source: Various online news sources. Uptime Institute Intelligence, January 2019

Severity and impact of outages

Although there are various ways to categorize the “mission criticality” of various systems as a planning tool for disaster recovery and availability/redundancy investments, there is no “Richter Scale” for measuring the severity/impact of outages.

Such a rating would clearly be useful. For example, the effect of losing access to a human resources system for 2 weeks might be frustrating, but negligible, even in large organizations, while a 5-minute loss of a currency trading system could be near catastrophic.

For classifying the impact of public outages, Uptime Institute has created a five-level Outage Severity Rating (see Figure 3):

In 2018, most publicly reported outages fell in the low to middle end of the scale (see Figure 4). But looking back over the past three years,

Figure 3. **Uptime Institute Outage Severity Rating**

The Uptime Institute Outage Severity Rating (OSR) describes the business/service of a major IT service outage, regardless of the cause.

Outage Severity Rating	Description	Impact of Outage
Category 1	Negligible	Recordable outage, but little or no obvious impact on services.
Category 2	Minimal	Services disrupted. Minimal effect on users/customers/reputation.
Category 3	Significant	Customer/user service interruptions, mostly of limited scope, duration or effect. Minimal or no financial effect. Some reputational or compliance impact(s).
Category 4	Serious	Disruption of service and/or operation. Ramifications include some financial losses, compliance breaches, reputation damages, possibly safety concerns. Customer losses possible.
Category 5	Severe	Major and damaging disruption of services and/or operations with ramifications including large financial losses, possible safety issues, compliance breaches, customer losses, reputational damage.

Uptime Institute Intelligence, January 2019

The Outage Severity Rating was developed by Uptime Institute © 2019, All Rights Reserved

there are some significant changes underway: the proportion of Level 5 outages (severe, business-critical outages) is falling, while the number of less-serious recorded outages grew. Uptime Institute has two explanations for this:

- the reporting of outages, in social media and mainstream media, is increasing as more people are affected (due to higher adoption of public cloud, SaaS, and managed hosted services), and it is easier to spread the news, even about smaller outages.

- IT-based outages, which are now more common than full data center outages, are more likely to be partial and, while certainly disruptive, can often have less impact than a complete data center outage, which may affect all applications and create cascading effects.

