



The SRE Report 2023

Service **Reliability** and Internet **Resilience** for All

With special contributions by **Adrian Cockcroft** and **Steve McGhee**



Foreword

As with most change, there was some resistance to DevOps when the movement first started. At Netflix, when it was a small company and we were trying to outrun our large competitors, agility was essential to winning the race to building personalized video delivery and enabling streaming at a global scale. We listened to Amazon CTO Werner Vogels' advice to "run what you wrote" and arrived at organizing ourselves according to DevOps principles because in order to be fast, we had to remove the handoffs between development and operations, and provide developers with a cloud based platform that they could use to automate operations tasks directly.

Site Reliability Engineering is a flavor of DevOps, which famously got started in 2003 at Google as an early means of bridging engineering and ops with the goal of maintaining extremely reliable software systems. As large-scale systems stabilize their functionality and the focus changes to reliability, it makes sense to harden the code and automation, hand them over to a central SRE team, and to free up developers to explore new areas.

Adopting a posture of learning, testing, and continually adding to one's toolkit is critical to building the kind of continuous resilience mindset necessary to producing reliable, resilient systems. Equally

important is being part of a community of SRE and DevOps practitioners in which peer-to-peer communication is freely encouraged.

That's why reports like this one, The SRE Report, are important. Outside of Google, this is the most data-backed report of its kind and the longest-running, now entering its fifth year. In addition to capturing benchmarking data year over year, it delves into new fields of practice each year, and most importantly, draws on a global community of reliability practitioners to provide the backbone of its research findings.

In an era of rapidly changing digital initiatives and continually shifting work paradigms, we need to keep a lookout for the next wave of disruption and the new approaches and toolkits that will enable us to surf it, rather than drown. It's my belief this year's SRE report will help you do just that.

Adrian Cockcroft, Catchpoint Advisor



Introduction

Why do we care about reliability as a feature? Is it because we can? Because we want to? Or because we have to?

We care about reliability as a feature because we are human, and we are digital resiliency hunters. We use tools to capture resiliency, and in our downtime, we may strum a few chords, relax with friends, or enjoy a good read... at least until the next release, next launch, or next learning from a near miss.

Our tools are ever evolving. A broad range of them are encapsulated as part of a DevOps lifecycle.

Welcome to The SRE Report – written in partnership with Blameless. As you delve into our five critical insights on the field’s most significant findings, SREs’ unfettered responses to open-ended questions, and the benchmarking data we gather year over year, we hope you find a great deal to stimulate, provoke and enlighten.

Each year of The SRE Report, we choose to highlight an emerging area in the field. After focusing on the need for Platform Operations last year, this year, we have examined the concept of Total Experience to understand how companies are balancing their focus between delivering customer experience and ensuring employee productivity and satisfaction. The results are not what we expected and present some fascinating hidden findings for companies and managers, in particular, to consider.

Let’s cut to it. We know you’re busy... So why should you read this report?

01.

If you believe writing integrity and data-driven decisions should not be influenced by market hype.

02.

If you understand that averages can be misleading and there may be anti-patterns/paradoxes in subsets of data worth examining.

03.

If you wish to see how operating with a “just culture” has quantifiable benefits.

As you dig in, bear in mind that over 550 global reliability practitioners (ranging from individual contributors to executives with all manners of title in between) have made this report possible.

Thanks to everyone who gave their time and insight in taking part.

SRE Report Highlights


Here's a sneak peek of some of the insights we'll dig into:

AI is not replacing human intelligence anytime soon - it should be considered within larger observability contexts

This year's data tallied with last year's - with the majority of reliability practitioners continuing to state that the value received from AIOps is low. Given that AIOps can be considered as a means of enabling IT or business-level observability, we suggest that practitioners consider which data sources currently feed those frameworks and consider a reappraisal. We also recommend SREs (and their business counterparts) break down the broad category of AIOps into individual capabilities as they work toward larger observability implementations.

A stark dichotomy exists between Individual Practitioners and Executives

In multiple critical areas, from the value of AIOps to whether tool sprawl is a problem, we found a gulf in the point of view between Individual Practitioners and Executives. An opportunity clearly exists for better conversations, new ways to create alignment around goals and objectives, and decision-making to be informed by project ownership and closeness to everyday activities versus being made along simple authority lines or largely determined by budget considerations.



Elite performing organizations report a very or extremely blameless culture

Blamelessness is not just a matter of engendering a just culture for the sake of itself, but our data shows that an overwhelming majority of Elite organizations (per the DORA definition) report they are "very" or "extremely" blameless. Further, the data shows that the value of PIRs (post-incident reviews) increases with greater levels of blamelessness. Interestingly, when there is a higher degree of just culture implementation, the impact of PIR work is more widely seen across the organization.

Total experience strategies need to be considered within a wider context

The reliability focus for external, customer-facing products or services eclipses that of internal, employee-facing tools or systems. Despite this skew currently being aligned with Elite performing organizations, there is evidence that innovation velocity is not sustainable, particularly amidst sustained work-from-home policies and an environment in which the only constant is change.

Toil numbers continue to fall

The SRE Report benchmarking data gleaned over the last five years reveals longstanding trends and patterns in terms of how SRE time is spent, providing a comparative gauge for SREs and businesses to determine adjustments as needed. This year, we saw levels of toil continue to fall while time spent working exclusively on engineering activities and time spent on call remains the same.

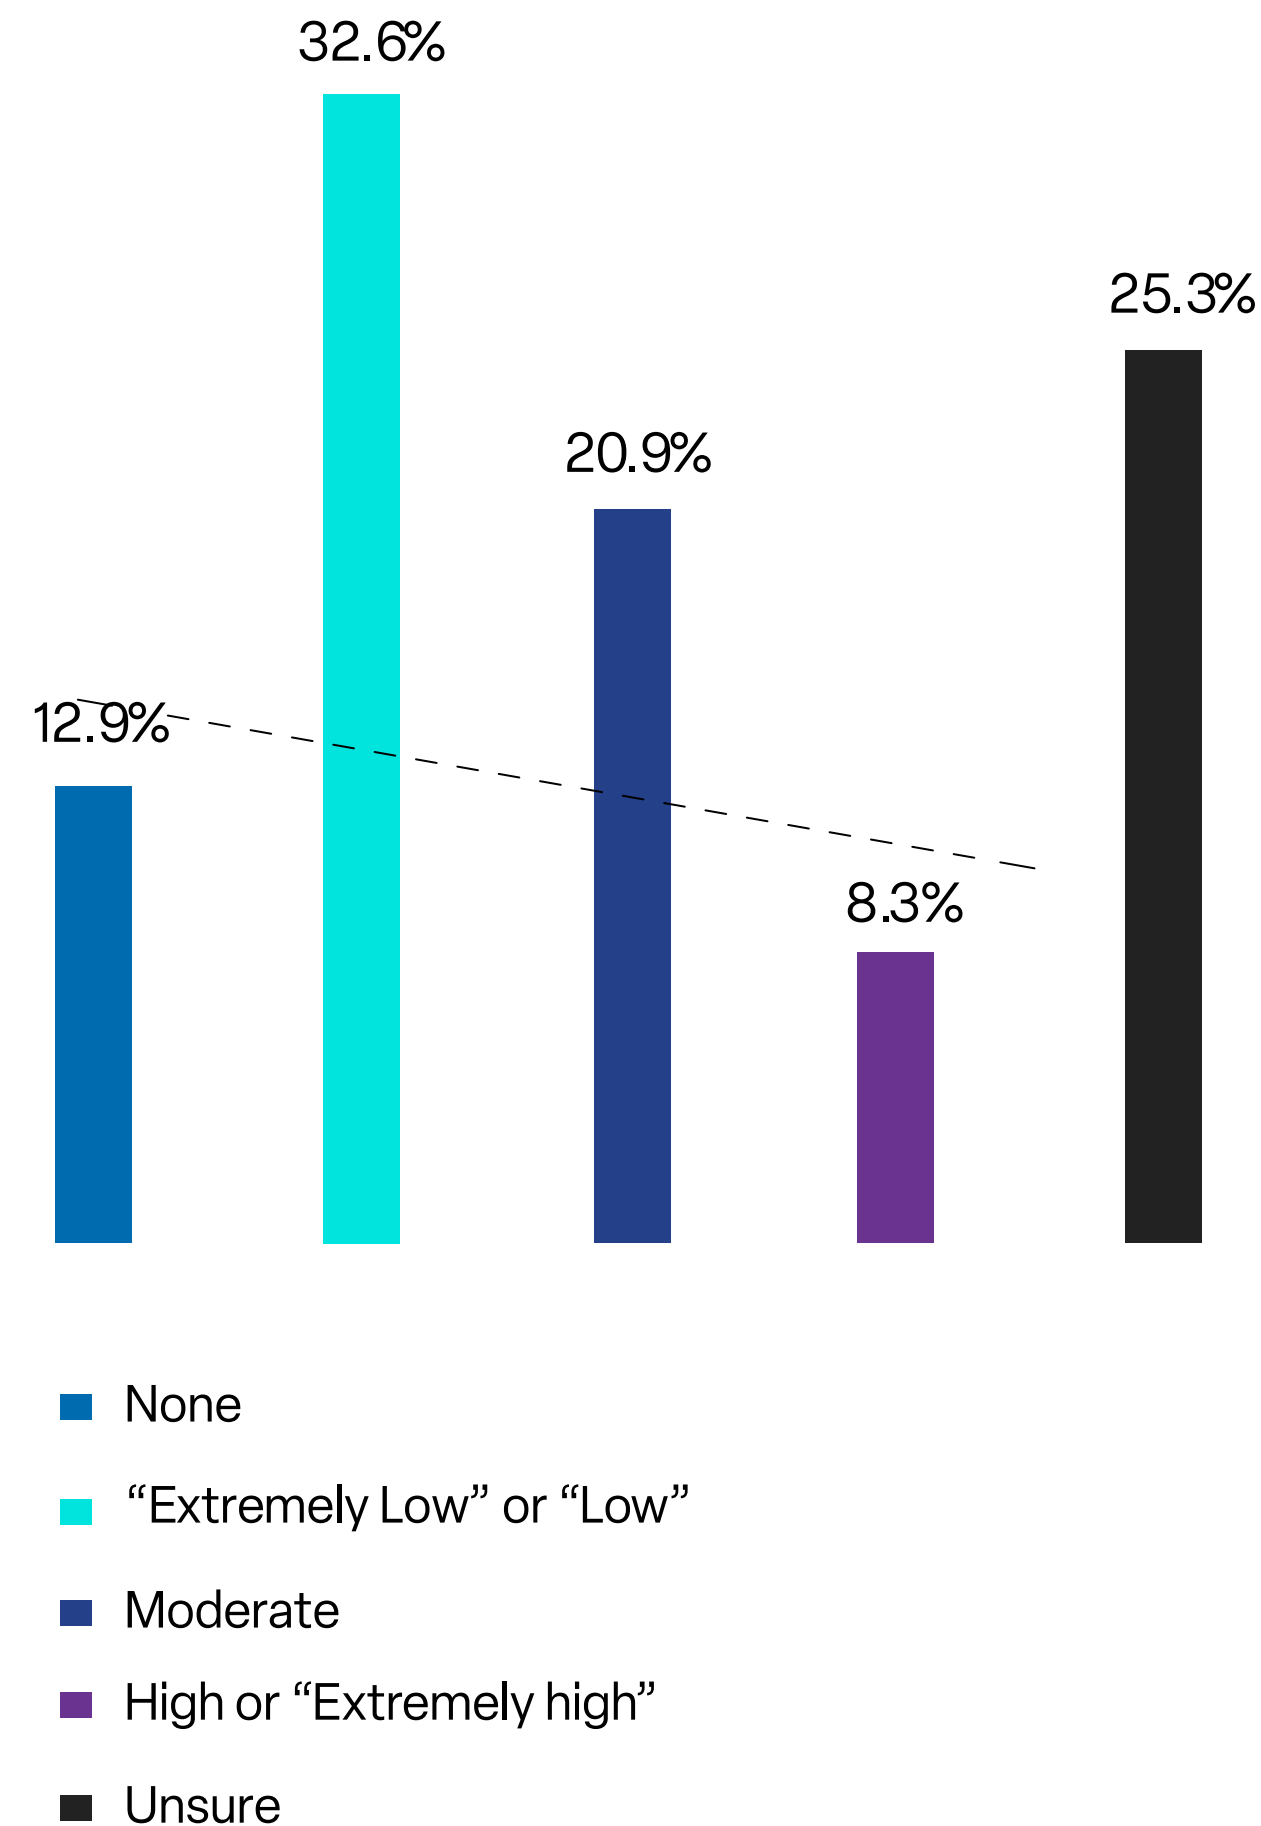
INSIGHT I

Insight: AI is not replacing human intelligence anytime soon - it should be considered within larger observability contexts.

Recommendation: Evaluate individual AIOps capabilities (e.g., the ability to correlate data or perform root cause analysis) against existing business use cases. Then audit your available resources (e.g., subject matter expertise, tool stacks, and processes) to determine which capabilities can be incrementally fulfilled on the path to larger, fuller observability implementations.




Please rate the value received from Artificial Intelligence for IT Operations (AIOps)



Is AIOps the panacea to wider business challenges?

Businesses have always faced a similar set of enduring challenges, such as battling fierce competitive landscapes, while trying to increase revenue streams, preserve existing customers, or protect the integrity of their brand. Increasingly, businesses are having to deal with new challenges, such as coping with a changing economy, the continued impact of The Great Resignation on retaining talent, and managing intellectual knowledge property (more on these further in the report). As businesses continue to undertake activities to improve efficiencies, achieve new economies of scale, or mitigate transformational risk, AIOps has been touted as the next-generation panacea to help achieve these goals. However, for the second year in a row, The SRE Report shows received AIOps value remains low - what should be done about this?

Individual capabilities are the gateway to positive business outcomes




We caution to ignore the hype around AIOps, but not to ignore AIOps itself. Break it down to smaller capabilities, and then evaluate each one against defined use cases, goals, or objectives. By breaking AIOps down to its smaller capability components, incremental value may be realized through a structured set of steps. Example capabilities include the ability to: detect anomalies, correlate events across different data sources, suggest (or even execute) remedial action, or perform root cause analysis. Usually, these capabilities are meant to improve customer experience or increase levels of employee productivity, with the ultimate goal of adding value to the business.

When considering these capabilities, there are various critical components to consider as foundational to pave the path to larger, fuller observability implementations.

