REPORT REPRINT

Is 99.99 an industry myth? Uptime Institute study shows outages are common and costly

ANDY LAWRENCE, RHONDA ASCIERTO

6 AUG 2018

New research shows that costly unplanned downtime is common, and may even be increasing. Advances in availability have been offset by complexity and interdependencies.

THIS REPORT, LICENSED TO SUNBIRD, DEVELOPED AND AS PROVIDED BY 451 RESEARCH, LLC, WAS PUBLISHED AS PART OF OUR SYNDICATED MARKET INSIGHT SUBSCRIPTION SERVICE. IT SHALL BE OWNED IN ITS ENTIRETY BY 451 RESEARCH, LLC. THIS REPORT IS SOLELY INTENDED FOR USE BY THE RECIPIENT AND MAY NOT BE REPRODUCED OR RE-POSTED, IN WHOLE OR IN PART, BY THE RECIPIENT WITHOUT EXPRESS PERMISSION FROM 451 RESEARCH.

451 Research®

©2018 451 Research, LLC | WWW.451RESEARCH.COM

In July, TSB, the UK bank owned by Sabadell of Spain, announced its half-year results. The bank swung from a profit of £108.3m ($142.17m) in the first half of 2017 to a loss of £107m in the same period for 2018. The main cause: a £176m hit caused by a botched IT system upgrade, which from April to June resulted in data corruption, security flaws, cascading failures and tens of thousands of customers experiencing problems with their accounts.

TSB's problems meant it joined a growing list of big companies that have suffered big, costly public outages or serious IT service problems in the past two years. The causes range from issues with IT systems, network equipment or connections, data power or infrastructure, and (increasingly) problems beyond their control at colocation, hosting or cloud partners. United Airlines, Delta Airlines, British Airways, OVH, AWS and Nasdaq Europe are some of the others.

These big cases make headlines, but they are far from rare, according to research published by the Uptime Institute (an independent division of The 451 Group). The research paper 'Uptime data shows outages common, costly and preventable,' which draws on a two-year study of publicly reported outages and a global survey of IT and datacenter operators, finds that failures and downtime are not only common, but 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 the ever-greater reliance on IT in every area of society and business. Distributed cloud architectures solve many problems, but can lead to software issues that may be difficult to resolve.

Uptime also found that many organizations had not properly analyzed risks and consequences of failure in their organizations, and that most failures could have been preventable (four-fifths of survey respondents said their biggest/most recent outage was 'preventable'). Often, significant investment in risk-reduction technologies and processes could be misplaced, the research finds.

THE 451 TAKE

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 datacenter operations and well-drilled process-oriented facilities staff. But this research proves that the levels of complexity and sensitivity in modern datacenter and IT operations, coupled with the high level of interdependencies, may be working against that trend. 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 likely lead to improvements, but diligence, investment and planning are required; it is, above all, a management issue.

ARE OUTAGES INCREASING?

IT equipment has become much more reliable over the past several decades, and architectures and management practices have generally improved resiliency – at the best-run organizations. This should mean that the number of downtime incidents is falling steadily.

However, Uptime has found no evidence for this. In fact, in its recent survey of about 1,000 operators, almost one-third (31%) of those responding (n=664) had experienced an IT downtime incident or severe degradation in the past year. Moreover, about half (48%) said they had experienced at least one outage either at their site or that of a service provider in the past three years.

This is a very high number, and is somewhat at odds with most publicly proclaimed availability figures (usually 99.99% or higher, which implies about 50 minutes a year of downtime, including planned maintenance). The survey data suggests that SLAs (service-level agreements) are commonly broken, and furthermore, that datacenter/IT service providers cite optimistic design-based data for resiliency, rather than realistically projecting forward data from actual records.

Uptime notes that the high profile of major providers and the ubiquity of IT today make most IT failures more visible and damaging, but it still finds evidence that failures are not declining as an overall portion of the whole. In Uptime's 2017 survey, 25% had experienced an outage during the previous 12 months – a lower figure than the 31% of 2018.

