Uptime Institute® INTELLIGENCE

RISK & RESILIENCY

UI Intelligence report 35

Annual outage analysis 2020

The causes and impacts of publicly recorded IT service and data center outages from 2016-2019

Authors

Andy Lawrence, Executive Director of Research, Uptime Institute

Rhonda Ascierto, Vice President of Research, Uptime Institute

Uptime Institute's examination of the outages that made media headlines in 2019 shows that failures are affecting all sectors and types of IT services. The frequency of outages is higher than ever, while the primary causes are increasingly shifting from physical data center infrastructure to software and IT systems.

UII-35 v1.1 published February 3, 2020 last updated February 26, 2020

This Uptime Institute Intelligence report includes:

ABOUT UPTIME INSTITUTE INTELLIGENCE

Uptime Institute Intelligence is an independent unit of Uptime Institute dedicated to identifying, analyzing and explaining the trends, technologies, operational practices and changing business models of the mission-critical infrastructure industry. For more about Uptime Institute Intelligence, visit uptimeinstitute.com/ui-intelligence.

Introduction

Critical IT systems, networks and data centers are far more reliable than they once were. Not only has equipment improved over time, with higher quality builds, self-management, and redundancy, but also management processes are now in place to anticipate failures or limit the consequences of failure. In an era of cloud, distributed architectures, traffic management and low-cost replication, IT can reroute around many failures, often automatically.

Despite all this, Uptime Institute Intelligence finds that major failures are not only still common, but also that the consequences of such failures are high — undoubtedly higher than in the past. In 2019, there were — once again — major failures of financial systems, hours-long outages of 911 emergency service call numbers, aircraft unable to fly, and healthcare systems lost during critical times. While visibility into the causes of many of these outages has become more opaque, the media headlines regarding their impact are there for all to see. These headlines may relate not only to the immediate impact, but increasingly, to the lawsuits that follow months later.

Reliable research on the causes and impacts of IT service failure is difficult to find, with good reason: Such research is difficult to do. Most of those involved in outages do not want to share their experiences and are often advised not to. Most studies have small samples and have often concentrated on areas of interest to commercial sponsors, such as security, denial of service and data center power. This approach does not fully address the fact that most failures have multiple causes, and that the impact of one failure can cascade between data centers and across networks, triggering secondary failures. Assessing the impact of failures — or partial failures — presents yet further challenges, as discussed below.

> Major IT service outages are not uncommon; a senior manager could expect to experience five to 10 in the course of their career.

The growing move to cloud services and the extensive use of colocation makes the need for a good understanding of outages and their consequences more important at the executive level. While it is possible to outsource the work, it is not possible to outsource the responsibility. Similarly, the use of outsourcing for all maintenance and management can create problems. Lawsuits have resulted from at least five outages that Uptime Institute has tracked — in colocation, banking, transportation and telecoms.

Society's increased dependency on IT generally — and on outsourced IT, in particular — is creating a need for a firmer contractual basis relating to preventing, managing and accounting for IT failures; for shared responsibility; and for greater visibility and legal accountability. Cloud and service providers can sometimes be

disarmingly open in discussing their failures, but more commonly they provide little or no commentary, and sometimes they do not admit to full outages at all.

This report, focusing primarily on high-profile, public outages, is one of a series that Uptime Institute produces analyzing IT service resiliency. Although supported by our own survey data, the report is based primarily on publicly available data, such as media sources, which is of mixed quality. Even so, it provides some key insights that highlight how the evolving world of complex, hybrid IT is creating new challenges. For more on how the data was collected and analyzed, see **Appendix: Sources and methodology**.

KEY FINDINGS

- **Major and damaging outages continue to trouble IT service operators of all kinds, despite improvements in technology and management. As in previous years, there is clear evidence that availability does not match marketing claims (service level promises).**
- **Data from publicly reported outages and from Uptime Institute's annual surveys suggest that not only are the causes of outages changing (more are due to IT and networks, fewer to power), but so also are the recovery times (they are longer) and the costs and consequences (they are greater).**
- **Reporting of IT service outages in the media (including social media) is now routine and often immediate, as well as technically vague or inaccurate, and sometimes damaging. Many organizations are now wary of giving precise root causes, limiting what can be learned.**
- **Public cloud-based services account for a significant number of reported service outages. Although these operators have generally good reliability/availability, the scale and complexity of their architectures creates its own challenges. Their high profile means their outages are always detected and reported.**
- **The distributed availability model used by cloud service and many other large operators does, according to Uptime Institute's data, provide some improved resiliency. But the frequency of outages affecting organizations with such architectures is only marginally better (i.e., the frequency is lower) than those with more traditional setups.**
- **The growing dependency on IT generally — and on outsourced IT, colocation and cloud in particular — is creating a need for a firmer legal basis relating to preventing, managing and accounting for IT failures. Uptime Institute is aware of at least five lawsuits resulting from outages, as well as multiple regulatory initiatives.**

Reporting and counting — a conundrum

There is a growing trend for managers at organizations affected by outages not to publicly disclose or discuss its cause. Indeed, in some cases, suppliers and service partners must sign confidentiality agreements to that effect. (Uptime Institute often produces or is given access to exhaustive root-cause analysis reports — but these are invariably covered by nondisclosure agreements and so cannot be discussed here.)

There are probably many reasons for this wariness to share the causes of outages, but a rise in legal cases and the growing losses resulting from outages provide context: a rushed or incomplete media statement may backfire, and many conclude it is better to say nothing. The root-cause analysis statements published by many of the major cloud operators are noteworthy exceptions.

This growing unwillingness to reveal causes presents us with a methodological conundrum when it comes to analyzing the data on publicly reported outages. In previous years, the number of outages for which the cause was not known was very small; some were excluded from our sample. For others we were able to hunt down the causes or we made an intelligent estimate based on known information. But in 2019, this figure rose to 11% of even the most serious outages (category 4 and 5 outages; see **Severity of outages** for more on our classification system).

Until or unless there is some form of mandatory reporting of the causes of service interruptions, the data will always be patchy. Such legal changes may come: In certain industries (the financial services sector in the United Kingdom [UK] is one example), outage reporting is now mandatory. However, full details of causes are not publicly revealed. Although efforts have been made to encourage detailed, anonymized reporting of outages (such as the Data Center Incident Reporting Network), such initiatives have so far had a limited impact.

The rise of ransomware

Uptime Institute's primarily goal in this research is to better understand the extent and causes of service disruptions. Data security and cybersecurity are major issues that require attention and investment but are not currently areas on which Uptime Institute advises. Most security-related issues are data breaches, and although these can have serious consequences, most do not lead to a service interruption.

However, two forms of malicious attack often do lead to outages or severe service degradation. The first is a distributed denial of service (DDoS), where a coordinated attempt is made to overwhelm a site with traffic. We have tracked some of these in recent years.

The second is ransomware, an emerging cause of outages, with some recent examples at Travelex, the currency trading company; the UK's National Health Service (NHS); and numerous municipal authorities in the United States (US). Ransomware attackers use malware to deny authorized users access to their own data (by encryption) and may threaten to destroy or corrupt it unless a ransom is paid. Often, operators have no choice but to take down all involved IT services in an attempt to recover access and purge the systems of viruses.

In the past two years, the number of ransomware attacks has increased dramatically. Government offices are a particular target. Kaspersky Research Labs, operated by security software supplier Kaspersky, reported 147 municipality attacks in 2019 (up 60% over 2018), with up to $5.3 million in ransom demanded. The IT Governance blog, based in the UK, recorded 19 major ransomware attacks globally in the month of December 2019 alone.

