The Changing Landscape Of IT Incident And Crisis Management

Safety-Critical Practices And Collaborative Platforms Are The New Normal

by Charles Betz

February 16, 2018 | Updated: February 20, 2018

Why Read This Report

Digital services are ever more critical to human well-being. Customers, as well as other stakeholders, are increasing the pressure on providers to respond quickly and effectively to operational disruptions. Digital transformation means that computer failures can threaten life and safety, as the WannaCry ransomware's impact on the UK's National Health Service (NHS) demonstrated. Infrastructure and operations (I&O) professionals have much to learn from large web-scale properties' adoption of practices from safety-critical domains like fire and police services and their use of chat platforms to supplement traditional ticketed workflows.

Key Takeaways

Digital Incidents Require Fast, Disciplined, Coordinated Responses
Identifying an incident condition, assembling a team, and tracking the response to the incident remain critical practices for digital organizations. Responding capably to an incident requires frictionless, rapid dispatch and close coordination.

Digital Managers Are Learning From Safety-Critical Practices
Web-scale properties have found that incident management practices from fire and police services are valuable in a digital context. The influence of these practices continues to spread.

Organizations Are Turning To Chat-First, Highly Integrated Collaborative Platforms
Chat platforms are increasingly important. These platforms contribute a dynamic, low-bandwidth mechanism that integrates with other systems in various ways, including direct interaction with operations such as ChatOps.

FOR INFRASTRUCTURE & OPERATIONS PROFESSIONALS

The Changing Landscape Of IT Incident And Crisis Management

Safety-Critical Practices And Collaborative Platforms Are The New Normal

by Charles Betz
with Eveline Oehrlich, Chris Gardner, Robert Stroud, Josh Zelonis, Julia Caldwell, and Diane Lynch

February 16, 2018 | Updated: February 20, 2018

Table Of Contents

2 Complex Systems Fail

Outages And Interruptions Are Costly And Catastrophic

Traditional ITSM Approaches Are Only Part Of The Continuous Resolution Puzzle

There's More To Incident Management Than Ticketing

6 Continuous Resolution Needs Integrated Platforms And Practices

Recommendations

7 Update Your Incident And Crisis Response For DevOps

What It Means

8 The Rise Of Chaos Engineering

9 Supplemental Material

Related Research Documents

Change Management: Let's Get Back To Basics

Reshape Your Application Support For Digital Operations

Vendor Landscape: Cognitive Operations

FORRESTER®

Forrester Research, Inc., 60 Acorn Park Drive, Cambridge, MA 02140 USA
+1 617-613-6000 | Fax: +1 617-613-5000 | forrester.com

Complex Systems Fail

Automation, web-scale engineering, and continuous delivery have all improved the stability of digital systems. But organizations inevitably tend to push these systems to their limits — and as digital systems become more pervasive and critical, they also become a major point of vulnerability. Experts in complex systems failure and human response emphasize that completely preventing failure is an unrealistic goal.[1] Exacerbating the challenge is the fact that it's rare for companies to entirely replace old systems with new ones; instead, they layer complexity upon existing complexity, making it ever more difficult for I&O pros to diagnose problems.

Outages And Interruptions Are Costly And Catastrophic

In the past three years, organizations as diverse as Barclays, JetBlue, the New York Stock Exchange, the Royal Bank of Scotland, Southwest Airlines, and Verizon have endured well-publicized outages.[2] Digital disruptions at this scale increasingly mean operational losses, brand damage, and even increasing concerns for public health and safety. Victims in the past two calendar years alone include:

- **Delta Air Lines.** In August 2016, Delta was forced to cancel 2,300 flights when a single piece of malfunctioning equipment caused a power outage at its operations center in Atlanta.[3] The problem was compounded by insufficient protections, which led to a complete shutdown of operations at the airline's headquarters. Total cost to the company: a reported $150 million.
- **GitLab.** In January 2017, a tired systems administrator at GitLab accidentally deleted hundreds of gigabytes of customer data.[4] Multiple redundant backup protocols failed, and the firm was ultimately unable to recover its production data. GitLab lost customer projects, comments, and other data; although its core source code repositories weren't affected, the failure is particularly troubling for a company in the business of stewarding such data.
- **Amazon S3.** In February 2017, an operator incorrectly entered the input to a runbook that had no safety controls, resulting in a 4-hour disruption to Amazon's core storage services. This had multiple impacts on significant web properties, including Alexa, IFTTT, Quora, and Trello.[5]
- **FedEx and the NHS.** In May 2017, WannaCry ransomware encrypted computers around the world. The UK's NHS suffered operational disruptions, one of many crises the malware caused. The cyberattack hindered urgent NHS services by blocking access to its computers, locking out vital medical equipment such as MRI scanners and devices for testing blood and tissue samples, and forcing some hospitals to divert ambulances to other locations.[6] FedEx was another victim; one FedEx subsidiary reported $300 million in losses resulting from the ransomware.[7]

