



2019 SRE Report

Table of contents

Acknowledgements / Executive Summary	Page 3
Key Findings Summary	Page 4
Survey demographics and firmography	Page 6
Key Finding 1: SRE still emerging practice	Page 8
Key Finding 2: Incident Resolution massive part of job	Page 13
Key Finding 3: Resolving incidents is stressful	Page 16
Key Finding 4: Team support reduces stress	Page 20
Methodology	Page 22

Executive Summary

For the second year, we surveyed Site Reliability Engineers or those that identify as an SRE to understand more about this emerging role. Last year we focused on who SREs are, where they work, what they do, and how they do it. The 2018 report explored the skills, toolset used, and the corporate culture to determine if there is a core set of principles across teams and organizations.

This year's survey examined team structure, outages, incidents, and post-incident stress. We looked to answer the question of "What impact do incidents have on organizations and the people responding to them?" Organizations are focused on building resilient systems and recovering quickly, but does this focus extend to employee resilience and recovery from post-incident stress?

The 2019 report analyzed responses from 188 SREs globally across a range of industries and company sizes. This report provides a unique view of trends and issues facing site reliability engineers and the organizations that employ them.

Acknowledgements

This report would not exist without the contribution and support of many people inside and outside of Catchpoint.

The inspiration for this year's survey came from a talk given by Jaime Woo from Incident Labs at SREcon. Jaime was instrumental in defining and refining the questions on post-incident stress. Thank you for the inspiration and assistance.

Not being an SRE myself, I wanted to make sure the questions made sense and weren't missing important pieces. Seth Vargo from Google and Liz Fong-Jones from Honeycomb provided valuable feedback on the types of questions to ask.

Phrasing survey questions to get meaningful answers is harder than it looks. Nicole Forsgren from Google gave me guidance on avoiding loaded questions/answers, providing clarification and definitions, and using best practices for survey questions.

Thank you to the Catchpoint employees that are instrumental in the creation of this report:

- Taylor Meluch for her design wizardry.
- Kayla Lee for her detailed editorial eye.
- Peter Saulitis for his guidance and insights on branding.
- Sarah Sanders for handling the logistics, communications, and overall program management.

Finally, this report would not be possible without those that took the time to respond to and share the survey. We took a bit of a risk asking people to share personal information on how they handle post-incident stress. Thank you for sharing your thoughts and feelings.

Key finding 1

Site Reliability Engineering is still **emerging** as a practice.

64% of respondents indicate the SRE role or team has been in existence for three years or less.

Key finding 2

Incident **resolution** is a massive part of the job.

49% of respondents indicated they had worked on an incident in the last week.

Key finding 3

Resolving incidents is **stressful**.

79% of respondents have stress.

Key finding 4

A **supportive** team reduces post-incident stress.

67% of SREs who feel stress after every incident do not believe their company cares about their well-being.



Survey demographics & firmographics

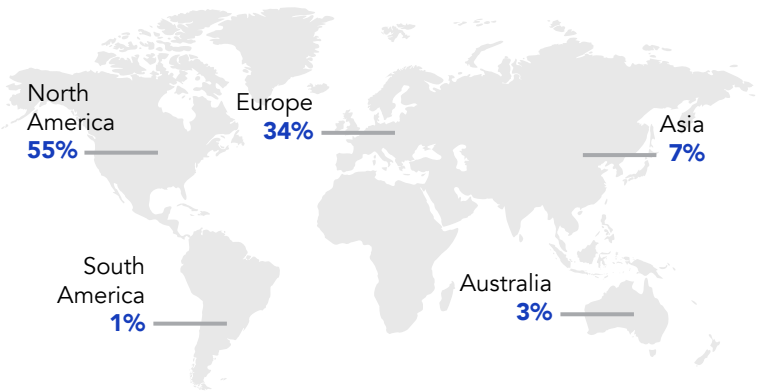
Industry

62% of respondents work in technology related industries

14% of respondents work in retail/consumer ecommerce

No other industry had more than 5% of respondents. As a result, we did not analyze whether industry impacts an answer.

