RISK & RESILIENCY

Uptime Institute® | RESEARCH

Uptime Institute data shows outages are common, costly, and preventable

FOCUS | JUNE 2018 | UII report 11
Andy Lawrence, Executive Director of Research, Uptime Institute

New Uptime Institute research data shows that downtime is common and may even be increasing, despite many advances. Complexity and extensive use of third parties has made life more difficult for management.

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. In the new cloud world, distributed architectures, traffic management, and low-cost replication mean that IT can re-route many failures.

All of this appears to represent significant progress. But new research by Uptime Institute suggests far from a simple picture, with failures and downtime still common and possibly even increasing. When problems do occur, recovery times can be lengthy, fault diagnosis can be complicated by interlocking or interdependent systems, and costs can be significant because of ever-greater reliance on IT in every area of society and business.

Uptime Institute also found that there has been little research into the causes and costs of downtime in IT, both at a macro level (largely due to a reluctance by organizations to report failures externally) and at a micro level. This finding suggests that significant investments may be being made without real assessments of the risks and that some of these investments are misplaced. Moreover, many organizations clearly have management problems, even when diligent efforts have been made. Most failures, the new research finds, could have been prevented.

KEY FINDINGS

- IT service and data center outages around the world are not only common, suggesting most SLAs are very often broken, and that outages may actually be increasing.
- The biggest cause of IT service outage is a data center power outage, closely followed by network problems, and then by an IT system failure.
- Failures at third-party cloud, colocation, and hosting providers, when aggregated, are now the second most commonly cited reason for IT service failure.
- Power failures accounted for 36% of the biggest, global public service outages tracked by Uptime Institute since January 2016.

RISK & RESILIENCY

Uptime Institute® | RESEARCH

Uptime Institute data shows outages are common, costly, and preventable

KEY FINDINGS

- 41 survey respondents reported an outage that had cost over $1 million. One outage cost over $50 million.
- Around a third of all reported outages cost most than $250,000.
- Most data center managers think that extending the life of a data center increases both risks and operating costs.
- Most IT service outages are preventable. 80% of respondents say that their most recent service outage could have been prevented.
- Many organizations have little understanding of the likely financial and overall business impact of particular IT service failures, nor have they carefully assessed the particular risks they face.

Uptime Data

Uptime Institute has three sources of data on the causes of 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 under NDA. No information from this database is included in this analysis.
- Uptime Institute Research has collected data relating to over 100 major publicly recorded service outages since January 2016.
- The Uptime Institute 2018 Data Center Survey. The Eighth Annual Uptime Institute Industry Data Center Survey provides an overview of the major trends shaping IT infrastructure delivery and strategy. The survey was conducted via email between February and May 2018, and includes responses from 867 data center operators and IT practitioners globally, from enterprise and service provider facilities. For the first time, this survey included many detailed questions about outages.

This is the most comprehensive data set that has ever been collected on outages at an industry-wide level. Although each of these research methods and the types of information they collect are different, they combine to create a complex picture that suggests that outages are far from rare and continue to be a major problem for operators. And as the publicly reported outages show, the consequences of failure can be more expensive and damaging than ever.

Uptime Institute® | RESEARCH

Uptime Institute data shows outages are common, costly, and preventable

Outages are neither rare nor declining

Uptime Institute's survey found that almost one third (31%) of those responding (n=664) had experienced an IT downtime incident or severe degradation of service in the past year. Moreover, about half (48%) said they had experienced at least one outage in the past three years either at one of their own sites or that of a service provider.

This is a very high number and is entirely at odds with most publicly announced availability figures (usually > 99.9%). The result suggests that most service level agreements (SLAs) are commonly broken (although they are usually phrased so that the operator pays only a nominal penalty). It is also out of line with the number of outages reported by Uptime Institute Network members, who very often report sustained periods without any site outages at all.