Perhaps the two most serious attacks in 2019 were on the city of Baltimore (US), which refused to pay the ransom and budgeted $18 million to fix its problem; and the city of Atlanta (US), which also refused to pay the ransom and paid more than $7 million to fully restore operations. The WannaCry virus attack on the UK's NHS in 2018 reportedly cost the NHS more than $120 million (£92 million). Starting on New Year's Eve 2019, Travelex's currency trading went offline for two weeks due to a virus attack, costing it millions.

Uptime Institute does not currently record security issues as a cause of outages, even though ransomware will likely require that systems be immediately taken down. This is partly because outages in this category are best tracked by security specialists, and partly because, given the trends, security issues would swamp the dataset. We may revisit this topic, and this decision, in the future.

Outage frequency

How common are outages? Is the number of outages increasing? The answers to these questions require care if misleading conclusions are to be avoided. Yes, outages are common and becoming more so. However, the number of outages relative to the overall rise in IT, and our growing dependence on it, is not demonstrably rising.

During 2019, we published the results of the Ninth Annual Uptime Institute Data Center Survey, with a focus on the prevalence of outages (globally). For the second successive year, about a third 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).

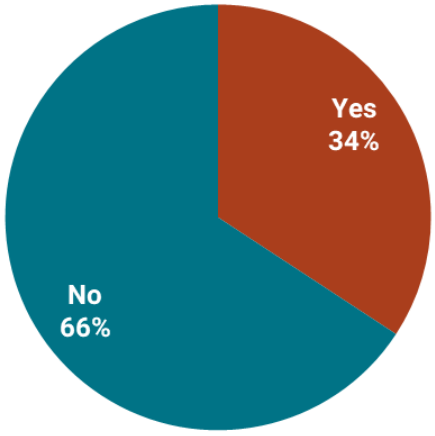

Has your organization experienced an IT service outage or severe service degradation in the last year, either in your own site or a third-party provider?

Source: Uptime Institute Global Survey of IT and Data Center Managers 2019, n=479

Uptime Institute | INTELLIGENCE

Figure 1. **A third of 2019 Uptime survey respondents experienced an outage in the past year**

Looking further back in time, more than half of the respondents, again for the second successive year in our survey, reported they had a serious downtime in the past three years.

As we have previously noted, these figures run contrary to the published service level agreements of most data center and IT service providers, whether they are enterprises with internal customers, colocation companies or cloud providers, which routinely offer at least 99.95% availability. Put simply, major IT service outages are not uncommon; a senior manager could expect to experience five to 10 in the course of their career.

Publicly reported outage frequency

When we look only at the publicly reported outages for 2019, we see a sharp jump (see Figure 2). This does not necessarily mean there were more outages; it could be that there is greater visibility of outages, greater reporting by the media and website trackers, and better data collection by the Uptime Institute Intelligence team. However, given that there is ever more IT, and that our (separate) survey data shows the rate of outages is still high and shows little sign of falling, we may conclude that the number of big, publicly reported is rising steadily too — and that this trend is not abating.

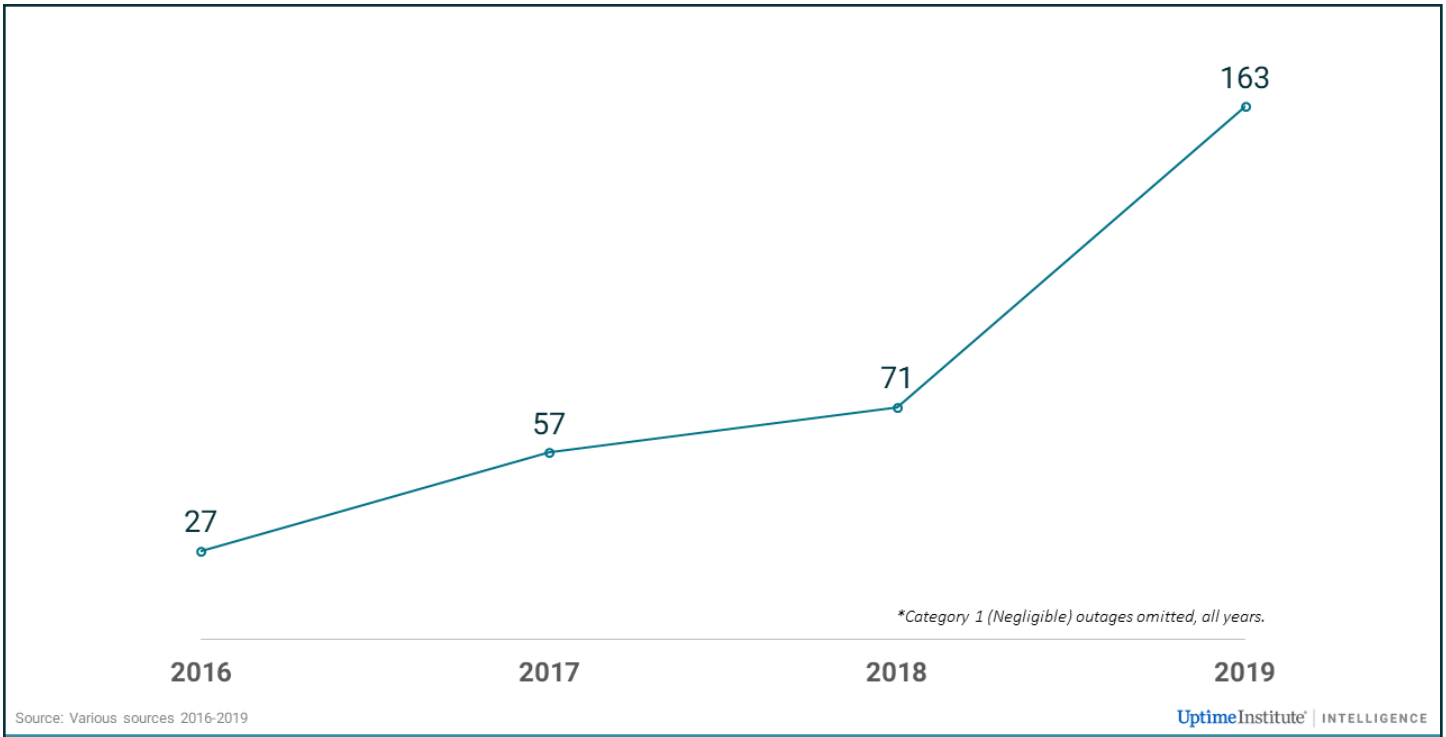

Figure 2. **Publicly reported outages tracked by Uptime Institute rose significantly in 2019***

In 2019, we eliminated from our dataset all publicly reported outages that were classified as category 1 for all years (2016-2019), according to the Uptime Institute Outage Severity Rating (see* ***Severity of outages* *for more on our classification system). This was due both to the volume and unreliability of data for this type of outages.*

Severity of outages

For those operating IT services, there are two main concerns:

- Increasing overall service availability/reliability, which often means eliminating minor and common causes of small outages in a rigorous fashion, through monitoring and analysis and disciplined process improvement.
- Preventing and reacting to serious/severe outages, which can be costly and reputationally damaging. This involves rigorous analysis and processes (as above), as well as a rehearsed process for dealing with different types of escalating incidents.

To differentiate between severity levels, we classify outages according to a simple rating system, the Uptime Institute Outage Severity Rating (Figure 3).