Examples are:

Telemetry from the entire service production stack	Multi-directional, robust API ecosystem integrations
Telemetry from the entire service delivery stack	Telemetry from critical business KPI or experience-based perspectives
Multi-dimensional analytics and high data fidelity	Data repository




Is tool sprawl a problem?

AIOps and observability are high-level capabilities. They will probably require multiple tools in order to fulfill. While other research asks, “How many monitoring tools do you use?” and automatically equates that to tool sprawl, we decided to ask the question plainly, directly, and simply.

Despite what we regularly hear in the industry press, the majority of reliability practitioners said that tool sprawl was a minor or non-existent problem for their company. On the one hand, no one should simply say, “We need to add more tools to the stack because we don’t have enough.” Conversely, no one should say, “We need to remove tools from the stack because we have too many.”

When evaluating whether there is tool sprawl, the primary consideration is the value received from the tools versus their cost, where the cost takes many forms. It is crucial to consider that tool sprawl is not just, “How many tools are in the stack?” Rather, if the overall value contribution is net positive, then there is no tool sprawl problem.

How large of a problem is tool sprawl for your company?



59.1%
Percent of respondents who said tool sprawl was a minor or non-existent problem for their company

40.9%
Percent of respondents who said tool sprawl was a moderate or serious problem for their company

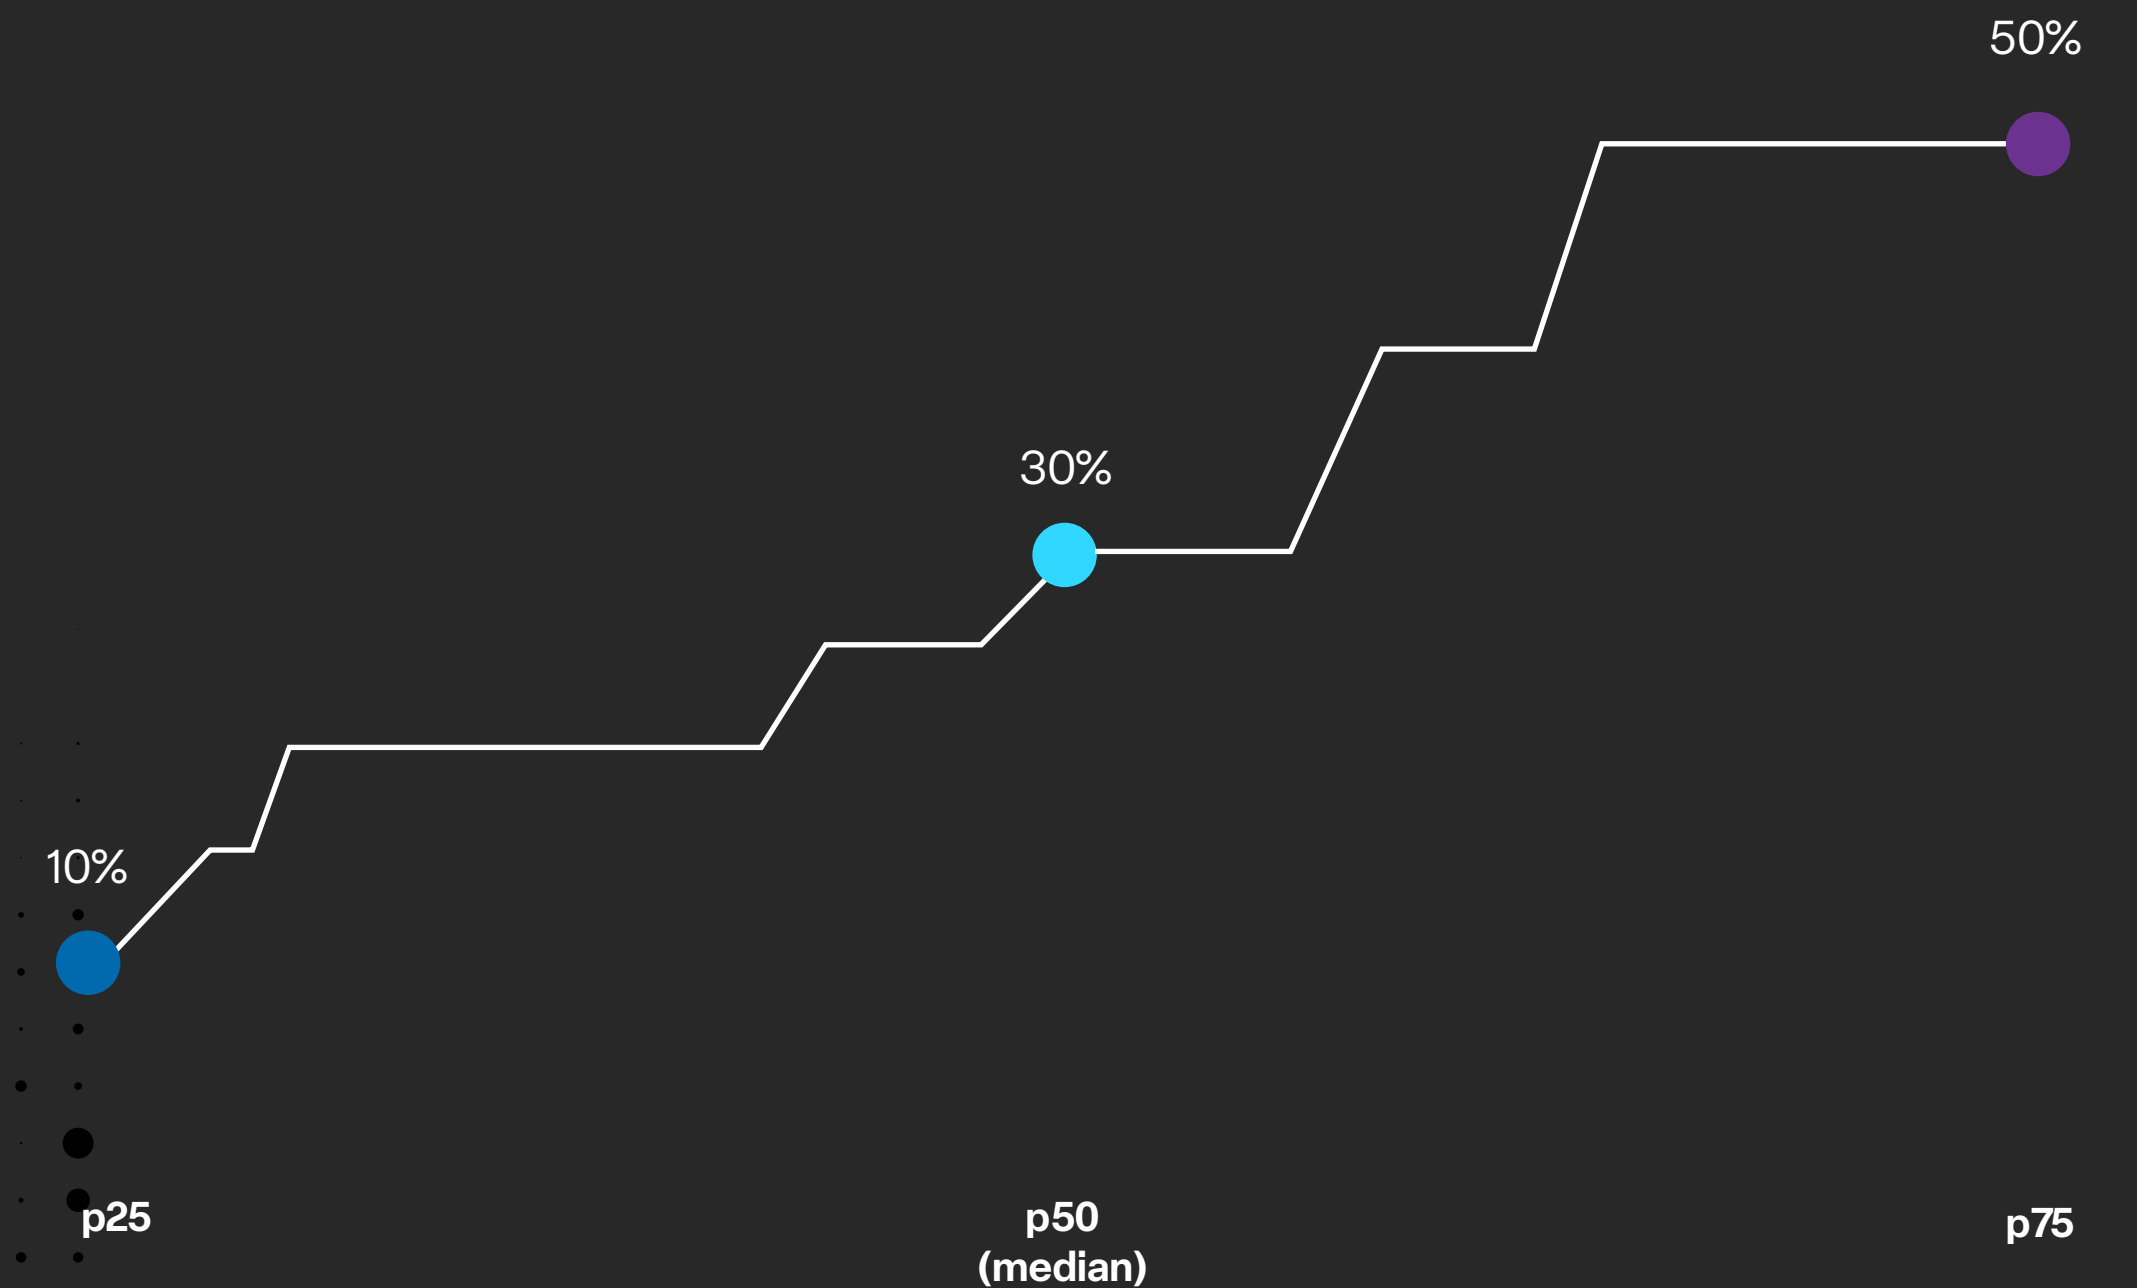
To build or buy?

30% *Half of the survey respondents said they build up to 30% of their tools*

When determining which capabilities should be built versus bought, consider existing risk profiles and how they may change over time. Whether building or buying, there is always an associated cost in the form of grunge (work that adds value, but that cannot be necessarily automated). This cost is often overlooked when compared to money cost, but what happens when your primary tool-building grungologists leave your company? Or what happens when your tool vendor cannot meet your customization needs (and you have to customize yourself)?

The above grunge cost is just one risk example to consider when having the build versus buy discussion. Other risk factors include harnessing and retaining tacit undocumented knowledge, building relationships, and retaining talent - all of which are consideration factors being impacted by The Great Resignation, for example.

What percent of your tools are exclusively built in-house (versus commercially purchased)?




Deep and wide data intake is critical

53.5%

[Derived] Percent of respondents who have three, or more, telemetry sources feeding their monitoring or observability frameworks

When considering the AIOps capability examples listed here, it is logical to say that deep and wide data intake is critical. This data should come from best of breed sources to avoid a “garbage in, garbage out” scenario. This means different data sources may come from different tools. If AIOps is supposed to enable IT or business observability, consider the current sources used to feed those frameworks. Since Observability and AIOps are a collection of capabilities, we expect their boundaries will continue expanding to more than just an IT realm - into the wider business.

Which telemetry feeds your monitoring or observability frameworks?



Conclusion: Consider how AIOps works in tandem with larger observability frameworks

SREs face a world of challenges when implementing reliability as a feature. Finding sufficient talent, realizing business value, complexities of architectures, and end-to-end visibility were the top challenges revealed by this year's survey data. When we think about these challenges in the context of AIOps and how AIOps works with larger observability frameworks, SREs have a golden opportunity to have new or better business conversations. Else, they will further perpetuate the IT-to-business gap, which is discussed in the next section.

We have seen the concept of AIOps and Observability expand to more than just the IT stack. Observability is not a tool. Rather, it is a capability of being able to infer relationships between inputs and outputs. Hence, terms such as Data Observability, Business Observability, or AI-Bops ("Artificial Intelligence for Business Operations") may become an SRE 'necessary evil' as non-tech members learn about these new capabilities to achieve positive business outcomes.

For SREs, this means the definition of observability would need to encompass more than just the ostensible "events, metrics, and traces" as has been discussed in an application context. Therefore, we recommend SREs do not treat these expanded observability boundaries as *word salad*.



View from the field

“Since AIs are not sentient, they cannot understand what happens inside your company and your business or between people. AIOps can definitely help with data, but Observability needs SRE eyes and ears - to be fully embraced.”

Simone Cibba **Site Reliability Engineer, Kyndryl**



“The quantity of information available from observability systems today can be really overwhelming. There is certainly potential for AIOps to help, by surfacing relevant information - anomalies, outliers, correlations, clusters of related information. However, we’ve also seen AIOps projects that simply generate a lot of meaningless noise, and it doesn’t seem that AIOps is currently able to deal with truly complex problems in distributed systems.

This is an area with a lot of hype, and there are a lot of companies selling AIOps solutions right now, but there are still relatively few success stories from practitioners. As with everything else, whether to invest right now depends on whether your organization is experiencing the kind of challenges that are currently amenable to AIOps. For now that means generic problems like detecting infrastructure issues such as resource exhaustion or saturation, or forecasting future capacity usage - problems that some organizations have solved well with traditional approaches.”

Laura Nolan **Principal Software Engineer, Stanza**

What is the number one challenge hindering successful reliability implementations?

Talent (hiring, retention, assimilation)	7.9%
Complexity of architecture	7.5%
Business value is hard to realize	6.7%
Lack of end-to-end visibility	6.3%
Alignment or prioritization	4.2%
Time management	3.8%
Communication or collaboration	3.8%
Knowledge, training, or education	3.3%
Lack/misuse of resources	2.9%
Cost or budget	2.5%
Perpetual evolution or change	2.5%
Balance - velocity versus reliability	2.5%
Lack of buy-in	2.5%
Sprawl - tools	2.1%
Culture	2.1%

[See Appendix I for full list](#)

Editor's note on challenges

This year we decided to ask, "What is the number one challenge hindering successful reliability implementations?" as an open-ended question. This gave survey respondents the ability to type in anything they wanted - and the data invoked a tremendous amount of emotion when we parsed it.

While survey respondents were presented with example answers to give an idea of what we were looking for (a type of anchor bias), multiple respondents did in fact type in their own. In the data presented here, we have bolded the anchor bias answers to provide a richer, more powerful set of insights. For example, even though, "tool sprawl" was an anchor bias, notice how it ranked lower on the list, even lower than some non-anchor bias answers. And notice the top challenge was not an anchor bias at all! Enjoy.

INSIGHT II

Insight: Executives are from Mars. Individual Practitioners are from Venus.