Figure 4. **Severity of publicly reported outages**

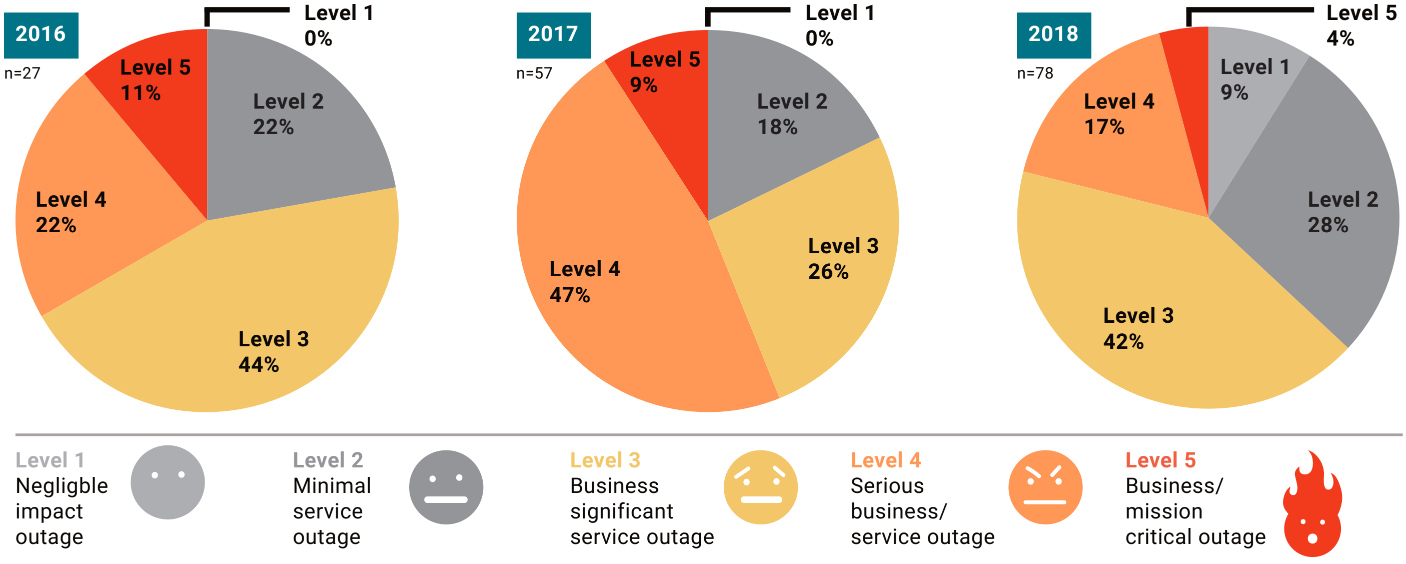

(% of each year's total) Uptime Institute Intelligence, January 2019 n=162

During the three years, there were 57 outages in total that were classified as either Level 4 or Level 5:

- In 2016, there were three Level 5, severely disruptive business-critical outages (caused in two cases by power outages, and in one by a network equipment failure)
- In 2017 there were five Level 5 outages (two caused by IT systems, one by a network issue, one by a power failure, and one by cooling/mechanical systems failure)
- In 2018, three Level 5 outages (one by a network failure and two by IT systems)

The difference between a serious Level 4 and a severe Level 5 outage may be minimal and in some cases debatable—both are costly. When we take all of the 57 serious/severe outages together, the three biggest primary causes were roughly equal: power, network, and IT system (see Figure 5).

Figure 5. **Level 4 and 5 outages in 2018: Primary causes**

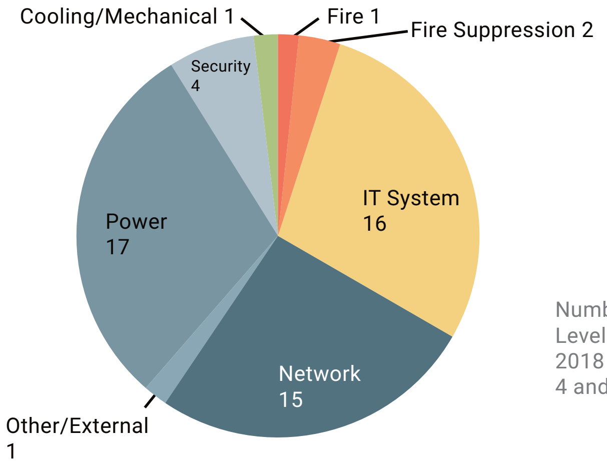

Number of Level 4 and Level 5 outages 2016 to 2018 (total number of Level 4 and 5 outages = 57)

Source: Various online news sources, Uptime Institute Intelligence, January 2019

The cost of outages

Uptime Institute has always been cautious about estimating the costs of outages because they vary widely from company to company and from industry to industry. In five of the 11 severe business critical outages (Level 5) from 2016 to 2018, airlines were involved—demonstrating their very high dependence on IT and the way in which IT disruptions can halt bookings, check-ins, and flights. During the past three years, Delta Air Lines, United Airlines, and British Airways have taken multimillion-dollar hits from outages, as has UK-based TSB Bank.

Colocation and cloud providers, however, are not only less vulnerable to severe failures but also have commercial agreements with customers to protect themselves against losses. Even so, in several cases, colocation and hosting companies have lost business and suffered heavy reputation damage as a result of failure.

The high cost of outages was demonstrated in the 2018 Uptime Institute Annual Survey. Half of the outage incidents reported by survey respondents cost under $100,000, but there were 39 outages that cost more than $1 million (15% of the total reported). Around one third of outages reported by respondents cost over $250,000. These figures should help persuade CIOs to maintain spending on availability at all levels.

(A surprising finding in Uptime Institute's 2018 Annual Survey was that, in cases where a significant incident did occur, 43% of respondents did not actually calculate the cost at all.)

A further measure of impact is the time that users/customers were affected. Just over half of the publicly reported outages in our 2018 sample lasted over 4 hours—higher than in earlier years (see Figure 6). This may align with the data, which suggests that IT/network related disruptions can take longer to fix than power/data center facilities disruptions.