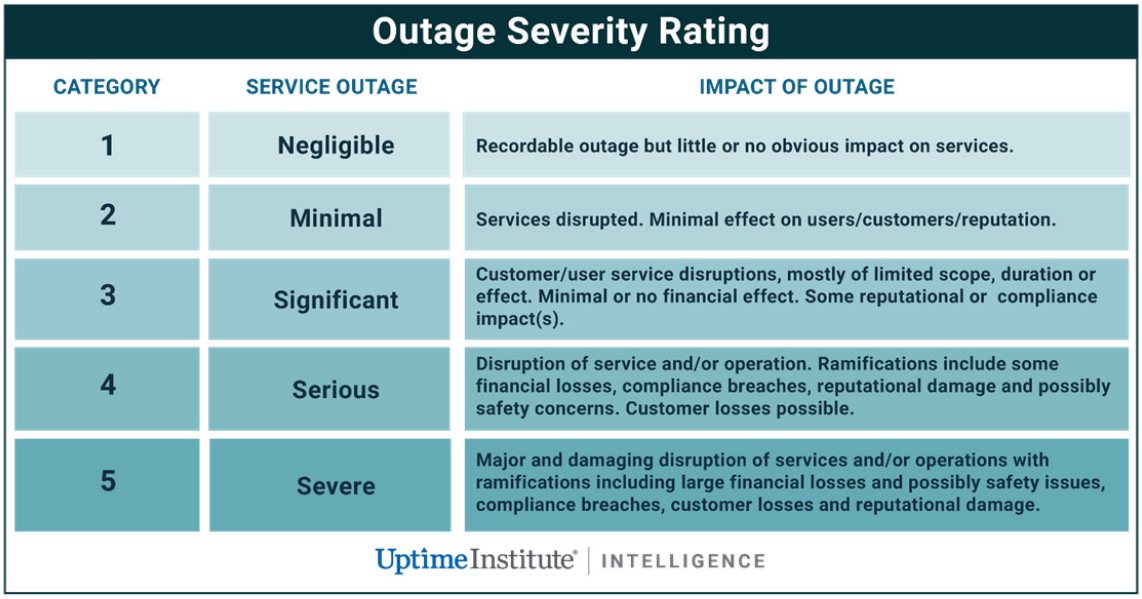

Figure 3. **Uptime Institute Outage Severity Rating: A way to compare outage impact**

For most operators, small service outages are irritants but are also clear signs that attention and investment is needed: serious outages — category 3 and above — have major ramifications, as discussed below.

During 2019, as part of its Digital Infrastructure Assessments, Uptime Institute asked many of its clients to provide data on the number of application outages they had suffered in recent years. The number of smaller outages was higher than we expected, leading us to conclude that there is an inverse pyramid of outages, as shown in Figure 4.

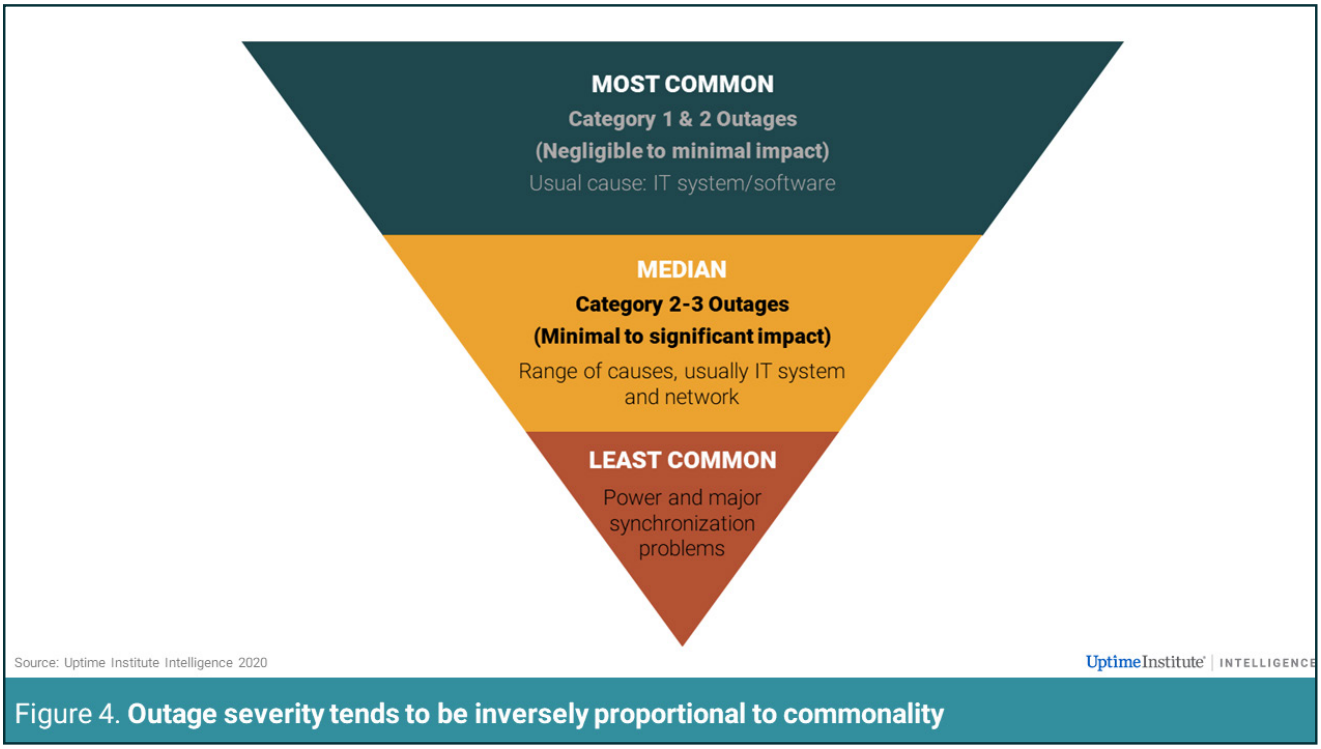

Figure 4. **Outage severity tends to be inversely proportional to commonality**

In the past, many smaller outages have often gone unnoticed — but compliance issues, which may involve mandatory reporting, and the growing use of bots to detect even momentary unavailability have rendered these outages more visible. For software issues, the disciplines of DevOps and site reliability engineering also require that failures are captured and that developers/production staff are held accountable. There are fewer places to hide, and as a result, reliability can be expected to improve over time.

Severe outages rising?

According to our public outage tracking, as shown in Figure 5, there were more severe outages in 2019 than in any of the previous three years, both by number and as a percentage of that year's total. This is troubling, since the data on serious/severe outages, due their high profile, is more reliable than that of lesser outages. Effectively, almost four in 10 of the outages in the 4-year sample were category 4 or 5, involving serious financial losses, reputational losses and problems with compliance and customers.

