

2021
SRE Report

FREE
TO BE AN
SRE

Contributing
Partners



vmware

catchpoint[®]

Foreword

How can leaders grow digital business services without incurring the problems associated with rapidly increasing scale, such as increased cost, risk, and complexity? The obvious part of a larger answer (especially given the name of this report) is to implement site reliability engineering (SRE).

SRE is a discipline which fundamentally asks software engineers to design operations teams. It applies aspects of software engineering to infrastructure and operations problems. The result: ultra-scalable and exceptionally reliable distributed software systems.

SRE complements DevOps by measuring and achieving the reliability of applications and services working in production and DevOps infrastructures in a prescribed manner. It combines the use of error budgets, team relationships brokered by collaboration, ops-as-code, and reliability control practices to ensure deployments meet Service Level Objectives (SLOs).

In analyzing the data from almost 300 practitioners and experts who participated in our survey on the state of SRE, we found some familiar trends and some provocative anti-patterns. As you go through

the results, you will also discover interesting changes from the previous year's survey, not to mention some noteworthy correlations.

Why should you read this report?

- Find out about the expected patterns and surprising anti-patterns of SRE.
- Gain insight into the industry's most comprehensive SRE tenet baselining data.
- Understand usage of the five major monitoring components (digital experience, application, infrastructure, AIOps, and network).
- Learn from an example value stream to bridge the IT-to-business conversation gap.

Before continuing, you should relax your baselining efforts when trying to match the Google definition of SRE. There are too many dimensions for a one-size-fits-all approach.

You are hereby free to be an SRE.




Mehdi Daoudi
CEO, Catchpoint



Contents

- 2 Foreword**
- 4 Introduction**
- 5 SRE Report Highlights**
- 6 Key Finding 1**
Levels Of Toil Lower Around the World;
Budget Usage Less Than Expected
- 14 Key Finding 2**
Multiple Providers Highlight the Need
For Platform Ops
- 20 Key Finding 3**
The Shift Toward AIOps Is Slow
- 25 Key Finding 4**
Observability Must Include Digital
Experience Metrics and Business KPIs
- 30 Conclusion**
- 31 Demographics and Firmographics**
- 32 Bios**



Introduction

We are grateful to all the SREs worldwide who responded to this year's SRE Survey. We had almost 300 respondents from site reliability engineers working in tech, the financial sector, media and entertainment, telecom, manufacturing, and other industries. Thank you so much for sharing your precious time and making this year's SRE Report the largest of its type in the field.

As we consider the past year's global events, current emerging trends, and future thoughts on how SRE implementations will change the world, we hereby declare:

SREs are beautiful.

SREs build reliable services.

SREs build resiliency in themselves.

The SRE community is strong.

ACKNOWLEDGEMENT AND CONTRIBUTIONS

SRE Report Passioners

Thank you to [Leo Vasiliou](#), [Eveline Oehrlich](#), and [Anna Jones](#). This report would not have been written without your passion and inspiration.

SRE Report Contributors

Thank you to [Holly Chessman](#), [Laura Meyer](#), [Neil Myers](#), [Gordana Neskovic](#), and [Stela Udovicic](#). Your contributions shaped this entire report.

Supporting Partners


Thank you from all of us at [Catchpoint](#) to [DevOps Institute](#) and [VMware Tanzu](#). The sum of our partners is greater than the whole.

Special Thanks

Thank you to our industry experts: [Kurt Andersen](#), [Simon Aronsson](#), [Helen Beal](#), [Johnny Boursiquot](#), [J. Bobby Dorlus](#), [Tamara Miner](#), [Sanjeev Sharma](#), [Gaurav Shukla](#), and [Jaime Woo](#). Your wisdom is an invaluable part of this year's report.


SRE Report Highlights

This year's report reveals a wide range of fascinating insights. Although you will definitely want to read the report in its entirety, here is a sneak peek into a few of the insights we'll dig into:



Many companies don't establish comparative baselines for SRE conditions

If you don't have a baseline, it's hard to tell what improves. For example, our survey showed that the average drop in toil from last year was 15%. Was this due to COVID-19, people doing less busy work at home, and the work feeling more meaningful? Will self-reported toil rise next year as offices open and the pain, problems, and challenges of the year before are reintroduced?



Companies are using a multi-provider strategy across the board

From cloud to DNS to API to CDN, it's all about the numbers. But managing a variety of providers is complicated, making the need for Platform Ops rise to the fore. While 53% of SREs rank "providing trainings on third-party platform capabilities" as a minor or not applicable activity, maybe it's time to re-think that stance.



AI is exciting but it has not yet been fully adopted or embraced

53% of SREs say the number one cloud application monitoring challenge is unified visibility across the stack. When monitoring and management strategies get to the point that they can depend on artificial intelligence and machine learning to improve decision making and automation, that problem will be vastly reduced. But for the most part, we're not there yet.



SREs should look to expand the boundaries of observability so they can become more customer-focused

Blending business, performance, and operations insight with monitoring is key to growing the SRE role. To make SRE successful, SREs need to tie their success back to the business and get business leaders onboard with SRE practices. That will be of huge benefit to SREs, businesses, and customers alike.

BASELINE

KEY FINDING 1

Levels Of Toil Lower Around the World; Budget Usage Less Than Expected

Are key SRE tenets of toil reduction, time on call, or the dev versus ops split increasing or decreasing? Without an established comparative baseline, SREs and businesses will have a difficult time knowing whether conditions are getting better or worse.

We recommend companies ensure they baseline the amounts of time being spent by SREs on any single category of activity. Then, compare internal baselines with industry data, such as this SRE Report. If the baseline is leaning too far from the “50/50 split,” with ops predominating, the nature of SRE work needs to be rethought. This starts with the measurement of how SRE time is spent (since what is measured is what will improve).