Recommendation: Underneath the averages (which are notoriously misleading), it's clear a dichotomy exists between individual practitioners and executives on the received value derived from AIOps and other researched topics. We recommend these personas find new opportunities to communicate and collaborate: reevaluate feedback loops, align to define shared goals and objectives, and drive accountability for data-driven decisions versus making them based on authority or funding alone.



It's time to have new or better conversations

Since received AIOps value varies considerably by role, we recommend reevaluating communication and feedback loops. This dichotomy may exacerbate preexisting frustrations between practitioners and executives, especially when it comes to who has decision-making authority versus who is held accountable if little or no value is ultimately received.

When having new or better conversations, remember that individual capabilities are the gateway to positive business outcomes - so focus on which capabilities are required and why. For example, on one end of the spectrum, an individual practitioner discussing "speeds and feeds" may make it difficult for a business leader to understand how value may be realized. On the other hand, simply saying, "We need to show value" is nefarious nothingness and may make it difficult for practitioners to know which "speeds and feeds" are important.

Bridge this gap by discussing capabilities and getting aligned on what greater use they deliver to the business, such as contributing to lowering customer churn or increasing average client spend. It's important to be aware that any conversations across a power gradient can be fraught with risk for the "less powerful" person who will be potentially risking their career by giving bad news about a "pet project". There may also be deeper, darker reasons for this dichotomy, but nothing ventured, nothing gained. Consider an "agile conversation" approach, which is built around defusing psychological land mines by using conversations to foster high trust, low fear through understanding why, making commitments, and being accountable. Overall, when reevaluating communication and feedback loops, ensure a just culture of openness, sincerity, and transparency. As seen later in this report, a just culture has substantial business value opportunities.

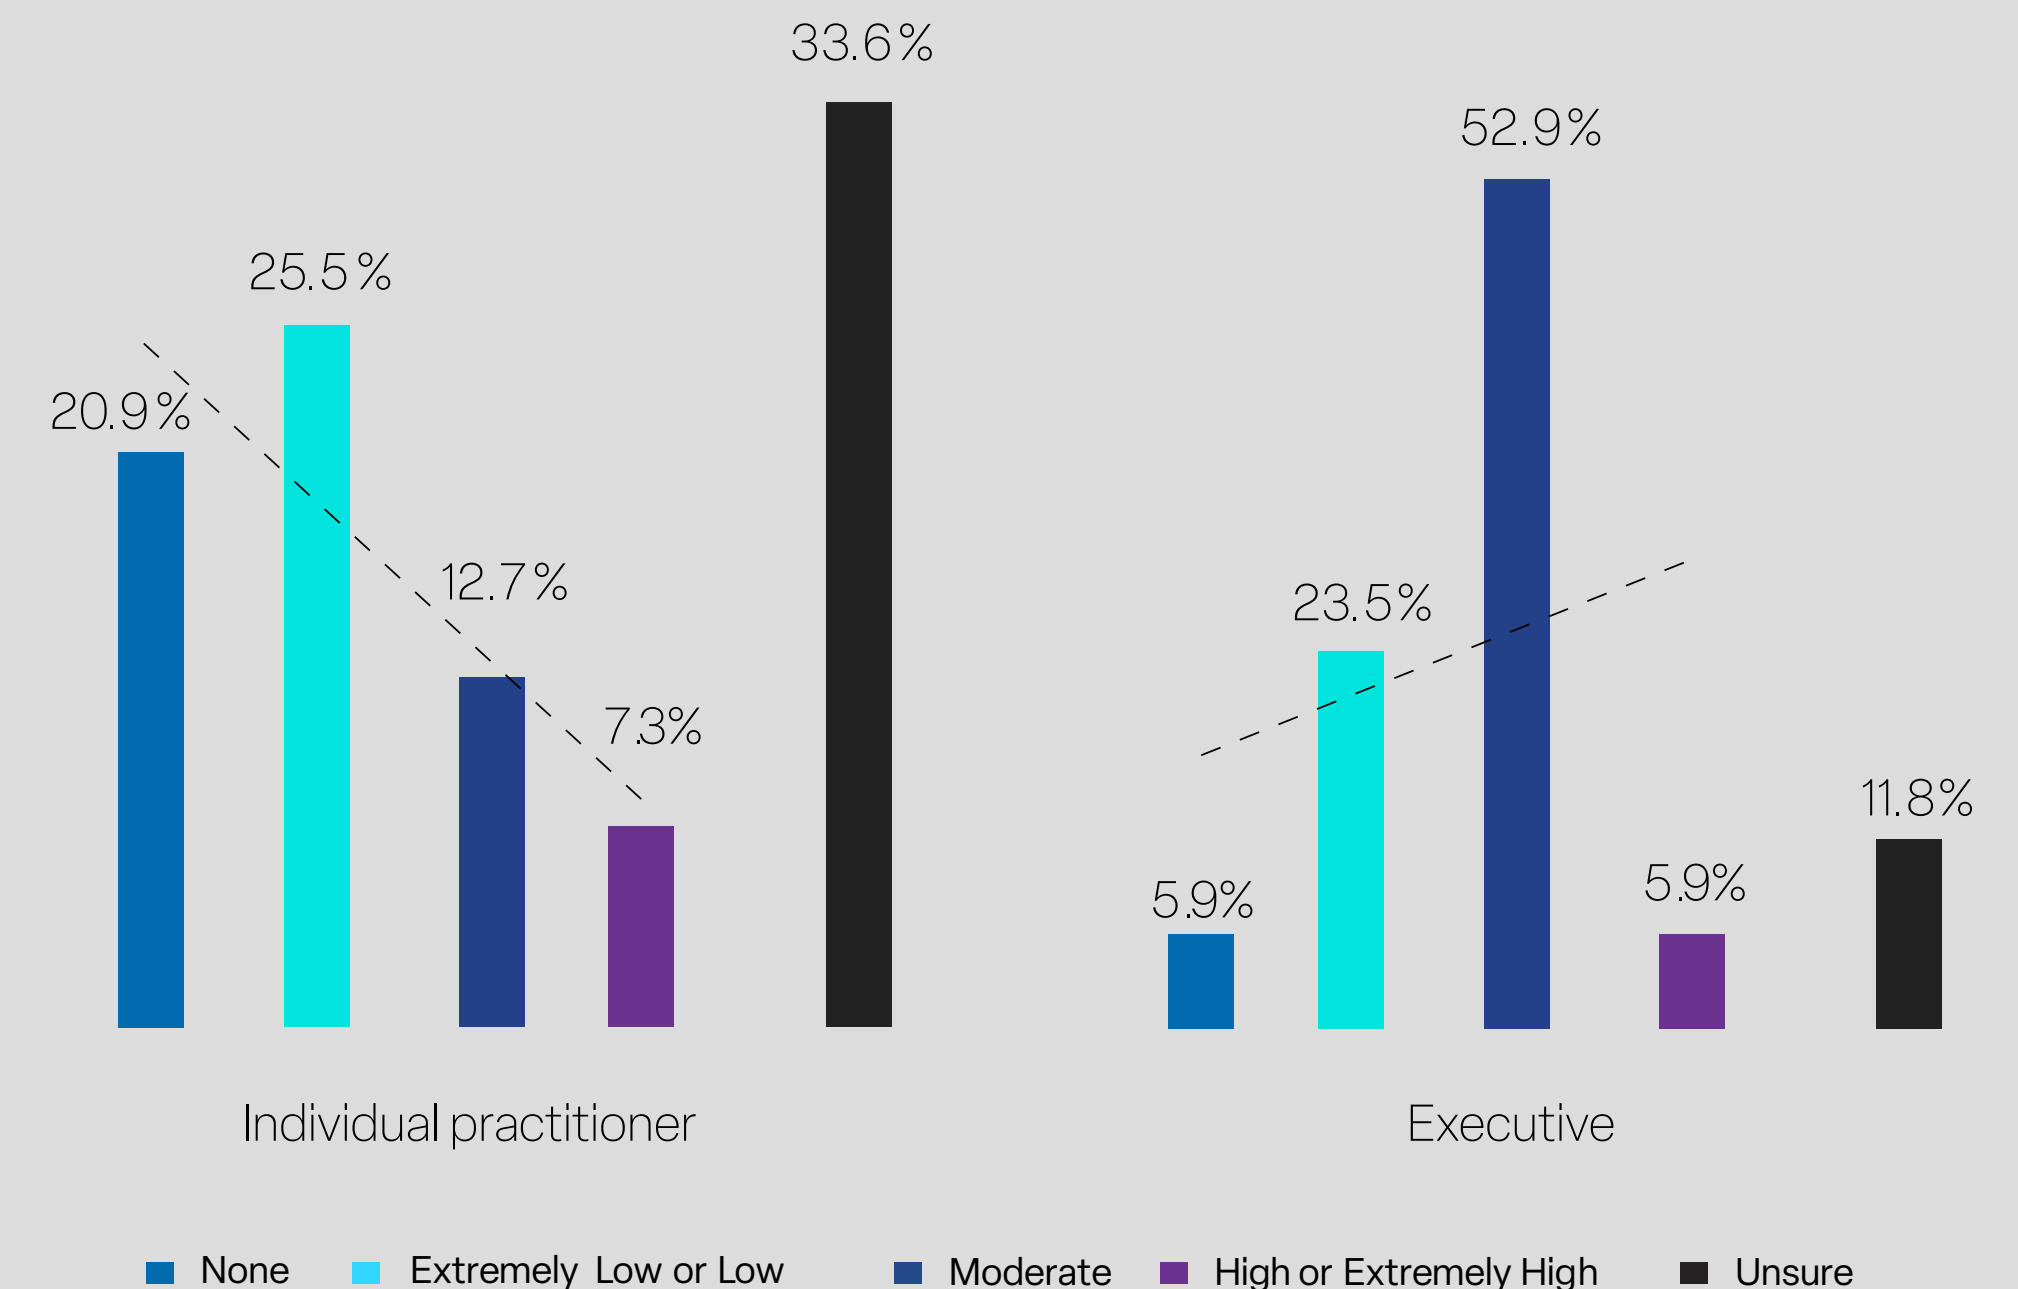
20%

Percent of Individual Practitioners who receive moderate or *high* AIOps value

58.8%

Percent of Executives who receive moderate or *high* AIOps value

Please rate the value received from Artificial Intelligence for IT Operation (AIOps)