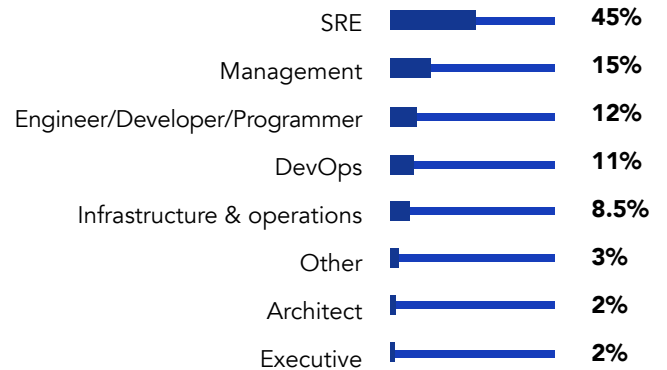
Geography



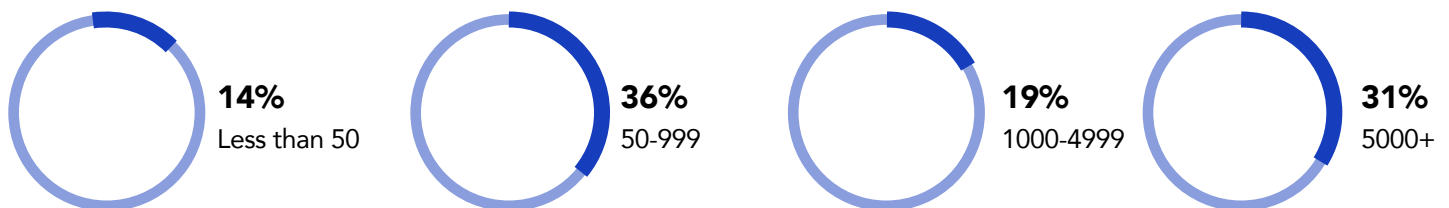
Role

A wide variety of titles are used to refer to people doing SRE work – 45% had a title of SRE. But the remainder self-identified as doing SRE work. When including management with an SRE title (SRE Manager, SRE Director, etc.), the percentage increased to 49%.

29% held senior positions (this includes people with the word lead, architect, or senior in their title) 16% are in leadership positions (manager, director, VP, or executive). The remaining are junior or mid-level.



Size of organization



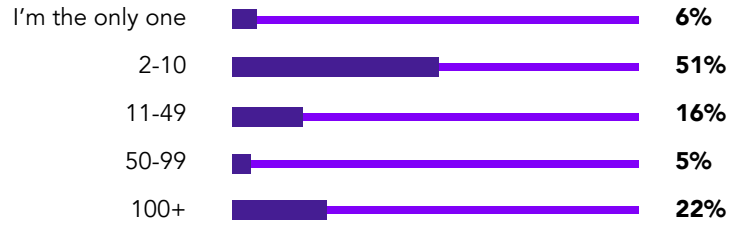
Key finding **1**

SRE is new and still being defined

SREs

How many SREs are in your organization?

The majority of SREs work in organizations with fewer than 10 SREs. 6% are the sole SRE in the organization.



How long has the SRE team been in existence?

The SRE concept, while in existence for over 15 years, is still in its infancy.



How was the SRE team built?

Given that SRE is relatively new, we were interested in how the team/role came into being across organizations. Some answers sounded similar, but there were nuances to them. For example, renaming a team is not the same as evolving a team or it growing organically.

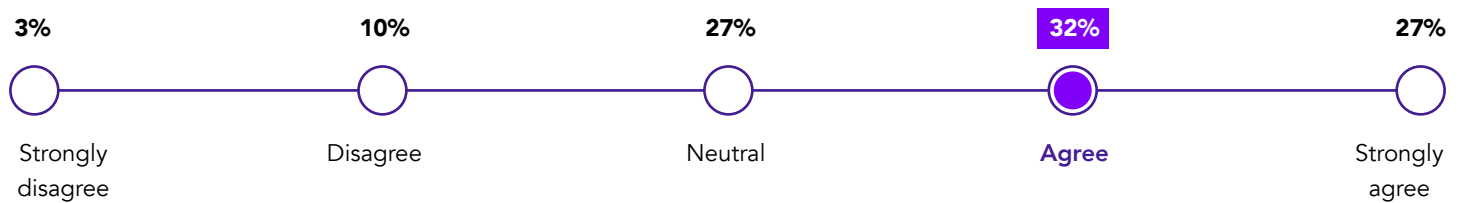
31% of respondents said it grew organically

- 29%** Evolved from operations/systems administration
- 13%** We renamed an operations/engineering/system administration team the SRE team
- 13%** Select people were chosen for a team
- 9%** Executive sponsor said "we are now doing SRE"
- 2%** We hired junior level people and trained them

Impact of toil

Toil is manual, repetitive, automatable, tactical work that scales linearly and is the main source of concern for SREs. 59% believe there is too much toil in their organization and not enough has been automated to reduce that toil. Nobody strongly agreed with the statement “We have used automation to reduce toil” while 48.5% disagreed or strongly disagreed. The two main sources of toil for SREs are in maintenance tasks and non-urgent, service-related messages. Those maintenance tasks are an automation opportunity to help reduce toil.