Explanations for this high rate vary. A possible reason is that the increased complexity and interdependencies of different systems and different datacenters, using increasingly complicated management systems, may be increasing the number and impact of failures (remembering that most datacenters handle more work each year). There has also been evidence, in earlier research by Uptime, that failures tend to occur most commonly during periods of technology change and investment, and secondly, at sites where there has been underinvestment and legacy assets are not upgraded. Many sites fall into one of these two categories. The move to cloud may be resulting in underinvestment of enterprise infrastructure.

Management expertise and attention play a significant role: In an earlier study, Uptime found that, among members of the Uptime Institute Network, where expertise about design, operations and resiliency is shared among managers and engineers, the amount of downtime was halved.

POWER STRUGGLE

What is the main cause of IT service downtime? Uptime's survey data shows that the biggest single cause of failures is the one that the industry has always invested so much in to prevent – a loss of on-site power. This was cited 93 times by 285 respondents who had 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 group that is getting more of the blame for outages: other organizations. When all the various failures at third-party service providers (colocation, hosting or cloud) were aggregated, it accounted for 87 of the 285 incidents reported (31%), only slightly fewer than primary on-site power failures under local control. Third-party failures have become a critical issue, and a particularly troubling one, because CIOs and service providers have limited visibility into, and limited control over, third parties.

FIGURE 1

Source: 451 Research, LLC

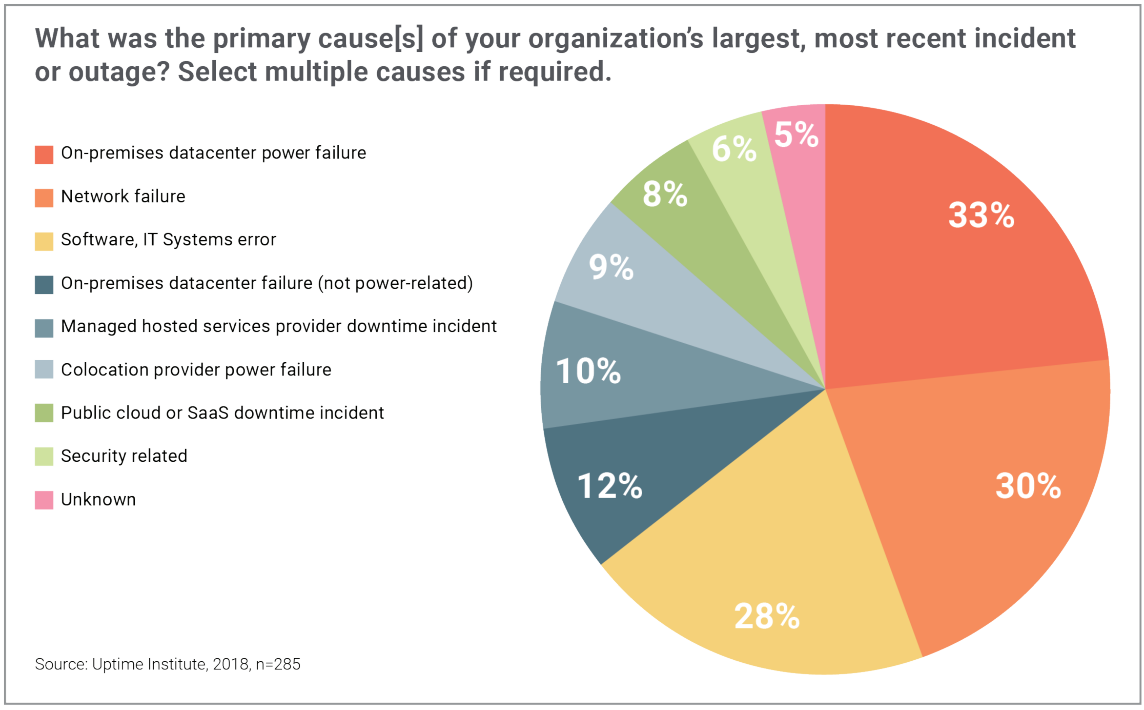

This survey data is closely aligned with the causes of big, public failures that Uptime has tracked since early 2016. In these big incidents, reported in the media and social media, power accounted for 36% of failures, followed by 25% for network issues and 22% for IT/software issues.