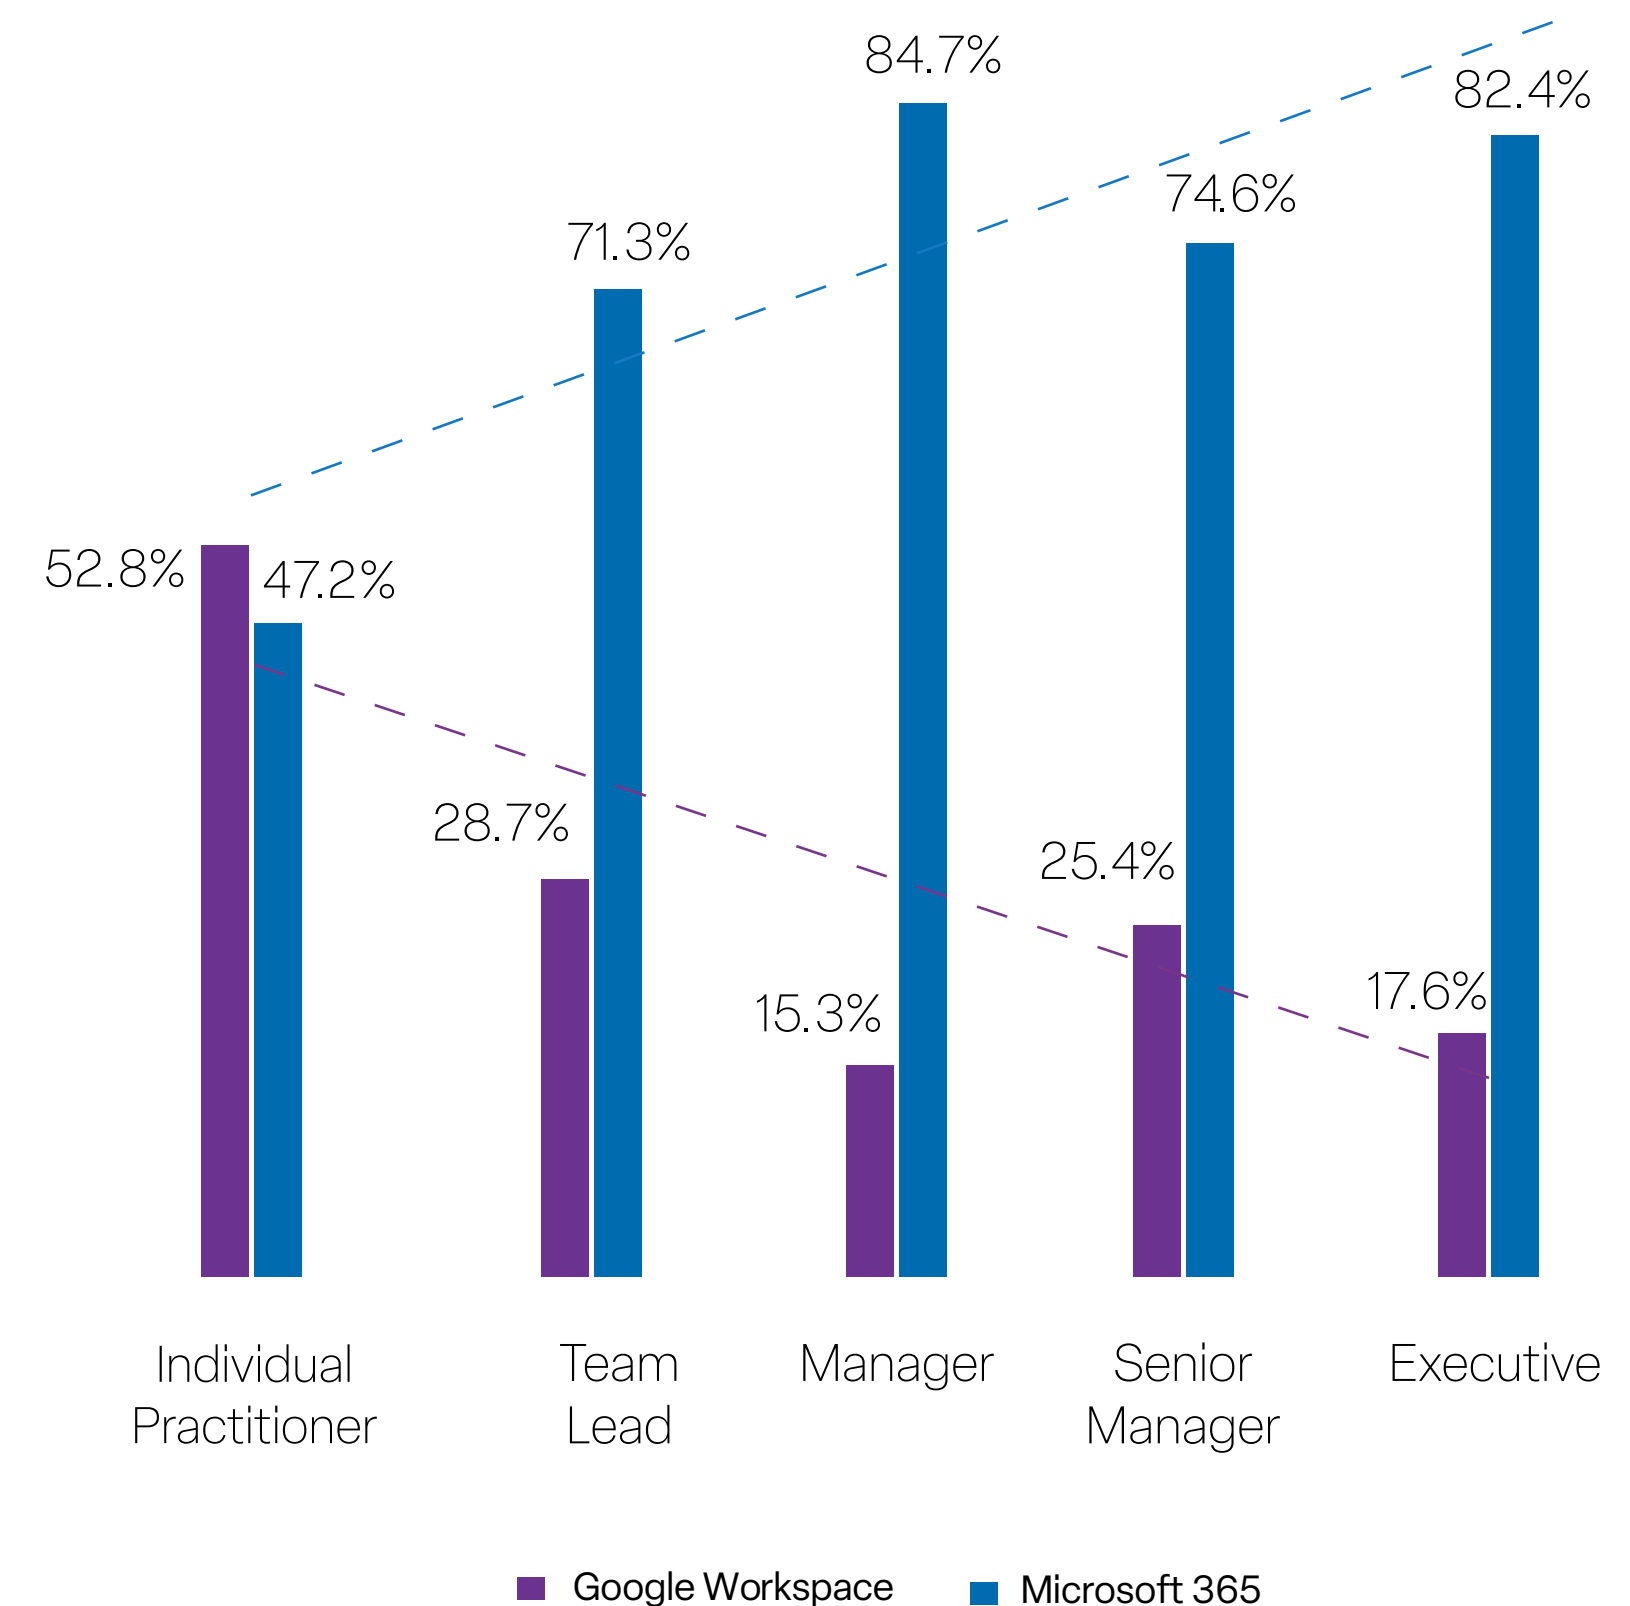
Consider capabilities instead of bias

Individual practitioner: the only respondent persona to prefer Google Workspace versus Microsoft 365


The AIOps value dichotomy between practitioners versus executives is an example of a larger communication challenge; where do other potential gaps exist? This disparate view also probably applies to more than just IT or reliability practitioners – businesses will want to use this survey data to initiate meaningful, purposeful conversations across all their organizational units.

As mentioned, individual capabilities are the gateway between “speeds and feeds” to positive business outcomes. An example conversation may be around the ability to collaborate with external vendors versus choosing one tool that aligns to your personal beliefs, bias, or convenience of licensing.

Prefer Google Workspace or Microsoft 365?



Tool sprawl is not a problem



43%

Percent of Individual Practitioners who said tool sprawl is a moderate or serious problem

22.2%

Percent of Executives who said tool sprawl is a moderate or serious problem

How large of a problem is tool sprawl for your company?

Individual Practitioner



Executive



■ "Not at all" or Minor ■ Moderate or Serious

While both personas had a net "Not at all" or "Minor" response to whether tool sprawl is a problem, the delta between the personas is noticeable. Consider the outcomes you are trying to achieve and remember that tool sprawl is **not** just, "How many tools are in the stack?" Rather, it is an evaluation of the received value versus cost - where the cost takes many forms (for example, by virtue of having multiple tools, it may cost more to do training). Remember, if the overall value contribution is net positive, then there is no tool sprawl problem.

But the story goes deeper. On the one hand, individual practitioners may need multiple tools and the ability to quickly throw something together to achieve a spot result. In this case, these types of tools may be completely invisible to executives. On the other hand, executives may be considering only tools for which there are line items on a balance sheet (versus those which are silently built). Therefore, both personas have an opportunity to take stock of all the tools in the stack where the outcome is improving conditions for all roles in the organizational chart (in part by valuing the tools different groups use to perform their jobs)

Find balance and ensure stack resiliency

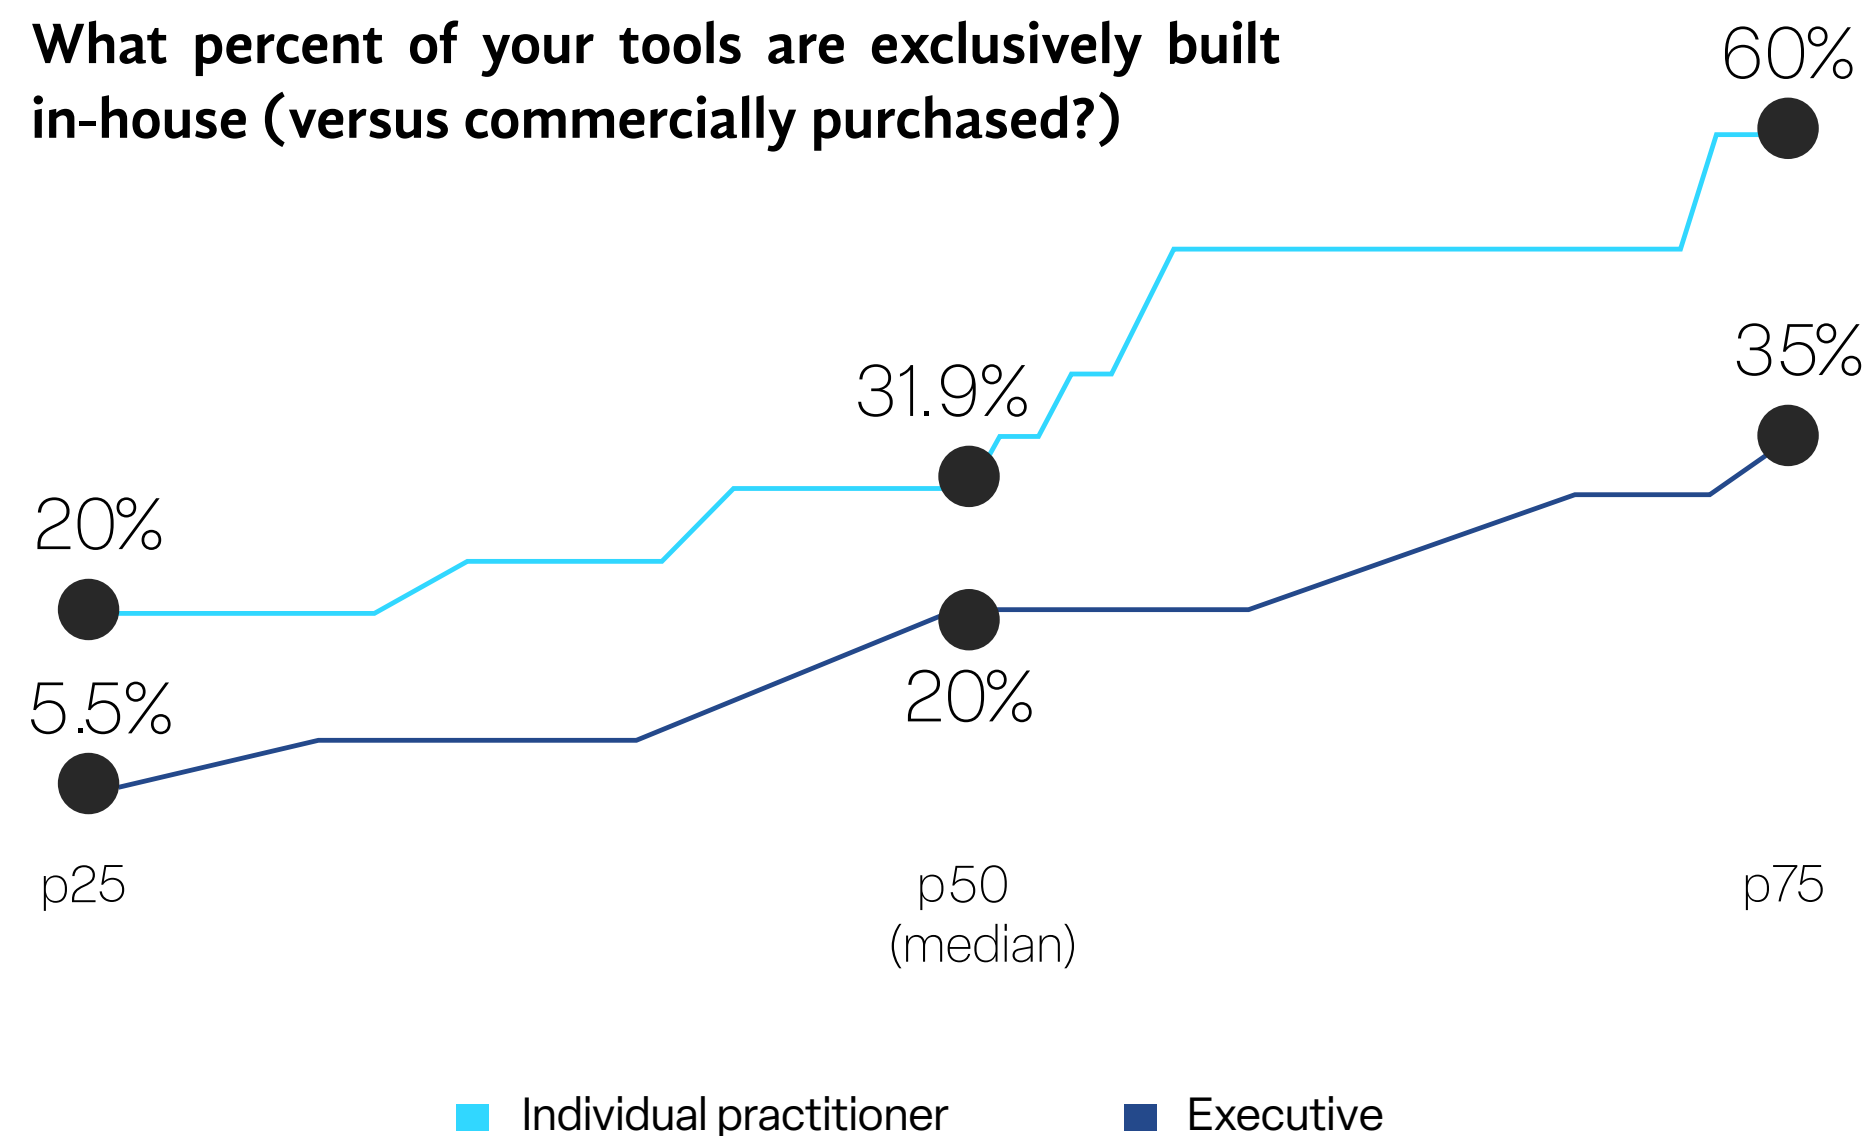
The build versus buy debate usually occurs only when companies have both (1) the money to buy and (2) competent staff to build. If companies have only one versus the other, then there really is no debate. When deciding what build or buy combination to use, practitioners will want to have open conversations with executives around where debt or risk will potentially lie. Find balance by having enough diversity in order that the entire resource stack (to broadly encompass e.g., people, places, and processes - in addition to tools) has enough resiliency to withstand change or volatility, regardless of whether tools are bought or built.

31.9%

Half of Individual Practitioners said they build up to 31.9% of their tools

20%

Half of Executives said they build up to 20% of their tools



Conclusion: Have new, better or more agile conversations across the divide

All persona ranks of the professional ecosystem - from individual practitioners, to architects, to executives, and to external partners - have the opportunity to use this report for data-driven decisions. When having new, better, or more agile conversations, remember that specific capabilities are the gateway to greater alignment and positive business outcomes. They bridge the gap between “speeds and feeds” on one end of the spectrum versus “business value” on the other.

But what form/s does “business value” take? And is everyone viewing this abstract value concept through the same lens?

In this next provocative dataset, notice how the number one answer to, “How do reliability engineering practices add business value to the company?” was a tie between an anchor bias and non-anchor bias answer. In other words, survey respondents were able to type in anything they wanted (see Editor’s note). The other startling revelation is to notice how several of the answers are arguably not business value contextual; this also aligns with the “business value is hard to realize” answer to top challenges.

We suggest not to let bias, politics, or emotion influence the reading of these responses. Some answers may seem painfully self-evident to some, while others may seem painfully without business context.



How do reliability engineering practices add business value?