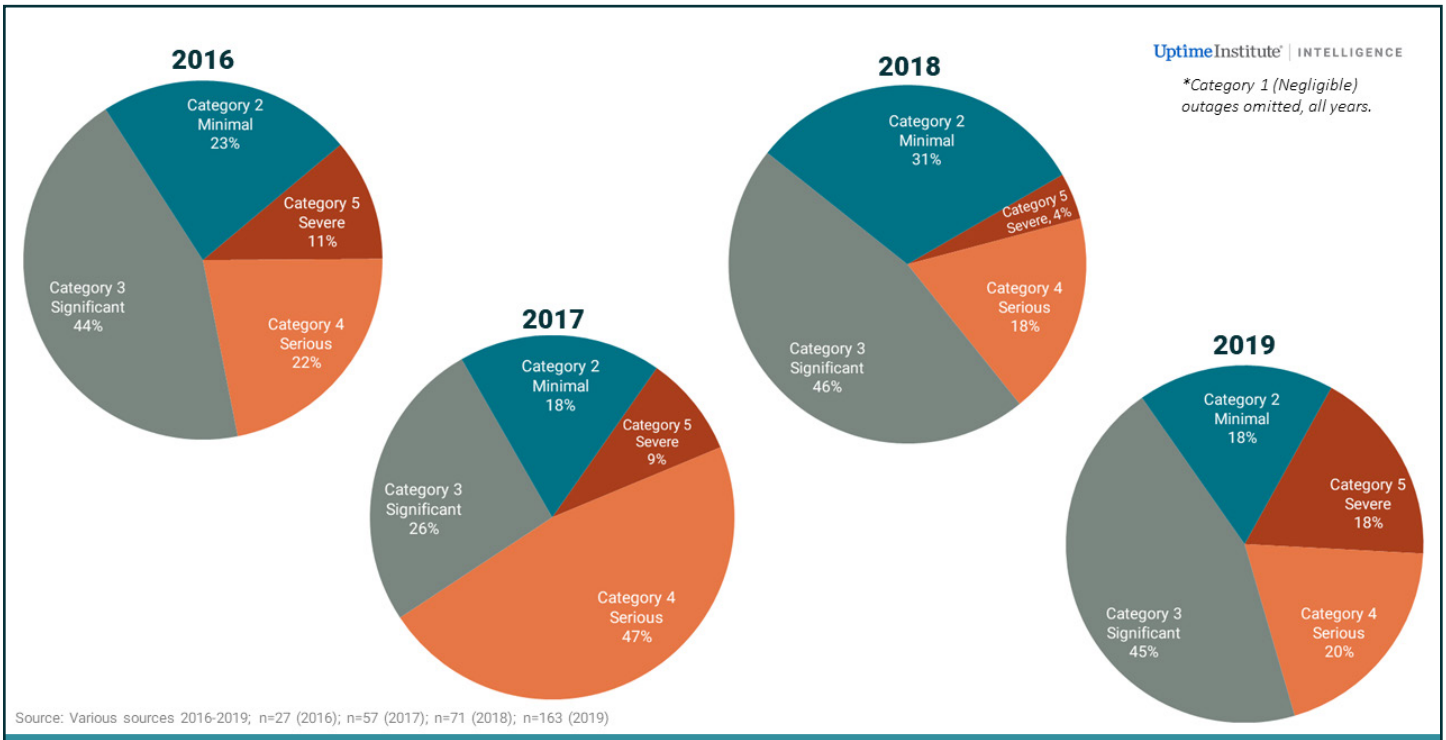

Figure 5. **The proportion of severe (category 5) publicly reported outages grew from 2016 to 2019***

Given that the outages had sufficient impact to be reported in the media, it is not entirely surprising that a significant proportion should be categorized as serious. But we also compared the severity of outages with those reported by operators in Uptime Institute's annual survey of 2019. This showed that about one in four of survey respondents' most recent service outages were categorized as serious or severe, as shown in Figure 6.

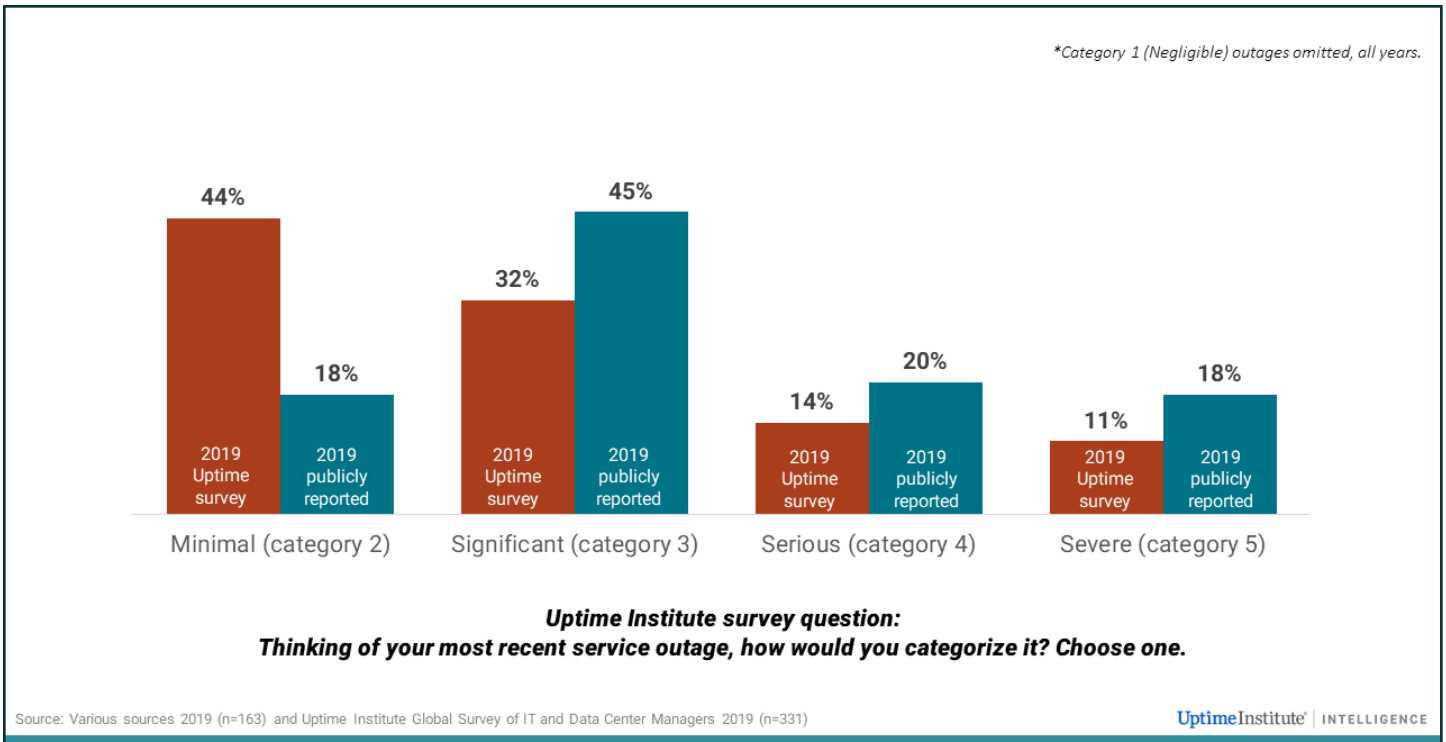

Figure 6. **How serious are the outages? Uptime Institute's 2019 survey vs. public reports***

The two sources (survey and public outage data) tell slightly different stories, but a clear conclusion can be reached: Any organization operating a significant IT service carries a substantial risk of not only a service outage, but also of one that has severe consequences.

Outages — the causes

Serious outages make news headlines and trend on social media but often are not well-explained or, increasingly, are not explained at all. With few exceptions, detailed information about the cause of an outage is rarely publicized. (See **Reporting and counting — a conundrum** for more on this.)

Where explanations are given, they can usually be categorized as some problem in a few high-level areas: networking; IT system (including IT hardware or software, or both); a third-party provider or partner; IT capacity or demand; or data center power, cooling, fire or fire suppression.

Of course, there is rarely a single cause of a major accident or failure; problems commonly cascade and very often there are issues with operational process, personnel training and management that compound an initial failure. In this report, we focus on the primary cause of an outage — although most incidents have a complex backstory.

During the three years from 2016 to 2018, there were 53 outages in total that we classified as serious or severe (either category 4 or 5). We summarize the most severe here:

- In 2016, there were three category 5, severely disruptive business-critical outages. Two of these were caused by data center power outages and one by a network equipment failure.
- In 2017, there were five category 5 outages. Two were caused by IT systems, one by a network issue, one by a data center power failure, and one by cooling/mechanical systems failure in the data center.
- In 2018, there were three category 5 outages. One was caused by a network failure and two by IT systems issues.

The difference between a serious category 4 and a severe category 5 outage may be minimal (and in some cases, debatable) — both are costly. When all of the 53 serious/severe outages between 2016 and 2018 are taken together, the three biggest primary causes are roughly equal, with power, network and IT system each accounting for about a third of the total (see Figure 7).