Traditional ITSM Approaches Are Only Part Of The Continuous Resolution Puzzle

Traditional ITIL-based IT service management (ITSM) tends to assume that there's a hard boundary between the development of digital systems and the deployment, operation, and support of those systems. Development and operations (DevOps) approaches integrate operational teams into the development life cycle and foster a team approach to addressing defects — although many organizations find that these two sides aren't well aligned when it's time for major incident remediation.[8] If I&O teams are to fulfill their responsibility for speedy incident remediation, they must bridge this gap by embracing their role in continuous resolution (see Figure 1). Forrester suggests that I&O pros view their tasks as falling into three major areas:

- **Service support.** These day-to-day tasks include daily operations and management, the incorporation of event management and telemetry, and the maintenance of key digital management systems and personnel data such as assignments and schedules. It also includes ongoing resolution of noncritical incidents and activities such as end user provisioning and support.
- **Proactive remediation.** Bad things happen, so I&O pros need to prepare automated operational responses; correlate performance analytics to predict and avoid incidents; support preventive assessments, drills/game days, and chaos engineering; and use postmortems to improve systems. These processes should include a feedback loop from operations to development.
- **Rapid response.** To gear up and fight the fire, organizations must be able to quickly bring a distributed team together and correlate multiple sources of information to restore service or prevent an outage. Discipline, culture, and standard response protocols are essential. Designate an incident commander, and ensure that responders coordinate their activities through that person. Responders also need tools that don't bog them down but enable them to carry out their duties and track their discussions and actions transparently.