Lower cost	12.5%
Customer experience, sentiment, or satisfaction	12.5%
Maintain reliability, performance, or uptime	11.1%
Retain existing customers	6.5%
Avoid SLA penalties	6.0%
Increase operational efficiency	5.6%
Increase new logo count or revenue	4.6%
Talent attraction/retention	3.7%
Self-evident business value	2.8%
Improve time to repair	2.8%
Trust or integrity	2.3%
Increase or preserve brand integrity or reputation	1.9%
Increase innovation velocity	1.4%
Increase business competitiveness	1.4%
Remove technical debt	1.4%

[See Appendix II for full list](#)



View from the field

“My career to date has spanned a number of different roles – from CEO to Solution Architect to Manager of SRE/DevOps. In looking at the disparity between how individual practitioners (SREs) and executives regard several core SRE tenets, such as the value of AIOps and whether or not tool sprawl is a challenge, it illustrates not only a lack of communication between the groups, but also their different priorities and perspectives. Greater communication to improve alignment is essential, as is improved clarity in setting business goals across teams. Not to mention that accountabilities regarding tools also needs to be clearly defined so that maximum value is derived from every tool.”

Adriano Velasco Nunes

CEO and Co-Founder, Stagefy



“The survey data shows a disparity between Individual Practitioners and Executive Leaders that will continue to grow until there is more wide-spread adoption of tools that help unify process, communication, and decision making. To put it plainly: we need to adopt more DevOps practices!”

Keri Melich

SRE, Nobl9

INSIGHT III

Insight: The power of high Blamelessness and valuing postmortem learnings are characteristics of Elite performing organizations (compared to non-Elite organizations) and are not tied to company size.

Recommendation: In an environment where the only constant is change, we recommend companies reevaluate both written and unwritten rules of the road, especially with respect to embodying a “just culture”. SREs have an opportunity to increase and improve both the amount - and reach - of value from postmortem incident reviews (PIR).




Why blameless?

Over the last 10-20 years, resilience engineering (RE) - and the human factors contributing to effective operations in complex environments - has been moving from an academic research area into daily operations across a wide variety of industries. Aviation, medicine, and even mining have benefited from these insights. Online services have changed many of the “ground rules” that applied to “shrink wrap” software. With the ability to rapidly deploy new versions, the software services industry has also begun to learn from resilience engineering. With the dramatic social upheaval of the worldwide COVID pandemic and its ongoing impact on our lives, including sociotechnical systems, we thought that it would be beneficial to retest some of the cultural aspects of high performing organizations.

[Erik Hollnagel](#) characterizes four key capabilities of a resilient system or organization:

- Responding (to opportunities and threats)
- Monitoring (how to recognize environmental changes)
- Learning (adjusting both monitoring and responses, post hoc)
- Anticipation
 - ◇ Planning on the basis of leading indicators for environmental shifts
 - ◇ Recursively - with both internal and external inter-related changes



The value of a “just culture”

The concept of “just culture” evolved during the 1990s in the context of safety science. [Pioneered by James Reason and David Marx](#), “just culture” is the concept that organizations should be responsible for the sociotechnical systems they put in place and that employees are generally intent on doing the best they can within the wider context of demands placed upon them. The purpose of a “just culture” is to use learning from incidents to lead to forward-looking accountability, which changes ineffective system dynamics. As such, a “just culture” organization seeks to ensure that reporting of incidents (monitoring) is not met with punitive measures since such a response tends to dampen willingness to recognize and report incidents. Without recognition, none of the latter stages of resilience can take place.

Psychological safety is critical

Cognitively, for individuals, learning takes place as a higher-level function in the prefrontal cortex involving the use of judgment, reason, and memory. Psychological safety is essential to enabling learning to happen. When blame enters the picture, neuroscience has shown that this invokes the amygdala and the resulting defensive reaction is known as an “amygdala hijack” - which cuts off the higher mental functions and prevents learning. [Research by Amy Edmonson](#) since the 1990s has shown a strong connection between psychological safety and teams being able to learn. [Google’s Project Aristotle](#) also confirmed the strong connection between team effectiveness and psychological safety.

How are organizational effectiveness and blameless interlinked?

Dr. Ron Westrum who, like Edmondson, is a researcher in human factors related to safety and organizational performance, discovered that an essential component of organizational effectiveness is related to the free flow of good information (good = pertinent, timely, usable). To form a typology of [different organizational cultures](#), Westrum correlated characteristics such as:

- Messengers are not punished when they deliver news of failures or other bad news.
- On my team, responsibilities are shared.
- On my team, cross-functional collaboration is encouraged and rewarded.
- On my team, information is actively sought.
- On my team, failure causes inquiry.
- On my team, new ideas are welcomed.

Similarly, in the research by DevOps Research and Assessment ([DORA](#)), led by Dr. Nicole Forsgren, the “generative” culture of the Westrum typology, which optimizes information flow, is one of the indicators for high performing organizations. A generative culture construct has six distinctive characteristics: (i) high cooperation, (ii) messengers are trained - to deliver good information, (iii) risks are shared, (iv) bridging is encouraged, (v) failure leads to inquiry, and (vi) novelty is implemented. This **is** a blameless culture.


A brief note for those who might pursue further reading in this area: Over the last few years, some practitioners have advanced the perspective that it is not possible to completely avoid blame from a psychological point of view so they have advocated the use of the term “blame aware” - as in mitigating the natural tendencies toward blame. We have kept with the more common parlance of “blameless” in both the survey questions and this writeup.



How do blamelessness and post-incident reviews relate?

In this year's survey, we asked about blamelessness, the value received from post-incident reviews, and how broadly distributed that value was within the organization. We also asked the "DORA Four" questions (see Appendix III). We then used their grading rubric to classify organizations as either low-performing, medium-performing, high-performing, or elite-performing to examine correlations between these different measures.


The principle of learning from failure (and near misses) which has been widely implemented in other fields such as aviation and medicine leads to carrying out post-incident reviews (PIRs). In some organizations, PIRs are known variously as "retrospectives" (though not the same as "agile retrospectives") or "post-mortems" (as in "the system died, what can we learn from this?").



Organization size is not predictive (good or bad) of blamelessness³


Regardless of the level of blamelessness reported by respondents, the curve tied to company size is fairly consistent and matches the distribution of respondents across company size ranges. This is good news because it says that you can be highly blameless whether you are based at a tiny company or a huge mega-corporation.

How effectively do you achieve the goal of being “blameless” in your postmortems (by company size)?



³A note on the categories of “blamelessness”: We have kept with the terms used in the survey responses in this report, but it may be easier to understand the range of responses as “Very blameful”, “Somewhat blameful”, “Slightly blameless”, “Very blameless”, “Extremely blameless”.

Different roles perceive blamelessness differently



44.7%


Percent of Individual Practitioners saying extremely or very blameless

9.6%

Percent of Executives saying extremely or very blameless

Looking at the level of blamelessness attributed to their groups, we notice another perception gap across different roles. Engineers consider their organizations tend toward the higher end of the blameless spectrum while managers and executives less frequently put their companies into the highest levels of blamelessness.

What does this dichotomy reveal about blameless culture within an organization? Do managers and executives prefer a more blame-based culture? Or do they hold themselves and their organizations to a higher (relative) standard than individual contributors?



Value of PIRs increases with blamelessness

57.9%

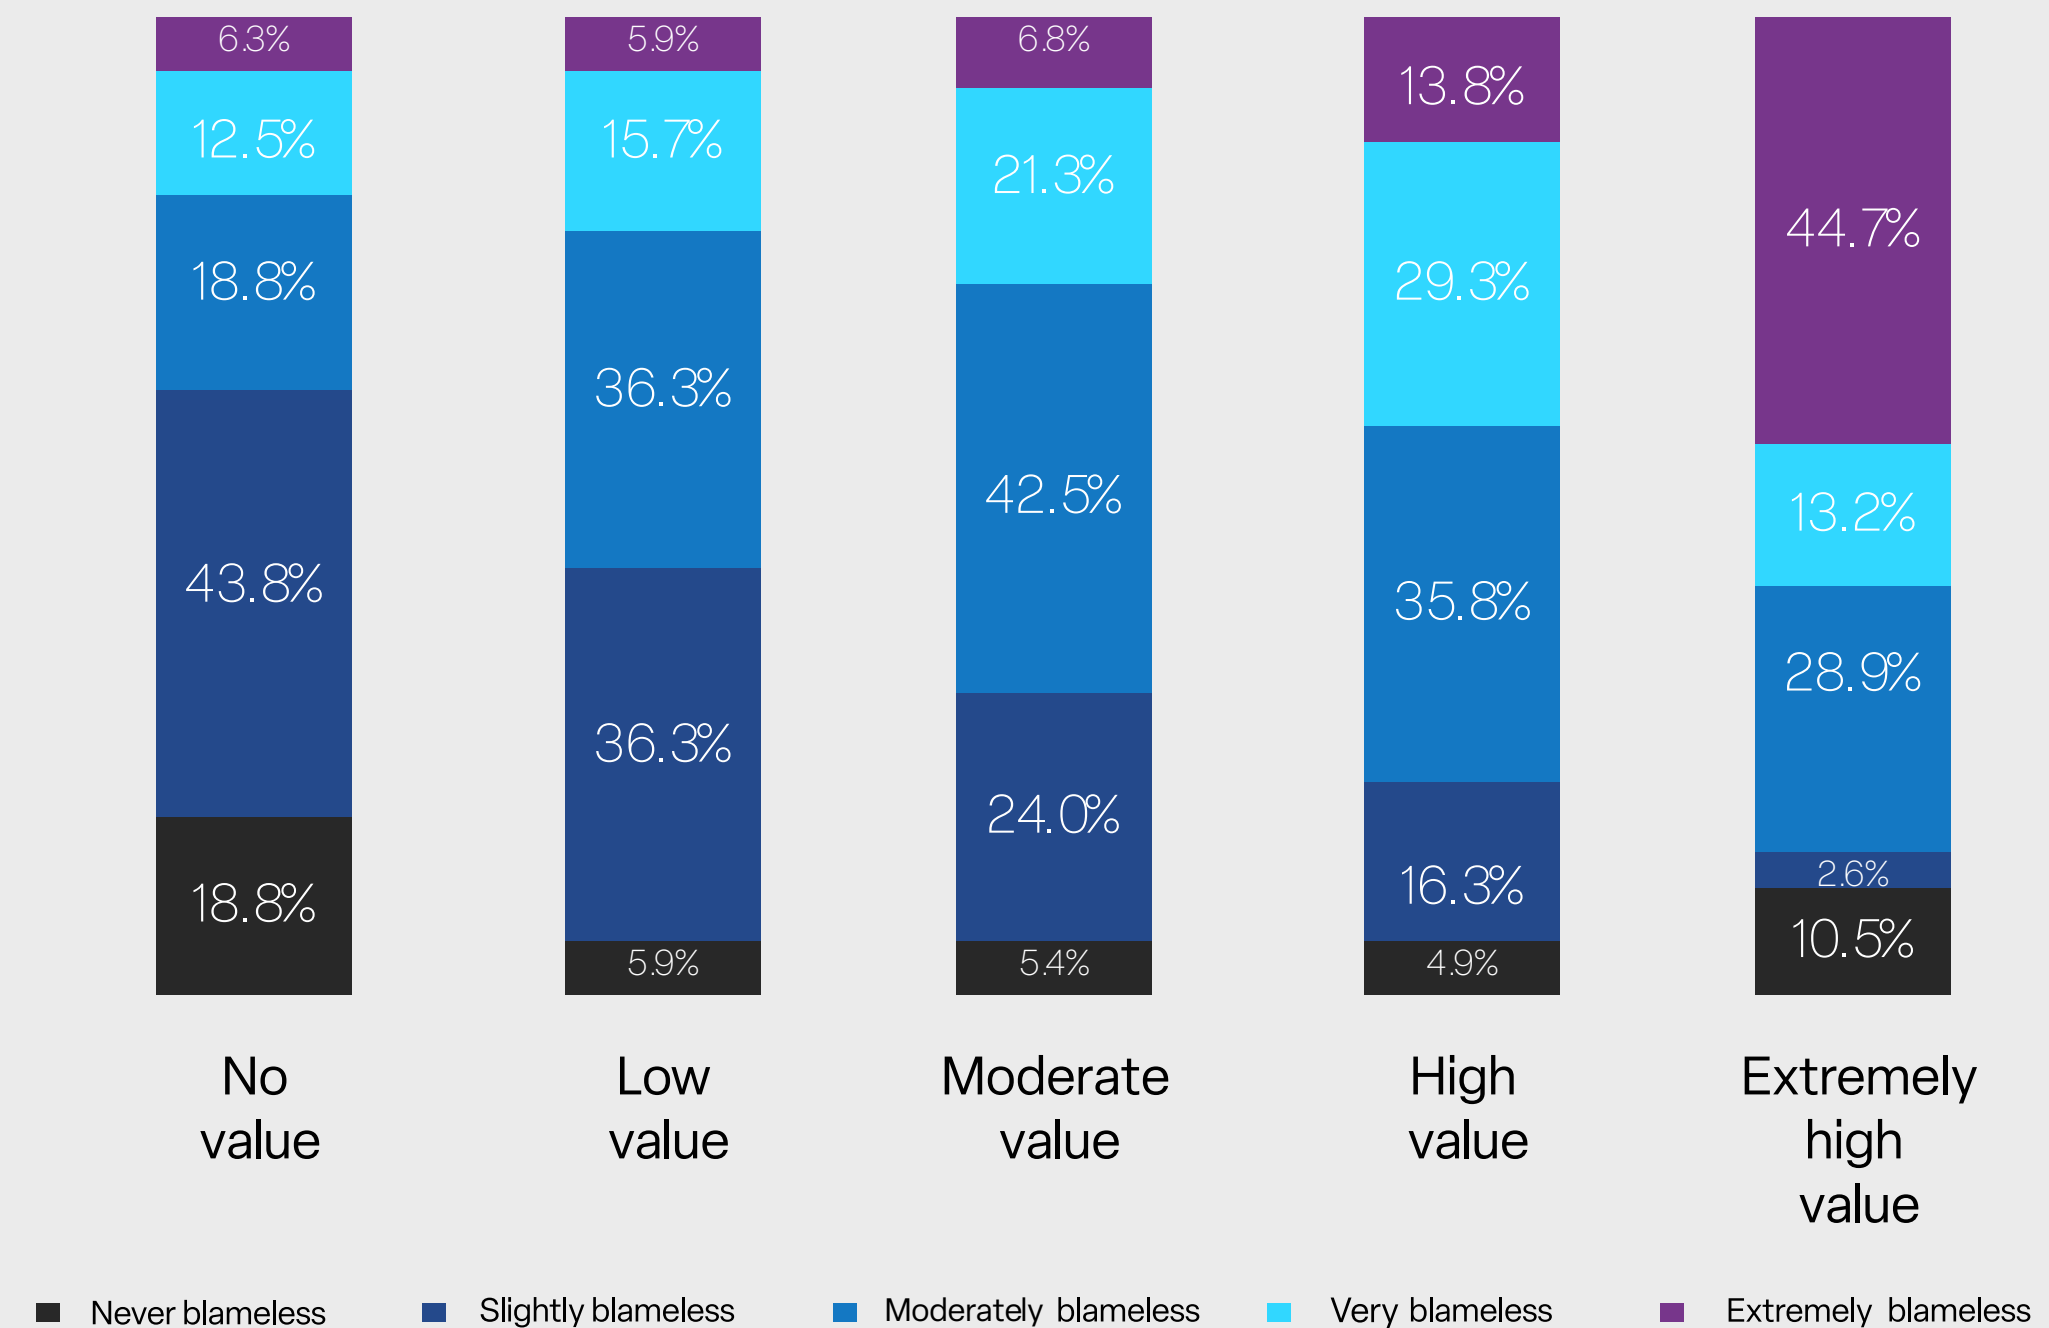
Percent of very or extremely blameless respondents who say they get extremely high value from PIRs

13.1%

Percent of never or slightly blameless respondents who say they get extremely high value from PIRs

Interestingly, companies that perceive the value received from conducting PIRs to be significantly higher tend toward “extremely blameless” cultures when compared to lower levels of blamelessness.