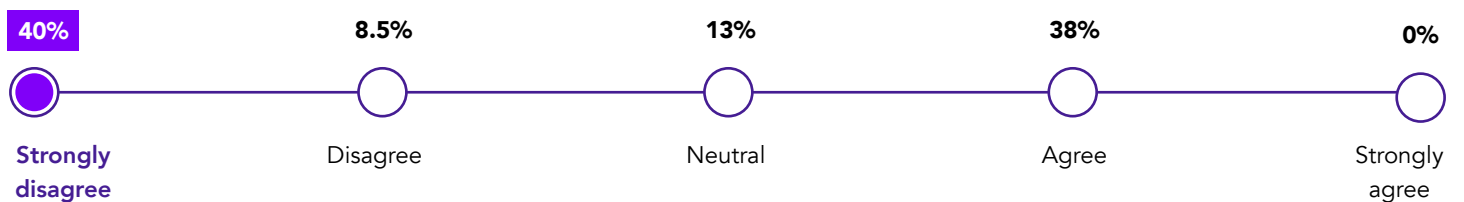
There is too much toil in the organization



What is your top source of toil?



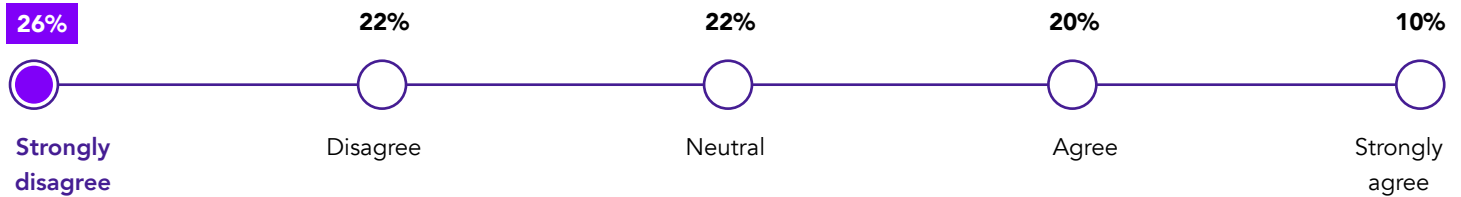
We have used automation to reduce toil



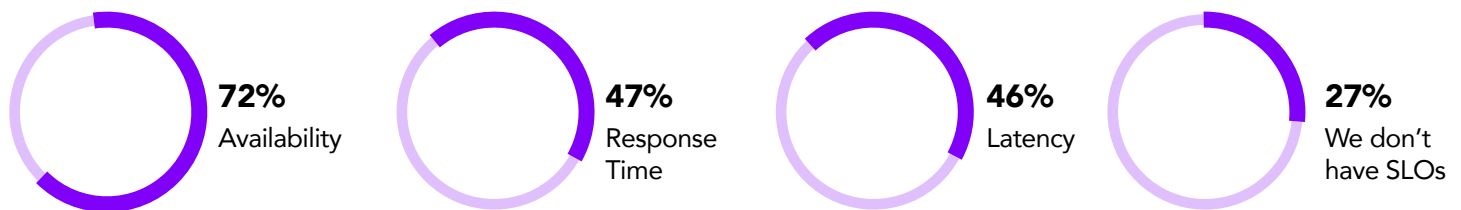
Service level objectives

Setting and monitoring service level objectives is a key aspect of the SRE role. The most widely tracked SLO is availability. Considering that 27% of respondents indicated they do not have SLOs in their organization every SRE that has SLOs tracks availability.