FIGURE 1 Urgency Distinguishes Incident And Crisis Practices

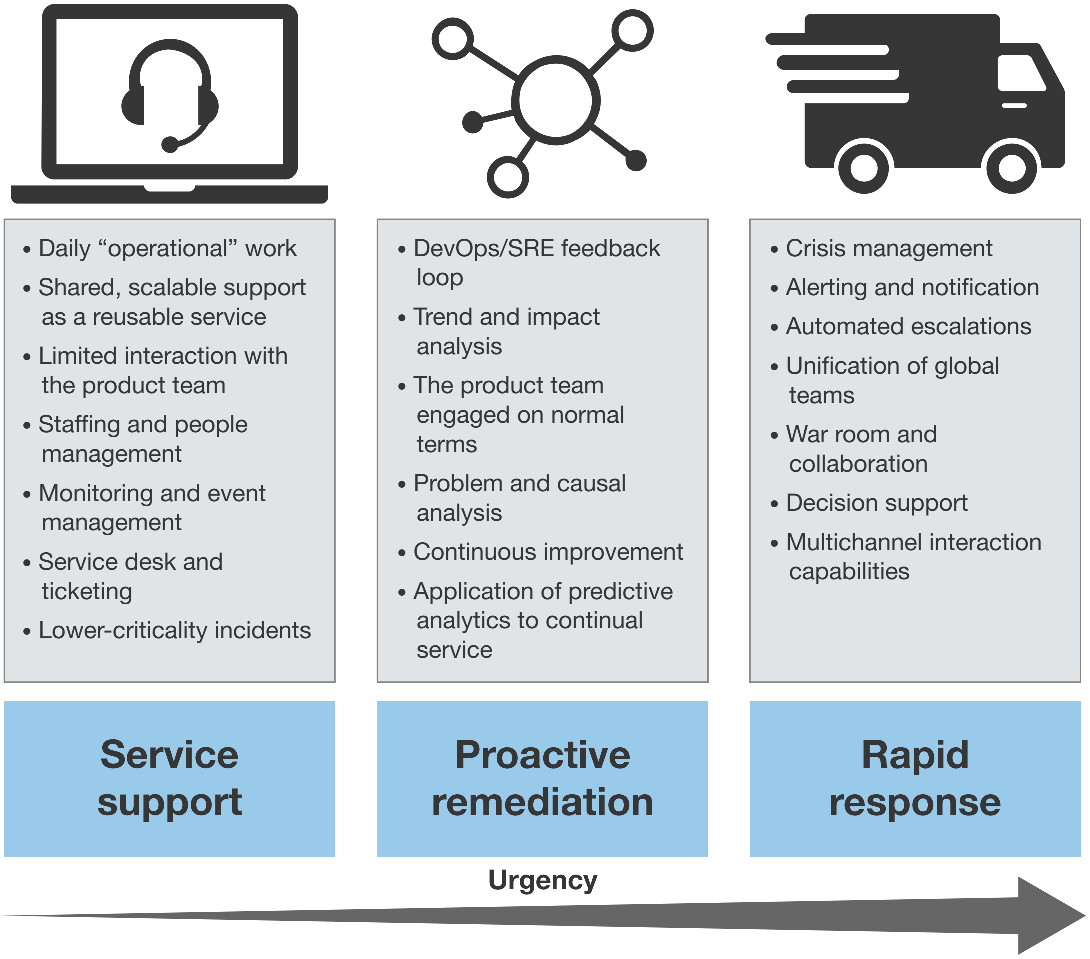

- Daily "operational" work
- Shared, scalable support as a reusable service
- Limited interaction with the product team
- Staffing and people management
- Monitoring and event management
- Service desk and ticketing
- Lower-criticality incidents

Service support

- DevOps/SRE feedback loop
- Trend and impact analysis
- The product team engaged on normal terms
- Problem and causal analysis
- Continuous improvement
- Application of predictive analytics to continual service

Proactive remediation

- Crisis management
- Alerting and notification
- Automated escalations
- Unification of global teams
- War room and collaboration
- Decision support
- Multichannel interaction capabilities

Rapid response

There's More To Incident Management Than Ticketing

Modern organizations are moving away from the traditional ITIL focus on incident management as a process that emphasizes a high level of consistency in how incidents are captured and classified.[9] Instead, the concern is for the collaborative experience of assembling and coordinating a team. We interviewed multiple organizations that indicated that crisis management requires maintaining discipline and focus during high-stress interactions while ensuring that communications and collaboration are appropriate and friction-free. Over the past decade, leading digital organizations like Amazon and Etsy have developed a new school of thought about how best to manage critical situations in digital operations:

- **Target both speed and quality.** The ITSM architect of a major agricultural products company that uses Everbridge told us, “Our business wants and expects a faster and more flexible, scalable, and available set of technologies.” Major incidents are more disruptive and more likely to directly affect revenue and brand.

- **Adopt and learn from the field of safety science.** Human factors and the safety sciences are increasingly shaping thinking about how to manage today’s technology incidents and crises. Industry leaders like former Etsy CTO John Allspaw are informing their thought by bringing in academically rigorous discussions on joint action. Allspaw observes that, while incident management in digital operations is understudied, there’s much applicable safety science research on the subject of “teams engaging in understanding and resolving anomalies under high-tempo, high-consequence conditions such as healthcare, aviation, space operations, and the military.”[10] Prominent safety science experts, including Dr. Sidney Dekker, appeared on a high-profile panel at the 2017 DevOps Enterprise Summit and called on business technologists to adopt a higher level of professionalization.[11]

- **Learn from the incident management system (IMS) protocol.** In 1970, a number of catastrophic wildfires highlighted poor coordination among fire services.[12] US government agencies created the IMS as an “all-hazard, all-risk framework designed specifically for emergency operations” that stipulates common terminology, standard roles and practices, and behavioral and cultural expectations. In the mid-2000s, Jesse Robbins, Amazon’s “master of disaster,” introduced the IMS concept as the basis of Amazon’s incident response, and its formal and informal influence on digital operations continues to grow.[13] PagerDuty, a vendor focused on automating incident response, brings presenters to its regional summits to talk about incident command so practitioners can learn and adopt best practices.