Perhaps more remarkably, the portion of operators suffering an outage appears to be increasing. In the 2017 survey, "only" 25% had experienced an outage during the previous 12 months (among members of the Uptime Institute Network, this level was halved).

Explanations for this high outage rate are not clear. Possibly, increased complexity and interdependencies of different systems and different data centers using ever more complicated management systems may be increasing the number and impact of failures (remember that most data centers handle more work each year). Uptime Institute has also found evidence that failures tend to occur during periods of technology change and investment but also at sites where there is under-investment and legacy assets are not upgraded. Some survey evidence suggests that many sites fall into one of these two categories.

Power still the big challenge

What is the main cause of IT service downtime? Uptime Institute's survey data show that the biggest single cause of failure is the one that the industry invests so much to prevent – a loss of on-site power, which was cited 93 times by 285 respondents who suffered at least one outage (33% of all respondents suffered at least one such incident). This was closely followed by network failure (30%) and an IT/software error (28%).

But there is a new factor reducing IT service availability. The failures experienced at third-party service providers (colocation, hosting, or cloud) account for 87 incidents (31% of those reporting an incident), which is only slightly fewer than on-site power failures at enterprise data centers. Third-party failures have become a critical issue; in hybrid environments, CIOs need to be as mindful of their data center suppliers as they are of their data center operations.

Uptime Institute® | RESEARCH

Uptime Institute data shows outages are common, costly, and preventable

ANSWER CHOICES	RESPONSES	
On-premise data center power failure	32.63%	93
Network failure	29.82%	85
Software, IT systems error	27.72%	79
On-premise data center failure (not power related)	12.28%	35
Managed hosted services provider downtime incident	9.82%	28
Colocation provider power failure	9.47%	27
Public cloud or SaaS downtime incident	7.72%	22
Security-related	6.32%	18
Unknown	5.26%	15
Colocation provider failure (not power related)	3.51%	10
Total Respondents: 285		

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

Uptime Institute's data aligns closely with the causes of big, public failures recorded since early 2016. Power outages accounted for 36% of failures, followed by 25% for network issues and 22% for IT/software issues.