We have defined SLOs for all essential services



Our service level objectives cover:



Business impact of incidents

A missed SLO can have a noticeable impact on the business. One SRE rightfully indicated that a consequence of an incident is “the world turning into a mess.” Not all SREs work on external-facing applications, some SREs support internal applications—which is why we asked about a drop in employee productivity. The drop can be related to employees not being able to access systems, or from employees having to resolve the incident.



Key finding **1**

Catchpoint's take on Finding 1

SRE disciplines are still nascent. SREs ensure applications and services are reliable. This includes defining what reliable means in terms of service levels. If an API is available, but it takes 5 seconds to respond to a request, will that meet users' expectations? Before deciding that your organization is ready to take on SRE work (or if you already have) consider what are acceptable service level objectives. Establish benchmarks of current application and service performance from multiple perspectives and use these to guide the creation of your service level objectives.

For companies that are well-entrenched with SRE practices, find areas for improvement. What additional SLOs should be added? What toil currently exists in the organization? Are there things you can automate to reduce that toil? What new toil may be created when new SLOs are implemented, or new services are launched?

Consider the tools used by SREs. Are these adding to toil? Do they help you track service level objectives accurately?



Key finding **2**

Incident management is a massive part of the job

For the survey, we defined an incident as an unplanned interruption to an application or service that reduces the quality of the service. Incidents are assigned priorities based on the scope, impact, complexity, and urgency of the failure or interruption.

88% of SREs receive notifications about incidents via alerting and notification tools, but a handful are still being notified by coworkers or users contacting the helpdesk.

When was the last time you worked on a service incident?

Incidents are an unknown and can be difficult to prepare for; some are easy, some are not. Almost 50% of respondents have worked on outages lasting more than a day at some point in their career.

49% of respondents said within the last week

- 34%** Within the last month
- 10%** I'm working on one now
- 4%** I don't work on incidents
- 4%** I don't remember

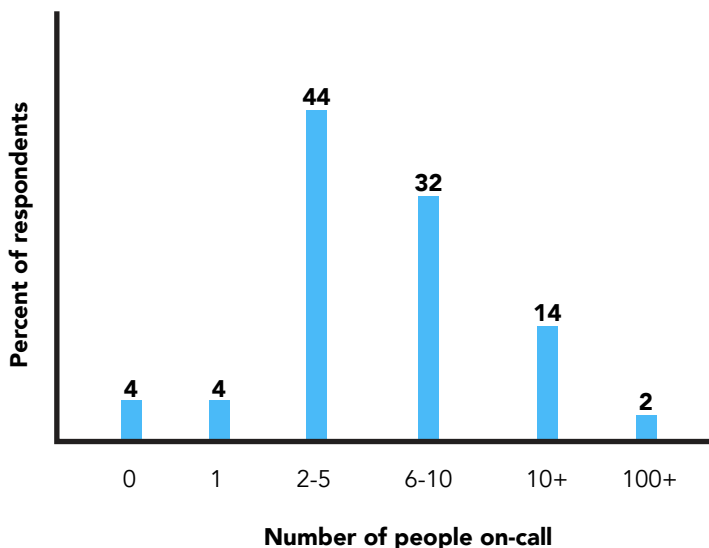
How many service incidents do you work on in a week?

Incidents are an unknown and can be difficult to prepare for; some are easy, some are not. Almost 50% of respondents have worked on outages lasting more than a day at some point in their career.

92% of respondents work on 5 or fewer incidents

- 48%** 1 or fewer
- 44%** Up to 5
- 4%** Six to 10
- 4%** Over 10

How many people are in your on-call rotation?



On-call rotation can vary. In some instances, there are rotations of rotations. Even companies with fewer than 50 people have varied sizes in their on-call rotation. 30% of respondents working at companies with fewer than 50 people report an on-call rotation of two people while other values have fairly even representation. One in 100 responses had 300 people in the on-call rotation, and another had 150.

One of the respondents who indicated zero people in the on-call rotation explained they are in pre-production, so the on-call rotation and responsibilities do not yet exist.



Key finding 2

Catchpoint's take on Finding 2

Consider how many SREs your team will have and whether they will be able to support the applications and services adequately. Include the appropriate people in the on-call rotation and ensure they have access to the alerting and notification systems. If you use Slack, integrate alerts into the appropriate Slack channels to reduce the number of times people find out about an incident from a coworker or users opening a support ticket.

Examine if there is a pattern when incidents occur. Do more incidents happen after code deploys? If so, consider if additional monitoring or testing in pre-production or development can reduce that.