- **Reduce the mean-time-to-assemble.** Marshaling and dispatching a team of responders is the only activity that the incident response team actually controls.[14] Modern incident management relies heavily on time-sensitive alerting and acknowledgement, which requires appropriate communication platforms such as PagerDuty or VictorOps. But many companies don’t clearly distinguish between notifications and dispatches — ambiguity that can have catastrophic results.[15]

- **Distinguish between internal and external communication.** Site reliability engineers (SREs) from CircleCI stress the importance of having an incident communication role distinctly isolated from and responsible for updating customers and thinking about the big picture. DreamLab, an IT hub for digital media in Central Europe and a part of the Ringier Axel Springer Media family, uses StatusPage for external communications and VictorOps for internal team collaboration during incidents.

- **Cue the multicausal, blameless postmortem after the dust settles.** Safety science calls for a systematic, thorough, and clinical examination of incidents. Rather than holding someone “accountable,” experts in human factors and safety like Dekker emphasize the need to examine the system that leads to a given individual’s actions. Human error has little or no place in safety science thinking, which focuses on the context giving rise to the error. As a product lead from VictorOps noted, the question is, “What made the engineer think that this was the right thing to do at the time?” Leading thinkers also deprecate the concept of a “root” cause; multiple factors typically cause complex systems failures.[16]

- **Treat postmortem outcomes as an equally important product demand.** Fixing the causes of incidents shouldn’t take a back seat to new product functionality. Leading firms implement a tight feedback loop from the postmortem to product backlog management. For example, from its database outage, GitLab identified 14 different issues requiring attention.[17] Both Statuspage customer CircleCI and VictorOps customer Skyscanner route any identified problem-solving tasks back to the master product backlog in Atlassian’s Jira Software.

Continuous Resolution Needs Integrated Platforms And Practices

Continuous resolution automation requires a suite of tools and capabilities. It’s rare that just one vendor can provide it, so integration is typical. Most of the customers that Forrester interviewed are using next-generation continuous resolution platforms such as BigPanda, PagerDuty, and xMatters for alerting, acknowledgment, and chat, while still relying on the ticketing and workflow of core ITSM systems such as Cherwell Software and ServiceNow. To automate intelligently and successfully, I&O teams must:

- **Let tickets fade into the background.** Tickets are cumbersome, especially when they’re created and updated manually. A PagerDuty product manager reported that some customers ticket only “major” incidents. Incident responders need to focus on solving problems, not on authenticating to multiple tools, navigating unwieldy user interfaces, and updating numerous data fields. Asked why they’re using tickets at all, SREs from Skyscanner identified consistent reporting and metrics as a key driver. “These big companies going through DevOps transitions, they’re working on hybrid environments, with one harmonized platform for incidents,” says a product manager at Big Panda. “People need to be able to quickly identify what’s going on and share notes with teammates, integrated with email, text, [and] Slack as well as smart ticketing solutions like ServiceNow and Jira.”

- **Support ChatOps — it’s ready for prime time.** Operations staff at ViaSat told us they’re “trying to remove the requirement for any resource working on an incident to interact with the ticket directly.” Instead, digital incident responders prefer chat platforms like HipChat and Slack, increasingly enhanced with direct operational integration via ChatOps.[18] ChatOps allows responders to use simple chat commands to instantly raise and update an official ticket. It can automatically attach documentation of discussions and actions from the chat channel to a traditional incident ticket, saving the toil of manual updates. Start with chat platforms such as HipChat and Slack, and integrate ChatOps plug-ins like Hubot to support your resolution process.

- **Evaluate smart chatbots and cognitive search solutions.** Chatbots bring natural language processing and real-time recommendations in the guise of virtual agents, accessible through instant messaging. Cognitive search solutions leverage artificial intelligence (AI) technologies such as natural language processing and machine learning to ingest, understand, organize, and query digital content from multiple data sources. Companies such as Moogsoft and Squirro are at the forefront of applying new cognitive and heuristic automation technologies to incident response and broader digital management objectives.[19]

Recommendations

Update Your Incident And Crisis Response For DevOps

As Christoph Goldenstern, an expert on problem management, incident postmortems, and causal management from Kepner-Tregoe, told Forrester, "You can't perform continuous improvement and integration in an atmosphere of continual firefighting." Ticketing and monitoring are table stakes; I&O pros need both, but in a broader context that also includes alerting, response, analytics, and frictionless collaboration. To evolve your incident and crisis response:

› **Evaluate the incident management system approach.** Train key staff as incident commanders and strive for a disciplined approach that 1) clearly distinguishes notification from dispatch; 2) ensures that staff agree upon problem-solving actions before attempting any; 3) designates a scribe to keep records; 4) uses formalized and consistent communication channels; and 5) protects individual responders from executive or media interference. Conducting and communicating the results of blameless postmortems is a well-accepted best practice that leading firms of all sizes employ. Every customer interviewed for this report indicated that this was part of its standard operating procedure for major incidents. Amazon has provided a prominent example with its public postmortems, emulated by many other digital firms.[20]

› **Implement and optimize ChatOps.** The use of chat platforms like Slack is widespread among organizations we speak with. However, the maturity of their usage varies. Fully leveraging ChatOps and chatbots is a journey. As a best practice, ensure that chat sessions related to an incident are readily accessible from that incident's record to enable postmortem analysis. While standalone chat platforms currently rule the day, more integrated solutions such as Atlassian, CA Technologies, and ServiceNow make sense. These ITSM vendors have already boasted messaging capabilities within their tools. For example, Atlassian's portfolio includes HipChat, Jira Service Desk, Jira Services, and StatusPage, covering chat, ITSM workflow, and product backlog management.

› **Carefully design which responsibilities to centralize.** It's currently fashionable to hand over more support responsibilities to product teams.[21] But at some point, it doesn't make sense to divert scarce engineering resources to resolving customer support issues or repeatable "known error" incidents (not all of which can be automated away). xMatters customer ViaSat is seeking to move away from using a centralized network operations center for triage, instead routing incidents automatically to high-level business domain owners. VictorOps customer Skyscanner pushes specific issues to feature teams but also maintains a global SRE team to coordinate responses across those groups as needed.

› **Experiment with different support models.** Try out an alternative to traditional level 1/level 2/level 3 support: swarming. In swarming, product teams respond to incidents dynamically, using less formalized routing and escalation.[22] According to its ITSM architect, an agricultural products firm "uses traditional tiered escalation for low-impact incidents but swarms for major incidents." On the other hand, ViaSat tried swarming and found that it often caused more problems than it solved, as too many people tried to work the same aspects of the incident and communication of activity taking place was inconsistent. Your mileage may vary, but you should try this alternative to the traditional level 1/level 2/level 3 support model.

› **Adopt advanced operational analytics.** Analytics vendor Squirro provides a command risk heat map. It assigns a risk profile to each command — such as move, copy, or remove — and categorizes it by the level of risk. The heat map shows if someone did something risky, who it was, and when it occurred, which speeds up remediation.

› **Prepare for a new generation of diagnostic tools.** Thought leaders increasingly cite the control theory concept of observability: How well can a system be understood from its external signals? Companies like DataDog and honeycomb.io offer tools to assist in the increasingly dynamic problem of analyzing failures across complex distributed systems. The CEO of honeycomb.io, Charity Majors, compares the resulting challenge to "searching for a needle in a planet-sized haystack."[23] Some customers report interest in distributed tracing frameworks like Envoy and Zipkin. Whatever the tool, start preparing now for these new operational challenges.

What It Means

The Rise Of Chaos Engineering

It's becoming more and more difficult to analyze severe crises in digital systems. Technology is increasingly advanced, resilient, and stable, and advanced practices like containers and infrastructure-as-code minimize the occurrence of simple failure modes, such as a corrupted package requiring reinstallation. Instead, failure modes become more complex to understand and troubleshoot, as they arise from interactions among multiple distributed system components (e.g., microservices). Ultimately, it's impossible to replicate a complex production system in full. Therefore, some I&O leaders are turning to the practice of injecting controlled amounts of failure into production infrastructure. Netflix's well-known Chaos Monkey randomly terminates production processes to find services that are insufficiently hardened to failure. Netflix's practices have inspired a new field of chaos engineering, intended to help "identify weaknesses before they manifest in systemwide, aberrant behaviors" by continually experimenting on real production systems.[24] Forrester is tracking the potential of chaos engineering to mitigate some of the risks of increasingly intelligent and automated systems.

Supplemental Material

Companies Interviewed For This Report

We'd like to thank the individuals from the following companies who generously gave their time during the research for this report.