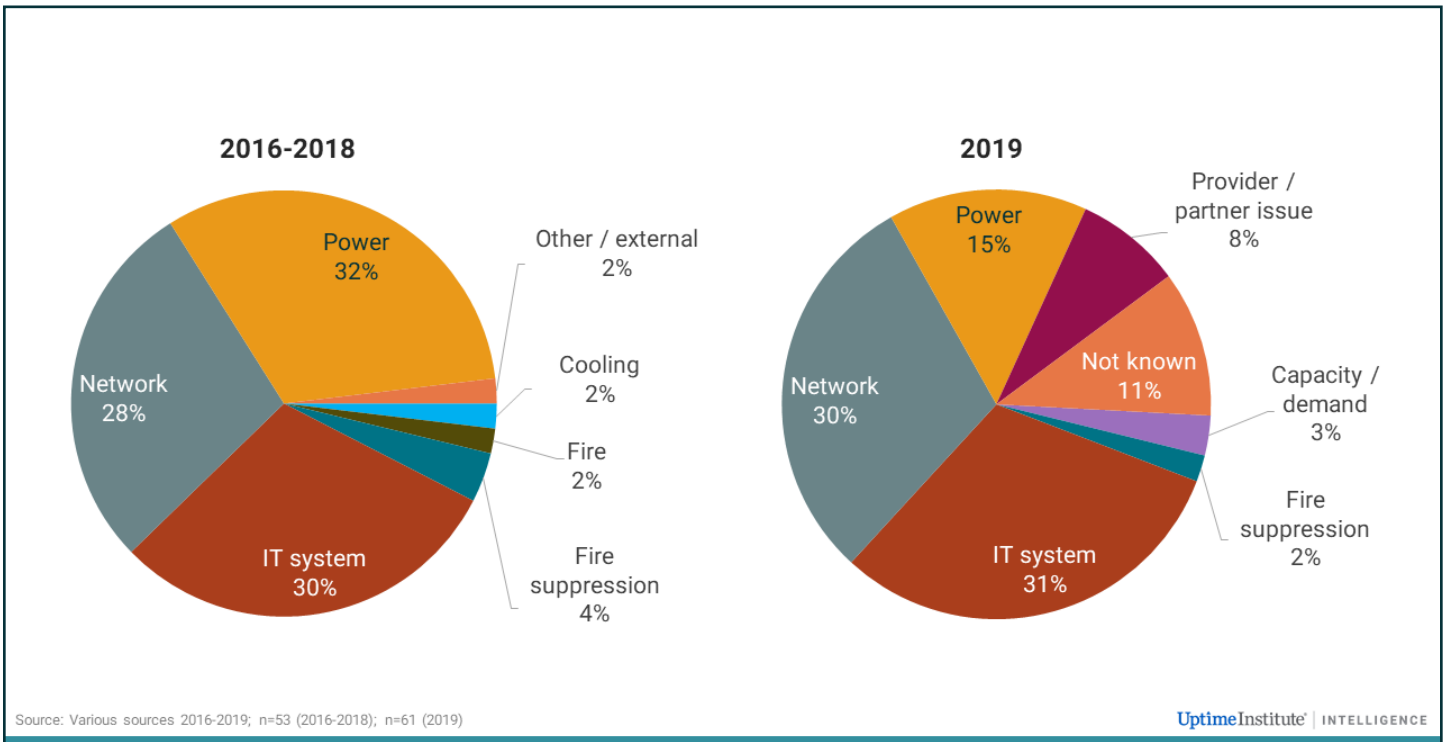

Figure 7. **Power is less frequently a cause of the most serious and severe outages**

In 2019, the number of serious or severe outages that made headlines rose significantly: there were 61 category 4 and 5 outages in total, with nearly half of these being severe (category 5). The majority were caused by IT systems (with software/software configuration issues being most common) and by network issues. Just 15% of the serious or severe (category 4 and 5) outages were data center power-related.

In the absence of good, detailed, root-cause analyses of failures, it is often difficult to draw conclusions about the causes of publicly reported outages. Often, "IT system" is used as a catch-all term by organizations' public relations teams when, in reality, the cause is either not yet known or well-understood, or when the organization is unwilling to divulge further details due to the risk of repercussions. It is not uncommon for the root cause to surface in a lawsuit filed against a third-party provider or partner months after a high-profile outage.

> While it is possible to outsource the work, it is not possible to outsource the responsibility.

As more organizations increasingly rely on colocation, cloud, hosting and other data center and IT service providers, we expect third-party provider or partner issues will become more common. In 2019, third-party providers/partners accounted for almost one in ten of all serious or severe outages (category 4 and 5).

Whether we look at just the most serious outages or all of the outages combined, it appears that data center power failures are causing fewer outages than in years past, while IT system and network issues are becoming prevalent. Figure 8 shows the top three causes of all publicly reported outages since 2016 in which a cause was given (in other words, it does not include outages where the cause was not disclosed).