Prioritizing Value-Based SRE Activities

To inspire innovation and solve complex problems, SREs perform all types of activities. This includes low-level operational tasks, tactical implementations, and high-level strategic initiatives. The derived value of some of these activities is quickly realized. Others may take time to show value or the value may be more difficult to measure.

This is why the desire to “automate ALL the things!” is a core SRE tenet. People want to give themselves the best chance of focusing on value-based activities. In other words, for SREs, the focus of automation is less about deploying code or releases as fast as possible. Instead, it is aimed at delivering value to the customer with a risk-mitigating, scalable approach that balances the need for agility against the need for stability.


For the first key finding, let’s explore the core tenets of SRE, including the amount of time SREs are spending on toil, the Dev versus Ops split, and the use of error budgets. After all, how can you measure improvements without a baseline?

SRE Activity Breakdown

Responding to incidents or outages



Post-mortem analysis and/or write-ups



Participating in on-call rotation



Developing applications or capabilities



Experimenting or receiving training to expand knowledge or skills



Authoring business processes, rules, or best practices




Performing audits of usage/cost allocation



Spinning up new hosts/instances



Planning release roadmaps



Performing chaos engineering exercises



Providing trainings on third-party platform capabilities



Load testing or other capacity management activities




● N/A ● Minor ● Moderate ● Major

Toil Drops In 2021... But Is This Temporary?

“Toil is the kind of work tied to running a production service that tends to be manual, repetitive, automatable, tactical, devoid of enduring value, and that scales linearly as a service grows.”


Site Reliability Engineering,
O'Reilly Media, 2016




How Is Toil Measured?

The average drop in toil year-over-year is 15%. This marked drop is important to note because toil numbers across the entire distribution (not just the respondent median) are down from 2020 to 2021.


That said, it is worth noting these are self-reported figures. Additionally, while around 45% of SREs said they measure toil, only 22% said it is measured scientifically. The other 54% said they either do not measure or track toil or that they are trying to do so, but are “finding it challenging.”




I/We do not measure or track toil today.



Toil is measured anecdotally (e.g., vent sessions in the hallway).




Toil is measured scientifically (e.g., detailed tracking via a system).



I/We are trying to measure toil, but are finding it challenging.

Causes Of Toil



Why Did Toil Numbers Drop?

Only 9% of SREs said COVID was a major factor for high amounts of toil. This leads us to theorize that the universal reported drop in toil is related to the global shift to working from home over the last year. Did the shift in related working conditions and priorities mean that work felt more meaningful and therefore less like toil? That could be the case, but we’re not sure.

So, will toil remain lower? Probably not. We expect the amount to rise again next year when a hybrid office is more of a reality. Note that the “hybrid office” could involve employees continuing to work from home, returning to the office full- or part-time, or some combination of the two.

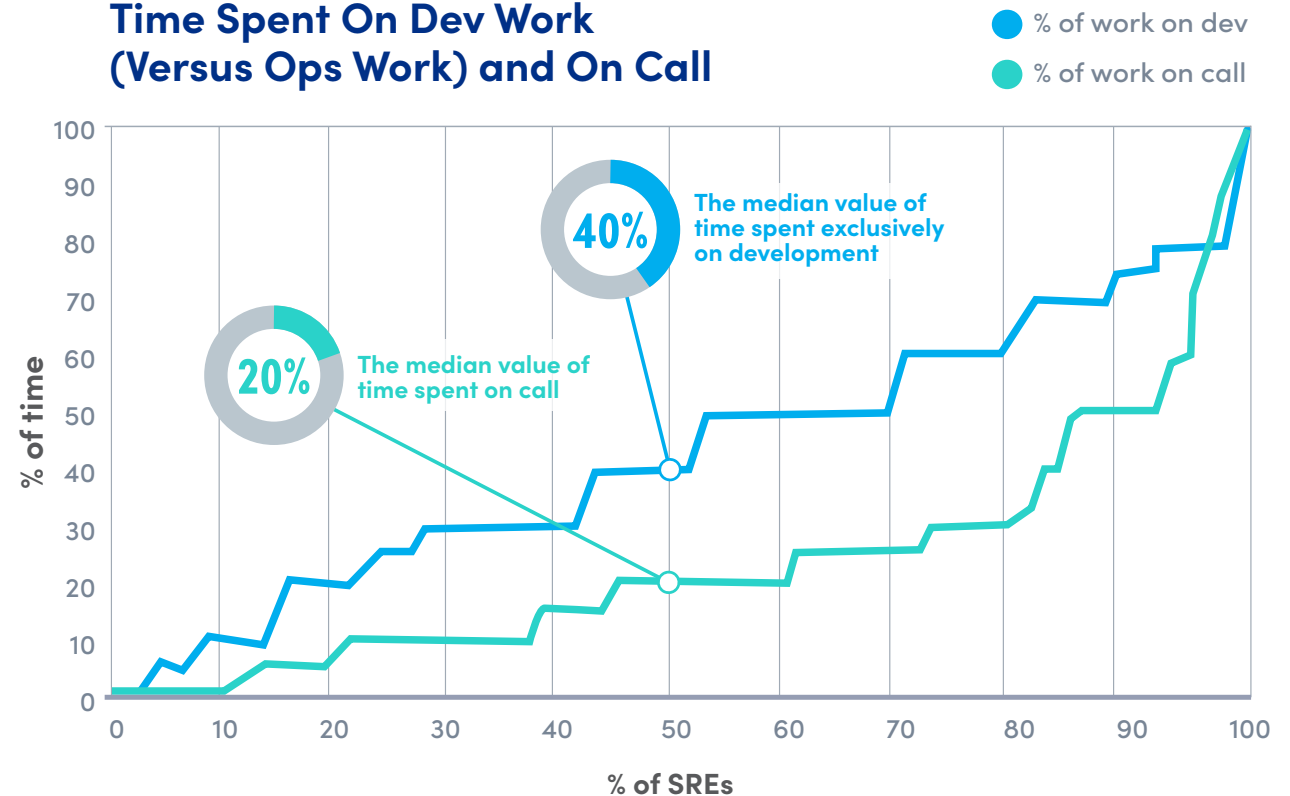
Does Google's Definition Of SRE Matter?