AlertOps

Atlassian

BigPanda

Cherwell Software

CircleCI

Everbridge

IBM

Ivanti

Kepner-Tregoe

Micro Focus

Moogsoft

Nexthink

OpsGenie

PagerDuty

ServiceNow

Skyscanner

Squirro

ViaSat

VictorOps

VMware

xMatters

Endnotes

[1] This is currently a hot topic in the DevOps community. Here are some of the pieces currently making the rounds. Source: Richard I. Cook, “How Complex Systems Fail,” Cognitive Technologies Laboratory, University of Chicago, 2000 (http://web.mit.edu/2.75/resources/random/How Complex Systems Fail.pdf); Sidney Dekker, *Drift into Failure: From Hunting Broken Components to Understanding Complex Systems*, CRC Press, 2011; and “STELLA: Report from the SNAFUcatchers Workshop on Coping With Complexity,” SNAFUcatchers, March 14-16, 2017 (https://drive.google.com/file/d/0B7kFkt5WxLeDTml5cTFsWXFCb1U/view).

Source: “DOES17 San Francisco — Convergence Of Safety Culture And Lean: Lessons From The Leaders,” YouTube video, November 30, 2017 (https://www.youtube.com/watch?time_continue=1&v=CFMJ3V4VakA).

[2] Source: G. Clay Whittaker, “Network Outages Like NYSE, United Airlines, Are The New Natural Disasters,” Popular Science, July 11, 2015 (https://www.popsci.com/network-outages-nyses-united-airlines-are-new-natural-disasters); Alexandra Zaslow, “Outdated Technology Likely Culprit in Southwest Airlines Outage,” NBC News, October 12, 2015 (https://www.nbcnews.com/business/travel/outdated-technology-likely-culprit-southwest-airlines-outage-n443176); Matthew Finnegan, “Bank IT meltdowns: Bank of England must do more to prevent bank IT failures, says Treasury Committee chair Andrew Tyrie,” ComputerworldUK, January 25, 2016 (https://www.computerworlduk.com/applications/hsbc-outage-is-first-bank-it-fiasco-of-2016-but-unlikely-be-last-3632901/); David Chernicoff, “Verizon data center failure causes JetBlue air travel delays,” Datacenter Dynamics, January 15, 2016 (http://www.datacenterdynamics.com/content-tracks/security-risk/verizon-data-center-failure-causes-jetblue-air-travel-delays/95538.fullarticle); and Scott Carey, “Amazon email review: AWS WorkMail vs Gmail vs Outlook - Amazon’s WorkMail app is just like every other email service on the web,” ComputerworldUK, January 21, 2016 (https://www.computerworlduk.com/applications/amazon-web-services-moves-into-email-but-is-it-any-good-3633814/).

[3] Source: Chris Isidore, “Delta: 5-hour computer outage cost us $150 million,” CNNTech, September 7, 2016 (http://money.cnn.com/2016/09/07/technology/delta-computer-outage-cost/index.html).

[4] Source: “Postmortem of database outage of January 31,” GitLab, February 10, 2017 (https://about.gitlab.com/2017/02/10/postmortem-of-database-outage-of-january-31/).

[5] Source: “Summary of the Amazon S3 Service Disruption in the Northern Virginia (US-East-1) Region,” Amazon Web Services (https://aws.amazon.com/message/41926/) and Jacob Kastrenakes, “Amazon’s web servers are down and it’s causing trouble across the internet,” The Verge, February 28, 2017 (https://www.theverge.com/2017/2/28/14765042/amazon-s3-outage-causing-trouble).

[6] Source: Owen Hughes, “WannaCry Impact on NHS considerably larger than previously suggested,” Digital Health, October 27, 2017 (https://www.digitalhealth.net/2017/10/wannacry-impact-on-nhs-considerably-larger-than-previously-suggested/).

[7] Source: Scott Shane, Nicole Perlroth, and David E. Sanger, “Security Breach and Spilled Secrets Have Shaken the N.S.A. to Its Core,” The New York Times, November 12, 2017 (https://www.nytimes.com/2017/11/12/us/nsa-shadow-brokers.html).

[8] For more information on product teams, see the Forrester report "Organize And Staff I&O Pros For Successful DevOps Practices."

[9] Source: ITIL Service Operation, The Stationery Office, 2011. Note the information on pages 73-74, which strongly emphasizes standardization and consistency and doesn't speak to the collaborative aspects of incident resolution.