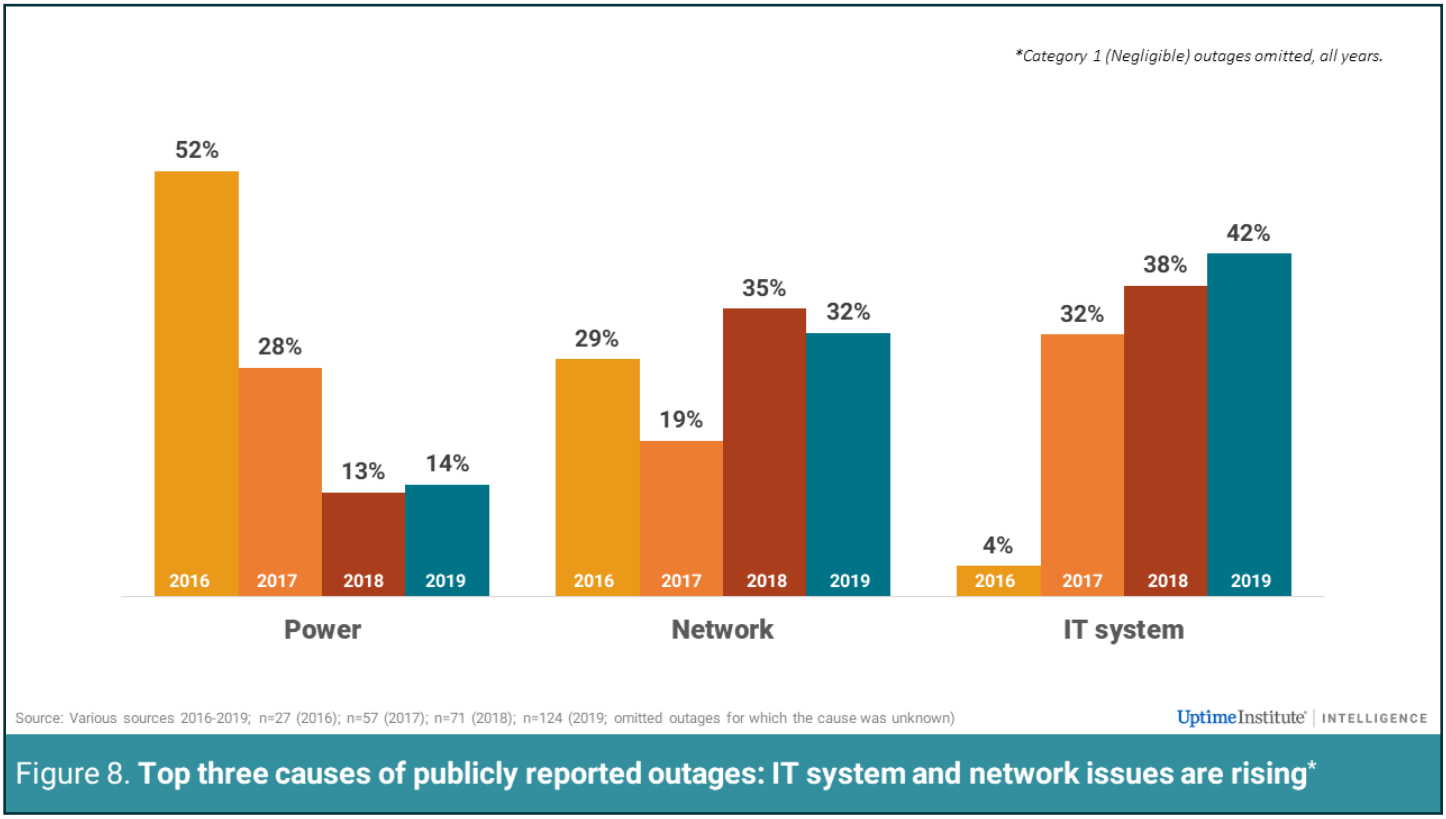

Figure 8. **Top three causes of publicly reported outages: IT system and network issues are rising***

The apparent rise in outages caused by IT systems and network issues may be due to the broad shift in recent years from siloed IT services running on dedicated, specialized equipment to more IT functions running on standard IT systems. Software-defined networking, for example, uses standard IT servers for network routing and other

functions. Another example of IT running on standard systems are application containers: containers can run modern IT workloads on most any type of IT system, including bare-metal servers, traditional server environments, virtual environments, or any type of cloud. As more organizations move to increasingly sophisticated virtualized network and IT approaches (driven by a desire for greater agility and automation), the underlying data center infrastructure is becoming less of a focus — the technologies are proven and stable, and there is an established discipline for management and maintenance.

Another interesting trend is the growing number of outages in which a cause was not disclosed at all. As Figure 9 shows, between 2016 and 2018 just 7% of all publicly reported outages were without explanation of the cause ("not known"); in 2019, this proportion jumped to 24%.

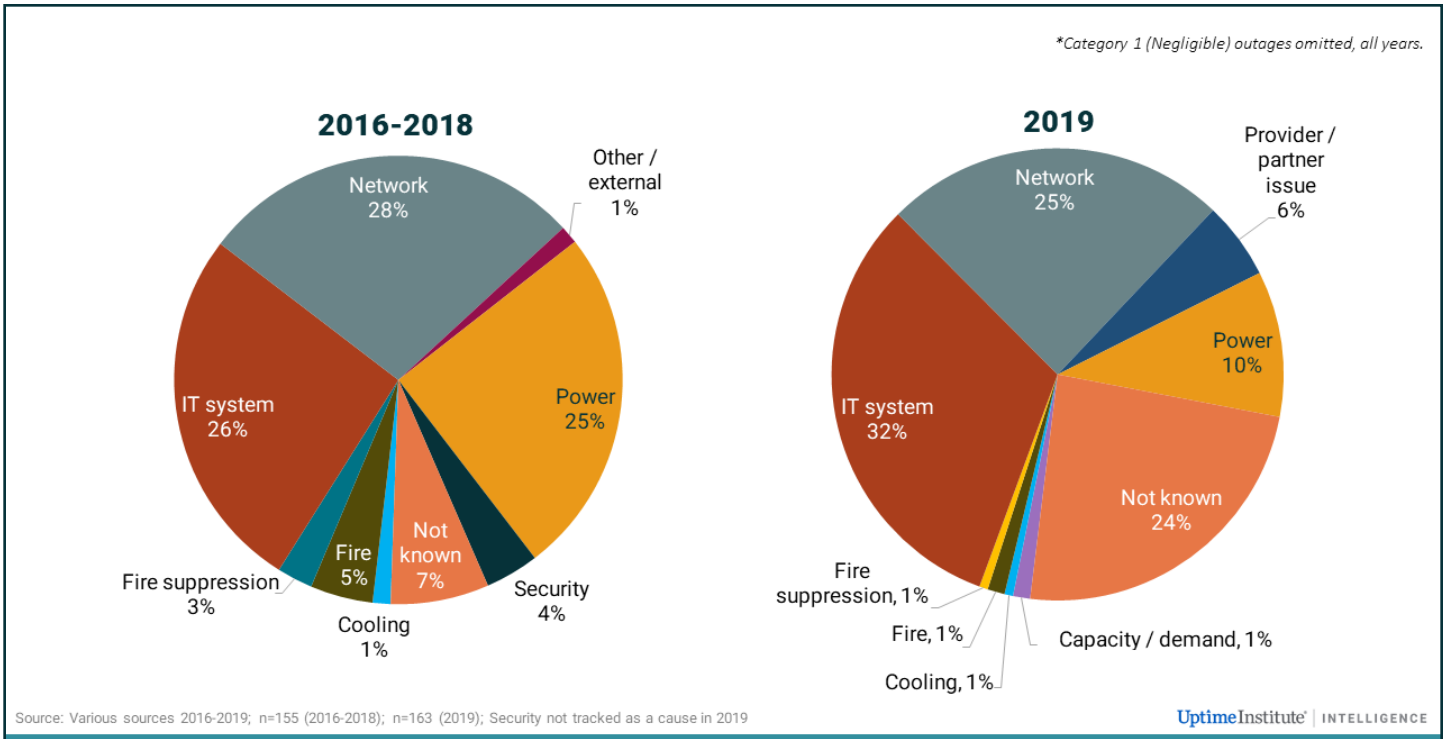

Figure 9. **In 2019, fewer organizations disclosed the cause of their publicly reported outage***

In addition to possible concerns over legal action, more organizations may be choosing not to comment on the cause of an outage as an attempt to "shorten the news cycle." In 2017, for example, when Amazon disclosed that an incorrectly typed comment was the cause of a major outage, there were countless news reports and social media posts about the "fat-finger typo" that took down big chunks of the internet. As a way to avoid a similar outcome, some organizations may not disclose the cause of an outage in the hopes of there being less for the media to write about and less for disgruntled users to comment about on social media. In other words, when the cause of a publicly reported outage is "not known," it is simply not known to the public — the organization will usually know the primary cause.

In 2019, we asked our survey respondents about the primary cause (or causes) of their most recent incident or outage. The results are shown in Figure 10.