Figure 6. **Length of user disruption for public outages in 2018**

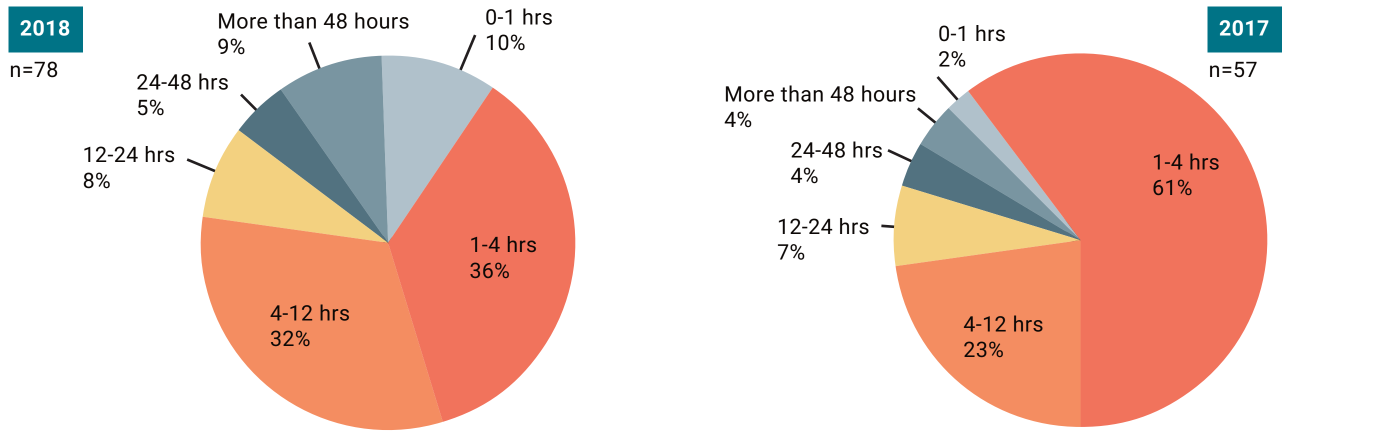

Source: Various online news sources, Uptime Institute Intelligence, January 2019

Causes of outages

Engineers often say that there is rarely a single of cause of a major accident or failure. Problems very commonly cascade, and very often there are problems with personnel training, operational process, and management that compound an initial failure. In this report, we focus on the primary cause, although most incidents have a complex back story. The results over three years show that, marginally, IT system failure is now the most likely cause of outages, followed by networks and then power (see Figure 7).

Figure 7. **Primary cause of public outages is changing**

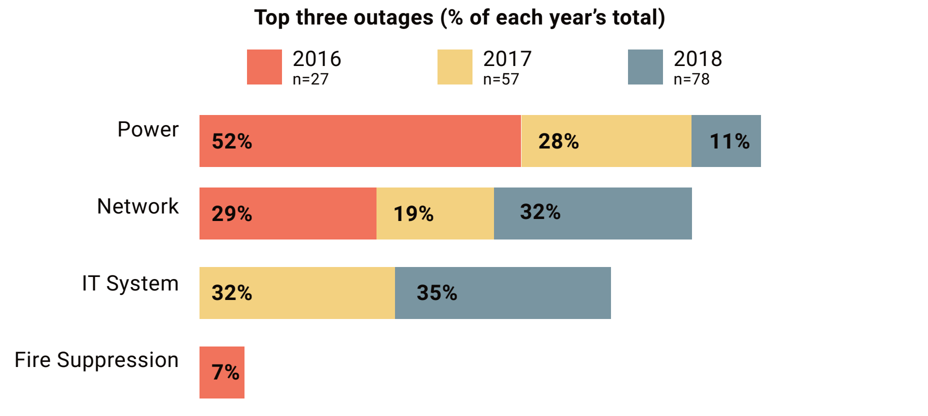

n=162
(% of each year's total)

Source: Various online news sources, Uptime Institute Intelligence, January 2019

However, taken as a whole, facilities issues (power, cooling, fire, and fire suppression) are still the biggest cause (at 32%) of outages (see Figure 8). In addition, many of the failures classed as IT and network were actually caused by a power-related problem at a single component, system or rack level—which we have not generally classed as a data center power problem.