[10] Source: John Allspaw, Trade-Offs Under Pressure: Heuristics and Observations Of Teams Resolving Internet Service Outages, Lund University, 2015.

[11] Source: "DOES17 San Francisco — Convergence Of Safety Culture And Lean: Lessons From The Leaders," YouTube video, November 30, 2017 (https://www.youtube.com/watch?time_continue=1&v=CFMJ3V4VakA).

[12] Source: Rob Schnepp, Ron Vidal, and Chris Hawley, Incident Management for Operations, O'Reilly Media, 2017.

[13] Source: Rob Schnepp, Ron Vidal, and Chris Hawley, Incident Management for Operations, O'Reilly Media, 2017.

[14] Source: Rob Schnepp, Ron Vidal, and Chris Hawley, Incident Management for Operations, O'Reilly Media, 2017.

[15] Source: Rob Schnepp, Ron Vidal, and Chris Hawley, Incident Management for Operations, O'Reilly Media, 2017.

[16] Some of the recent dismissal of "root cause" is overdone; seasoned investigators of complex systems failures have known for decades that a single "root" cause is unusual. Multiple "roots" are more common. Source: Sidney Dekker, The Field Guide to Understanding Human Error, Ashgate Publishing Company, 2006.

[17] Source: David Iffland, "GitLab.com Postmortem Digs into Root Causes of 18 Hour Outage," InfoQ, February 21, 2017 (https://www.infoq.com/news/2017/02/gitlab-outage-postmortem).

[18] See the Forrester report "ChatOps: The Missing Ingredient To Improve Customer Experience" and see the Forrester report "Cracking The Collaboration Conundrum: Accelerate Customer Focus With ChatOps."

[19] See the Forrester report "The Forrester Wave™: Cognitive Search And Knowledge Discovery Solutions, Q2 2017."

[20] Source: "Summary of the Amazon S3 Service Disruption in the Northern Virginia (US-East-1) Region," Amazon Web Services (https://aws.amazon.com/message/41926/).

[21] See the Forrester report "Reshape Your Application Support For Digital Operations" and see the Forrester report "Organize Your App-Dev Teams With Agile And DevOps."

[22] Source: Jon Hall, "ITSM, DevOps, and why three-tier support should be replaced with Swarming." Medium, December 17, 2016 (https://medium.com/@JonHall_/itsm-devops-and-why-the-three-tier-structure-must-be-replaced-with-swarming-91e76ba22304).

[23] Source: "'Observability for Emerging Infra: What Got You Here Won't Get You There' by Charity Majors," YouTube video, October 2, 2017 (https://www.youtube.com/watch?v=1wjovFSCGhE&t=1553s).

[24] Some companies report success with using chaos engineering principles in QA environments as well, when those environments have a high degree of fidelity to production and when load testing is well automated. Source: Principles of Chaos Engineering (http://principlesofchaos.org/).

FORRESTER®
CHALLENGE THINKING. LEAD CHANGE.

We work with business and technology leaders to develop customer-obsessed strategies that drive growth.

PRODUCTS AND SERVICES

- › Core research and tools
- › Data and analytics
- › Peer collaboration
- › Analyst engagement
- › Consulting
- › Events

Forrester's research and insights are tailored to your role and critical business initiatives.

ROLES WE SERVE

Marketing & Strategy Professionals

CMO

B2B Marketing

B2C Marketing

Customer Experience

Customer Insights

eBusiness & Channel Strategy

Technology Management Professionals

CIO

Application Development & Delivery

Enterprise Architecture

› Infrastructure & Operations

Security & Risk

Sourcing & Vendor Management

Technology Industry Professionals

Analyst Relations

CLIENT SUPPORT

For information on hard-copy or electronic reprints, please contact Client Support at +1 866-367-7378, +1 617-613-5730, or clientsupport@forrester.com. We offer quantity discounts and special pricing for academic and nonprofit institutions.

Forrester Research (Nasdaq: FORR) is one of the most influential research and advisory firms in the world. We work with business and technology leaders to develop customer-obsessed strategies that drive growth. Through proprietary research, data, custom consulting, exclusive executive peer groups, and events, the Forrester experience is about a singular and powerful purpose: to challenge the thinking of our clients to help them lead change in their organizations. For more information, visit forrester.com.