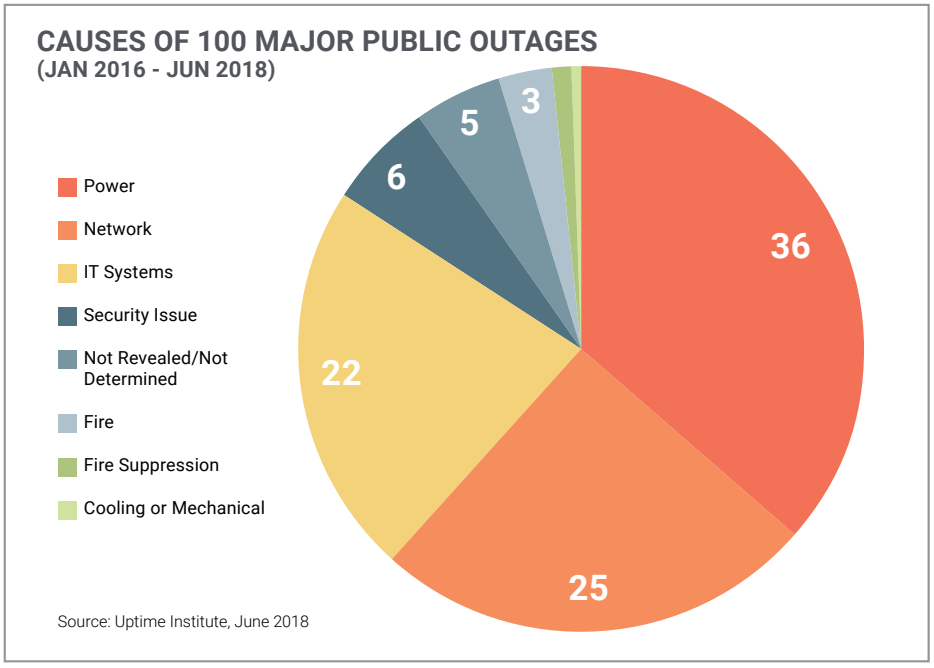

Uptime Institute® | RESEARCH

Uptime Institute data shows outages are common, costly, and preventable

These results support our view that power and facilities issues are most likely to cause problems but IT and network problems tend to cause the most issues, because of interdependencies, systems complexity, relatively difficult fault analysis, and longer recovery times. However, it is rare that the initial cause of a problem is contained, so a power problem may soon become an IT systems recovery issue, especially where multiple, interdependent databases are affected, as is usually the case. Older, single-site transactional databases can usually be recovered more quickly, for reasons of both design and relative simplicity.

Causes of 10 most serious outages

During the period from January 2016 to June 2018, Uptime Institute recorded 10 outages that were "extremely serious," meaning that they caused a very serious loss of revenue and brand and potentially an existential threat to the data center operator or its clients (for a formal definition and deeper discussion, see upcoming reports by Uptime Institute).

Primary cause of Ten big "category outages"	
IT Systems	3
Network	3
Power	3
Cooling	2

Preventable outages

Uptime Institute frequently sees percentages of failures ascribed to human error. In our experience, all major failures that happen in normal peacetime can be attributed to human error. Uptime Institute has not visited all the sites concerned in this report, but the hands of the error-prone operator, manager, and supplier are clearly evident.

A majority (80%) of the respondents to the Uptime Institute survey believe their biggest/most recent outage (for those who had suffered one) was "preventable." This result suggests, as Uptime Institute always advises, that the most common cause of problems lies in processes and practice, rather than architecture or equipment. But the survey also supports the view that cautious and careful design at the outset does reduce outages (see 2N vs. N+1 below, as an example).

Uptime Institute® | RESEARCH

Uptime Institute data shows outages are common, costly, and preventable

Costs of downtime

Historically, Uptime Institute has been wary of developing an average cost of downtime, largely because it would vary so widely according the service and type of data center. When applications are interdependent and spread across multiple data centers, it makes such estimates still less reliable or meaningful.

Even so, in the 2018 survey, we did ask respondents to estimate the cost of downtime for individual incidents they had suffered. The first and perhaps most surprising finding is that, in cases where a significant incident did occur, 43% of respondents (n=271) did not actually calculate the cost at all. As stated, this is not best practice, if only for assessing investment decisions.

The costs that were calculated and reported should worry any CIO. While half of the reported incidents cost under $100,000, there were 39 outages that cost more than $1 million (15%). Outages, it seems, are not common, but very expensive. Around a third of outages reported by respondents cost over $250,000.

Cost ($ million)	%	Number of Incidents
<1	50.55	137
1-2	6.27	17
2-5	4.43	12
5-10	0.74	2
10-20	1.85	5
20-50	0.74	2
>50 million	0.37	1
Did not calculate	35.05	95

Uptime Institute® | RESEARCH

Uptime Institute data shows outages are common, costly, and preventable

Overall, the costs breakdown looks like this:

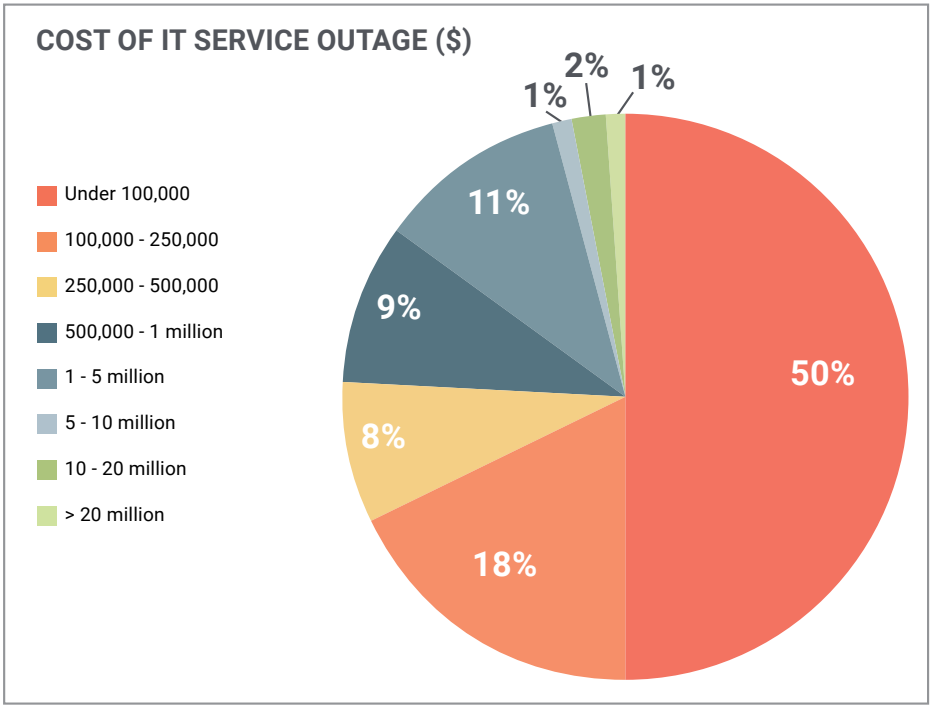

Reducing or increasing risks

In conducting this research, Uptime Institute wanted to know if certain strategies or architectures were likely to lead to more or fewer outages. We considered three: distributed resiliency, N+1 vs. 2N in power and cooling, and extending the life of (legacy) assets.

- **Distributed Resiliency.** This is often cited as the ultimate, and certainly the long-term, solution to outages. Cloud and hybrid architectures enable risks to be spread across many data centers. We found that 40% of respondents use high-availability, public cloud services, and that a similar number, 41%, say they replicate data across two or more sites. This, it seems, gives most a high degree of confidence in their solution. 60% say this has made them more resilient. However, roughly 10% believe it has not, and roughly 30% say they don't know.

- **Are 2N architectures more resilient?** By analyzing just those respondents with 2N (as opposed to N+1 architectures) for cooling and power, we are also able to see if N+1 architectures, sometimes viewed as riskier, are any less reliable.

Uptime Institute data shows outages are common, costly, and preventable

The results were clear: Of those with a 2N architecture, 22% had experienced an outage in the past year, rising to 35% in the past three years (remember that this is for all failures, not just on-site facility failures). But those with an N+1 architecture did not fare so well: 33% said they had an outage in the past year, rising to 51% in the past three years – higher than the overall average.

This response is clear and perhaps, given that 2N redundancy costs more and is designed to reduce failures, somewhat expected: 2N architectures, according the data, are more effective at preventing outages than N+1 solutions. The data, however, may change in the years ahead, as N+1 systems, aided by management software, become more effective.

- **Extending legacy asset life.** First, we asked if respondents were operating data centers beyond their expected life cycle. Of those responding (n-503), 34% said yes. Of this group, approximately three quarters (in each case) believed this practice increases the likelihood of an outage, increases operating costs, and reduces agility/introduces constraints.

Who's in charge?

Given the growing complexity of issues, especially where multiple service providers are involved, we were curious to know how organizations deal with risk/resiliency. Who is responsible for managing and assessing risk?

We found that in most cases it is the CIO/CTO. But in a few cases (7%), it is a chief risk or resiliency officer. This position may be more common in the future.