Google famously places a 50% cap "on the aggregate 'ops' work for all SREs — tickets, on-call, manual tasks, etc... Google's rule of thumb is that an SRE team must spend the remaining 50% of its time actually doing development."

Ultimately, the Google foundational definition of SRE is a guideline, not a rule. This is not to say that businesses should not baseline core SRE tenets against the Google definition. After all, entities performing both development and operational activities will be driven by the passionate desire to solve problems. However, businesses should focus primarily on inspiring and incentivizing SREs to achieve goals and objectives instead of trying to match the elusive 50/50 dev versus ops split. The drive to solve complex problems should always outweigh semi-arbitrary decisions, such as how an organization is structured or whether a business should mandate the use of open source for all tool selection.

SRE is fundamentally more about people and processes than technology. It's a good sign that SREs across the globe are spending time on development activities. The purpose of SRE is to bridge the gap between platform design, development, and operational execution. This means SREs have to shift-left (which they are doing) into development to share the wisdom of production to those teams delivering value-adding products and services to that production environment.

Time Spent On Dev Work (Versus Ops Work) and On Call



SRE Budgets: Hardly Outweighs Highly In All Categories

When asked how budgets are used — whether performance, toil, training, or error budgets — “hardly” outweighed “highly” in all categories. In fact, in terms of error budgets, only 20% of SREs aid they were “highly” using them.


Using error budgets demands that service level objectives (SLOs) and user journeys are designed and captured within service level agreements (SLAs). The next step in the journey toward wider usage of error budgets by SREs is to connect service level indicators (SLIs) and set SLOs.

While an error budget is calculated based on specific SLOs, you also need to establish specific error budget policies to determine what actions need to be taken if the error budget is depleted. When looking at the work done around SLOs, we find that half of SREs are continually refining their SLOs. This is a good indication that, even though error budgets are not used as highly as we anticipated, businesses are working to find the right balance between innovation and reliability.

Think of error budgets as something that force your SRE team to consider whether the right metrics are in place to meet company expectations. If you are not using them appropriately, the speed by which software is delivered and changes are implemented are likely impacting your company’s customers.

The goal is to shift left toward being able to address reliability problems sooner. The earlier you catch issues, the cheaper they are to fix.


Budget Utilization




“The business or the product must establish the system’s availability target. Once that target is established, the error budget is one minus the availability target. A service that’s 99.99% available is 0.01% unavailable. That permitted 0.01% unavailability is the service’s error budget. We can spend the budget on anything we want, as long as we don’t overspend it.”

*Site Reliability Engineering,
O’Reilly Media, 2016*


Service Level Attributes




We continually refine our SLOs.




We publish our SLOs to our users or customers to set expectations.




We look at SLOs in relation to system boundaries and then define tiers.




We know how each of our SLOs are treated as part of business SLAs.




We work to avoid triggering the consequences of missed SLOs.



We regularly compare SLIs and SLOs to decide what actions to take.



It is easy to find the right data to support our SLOs.



It is easy to choose SLOs.

SRE Spotlight

“

Most organizations and enterprises, no matter how small or large, are not like Google. Small and medium organizations don't have Google's scale. And large enterprises are not homogenous in their technology stacks or architecture like Google. So, yes, they all need SRE, just not in the same manner as Google does. They need to automate repetitive, typical tasks in operations in order to improve resilience and reliability and can learn from Google's take on SRE. They just cannot, and should not, emulate how Google implemented SRE.



Jaime Woo

Co-editor, *97 Things Every SRE Should Know*

@jaimewoo

SRE from: Toronto, Canada

Favorite: Haruki Murakami's *What I Talk About When I Talk About Running*

Sanjeev Sharma



SVP, Head of Automation and Platform Engineering, Truist Financial

@sd_architect

SRE from: Ashburn, VA

Favorite: All things SciFi and Superhero – from James T. Kirk to Loki

“

If Google is correct that too much toil leads to frustration, boredom, and burnout, could it be that during a pandemic, with its enduring waves of stress, SREs chose to spend their time elsewhere? Either way, did anyone even notice – and what does that say about toil?”

SCALE

KEY FINDING 2

Multiple Providers Highlight the Need For Platform Ops

Businesses that fail to implement Platform Operations teams will eventually hit the scalability ceiling. The rising use of multiple providers for same-service platforms (e.g., multiple cloud providers), will drive the need for SREs to develop and offer normalized capabilities even though underlying platform providers may have different interfaces and instrumentation.

We recommend businesses tactfully consider the assignment of SREs, or other development resources, to shared platform teams. The key result should be a higher level of scale. At that point, a single set of capabilities can then be treated as a product and re-used by many different teams.

Multi-Provider Use Increases and Brings With It Greater Complexity

Let's explore the number of multi-providers and levels of automation in use, the topology of SRE teams within the organization, and the nascent field of Platform Ops.


There is a vicious cycle of rising customer expectations causing businesses to manage ever-increasing complexities to meet those expectations. Migration to multiple clouds and using multiple third-party providers are two such examples of this trend. If the

cloud is your new data center and the Internet is the new network, than third-party services like DNS and APIs are your new racks and cabinets.