Figure 8. **Why? Primary causes: 2016-2018**

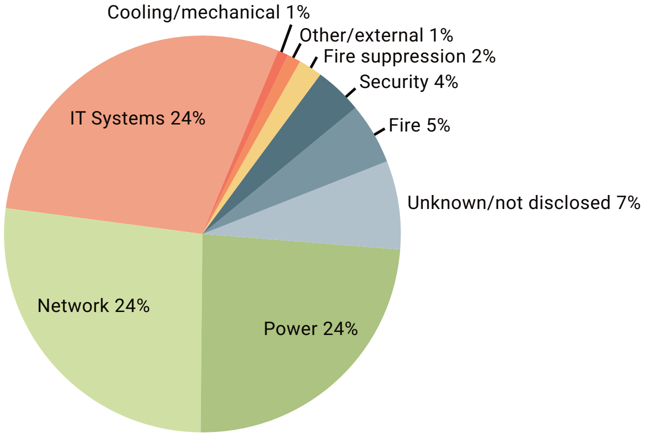

n=162
Source: Various online news sources, Uptime Institute Intelligence, January 2019

What accounts for most of these failures? Based on the sometimes-limited information made available about these outages (and frequently analyzed at Uptime Institute Network meetings), we are able to make the following observations:

IT systems: IT systems have long been designed to tolerate failure—at the component, system, and, most recently, the data center level. Technologies such as disk mirroring, RAID, clustering, real-time synchronous replication, and, more recently, wide area availability zones, are all effective. However, even so-called failsafe systems do fail.

Among the most disruptive causes of IT systems failures that resulted in a service interruption during the past three years:

- A poorly managed upgrade, on at least one major occasion, with insufficient testing at the software level.
- The failure and subsequent data corruption of large disk drives/storage area networks. This was likely caused by a hardware failure, exacerbated by configuration/programming errors.
- Failure of synchronization or programming errors across load balancing or traffic management systems.
- Incorrectly programmed failure/synchronization or disaster-recovery systems.
- Loss of power to non-backed-up single components (including servers and large disk drives).

Power outages. Power sits at the base of the "dependency pyramid," and so failures can have major implications. A power failure frequently reveals a failure in configuration of IT or in management software. Among the causes—all familiar to most data center managers and certainly to Uptime Institute Network members:

- Lightning strikes, leading to surges and lost power. Back-up software/configuration failed.
- Intermittent failures with transfer switches, leading to failure to start generators or transfer to a second data center.
- Uninterruptible power supply (UPS) failures and failure to transfer to a secondary system.
- Operator errors, turning off/misconfiguring power.
- Utility power loss and subsequent failure of generator or UPS.
- Damage to IT equipment caused by power surges.
- IT equipment not equipped with dual power suppliers switched to secondary feed.

Network. Over the past decade, the network has become as important as power to the data center. Without network connectivity, few applications can function. The move to cloud has made this yet more critical. It is no surprise, therefore, that network issues now account for the second highest number of outages in 2018.

Among the issues:

- Fiber cuts outside the data center, with insufficient routing alternatives (this is common).
- Intermittent failure of major switches, with secondary routers not deployed.
- Major switch failure without backup.
- Incorrect configuration of traffic during maintenance.
- Incorrectly configured routers/software-defined networks.
- Loss of power to non-backed-up single components (such as switches and routers).

Fire and fire suppression. Fire in data centers is extremely rare, but, our research revealed eight fires and four failures caused by fire suppression in outages that were made public. The fires were primarily electrical; two of the fire suppression outages were caused by accidental discharge of inert gas—an issue that Uptime Institute has warned about for several years.

Main sectors affected

It is sometimes said that "cloud services" (IaaS, SaaS, etc.) are designed to fail—meaning that there are layers of software that will allow failure and re-route workloads and traffic. In truth, the picture is more complex. Cloud services can be designed to tolerate very low failure rates, but the investment required on the part of the service provider is high, and the configuration on the side of the customer is complex. In addition, cloud providers are still building their infrastructure and honing their offerings—they have gaps in coverage and are grappling with huge scale and complexity. And, of course, no architecture is failsafe.

Furthermore, the cloud/internet giants have more infrastructure than anybody, more customers, and a higher profile. It is not surprising then, that, their failures score so highly (see Figure 9 below). This, we believe, fairly reflects their scale and position in the ecosystem, supporting so many other services. To a lesser extent, this is true of colocation and hosting companies, too. Taken together, cloud, commercial data center and managed service providers accounted for 34% of all outages: if telcoms/networks services are added, this is nearly half.

Two industries come out particular badly: Airlines, for the reasons already discussed (and widely discussed in a 2017 Uptime Institute paper) account for 6% of the failures and financial services (17%).