FIGURE 2

Source: 451 Research, LLC

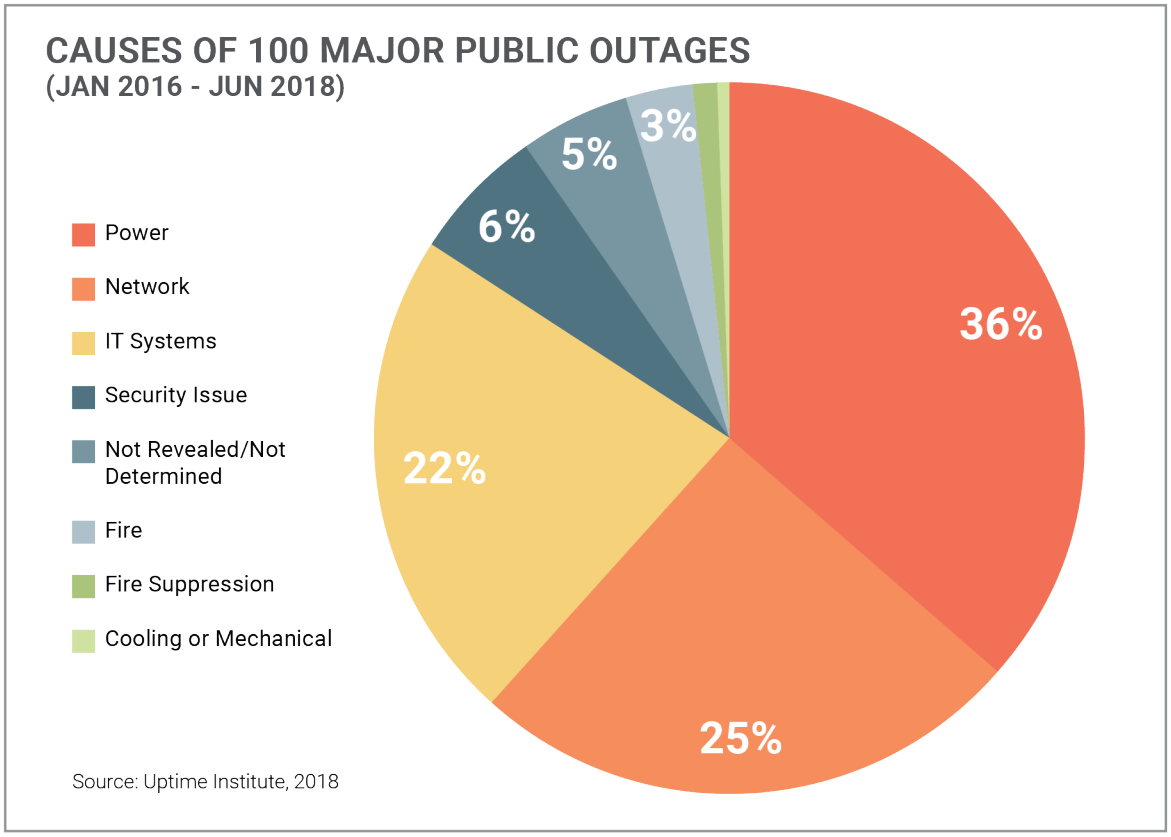

PLEASE TRY AGAIN LATER

From January 2016 to June 2018, Uptime recorded 10 outages that were 'extremely serious' and fell into the highest category of outage, which represents a material loss of revenue (and brand damage), or in extreme cases a threat to the sustainability of the business of the operator or its clients. Among these big failures, power and datacenter facilities problems are the biggest/most common cause of initial problems, but IT and network problems tend to cause most issues because of the interdependencies and complexities of such systems, relative difficulty in fault analysis, and longer recovery times.

COSTS OF DOWNTIME

Estimating the average cost of downtime is difficult and controversial because impacts vary widely. However, for its 2018 survey, Uptime asked 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 41 outages that cost more than $1m (15%). About one-third of outages reported by respondents cost over $250,000.

Uptime Institute helps organizations improve their IT and datacenter resiliency. Among its research programs, it tracks public outages, and conducts research on failures and outages. Data from Uptime's Abnormal Incidents Reporting Service (AIRS), a tracking system for Uptime Institute Network members to share and learn from datacenter infrastructure incidents in detail, was not used for this research for confidentiality reasons.