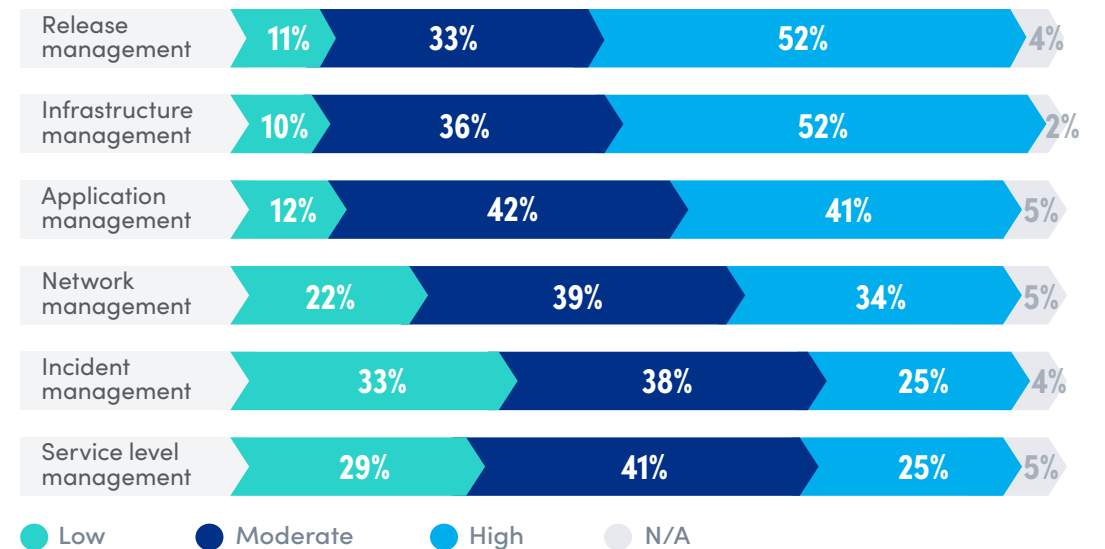
In response to the question, "Do you have a multi-provider strategy?", cloud led the way in all geographies. At the same time, it is clear that people are using a multi-provider approach across the board for DNS, API, and CDN. The reasons for this are numerous,

from improved resilience of systems to the ability to leverage different providers' strengths. However, managing a variety of providers brings with it greater complexity, from challenges posed by a reduction in visibility over the total infrastructure to confusion among teams over how best to take advantage of the different toolsets and rules of each vendor.

Multi-Provider Usage



Use Case Automation Levels



There Is No One Right Systems Reliability Engineering Topology

There are several different SRE models, each of which is organized and leveraged differently depending on the organization's leadership, needs, challenges, and goals. The following models are in use today:

CENTRALIZED SRE BUT SUPPORTING VARIOUS PRODUCTS OR PLATFORMS.


An SRE is typically associated with various products, applications, or platforms. These types of teams have SRE responsibility across a specific set of services and workflows with key boundaries. Here the SREs are functioning as an enabler to help the product or platform teams within the technical or product domain. The work of these SRE team members could include a variety of projects, such as the onboarding of specific technologies like Kubernetes, building out a mature infrastructure service mesh, working on continuous integration activities, or providing 24x7 support for specific applications and services.

DECENTRALIZED SRE ASSIGNED BY PRODUCT OR ANOTHER ATTRIBUTE.

These SREs are typically embedded into application/service teams to support the application team and are responsible for automating operations (runbooks) and deployment work, including monitoring and metrics population, as well as owning incident management around application issues. The advantage of this decentralized model is a solid integration within the business, with strong business value alignment and a new career path evolution for the embedded SRE. SREs can establish best practices, benchmarks, and automation architectures across the entire enterprise.

THE SRE AUTOMATION TOOLS TEAM ESTABLISHES AND GUIDES IT AUTOMATION BEST PRACTICES.

Another model is that of the SRE tools team, which supports its software development counterparts in improving system reliability through automation work, leveraging existing tools, or creating scripts. Possible work tasks could include focusing on monitoring, observability, and analysis across a particular application stack or the instrumentation of key infrastructure ecosystems as a best practice for other teams to leverage.




Greater Size Means Greater Decentralization


Related to the previous point is the finding that as the number of SREs and employees within an organization increases, SREs become more decentralized. Perhaps this path is unavoidable, given the different expertise of SREs. Since this pattern will likely increase, businesses must ensure they are avoiding a return to silos.

One way of preventing silos from occurring is by creating a Platform Operations team. Such a plan allows decentralized teams the autonomy and flexibility to do their work, while the centralized Platform Ops team provides capabilities to its decentralized colleagues. These normalized capabilities can then be integrated across workflows.

SRE Topology/Organization Size



SRE Topology/Number Of Employees



SREs, Meet Platform Ops


Platform Ops is fast gathering momentum among organizations that want to scale and need greater consistency (in tooling, security, training, etc.) across different pipelines and products. A Platform Ops team can focus on creating self-service capabilities by applying automation for application developers. This process can encompass everything needed to build an automated DevOps value stream.

The goal is to produce a system where developers do not have to worry about infrastructure and operations tooling. Instead, they can put their efforts towards developing, while the Platform Ops team implements guardrails around cost, monitoring, and compliance (and more) – without being a bottleneck. Platform Ops must collaborate with the development team, of course, to ensure that the designed workflows are improving the development team's experience. At the same time, Platform Ops needs to ensure operational boundaries and rules are established and followed.

In our question about most performed activities by SREs, “providing trainings on third-party platform capabilities” ranked the lowest. With 52% of respondents selecting this activity as minor or not applicable, this represents a key opportunity for scaling SRE implementations through a Platform Ops program.


To ensure a successful Platform Ops implementation, focus on three critical checkpoints:

1




A shared, self-service platform that offers an array of capabilities.

2



The ability for those capabilities to be enhanced within the context of the overall business, through the use of development/engineering resources.

3



The ability for those capabilities to then be productized and internally marketed for consumption, able to fulfill the needs of many different internal consumers.

SRE Spotlight

“

Spanning the gaps between the interfaces and the data that each provider offers increases the difficulty for SRE teams to automate across those multiple providers. These integrations are rarely simple except for the most superficial aspects. Effectively mapping disparate data models together may be the next frontier for SRE in a multi-vendor environment.”