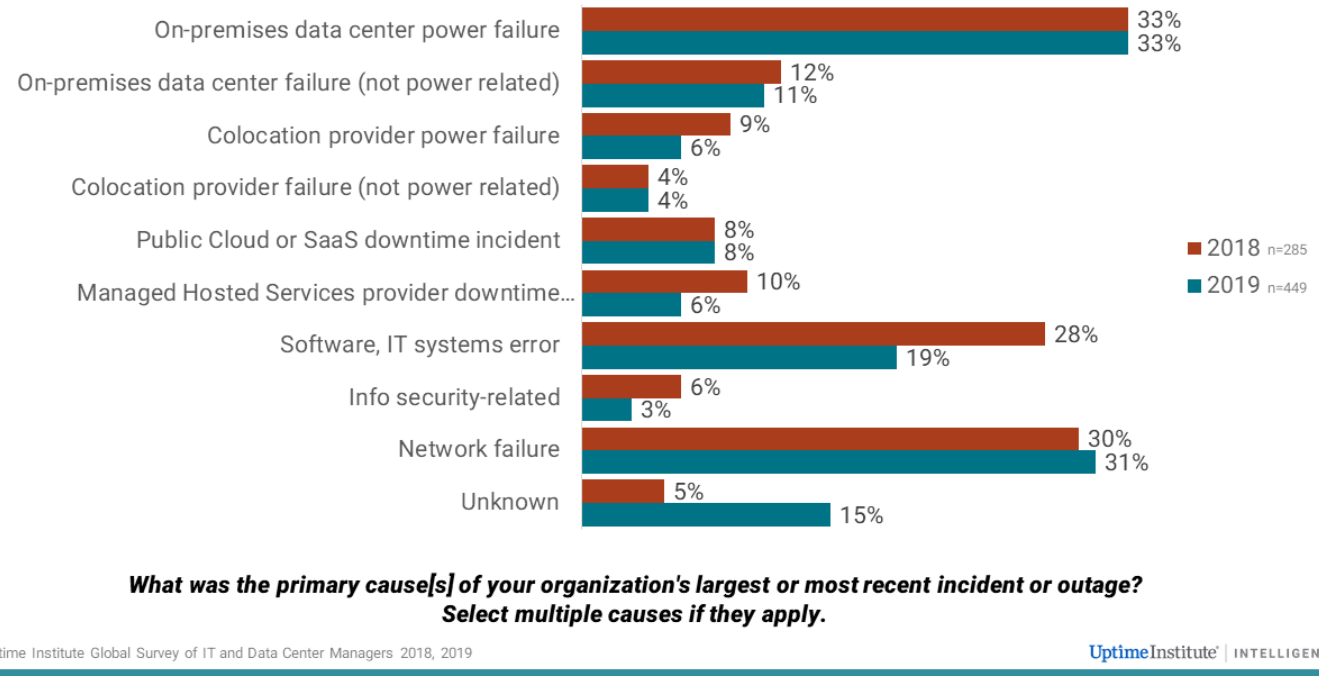

Figure 10. **Uptime survey 2019: Power, IT system and network issues remain top causes of outages**

In our survey, on-premises power failure was the single biggest cause of outages, accounting for one-third in 2019, as in 2018. Networking issues were close behind, at 31%; other IT issues (e.g., software and systems failures) fell significantly, from 28% in 2018 to 19% in 2019. Problems — some of which were power-related — at third-party suppliers (e.g., colocation companies, hosting and cloud companies) jointly accounted for 24%.

It is clear that data center power and power-related failures continued to be a major issue in 2019, even if power appears to have less influence in publicly reported outages.

The discrepancies between the major causes of outages in our surveys and that of publicly reported outages are likely due to several factors, including specificity and completeness of information. Uptime's survey questionnaire was written for IT and data center managers and included a relatively high level of granular information about outages and their causes; media reports, on the other hand, are written (typically under tight time constraints) for a broader, more mainstream audience.

There is also truth in anonymity. The identity of respondents to Uptime's survey are known to (and vetted by) Uptime, but the information they provide is shared only in aggregate and in a strictly anonymous form.

Outages — by sector

Cloud and IT services are designed to operate with very low failure rates. Layers of software and middleware, orchestrated by artificial intelligence and other big-data approaches, reroute workloads and traffic away from failure. On the whole, they provide high levels of service availability, at huge scale and growing complexity. Even so, no architecture is fail-safe; 2019 was a particularly bad year for cloud and IT services.

Taken together, cloud and digital services — which includes data center colocation, software as a service, IT hosting and managed services, and other information services — suffered more publicly reported outages (36% of the total) than any other sector, as shown in Figure 11. This, we believe, fairly reflects their scale and position in the ecosystem, supporting so many other services, as well as having more IT infrastructure and customers than any other group.

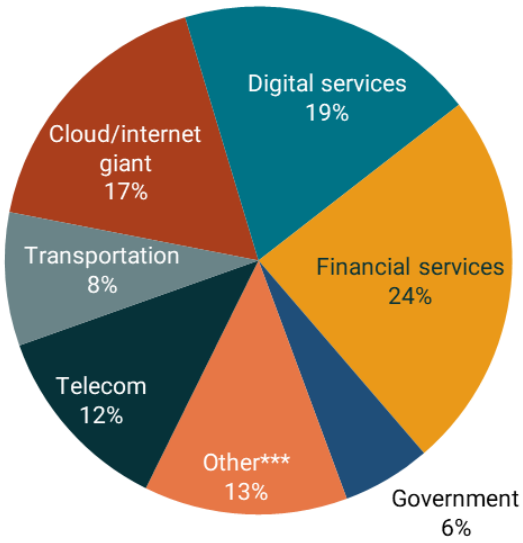

*Category 1 (Negligible) outages omitted, all years.

**Sector definitions used in 2016-2018 have been converted to 2019 categories.

***Includes the following sectors, each of which comprised <5% of all publicly reported outages during the reporting period: Education/Research; Healthcare/Pharmaceuticals; Logistics; Manufacturing; Retail; Utilities.

Source: Various sources 2016-2019; n=27 (2016); 57 (2017); 71 (2018); 163 (2019)

Uptime Institute | INTELLIGENCE

Figure 11. **Publicly reported outages by sector: Financial services, digital services and cloud dominate*,****

Another industry that comes out particularly badly is financial services. Even allowing for the financial industry's important and frontline role, it is clear from the known examples that many banks are suffering from a combination of complexity and lack of investment/modernization. A particular issue that has affected financial services, as well as other industries such as air transport and retail, is "asymmetric criticality" or "creeping criticality." This refers to a situation in which the infrastructure

and processes have not been upgraded or updated to reflect the growing criticality of the applications or business processes they support. Some of the infrastructure has a 15-year life cycle, a timeframe out of sync with the far faster pace of innovation and change in the IT market.

High-profile and persistent outages in the financial services sector is troubling regulators, which now plan to exercise more oversight in the US, UK, Europe and elsewhere.

Both finance and cloud make significant use of multi-site synchronous or asynchronous replication — copying all their data to a second or third site in real or near-real time. When groups of data centers are clustered (logically) to back up to each other, this is often termed an "availability zone." This approach is considered highly reliable — yet our survey research shows that those with such architectures achieve only a small advantage in availability. This is likely due to complexity and scale, and the level of innovation. Over time, these architectures are certain to prove more reliable.

Impact of outages

While the causes of an outage can often be obscured, the public impact is difficult to hide (especially if it appears on prime-time news). During 2019, all of the organizations that suffered a category 4 or 5 outage suffered significant financial and reputational damage.

The financial impact of outages is always difficult to estimate, both for individual incidents and in aggregate. For individual incidents, managers must decide if and how they will calculate consequential losses, lost opportunities and reputational damage. Sometimes, lawsuits add yet more to the bill. For aggregate figures, a few large incidents can distort the overall picture.

In the Uptime Institute 2019 survey, while most (60%) of the incidents reported had a relatively negligible cost impact, just above 10% cost more than $1 million (see Figure 12). There were six recorded incidents where the cost of the outage exceeded $40 million.

Figure 12. **One in ten major outages reported in Uptime Institute's 2019 survey cost over $1 million**

For the publicly cited outages, many of the category 4 and 5 incidents clearly cost tens of millions of dollars. Equally notable is the extent of the disruption that IT failures can cause. In 2019, there were examples of payment processing failures (affecting banks, retailers and consumers); major transportation service failures (affecting airlines, train services and consumers); telecoms failures (multiple instances, including 911 services); stock exchange outages; retail system failures; border service failures; and health systems failures (causing cancelled surgeries and treatments). And, of course, there were many cloud service outages, with a wide variety of consequences.