Please rate the value received from incident postmortems/retrospectives by level of blameless



How does DORA performance rating align with self-reported “just culture”?

Since we also asked respondents to provide information related to the four primary DORA constructs, we looked at the alignment between the DORA performance categories and the reported blamelessness of the organization and found a good correlation. Higher performing organizations (per the DORA definition) report they are “extremely” or “very” blameless:

72%

Percent of Elite performers who report very or extremely blameless

36.3%

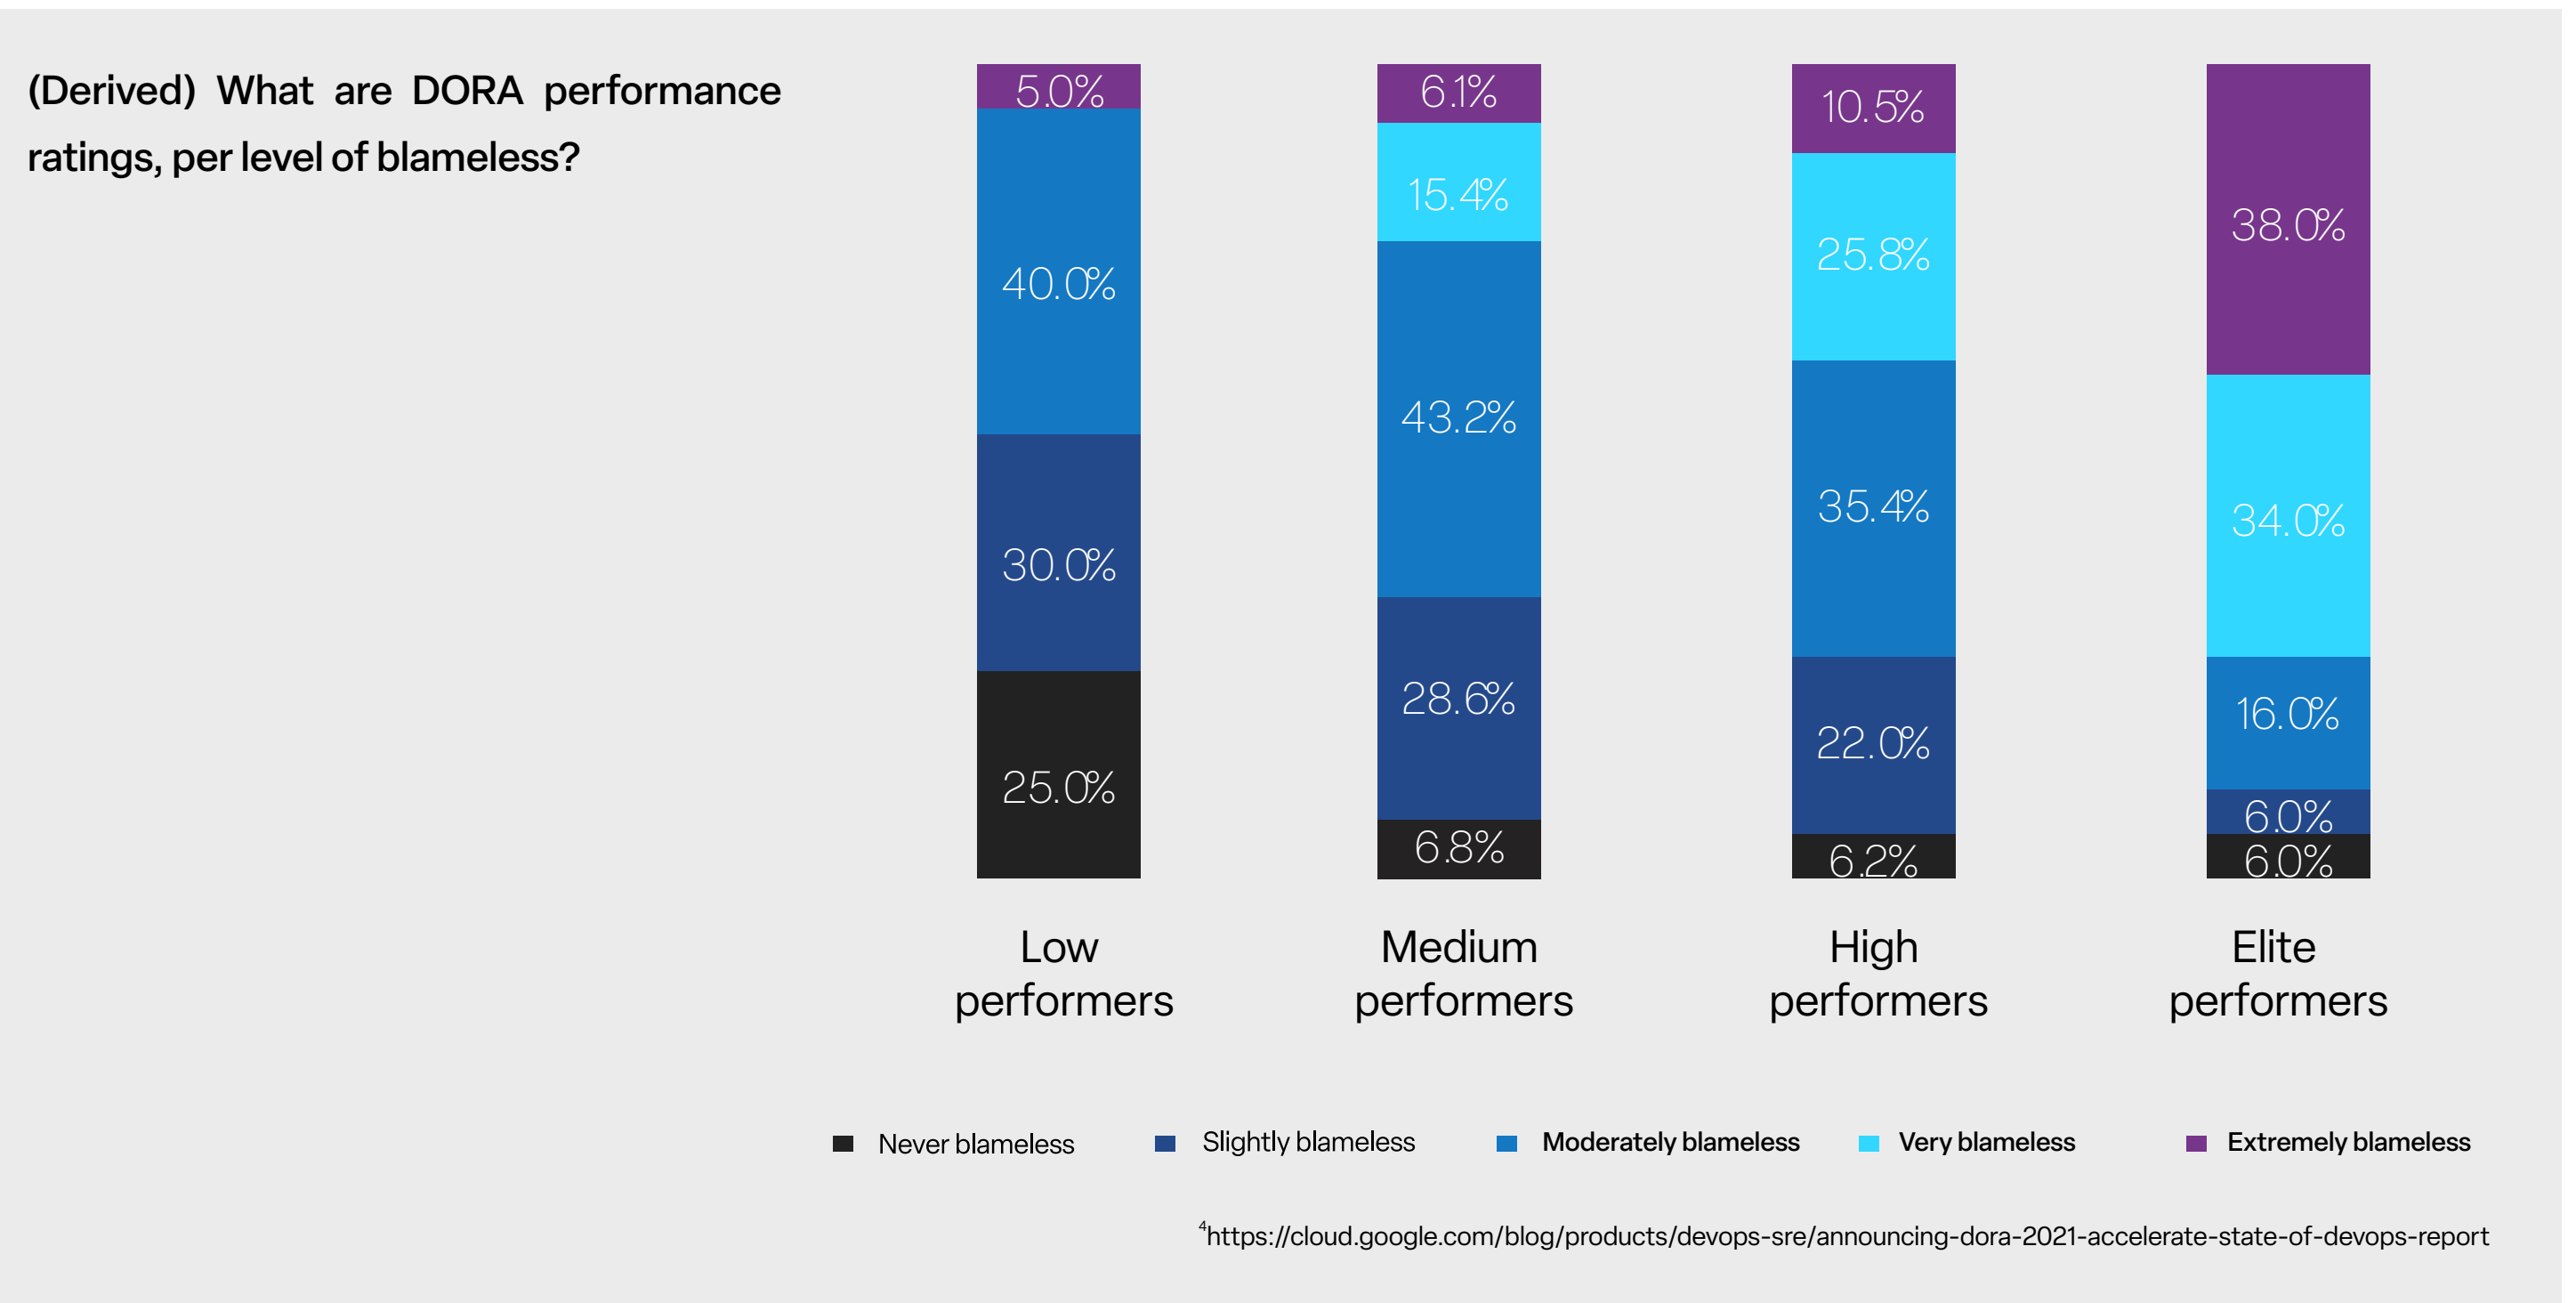
Percent of High performers who report very or extremely blameless

21.5%

Percent of Medium performers who report very or extremely blameless

5%



Percent of Low performers who report very or extremely blameless



Breadth of PIR value expands with greater blamelessness

As their levels of self-reported blamelessness increased, post-incident review participants said other parties more frequently received PIR value. In other words, a higher degree of “just culture” implementation also leads to a wider impact radius for the PIR work.

As the level of blameless increases, so too does the PIR value frequency.



Conclusion: Post-incident reviews can unlock significant value

Post incident reviews (PIRs) was the area where we focused this inquiry and analysis. As a tool and a practice, the value clearly correlates both in breadth and depth with the highest levels of organizational performance (as identified in the DORA research). Since “finger pointing” and/or scapegoating are both blameful activities which can happen in the wake of incidents, the data shows that avoiding such patterns also aligns with high performance.

From the open-ended questions about hindrances to reliability programs, PIRs have the opportunity to address at least four key categories:

Alignment or prioritization	4.2%
Communication or collaboration	3.8%
Knowledge, training, or education	3.3%
Perpetual evolution or change	2.5%

Driving a culture of open communication, curiosity, and learning through consistent behaviors in and around post-incident reviews can unlock significant value from existing teams and help to promote a wider awareness of customer-centric reliability.



View from the field


“The benefits of PIRs and ‘blamelessness’ are notoriously hard to quantify, it would be very useful for companies to institute some measure of ‘Culture Debt’ and track it transparently across various categories for them to truly understand the benefits they can see from this and see whether their efforts in this area are moving the needle.”

Premkumar C. Ingersoll Senior Manager-SRE, Verizon Business Group



“The last two years have seen a sea change in the way engineers and entire organizations perform their work. The pandemic and macroeconomic trends have highlighted the importance of further understanding the human contribution to both a functioning complex technological system and one that experiences an incident and requires the intervention of human expertise to repair. This presents an opportunity for teams to take a renewed look at the human aspect as they investigate incidents, and how to do so in a generative environment, one which requires blame awareness. Organizations who invest in both go beyond a laundry list of action items (which may or may never actually be completed), more deeply understand their incidents and increase their resilience and capacity to adapt to the demands of an ever-changing environment.”

J. Paul Reed Principal Consultant, Spective Coherence, Inc.




INSIGHT IV

Insight: Elite performing organizations emphasize customer experience reliability without ignoring the importance of employee experience reliability.