Johnny Boursiquot

SRE, Heroku

@jboursiquot

SRE from: Baltimore, MD

Favorite philosophy: Amor Fati

“

Chances are high you'll need a tailored SRE model for your organization. Be aware that the effort you put into building this custom SRE practice will only bear fruit when all parties affected have a shared understanding of existing shortcomings and how outcomes will be measured and reflected while on the journey. Nothing kills the effectiveness of SRE adoption faster than a misalignment on its aims and a lack of transparency on the value it provides specifically to your organization.”



Kurt Andersen

SRE Architect, Blameless

@drkurt

SRE from: North Idaho, ID

Favorite quote: “There is only one root cause: change.”
Dave Zweiback, *Beyond Blame*,
O'Reilly, October 2015



INCREMENT

KEY FINDING 3

The Shift Toward AIOps Is Slow

AIOps promises to turn mountains of data and inert action into molehills of information and precise actionability. But this year's SRE Report showed adoption is slow (when compared to other monitoring tools in use) and that "the value received from AIOps" covers a wide range of sentiment.

We recommend SREs break down the broad AIOps category into smaller components, such as event correlation, topology analytics, or smart alerting. Once these components have been identified, incrementally develop AIOps capabilities through an investment in training and the right tooling to help manage the continually increasing volume, velocity, and plethora of data sources.



The Value of AIOps

It's time to explore the value of AIOps and why its uptake is slow. Given the amount of data coming in from monitoring of all different data source types (from infrastructure to application to experience), the challenge of analysis or correlation by a human is enormous.

The promise of specific AIOps components – event correlation, for instance – is huge in helping to parse through big data to arrive quickly at the important information or (even more critical) the prescribed action. However, while you will need AIOps to make sense of all this data and to reduce toil, not all components are equally helpful. For instance, vendor promises of self-healing and automatic remediation should be heavily scrutinized if you are considering them to be part of AIOps' received value.

AIOps Is Still An Underused Field

SREs are always performing two types of activities: development and operations. This is a form of bi-modal IT, where mode I involves maintaining existing activities and mode II involves exploring and researching the unknown. The result is that SREs are double struck with challenges as they are required to handle constant context switching. This constant switching can make it difficult to accomplish goals in either domain.

One key way to help mitigate the challenges of being multi-mode for mode II (the new and unexplored mode) is AIOps. However, it is clear from this year's survey results that AIOps is still very much an underused capability.

Monitoring Tool Usage

Infrastructure monitoring (e.g., disk or CPU)



Network performance monitoring or diagnostics (e.g., latency or saturation)



Application performance monitoring (e.g., tracing or events)



Digital experience monitoring (e.g., Synthetics or RUM)



Artificial intelligence for ITOps (e.g., anomaly detection or self-healing)



Public/social sentiment monitoring




Competitive benchmarking intelligence



● Never ● Rarely ● Sometimes ● Always

Received AIOps Value



● Low (1-3)
● Moderate (4-6)
● High (7-9)

*Based on a 1-9 value scale.

AIOps Is At An Inflection Point Toward Becoming Valuable

Availability and reliability are key aspects for ensuring customer satisfaction and avoiding a negative impact on customer confidence and your company's brand. There are a number of monitoring tool types that can ensure both availability and reliability. Digital experience monitoring (DEM) solutions are an essential part of most companies' efforts to reduce downtime and ensure ongoing services. Web analytics and other IT monitoring technologies (APM, ITIM, NPMD, and AIOps) also play a role in a broader monitoring strategy, helping to offer a complete view of the end user experience.


At the same time, clearly not every company is able to get what they need from the monitoring tools they have. After all, 53% of SREs said the number one cloud application monitoring challenge they face is a lack of unified visibility across the stack. What is causing this disconnect?

Large volumes of data are draining the value of traditional APM and other monitoring tools and making it impossible for SREs to become proactive. The demands of the digital business require a modern way of managing incidents and service health, but many companies aren't there yet.


Today's monitoring and management strategies need to leverage artificial intelligence (AI) and machine learning (ML) to improve decision making and automation. Observational and engagement data must be analyzed to allow businesses to react in real-time. When combined with automation either during or after analysis, this data will enable continuous ongoing improvement and help shift IT operations to working in a proactive and predictive way. This is exactly what AIOps automation solutions are promising.

However, its adoption from our survey respondents is not looking optimistic. Why is this?


Cloud Application Monitoring Challenges




Lack of unified visibility across the stack.




Huge effort in maintaining existing tool(s).




Lack of analytics.



Monitoring tools do not scale.



I/We face no cloud application monitoring challenges.



Other

Break AIOps Into Individual Components To Understand Its Value

Unrealistic expectations of AIOps components, such as self-healing and remediation, have led to disappointing results in production environments and an undervaluing of AIOps as a whole. To develop a clearer view of what AIOps can offer, we recommend breaking down its individual capabilities, then training SRE teams in AI and ML. This will make it easier to realize AIOps' value. By looking at the individual components and incrementally developing capabilities, you can realize the potential of those that are most relevant to your business.

Important AIOps use cases for monitoring include:

- Event correlation
- Anomaly detection
- Proactive analysis
- Root cause analysis

When it comes to monitoring, a good deal of value can be derived from deploying AIOps and leveraging these use cases in combination, to determine unexpected patterns in large volumes of data. You can achieve this value using a real, automated inference capability.

AIOps is broader than a single monitoring tool. With that in mind, it's worth taking the time and investing in training in AI and ML for SRE teams to understand the different domains and how they might work best together. Note, however, that while there may be no immediate perceived value. The real worth comes in the medium to long-term gains.

SRE Spotlight

“

During the last year, most of us have had to adjust to spending a lot more time in our homes. We believe this transition to be one of the primary explanations for the significant decrease in how much of our work we consider to be toil.

As we slowly start recovering from these trying times, some of the work we've considered as good alternatives to being isolated will likely begin to feel like toil again.

We expect site reliability to stay mission-critical moving forward, justifying continued efforts in eliminating toil. Organizations that properly manage to identify activities suitable for automation and AIOps will have an excellent opportunity to keep the perceived level of toil low even as we readjust.”