The examples below from 2019 were not selected for their size or level of impact, but rather to demonstrate the wide variety or problems caused by IT outages:

- **UK Ministry of Justice (UK)** - All the UK's prosecution and court administration IT services were down for more than 3 days, crippling the legal system. The outage was discussed separately in both houses of Parliament. According to published reports, the cause was an infrastructure failure in the Ministry of Justice's suppliers' data center.
- **Amazon Web Services (US)** - A power surge in a Northern Virginia (US) data center led to a failure of backup generators and exposed a fault with load switching software. Customers faced a service loss of many hours and for some, extensive data loss.
- **Facebook (US)** - The social media giant suffered an outage in March 2019 for more than 12 hours, then again on Thanksgiving

Day (November 28, 2019) for 6 hours, affecting Black Friday (November 29, 2019) trading and advertising. Facebook has an estimated 2.3 billion users.

- **Bank of America (US)** - Thousands of customers were locked out of their bank accounts and multiple services were disrupted. Although the outage was apparently brief in duration, Bank of America's stock fell 3.7% during the next day's trading.

Any organization operating a significant IT service carries a substantial risk of not only a service outage, but also of one that has severe consequences.

- **Commonwealth Bank (Australia)** - Multiple services were down for 12-18 hours, and customers of other banks could not receive payments. Up to eight million people were affected, and the cost was estimated at AU$5 million.
- **London Stock Exchange (UK)** - The Exchange suffered its worst outage in 8 years, lasting 1 hour 40 minutes. Many trades were delayed or switched to competitive exchanges.
- **Target (US)** - The retailer lost checkout services for up to 4 hours over 2 days, caused by a supplier's data center problem. Unconfirmed estimates of $50 million in lost sales were reported.
- **Telstra (Australia)** - Among the telco's many service outage problems in 2019 was a 4-hour outage in November that affected automated teller machines and electronic funds transfer at point of sale systems running on its network. Losses were estimated at AU$100 million.
- **Tele2 (Sweden)** - In Sweden and the Baltic countries, a serious telecom service outage affected 112 (emergency services) calls, landline, mobile and internet protocol-based calls for 2-3 hours. This was one of three crippling outages for Tele2 over a 12-month period. Other emergency service interruptions in 2019 affected AT&T (US), KPN (Netherlands) and Verizon (US).
- **Southwest Airlines (US)** - The airline, which was hit by major outage in 2017, suffered a multi-hour system-wide computer outage that left pilots unable to submit flight plans. An estimated 650 flights were delayed.
- **British Airways (UK)** - The airline has had numerous outages since its well-known and devasting data center outage in 2017. In August 2019, hundreds of flights were cancelled and delayed after an outage, and problems recurred in November. Failures have cost the airline's owner, International Airlines Group, a reported $103.6 million.

Duration of outages

Data center operators will testify that they are often able to resolve incidents, at a facilities level, in a few minutes or maybe a few hours. IT staff, however, may labor for hours or days beyond that to recover services. And even then, organizations may struggle to recover from a backlog or to move people and assets where they should be.

For outages in 2018 and 2019, there are signs of a trend: When we look at how long the incident lasts until the IT services are back to normal, a growing number of publicly reported outages last 12 hours or more, with a significant number lasting more than 24 hours — and the percentage of outages that were longer than 12 hours grew as well. This is true even though one of the biggest causes of lengthy outages — ransomware — was excluded from our sample.

Table 1. **Outages getting longer?**

Duration (hours)	2017 (n=57)	2018 (n=71)	2019 (n=140)*
0 - 1	1	4	20
1 - 4	35	25	49
4 - 12	13	25	26
12 - 24	4	6	14
24 - 48	2	4	14
> 48	2	7	17

** Outages for which the cause was not known were eliminated from the analysis.*

Note. *Times reported are time to service recovery, not time to full business recovery.*

Although complexity and interdependency are clearly common factors, the explanations for this vary, and there may certainly be problems of definition and reporting. But it is clear that resolving facility engineering issues is usually a relatively predictable matter — often recovery processes have been drilled, and spare parts are kept at hand. Software, data integrity and damaged/interrupted cross-organizational business processes can be much more difficult issues to resolve. Not only are the causes of outages changing, but so are the recovery times.

Summary

Almost 70 years have elapsed since the United States Department of Defense created the Advisory Group on the Reliability of Electronic Equipment. Uptime Institute introduced the Tier Standards for data centers over 25 years ago, and a decade has passed since Google introduced the concept of site reliability engineering. These developments have helped improved the reliability of complex, large-scale IT infrastructure and services. Despite all this intelligent endeavor and much investment, outages caused more problems in 2019 than in any previous year. Although we have not tallied the costs, the disruption and remediation run into billions of dollars globally.

Uptime Institute's research suggests that the industry is reaching a critical point. Regulators of financial services, emergency services, telecoms and central government are beginning to realize that greater visibility, accountability and control is needed. Many new regulations are in the pipeline, against a backdrop of increasing legal action. And the cause is simple: Outages keep occurring, and they are proving disruptive and damaging to industry and to the wider economy.

Many players and providers in the IT and critical infrastructure industry will view greater regulation as an unwelcome and intrusive development. That is debatable. What is not debatable is that resiliency requires ever greater investment, attention and discipline.

Appendix: Sources and methodology

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

- **Uptime Institute's Annual Data Center Surveys** - These long-running, global surveys, with over 1,100 respondents in 2019, ask detailed questions about outages; some of the findings are discussed here. This data represents the most statistically significant dataset relating to outages in the critical infrastructure industry.
- **Uptime Institute Intelligence's public outages database** - Since the beginning of 2016, Uptime Institute has collected data about major IT outages from public media reports and other public sources (social media, outage detection sites, etc.) on an ongoing basis. This database 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 the main subject of this report.

- **The Abnormal Incident Report (AIRs) database** - This is a long-standing confidential system for global Uptime Institute Network members to share details of incidents under a nondisclosure agreement. Most incidents recorded do not actually lead to outages — many are "near misses." We do not refer to these incidents in this report.
- **Professional services** - Uptime Institute conducts Digital Resiliency Assessments and root-cause analyses of failures on behalf of clients. Although these assignments are confidential, the experience garnered from these incidents informs our analysis.

The methodology used for the bulk of the findings in this report is limited and the data should be understood in this way — it is primarily useful for trending and, while we believe it is directionally accurate, it is not a representative dataset for all outages. There are several limitations:

- If a failure is not reported or 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 IT services in geographies with a well-developed and open media.
- We limit failures to those that had a noticeable impact on end users — a major fire during data center commissioning, for example, may never be registered. This year, we have also eliminated all category 1 outages — small, short failures where the business or reputational impact is negligible.
- The amount of information available varies widely from outage to outage, and sometimes there is very little information available at all (see **Reporting and counting — a conundrum**). It has regrettably been necessary, in some of the analysis, to include outages for which the cause is "not known" — meaning it was never disclosed.
- Finally, while we include IT system failures, we do not generally include security breaches, although these can lead to service interruptions (see **The rise of ransomware**).

Uptime Institute® | INTELLIGENCE

ABOUT THE AUTHORS

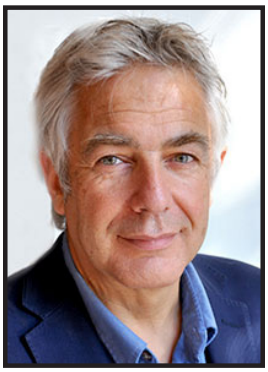

Andy Lawrence is Uptime Institute's Executive Director of Research. Mr. Lawrence has built his career focusing on innovative new solutions, emerging technologies, and opportunities found at the intersection of IT and infrastructure. Contact: alawrence@uptimeinstitute.com

Rhonda Ascierto is Uptime Institute's Vice President of Research. 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. Contact: rascierto@uptimeinstitute.com

ABOUT UPTIME INSTITUTE

Uptime Institute is an 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 US, Mexico, Costa Rica, Brazil, UK, Spain, UAE, Russia, Taiwan, Singapore and Malaysia. Visit 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