Key finding **3**

Resolving incidents is stressful: Understanding post-incident stress

The survey defined post-incident stress as changes to physical and psychological well-being experienced up to two days after an incident occurs. Post-incident stress can last for a few minutes or up to two days.

How often do you experience post incident stress?



67% of those that report post-incident stress worked on an incident in the last week, 14 % indicated they were currently working on an incident.

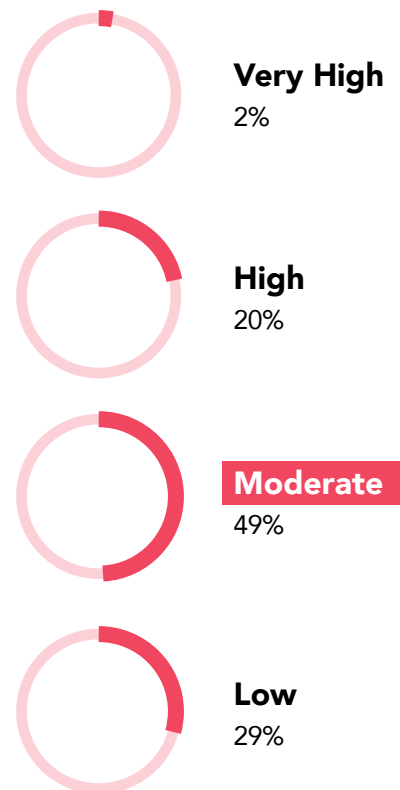
One way to not experience post-incident stress is to not work on incidents. 18% of those that never experience post-incident stress don't remember the last time they worked on an incident or reported that they don't work on incidents.

On the flip side, those who are the only SRE in an organization will always experience some stress. Post-incident stress happens after some or all incidents for the 12 people who are the only SRE in their organization. There is never an incident where they don't feel some level of stress.

On average, after an incident how would you rate your stress level?

Stress level is subjective. One person may classify their stress level as low while another as moderate or high. What matters here is this is an individual's perception of their stress level. Just because something isn't stressful for you, doesn't mean it isn't stressful for somebody else.

One respondent commented, "Culture and tenure have a lot to do with post-incident stress levels. I am fortunate enough to have a good culture and long tenure. Also, it is rarely my service that is the root cause of our incidents (knock wood, no brag, thug life)." Analysis of stress level based on seniority or title did not reveal any noticeable difference in the respondents.



Do you experience higher levels of stress during more severe incidents (i.e., widespread outage vs minor incident)?

82% of respondents
said yes

18% of respondents
said no

68% of those reporting low stress report the level varies based on the severity of the incident compared to 85-100% reporting moderate to very high stress.

After recent incidents, do you notice a change in any of the following?

Even those who reported never experiencing post-incident stress identified a change in one of the above after working on incidents. While they may not classify this as stress, they do experience a physical or psychological reaction to managing incidents.

52% of respondents
said mood

- 48%** Concentration
- 38%** Ability to sleep
- 38%** Desire to be social
- 32%** Ability to enjoy things
- 9%** Appetite
- 1%** None

Which of the following do you engage in to alleviate the effects of post-incident stress?

There are many things people do to relieve stress.

61% of respondents
said exercise / take a walk

- 52%** Spend time on a hobby
- 48%** Get a good night's sleep
- 43%** Spend time with people
- 35%** Drink alcohol

Catchpoint's take on Finding 3

The SRE role is stressful. There are steps organizations can take from a process and technology perspective to reduce that stress. One way to reduce the amount of stress employees feel is to deploy alerting and notification solutions. Of the respondents that always experience post-incident stress, 20% discover incidents from users contacting the help desk compared to 2-3% for those who never or sometimes experience post-incident stress. Stress levels may decrease if employees receive notifications before users start complaining.

If you already have an alerting and notification solution in place, but the stress levels are still high, explore whether this is due to alert fatigue. Are there too many false positives and false negatives occurring? Are you missing critical notifications because you aren't monitoring all the critical elements of the application or service?

Conducting game-day drills can help prepare the team for live incidents.