Simon Aronsson

Head of Developer Relations, k6

@0x12b

SRE from: Norrköping, Sweden

Favorite movies: *Hackers* (1995) and *WarGames*

SRE Spotlight

“

Today, AIOps solutions elevate SRE skills by automating incident responses. They can proactively monitor the SRE's golden signals and measure what really matters for customer experience, but if not introduced correctly, can backfire, ending up in a more complex system to be managed! So how can this adoption be simplified? To start with, we need to have a closer look at our current toolset and evaluate where we have gaps, taking into account common features that current AIOps products offer, such as baselining, RCA, anomaly detection, event correlation, simulation, etc. It's all about getting better at understanding and managing the complexity of your setup and then integrating automated helpers and actions.”



J. Bobby Dorlus

SRE, Twitter

@BobbyD_FL

SRE from: West Palm Beach, FL

Favorite: Independence Day, I, Robot, I am Legend

“

AIOps should be part of any SRE toolkit or any infrastructure managing large scale systems. Greater adoption is coming. It aligns with some of the challenges I've experienced at Twitter with managing large scale systems and in particular, the vast amounts of data, metrics, and logs we need to work through. How can we parse through all that as humans and be as effective as our customers need us to be? Most SREs working at scale are already leveraging machine learning, especially when it comes to efficiencies around data centers (locations, cooling, and all the things that happen inside it), for networks and building out infrastructure ... Evolving that into AIOps is the next logical step.”



Gaurav Shukla

SRE, Catchpoint

SRE from: Delhi, India

Favorite: The Pursuit of Happiness, Seven Pounds

EXPAND

KEY FINDING 4

Observability Must Include Digital Experience Metrics and Business KPIs

SREs that fail to deliver customer value run the risk of being stuck in an operational toil rut. Conversely, businesses failing to recognize the importance of SRE activities run the risk of losing talented employees and a competitive edge.

IT-to-business conversations should start around capabilities – the gateway to positive business outcomes and a middle point in the value stream. Talking about necessary capabilities allows the conversation to shift as needed, helping shrink the IT-to-business gap.


Monitoring Data Needs To Pivot Towards Details To Ensure Customer Outcomes

For the final key finding, let's explore what is driving the use of monitoring data, and how SREs can expand the boundaries of observability to include the digital experience and business KPIs. With this new angle, SREs can become more customer focused.

The responses to the question, "What drives the use of monitoring data?" resulted in few surprises this year. Augmenting troubleshooting and root cause analysis was the resounding leader. Ongoing changes continuously require better root cause analysis and troubleshooting, so this came as no shock. "Ensuring service level objectives (SLOs) are met" came in second. While only 49% of SREs said, "enhance the customer experience" was a major driver for the use of monitoring data, ensuring SLOs indirectly enhances customer experience. Thus, these two responses can be viewed together. A quarter of SREs chose "accelerate our ability to innovate automation" or "ability to support collaboration efforts" as major drivers for the use of monitoring data.

As can be seen from the survey results, SREs are pivoting towards focusing on customer experience, blending business, performance, and operations insights with their monitoring approach. For SREs to be successful, they must work with developers, customer intelligence (CI) pros, CX experts, marketers, and ecommerce professionals. This type of collaboration will enable them to understand what critical data needs to be observed and leveraged to predict and populate business-oriented KPIs.

Monitoring Data Usage Drivers



Expand Observability Boundaries To Enhance Customer Experience

In a world where businesses are rapidly implementing digital transformation initiatives, SREs must expand their observability boundaries to include digital experience and business KPI data. Since SREs are innovation providers, they are not bound by traditional definitions. Therefore, this is a business opportunity to take evolutionary concepts, in this case observability, and apply them to more than just an IT domain.

When asked what data sources feed SRE frameworks, again, there were few surprises. Application monitoring came out on top, with infrastructure monitoring a close second. What was interesting was the rare use of benchmarking intelligence and public sentiment/social media monitoring. Forward-looking


SREs will include digital experience data (both human and machine), such as social sentiment data, benchmarking data, and that which directly applies to business KPIs. Why?

IT needs data from the end-user point of view, to fully measure quality and consistency of experience delivery. Outside-in monitoring via digital experience monitoring (DEM) allows businesses to look at what a customer or end user is seeing. In this way, organizations can determine if the application is delivering what the customer wants and expects. A poor customer experience directly impacts a company's reputation and bottom line. That's why ensuring a good customer experience aligns directly with overall business goals.


By including a wider range of monitoring tools, such as those explicitly focused on business perception and ranking, SREs can better realize their role within the organization. This is true, not just in terms of maintaining reliable systems, but also in terms of helping realize greater value for the business.

On a similar note, more SREs than not are choosing to maintain traditional monitoring habits by focusing on inside-out monitoring versus outside-in. Measuring whether the servers are up and the lights in your NOC are green is a narrowly focused surface area. On the other hand, monitoring to ensure customer experience across the entire delivery chain is more broadly focused and will ensure a greater alignment with business goals.