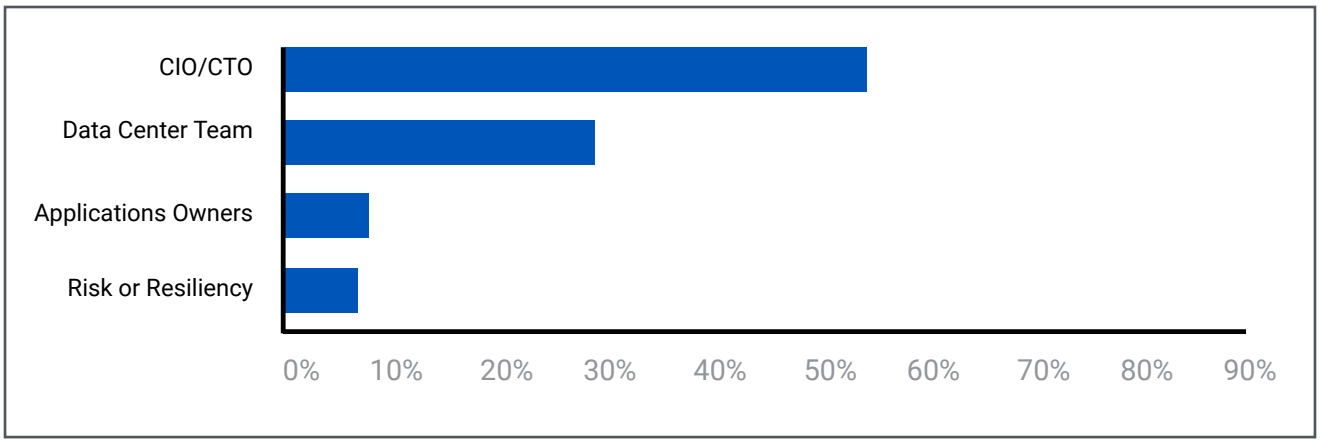

RISK & RESILIENCY

Uptime Institute® | RESEARCH

Uptime Institute data shows outages are common, costly, and preventable

Summary

There is a widespread belief that advances in IT systems and software have made IT services far more resilient, especially when coupled with highly engineered data center operations and well-drilled, process-oriented facilities staff. But this research proves that the levels of complexity and sensitivity in modern data center and IT operations, coupled with the high level of interdependence, may be working against that trend. Operators, this survey suggests, are struggling to keep pace. The result is that expensive and damaging failures keep occurring, and when they do, diagnosis and quick recovery can be challenging. Over time, experience and new technologies will no doubt lead to improvements, but diligence, investment, and planning are clearly required; it is, above all, a management issue.

SOME UPTIME RECOMMENDATIONS:

- Conduct regular system-wide resiliency analysis that spans data centers, power, cooling, connectivity, third-party services, planning, and management.
- Review failures at other organizations. Almost all failures reported or researched by Uptime Institute have happened before and are often well documented.
- Consider independent resiliency/risk analysis in the same way that your organization would carry out external security analysis.
- Understand and model the costs of downtime, whether complete or partial, which will help inform investments in resilient infrastructure.
- Ask service providers (cloud companies, applications service providers, hosting companies, colocation providers, and carriers) to provide detailed risk/resiliency reports.
- Remember that downtime costs are not necessarily unplanned. Planned downtime can also cause serious problems if not properly managed and with appropriate processes in place.
- Almost all downtime results from planning and investment decisions, coupled with poor processes or a failure to follow processes. They may therefore be termed management failures.
- When negotiating SLAs with colocation providers, include measures and multiple avenues of action that go beyond ‘standard’ contracts. These could include arbitration clauses and a limited number of unplanned events over a set period of time (rather than a limit of their cumulative duration, for example).

Note: Uptime Institute’s 2018 Data Center Survey was conducted between February and May 2018, with a total of nearly 1500 respondents. Not all the data collected in that extensive survey is in this report. Please contact Uptime Institute if you are interested in a more detailed breakdown.

Uptime Institute’s M&O Stamp of Approval for operating data centers has been shown to reduce failures (according to Uptime 2017 research). Please contact Uptime Institute for more details. More information: contact alawrence@uptimeinstitute.com.

Contact alawrence@uptimeinstitute.com

Uptime Institute® | RESEARCH

Uptime Institute data shows outages are common, costly, and preventable

NOTES