Even allowing for the financial industry's important and front-line role, it is clear from the known examples that many banks are suffering from a combination of complexity and lack of investment/modernization.

Figure 9. **Outages by industry sector: 2016-2018**

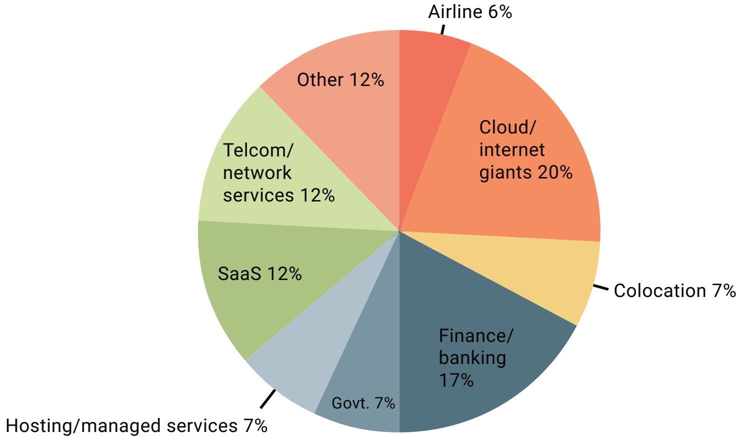

n=162

Source: Various online news sources, Uptime Institute Intelligence, January 2019

Conclusions

As we stated earlier, the findings from our research of publicly recorded outages need to be taken and understood in context: this is as much a record of media coverage of outages as it is of the outages themselves. However, when viewed in combination with other research that Uptime Institute undertakes (see Methodology), and the data about the outages that we have gathered, it is clear that some conclusions can be drawn.

First, outages continue to be a major, and expensive, problem for the IT industry. There is repeated anecdotal evidence, backed by the Uptime Institute's research, that management shortcomings play a major role in these failures. (As an example, only half of our 2018 survey respondents count the cost of a failure after it occurs, let alone model it in advance). When assessing these failures (as Uptime Institute Network members regularly do), it is clear than many could and should have been anticipated.

Second, failures have become complex, and affected services can span across multiple systems and data centers. This calls for an approach to understanding the causes and impacts that is more holistic, and less siloed. IT continues to operate in silos, a strategy that is successful because of the specialties required, but which can cause critical vulnerabilities to be overlooked.

Third, some of the failures of data centers and IT services have been very, very expensive—yet on many occasions, the organization most exposed to the risk is not legally or technically indemnified, nor sometimes even directly involved in designing, operating, or managing the service. For many large organizations exposed in this way, this seems unsatisfactory and there is likely to be more pressure to increase transparency and accountability across their hybrid infrastructures.

Uptime Institute® | RESEARCH

ABOUT THE LEAD ANALYSTS

Andy Lawrence is Uptime Institute's executive director of Research. Mr. Lawrence has built his career focusing on innovative new solutions, emerging technologies, and opportunities found at the intersection of IT and infrastructure.

Rhonda Ascierto is VP of Research at the Uptime Institute. She has spent nearly two decades at the crossroads of IT and business as an analyst, speaker, adviser, and editor covering the technology and competitive forces that shape the global IT industry.

Uptime Institute is an unbiased advisory organization focused on improving the performance, efficiency, and reliability of business critical infrastructure through innovation, collaboration, and independent certifications. Uptime Institute serves all stakeholders responsible for IT service availability through industry leading standards, education, peer-to-peer networking, consulting, and award programs delivered to enterprise organizations and third-party operators, manufacturers, and providers. Uptime Institute is recognized globally for the creation and administration of the Tier Standards and Certifications for Data Center Design, Construction, and Operations, along with its Management & Operations (M&O) Stamp of Approval, FORCSS® methodology, and Efficient IT Stamp of Approval.
Uptime Institute – The Global Data Center Authority®, a division of The 451 Group, has office locations in the U.S., Mexico, Costa Rica, Brazil, U.K., Spain, U.A.E., Russia, Taiwan, Singapore, and Malaysia.
Visit www.uptimeinstitute.com for more information.
All general queries:
Uptime Institute
5470 Shilshole Avenue NW, Suite 500
Seattle, WA 98107
USA
+1 206 783 0510
info@uptimeinstitute.com