Observability Data Sources




Application monitoring




Infrastructure monitoring




Network monitoring




Front-end user experience monitoring



Client-side device or endpoint monitoring



Collective benchmarking intelligence



Public sentiment/social media monitoring

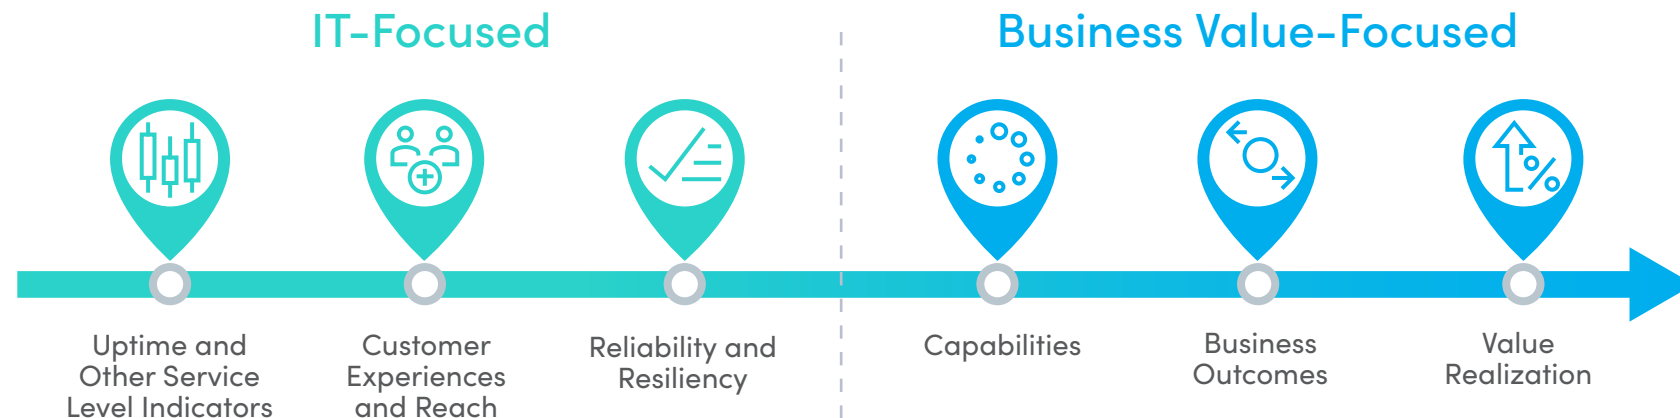
SREs Must Connect Monitoring Practices To Customer Value Realization

SREs are often inwardly focused on operational conversations and making sure services are reliable. SREs place a great deal of importance on the left three elements of the table on this page: inside-in monitoring of service level indicators, helping to extend reach (using third-party delivery chain components like CDN), and working toward reliable services (through use of high-availability, failover architectures, or chaos exercises).

What they are not doing, however, is sufficiently aligning themselves with business outcomes or realizing customer value. This is the direction in which business leaders typically take the conversation. They make more outward-facing, value-based decisions like how to land new logos or retain existing revenue. The result is an often large and frustrating gap or the less-than-optimal use of resources to accomplish goals or objectives.

To bridge this gap, SRE conversations should start with capabilities. For example, when asking for budget, instead of saying, “Can we have \$\$\$ to buy Google Docs?”, a more effective way to reframe the question might be, “Can we have \$\$\$ to increase our ability to quickly collaborate with external partners on large projects?”

In addition to reducing frustration, using a value chain path increases the ability to reproduce results. By having a path that leads to realizing value, SREs and businesses alike can scale their activities and avoid “blind luck” scenarios where they are not quite sure how value was realized (even though they must reproduce that value over and over).



Capabilities are the gateway to positive business outcomes.

Successful SRE Implementation Drivers

- 60%** | How quickly we resolve incidents
- 43%** | The amount of time between failures
- 41%** | How quickly we do root cause incident analysis
- 40%** | How quickly we push product updates
- 33%** | How quickly our business can expand to new markets
- 22%** | How quickly we can understand the cause of social media sentiment

SRE Spotlight

“

In addition to understanding how business value ties back to SRE SLIs, it is important to close the feedback loop by communicating actual customer sentiment and felt value back to the engineers working on improvements in these areas.

Recently at Improbable, we received some feedback from a customer around how happy they'd been with our team addressing key SRE metrics (e.g., shorter deployment times and reduced number of incidents). Our proactiveness in soliciting feedback via customer satisfaction surveys and ad hoc meetings was called out as a huge plus, allowing them to focus on improving their customer experience and increasing their scale significantly. Instead of that feedback stopping at the Account Manager, we ensured we had a closed loop process, so our engineers could see the impact of their work and motivate them to anticipate customer needs in future development.”



Helen Beal

Chief Ambassador,
DevOps Institute

@BealHelen

SRE from: Chichester, UK

Favorite: Yann Martel's
Life of Pi

“

People don't have time to analyze data all the time — they're busy dealing with incidents and trying to make improvements. And all those different monitoring tools feeding the observability frameworks... It seems there's a gap — we've got the data, we've got the performance model, but we haven't connected the two. The observability and AIOps framework vendors need to up their game for the SRE personas and create out-of-the-box patterns for the management of these all-important SRE metrics.”



Tamara Miner

Engineering Manager,
Improbable

@tammasaursrex

SRE from: London, UK

Favorite: Erin Meyer's
The Culture Map

Conclusion

With the exception of Google's research, we believe this to be the most data-backed report of its kind. We looked at the data for each individual question and then performed cross-question correlations, some of which did not make it into the report for the sake of brevity. Examples include:


- As people use more multi-provider third-party strategies, they buy more than they build.
- The more decentralized, the more time on call.
- Xbox users have a greater preference for Macs over their PlayStation counterparts.

In our desire to evolve the SRE field, we'll publish additional findings from the report over time. To stay up on our findings, please subscribe to our [blog](#) and follow us on [LinkedIn](#) or [Twitter](#).

One of the underlying themes for this year's Report was to look at the five major monitoring components (digital experience, application, infrastructure, AIOps, and network) and understand how the promise of AIOps will help them fit together. In order to run, you must first walk through drafting an established line in the sand, i.e., baseline at the outset so you know where your journeys begin.

As we've determined, SREs must attach to business value conversations. It's also important to consider the different perspectives of, and empathize with, different teams. Your IT counterparts, for example, may not have the development and engineering resources that a well-honed, well-tuned SRE may have. Consider offering talent sharing programs, because your development experience may exponentially increase their ability to improve the efficiency of their programs. In other words, an ounce of development to you is worth a pound of development to others.

This brings us full circle to one of our core SRE philosophies, which is the common, passionate desire to solve complex problems.