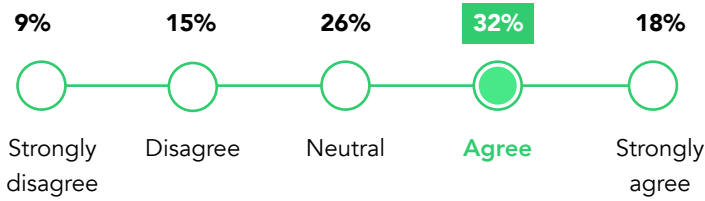
If all these are in place, the solution may be people related and not solvable with technology, read on to learn more.

Key finding **4**

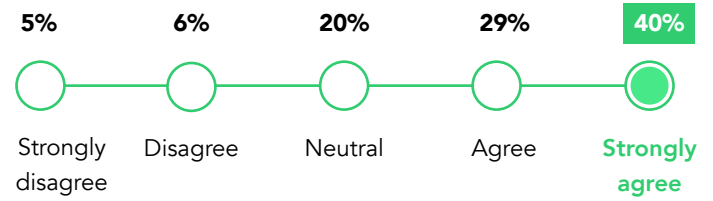
A supportive team reduces post-incident stress

Employees that feel their employer cares about their well-being experience less stress. 76% of the people who experience stress after every incident either feel neutral or disagree that their company cares about their well-being. SREs feel their teams care more about their well-being than their companies do.

My **company** cares about my physical and mental well being.



My **team** cares about my physical and mental well being.



What does your organization do to alleviate post-incident stress?

We opted not to include an option for “nothing” however 9% of respondents wrote that in. 10% that report stress after every incident indicate their company does nothing to help alleviate the stress. 16% of the people that reported high levels of stress after an incident say their company does nothing to help alleviate the impact of post-incident stress.

61% of respondents said reinforces a just/blameless culture

- 40%** Provide extra time off
- 38%** Checks in to see how you're doing
- 7%** Offers free massages
- 7%** Reduces on-call rotation

How supported do you feel by your team during/after an incident?

How supported an SRE feels by their team during and after an incident influences their stress levels. Overall 80% feel supported by their team after an incident. These numbers drop for those who experience stress after some or all incidents—64% for those feeling stress after some, and 43% feeling stress after every incident. 20% of those feeling stress after every incident feel little to no support after an incident compared to 14% of those experiencing some stress and 5% who never experience stress.

		During incident	Post incident
Not supported	1	1%	5%
	2	4%	7%
	3	13%	22%
	4	31%	23%
Supported	5	49%	41%



Key finding **4**

Catchpoint's take on Finding 4

Stress is considered "part of the job" but ignoring the stress is not healthy.

The notion of blameless post-mortems is good, but this doesn't eliminate the stress experienced when resolving incidents. Organizations need to deploy more concrete ways of reducing stress. Recognize that failure is stressful for many. Finding ways to reduce the number of incidents, and most importantly the number of high priority incidents, will go a long way towards reducing stress.

When people are on call, or when incidents occur, compensate those that worked on the incident. A few suggestions from respondents include paying people to be on-call or offering the option for extra time off. One company referred to this policy as "surge protection."

Conduct regular post-incident reviews. Document what went wrong, identify whether additional investments are needed to fix a problem. Share knowledge and information across the organization and teams.

Methodology

In January 2019 Catchpoint conducted an SRE survey promoted via email lists and social media. The survey questioned technical professionals from across a variety of industries about their role as a site reliability engineer, how they manage incidents, and the prevalence of post-incident stress. 188 people responded to the survey.

About the Author



Dawn is a Director at Catchpoint where she uses her storytelling prowess to write and speak about the intersection of technology and psychology. She makes technical information accessible avoiding buzzwords and jargon whenever possible. Dawn has spoken at DevOpsDays, Velocity, Interop, and Monitorama. Her articles have appeared in numerous technical publications. She uses her non-existent spare time to serve as a chapter organizer for Write/Speak/Code a non-profit organization to empower women and non-binary coders to become speakers, writers, and leaders.

About Catchpoint

Catchpoint is revolutionizing end-user experience monitoring to help companies deliver amazing digital experiences. Our platform provides complete visibility into your users' experiences from anywhere – and real-time intelligence into your applications and services to detect and fix issues faster. We are proud to partner with digital innovators like L'Oréal, Verizon, Oracle, LinkedIn, Honeywell, Priceline, and Qualtrics, who trust Catchpoint to improve their brand experience and drive their business success. See how Catchpoint can reduce your Mean Time to Detect at www.catchpoint.com/freetrial.