Recommendation: “[Total Experience](#) is a strategy that creates superior shared experiences by interlinking the user experience, customer experience, multi-experience, and employee experience disciplines source.”

We recommend conducting an inventory and audit of exclusive customer-facing reliability activities versus exclusive employee-facing reliability activities - and then determining whether there are economy of scale opportunities with a Total Experience approach. The goal is to ascertain whether a re-balance or alignment, within the business context, may have impact on business outcomes, such as retaining talent, improving morale, or increasing operational efficiencies.



What is Total Experience?

Delivering flawless customer experiences versus improving employee productivity have frequently been discussed as mutually exclusive sets of activities. Total Experience is a strategy aiming to change this by discussing these activities at the intersection of all involved parties.

19%


Percent of the reliability focus on employee-facing tools or systems

41.1%

Percent of the reliability focus on customer-facing products or services

Total Experience is about more than just reliability. However, in the context of this report, we narrow the criteria to focus largely on certain reliability aspects such as monitoring, observability and capacity management. Ultimately, however, what we are suggesting is to research an emerging trend claiming to achieve greater operational efficiencies - allowing organizations to grow exponentially versus linearly.

In your company, how balanced is the reliability focus of external (customer-facing) products or services, versus the reliability focus of internal (employee-facing) tools or systems?

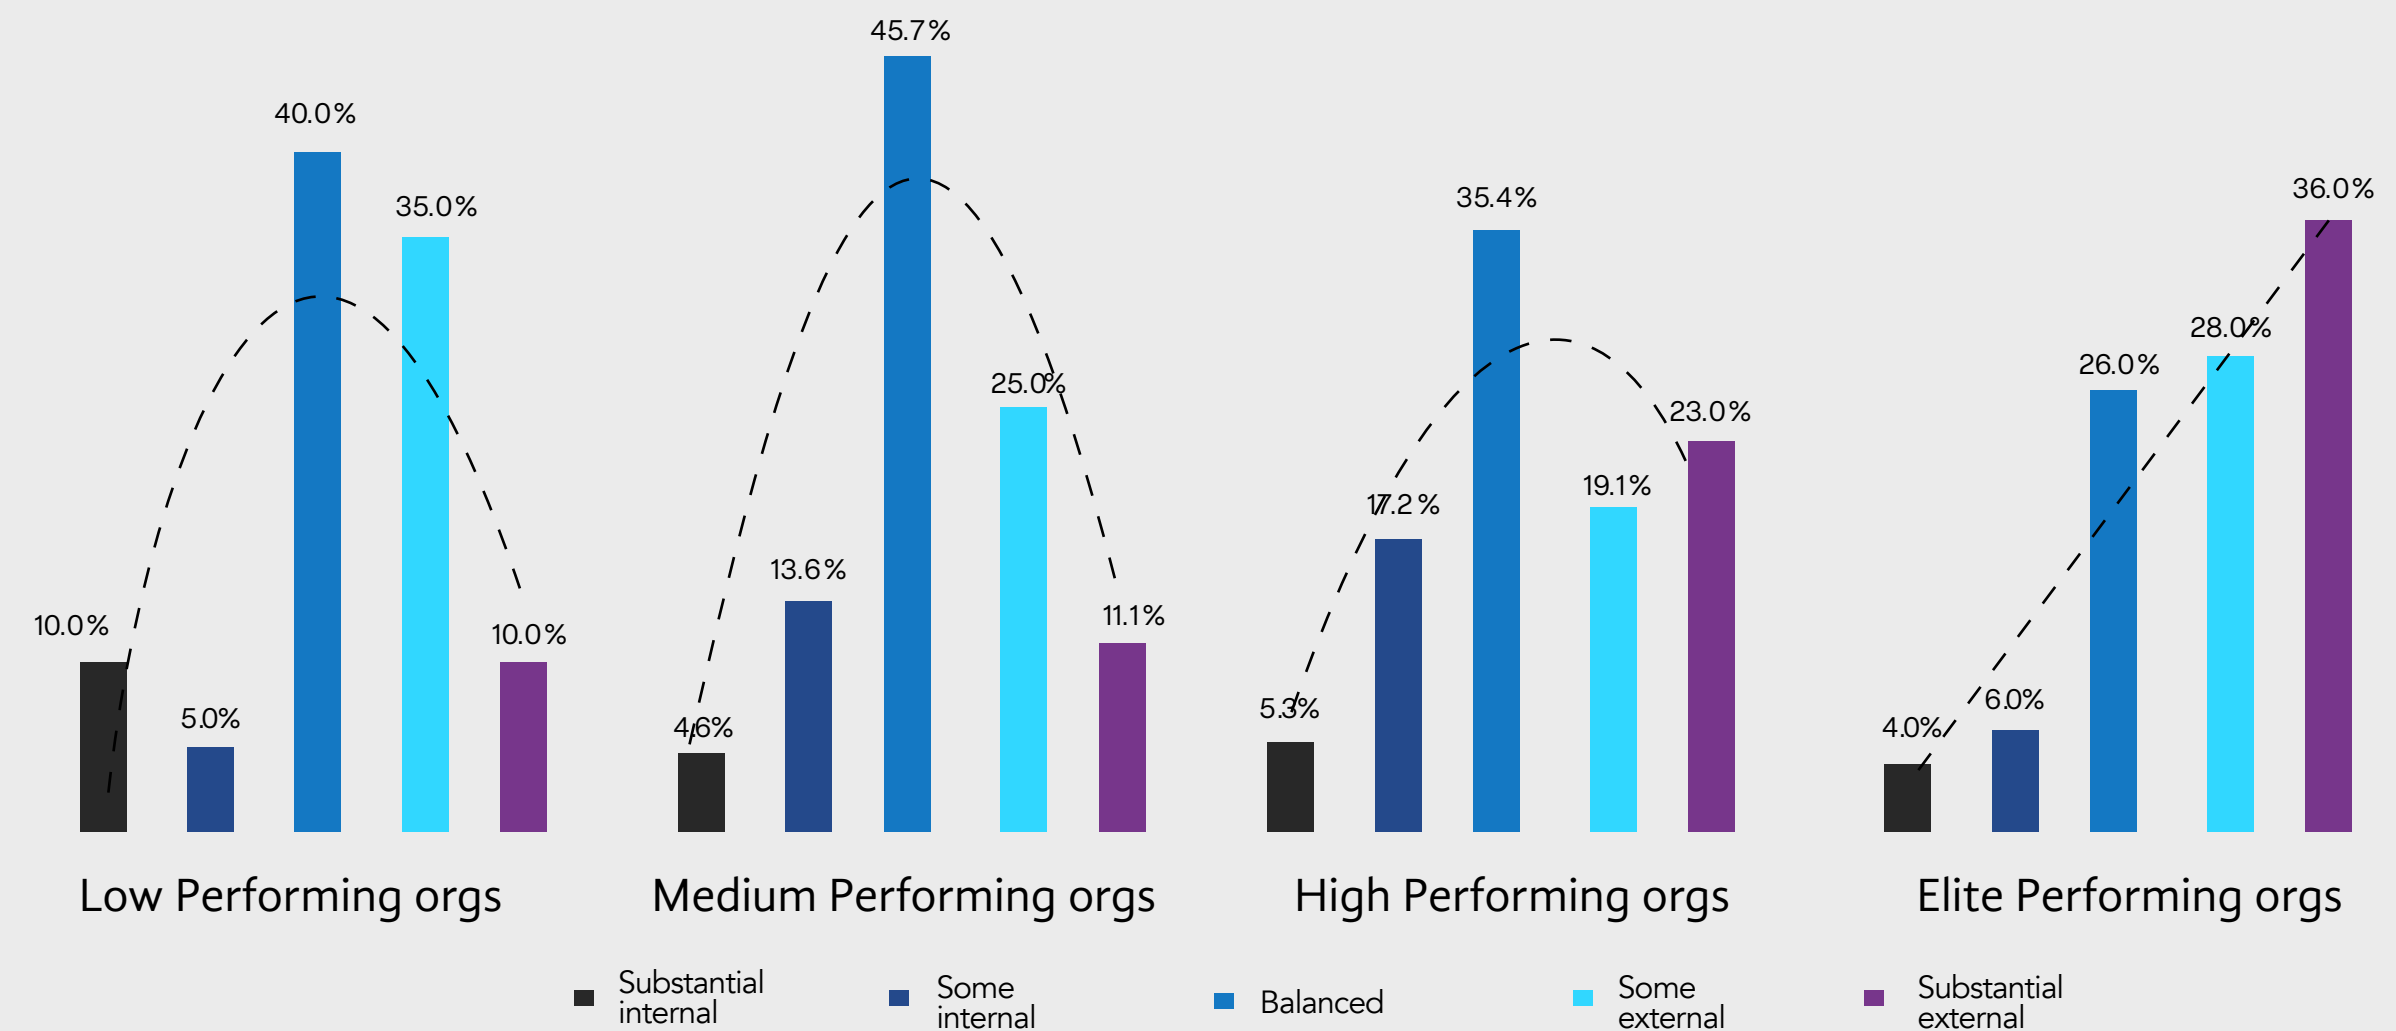


Elite organizations heavily focus on customer experience

By correlating responses in relation to DORA maturity metrics, we see that Elite performing organizations focus considerably more on external reliability (versus internal reliability), with 26% claiming to strike a balance between the two. While some industries or verticals (e.g., emergency services, medical, or government) may never be able to sincerely classify as Elite, care should be taken to nonetheless evaluate what this finding implies when evaluating how balanced is your external versus internal reliability focus.

Improving experiences at the intersection of all involved parties should outweigh the vanity desire to classify as a High or Elite performer. Within any organization's specific business context, attributes such as repeatability or sustainability should be continually reassessed as business and cultural landscapes change.

In your company, how balanced is the reliability focus of external (customer-facing) products or services, versus the reliability focus of internal (employee-facing) tools or systems?



36%

Percent of Elite performing organizations with a “substantial **external** skew”

10%

Percent of Low performing organizations with a “substantial **external** skew”

The hidden cost of innovation


A high deployment frequency is a KPI of Elite performers in the DORA metrics. The responses to this question showed that, more often than not, maintaining innovation velocity came at the expense of productivity or morale. There are good reasons, particularly for more innovative companies, to ask how sustainable this is. Employees working at companies at the cutting edge of innovation [report higher levels](#) of attrition and are more likely to quit than those working at a slower-moving organization. Ambitious targets and a heavy workload are often cited as reasons why people quit.

59%

[Derived] Percent of respondents who said maintaining innovation velocity “occasionally” or “often” negatively impacted productivity or morale

How is maintaining innovation velocity impacting employees at your company? Are you part of the 14.1% who said they were unsure? While measuring the benefit of innovation velocity may involve incorporating (possibly difficult to measure) business telemetry, we recommend companies consider how to mitigate the negative impacts of maintaining innovation velocity, such as offering employees more control over their schedules, evaluating project timelines, and considering the complete cost of achieving key targets.

How frequently does maintaining innovation velocity come at the cost of *negatively* impacting employee productivity or morale?



The impact of sustained work-from-home policies

44.7%

Percent of respondents who said “relationship building” was much or somewhat worse due to work circumstances

39.1%

Percent of respondents who said “talent hiring/retention” was much or somewhat worse due to work circumstances

29.5%

Percent of respondents who said “knowledge retention” was much or somewhat worse due to work circumstances

Talent hiring/retention, knowledge retention, and relationship building are the three categories ranking highest in the “worse” columns. Despite that, nearly 50% of respondents claimed that innovation velocity was “about the same” and productivity was 31.3% net “better” since working from home.

How have work circumstances (i.e., change in remote or hybrid work policies) since the start of the COVID pandemic affected the following?

	Much worse	Somewhat worse	About the same	Somewhat better	Much better	#N/A
Service reliability	1.1%	10.1%	51.2%	22.0%	9.4%	6.1%
Project completion	2.3%	15.9%	49.0%	19.6%	7.5%	5.7%
Innovation velocity	2.0%	18.9%	48.5%	15.7%	8.1%	6.7%
Knowledge retention	4.7%	24.8%	43.0%	16.6%	5.6%	5.4%
Morale	4.2%	21.0%	38.8%	22.1%	8.7%	5.1%
Productivity	1.8%	11.6%	36.7%	30.4%	14.3%	5.1%
Relationship building	11.6%	33.1%	32.0%	12.5%	5.4%	5.4%
Talenting hiring / retention	11.5%	27.6%	31.2%	14.9%	8.1%	6.6%

However, is maintaining this level of innovation and productivity sustainable amidst a backdrop where retaining knowledge, hiring talent and building relationships is harder to do? Movements like “quiet quitting” and “lying flat” suggest otherwise. Regardless, it’s clear that to successfully implement TX strategies, companies will have to overcome the challenge of building meaningful relations and ensuring a common focus in the hybrid work landscape

IT to business communication gaps exist

15.7%

Percent of respondents who collaborate with sales and marketing “often”

24.9%

Percent of respondents who collaborate with executives “often”

Speaking of relationship building, communication and collaboration are continually touted as core DevOps culture assets. However, our survey data shows that respondents communicate and collaborate markedly less with sales and marketing teams, followed closely by executives and customer support.

Unfortunately, this situation exacerbates an existing challenge in terms of management, in which executives, who tend to be the furthest away from the situation, end up making the closest decisions.

As already noted, there is an opportunity for SREs to find better ways of communicating with business leaders and generate greater alignment in order to have more control over the decisions that impact their roles daily. Apparently, another opportunity also exists for new lines of communication between sales, marketing and customer support. Given sales and marketing are effectively the intake of what’s happening in the industry

How frequently do you collaborate with the following groups?

	Never	Seldom	Occasionally	Often	N/A
Engineering	4.3%	8.6%	17.4%	64.2%	5.5%
Operations	2.7%	5.4%	22.9%	63.1%	5.9%
Infrastructure	5.4%	7.7%	22.4%	57.8%	6.8%
DevOps	5.9%	11.8%	26.1%	45.8%	10.4%
Network Engineering	5.0%	11.6%	30.1%	44.4%	8.9%
Product	7.2%	16.8%	30.2%	37.6%	8.2%
Customer Support	7.9%	18.8%	32.7%	34.9%	5.7%
QA/Test	7.5%	15.2%	32.4%	32.9%	12.0%
Executive Officer	16.3%	22.7%	31.5%	24.9%	4.7%
Sales or Marketing	21.8%	27.0%	25.9%	15.7%	9.5%

while customer support is the direct link to external customers, organizations must do more to ensure more effective collaboration to ultimately ensure that engineers are delivering what customers are asking for. Project ownership is an important part of this. If everyone in the business is given the opportunity to own their piece of the puzzle, a greater sense of responsibility, better communication - and ultimately business outcomes - will result.