Demographics and Firmographics

Respondents to the SRE Survey came from all over the world with almost half representing North America, a quarter representing Asia, and a fifth representing Europe. Other locations represented included Oceania and South America. Respondents came from a wide range of industries, with technology leading the pack, followed by financial services, then media or entertainment. Over half the respondents cited IT operations as their primary area of expertise, closely followed by application or software development/engineering and IT infrastructure. Sixty percent of respondents said they had more than one area of expertise.

What is your role?

SRE individual practitioner/subject matter expert	47%
Team leader/supervisor	16%
Manager	14%
Senior management (Director, Vice President)	9%
External consultant/contractor/coach	5%
C-Suite executive	3%
Other	6%

How many employees does your company have?

One to 100	18%
101 - 1,000	29%
1,001 - 10,000	23%
10,001 - 100,000	21%
More than 100,000	9%

What is your primary area of expertise?

IT Operations	52%
Application or Software Development/Engineering	50%
IT Infrastructure	47%
Architect	27%
Network Operations or Engineering	20%
Security	16%
Database Engineer	11%
Service Desk or Support	9%
CIO, CTO, CXO, or other C-Suite Executive	4%
Other	4%

How many SREs are in your organization?

One to ten	52%
11 to 100	34%
101 to 1,000	11%
More than 1,000	4%

In what industry is your organization?

Technology or Technology Provider	41%
Financial Services	13%
Media or Entertainment	8%
Telecom	6%
Manufacturing	5%
Consumer Packaged Goods or Retail	5%
Healthcare or Chemicals	4%
Professional Services or Consulting	4%
Energy	3%
Transportation	3%
Government or Non-profit	3%
Travel or Accommodation	<1%
Other	4%

Where are you located?

North America	48%
Asia	24%
Europe	21%
Oceania	3%
South America	3%
Africa	1%

Bios

Kurt Andersen is currently the SRE Architect at Blameless. His writings appear in various O'Reilly books including *What is SRE?* and he serves on the Board of Directors for USENIX and as part of the steering committee for the worldwide SREcon conferences.

Simon Aronsson is the head of DevRel at K6. He's a long-time reliability and DevOps nut helping software teams build reliable, performant software.

Helen Beal is a DevOps and Ways of Working coach, Chief Ambassador at DevOps Institute, and an ambassador for the Continuous Delivery Foundation. She is the Chair of the Value Stream Management Consortium and provides strategic advisory services to DevOps industry leaders such as Plutora and Moogsoft.

Johnny Boursiquot is a multi-disciplined engineer at Heroku with a love for teaching and community-building.

Mehdi Daoudi is the driving force behind Catchpoint's promise to help the Internet deliver on its potential for the human race.

J. Bobby Dorlus is a seasoned Systems Engineer with over 18 years of experience working for companies like Twitter and Citrix. In his eight-year journey at Twitter, he has worked on teams that are responsible for the foundational platform that most services @Twitter rely on. He is known for his historical knowledge of Twitter's technical journey.

Anna Jones is a fanatic about conveying technologists' pains and problems through the written word.

Tamara Miner has worked on infrastructure and developer tools for over 15 years in the US and Europe. She is currently the Engineering Manager of Improbable's Partner Engineering team in London. She is a recipient of the Forbes 30 Under 30 and the Microsoft Xbox Women in Gaming Rising Star awards.

Eveline Oehrlich is Chief Research and Content Officer at the DevOps Institute responsible for research and content strategy. Previously, she held the position of VP, Research Director and Principal Analyst at Forrester Research delivering research, advisory and consulting for software vendors, system integrators, service providers, and IT enterprise organizations.

Sanjeev Sharma is an internationally renowned digital transformation leader with a track record in the areas of DevOps, DataOps and Cloud Adoption, and Application and Data Modernization. Sanjeev is a former IBM Distinguished Engineer, and the former Global Field CTO of a VC backed startup. He is currently Head of Platform Engineering at Truist Financial.

Gaurav Shukla is an experienced and accomplished Linux Systems, Cloud and Site Reliability Engineer with a deep passion towards observability, automation, and designing of fault tolerant systems.

Leo Vasiliou is a former ITOps practitioner and current chart and graph whisperer.

Jaime Woo is a writer, SRE educator, and former molecular biologist and has studied the impact of stress and burnout in SRE since 2017. He continues to focus on mental health and well-being, working toward a Certificate of MBSR (mindfulness-based stress reduction) Facilitation from one of the field's foremost programs.

About Catchpoint

Catchpoint, the global leader in Digital Experience Monitoring (DEM), empowers business and IT leaders to protect and advance the experiences of their customers and employees. In a digital economy, enabled by the cloud, SaaS and IoT, applications and users are everywhere. Catchpoint offers the largest and most geographically distributed monitoring network in the industry – it's the only DEM platform that can scale and support today's customer and employee location diversity and application distribution. It helps enterprises proactively detect, identify and validate user and application reachability, availability, performance and reliability across an increasingly complex digital delivery chain. Industry leaders like Google, L'Oréal, Verizon, Oracle, Equinix, Honeywell and Priceline trust Catchpoint's out-of-the box monitoring platform to proactively detect, repair and optimize customer and employee experiences. Learn more at www.catchpoint.com.

Methodology

For our fourth annual SRE Survey, we received 278 responses. All responses were organic and non-paid.

The survey was open for the month of April 2021. It presented a list of 40 questions that covered a wide range of topics, including the split between development and operational activities, the approach taken to tool usage, and the type of monitoring tools and practices used. Qualifying questions such as, "What activities do you perform?" had both operational and development choices as answers, to build confidence in the quality of the received answers.

When analyzing the results of the survey, correlations between questions were actively made. Correlation does not necessarily mean causation, but it can yield fascinating insight into the nature of SRE approaches and activities. To ensure efficacy and integrity of correlated questions, the same mathematical base was applied.

For any questions about this report or the data, please contact pr@catchpoint.com.