The fallout of the Great Resignation

49.7%

Percent of respondents who said “talent hiring/retention” was moderately or seriously impacted by The Great Resignation

44.8%

Percent of respondents who said “knowledge retention” was moderately or seriously impacted by The Great Resignation

37.6%

Percent of respondents who said “productivity” was moderately or seriously impacted by The Great Resignation

37.5%

Percent of respondents who said “morale” was moderately or seriously impacted by The Great Resignation

Earlier in the report, we noted how talent hiring was the number one challenge hindering successful reliability implementations. This harmonizes with our findings on the impact of The Great Resignation. There is also some evidence of a sizable impact on knowledge retention, productivity, and morale.

How much negative impact has The Great Resignation had on the following?

	No impact	Minor impact	Moderate impact	Serious impact
Productivity	28.8%	33.6%	31.7%	5.9%
Knowledge retention	24.9%	30.3%	30.8%	14.0%
Project completion	32.2%	32.9%	28.2%	6.6%
Morale	26.7%	35.8%	27.4%	10.1%
Talent hiring/retention	22.7%	27.7%	26.5%	23.2%
Relationship building	30.5%	33.3%	26.0%	10.2%
Innovation velocity	33.0%	34.4%	25.9%	6.6%
Service reliability	36.8%	37.0%	21.9%	4.2%

Almost a quarter of respondents said talent hiring and retention were seriously impacted most by The Great Resignation, which should come as no surprise. What was eye-opening, however, were how different personas had different tendencies in their answers.

Middle managers feeling the pinch

Managers and Senior Managers bear the brunt of The Great Resignation’s impact on talent hiring or retention. Notice the contrast with executives who skew sharply toward “no” or “minor” impact.

55.9%

Percent of Managers who said The Great Resignation has moderately or seriously impacted “talent hiring/retention”

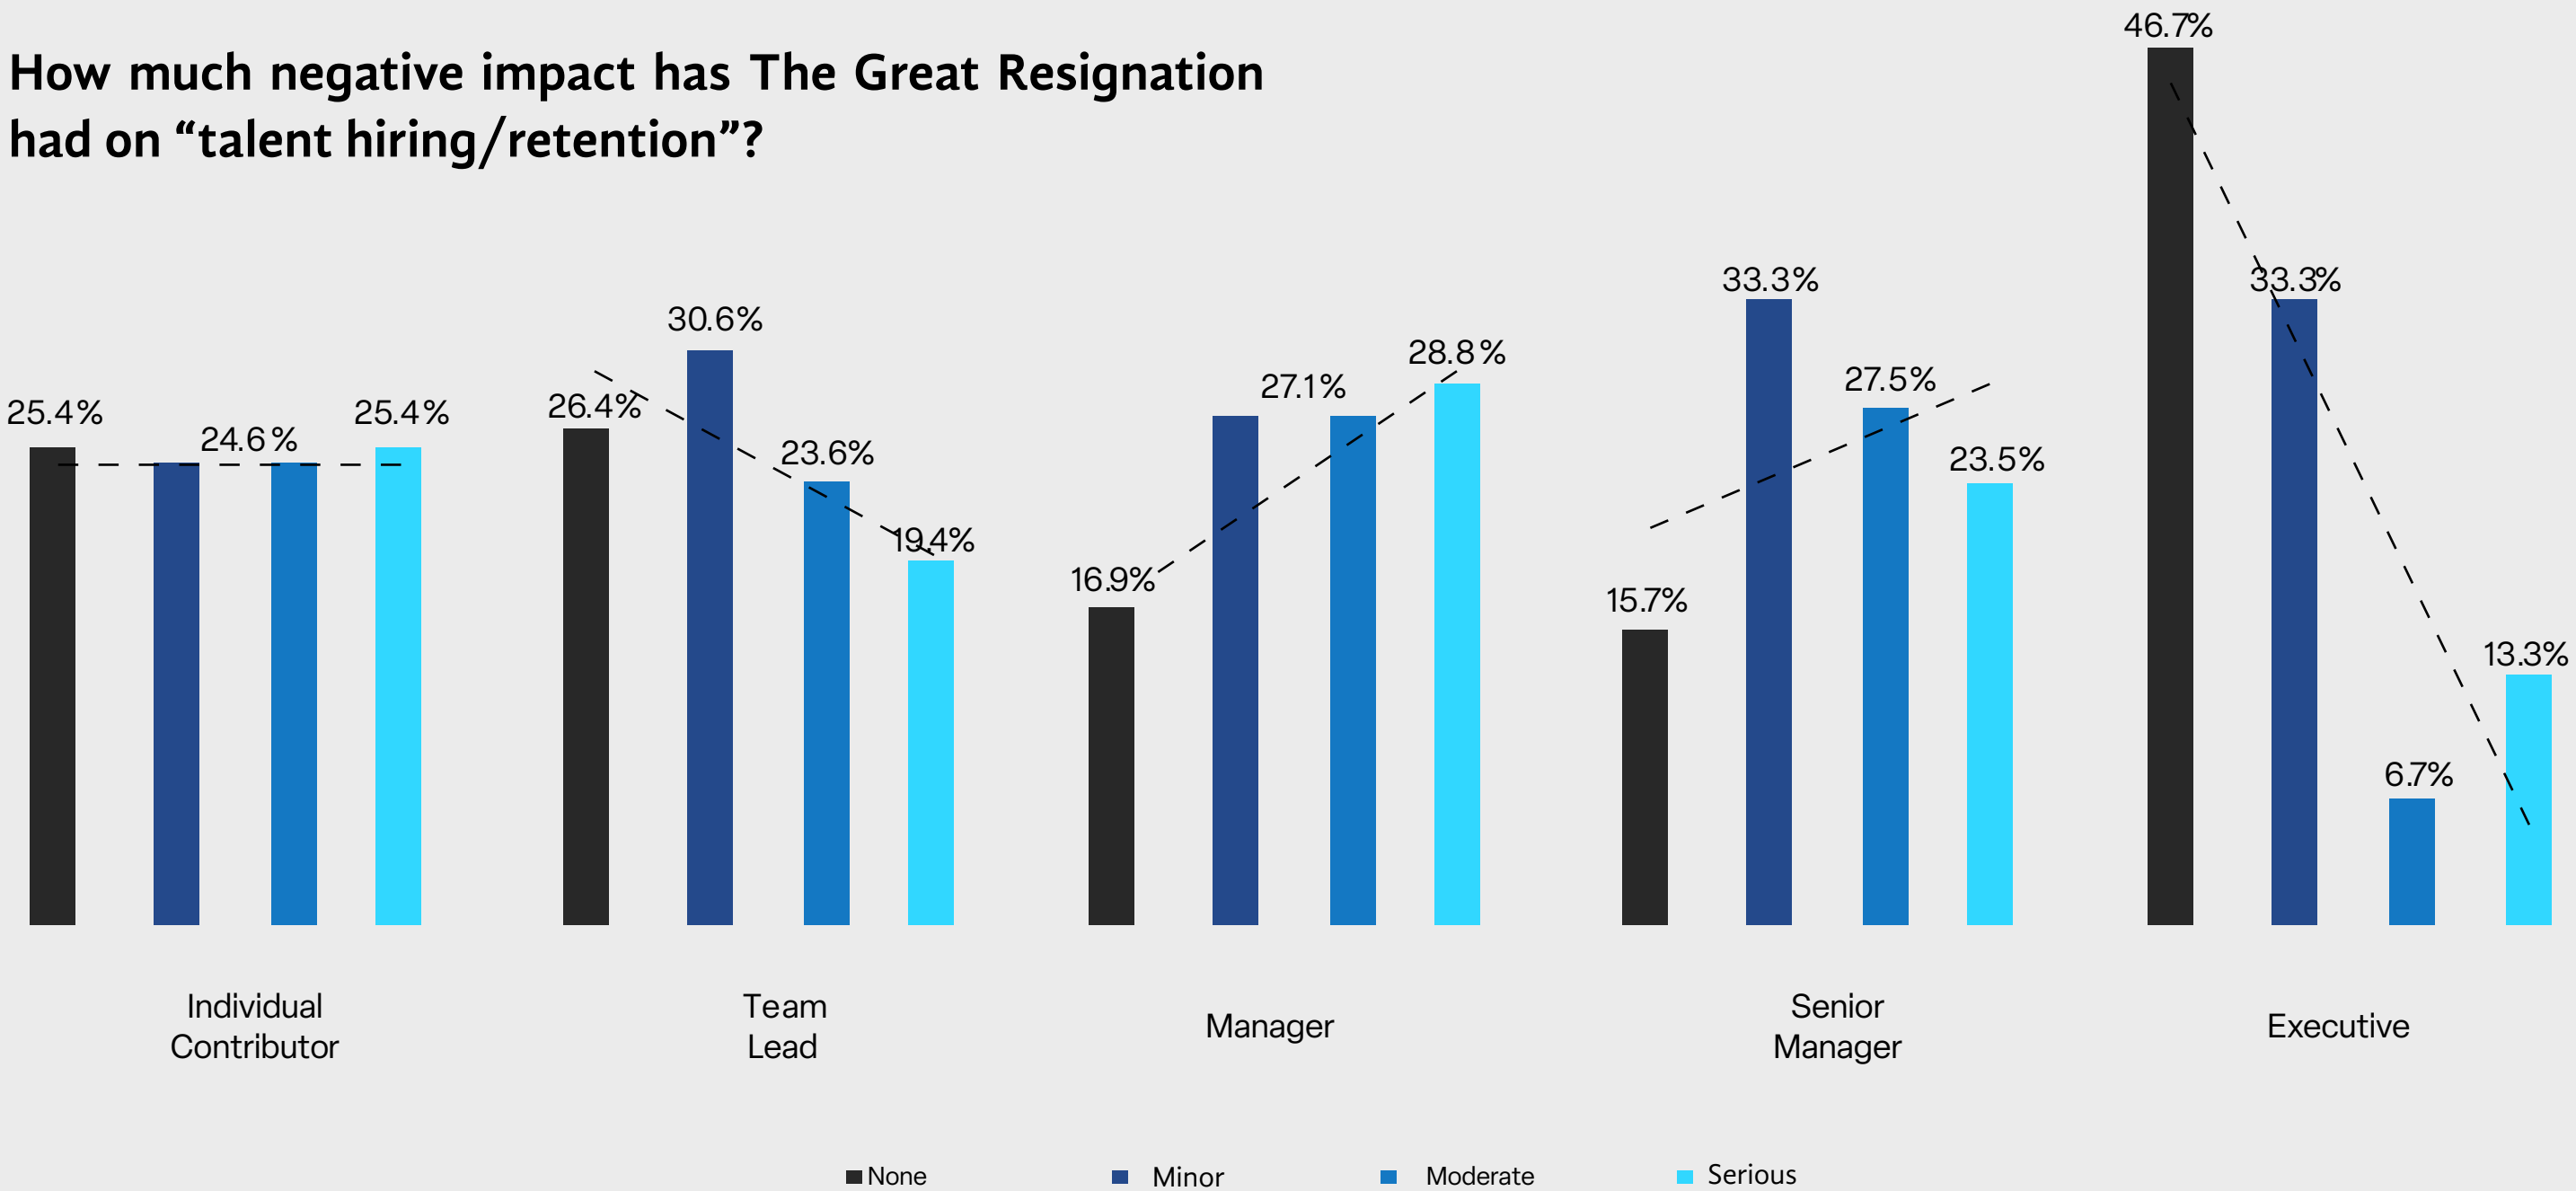
51%

Percent of Senior Managers who said The Great Resignation has moderately or seriously impacted “talent hiring/retention”

20%

Percent of Executives who said The Great Resignation has moderately or seriously impacted “talent hiring/retention”

How much negative impact has The Great Resignation had on “talent hiring/retention”?



While the data shows different tendencies for different personas with regard to the impact of The Great Resignation, recall this Insight is focused on whether Total Experience approaches will benefit all involved parties. This is another opportunity to share ownership, foster collaboration, and improve overall efficiency to achieve better outcomes.

Conclusion: Consider whether a TX approach will help you scale

In conducting the analysis around TX, we found some familiar trends and some thought-provoking anti-patterns. As “consumers” in our daily life, we’re all aware of how a single interaction with an employee can make or break our experience. Yet currently, the role employees play in generating exceptional customer experience is yet to be reflected in the level of focus businesses are giving to internal tools or systems, more so for Elite performing organizations. How sustainable this is given the factors we’ve noted is something we’ll look out for in next year’s report. Please consider that TX strategies may help to achieve better economies of scale and narrow the internal vs external gap. As is the case with all tech trends, it’s important to evaluate your particular situation and use caution to decide whether a TX approach is the right fit for your organization.



View from the field

“Total Experience approaches may or may not move the needle. Depending on your use of SaaS as part of your employee productivity suite, the effectiveness of the reliability focus is clearly limited by the extent to which you can actually exert control over those two very different production estates. However, the classic trade-off between “good for the team” versus “good for the company” won’t change unless we take some time to occasionally research (and implement) new approaches.”

Niall Murphy CEO, Stanza



“Whenever internal processes are streamlined - minimizing friction and stress amongst employees - it leads to improved customer experience and commercial success. This report shows that companies can look to improve internal ways of working to combat employee burnout. I often say, “happy devs write happy code”, which supports the Total Experience approach. If internal reliability isn’t taken into account, operational efficiency, and thereby productivity and morale suffer.

For example - in an ideal scenario your developers are making commits hundreds of times a day. However, if your CI/CD systems are less than continuous and take more than 10 min to compile a new build, you are losing thousands of minutes per day which translates into frustration (low morale), loss of flow (low productivity / high context switching), and wasted time that could otherwise be put to use working on external reliability. To improve your business’ scalability, efficiency, and productivity, you need to improve your ability for creative iteration.”

Tamara Miner VP of Engineering, disguise

INSIGHT V

Insight: Levels of toil drop marginally lower. *Time spent working exclusively on engineering activities and time spent on call remain the same.*

Recommendation: Now in its fifth year, The SRE Report is the longest running, non-Google, report of its kind. Therefore, we encourage you to use our benchmarking data to see where your company stands [per Google's recommendations](#) of what it means to implement reliability practices. Although there is no one-size-fits-all approach, Google's SRE guidelines are an accepted comparative gauge that allows SREs and businesses to know where SRE time is spent and make adjustments as needed.



Toil numbers continue to drop

5%

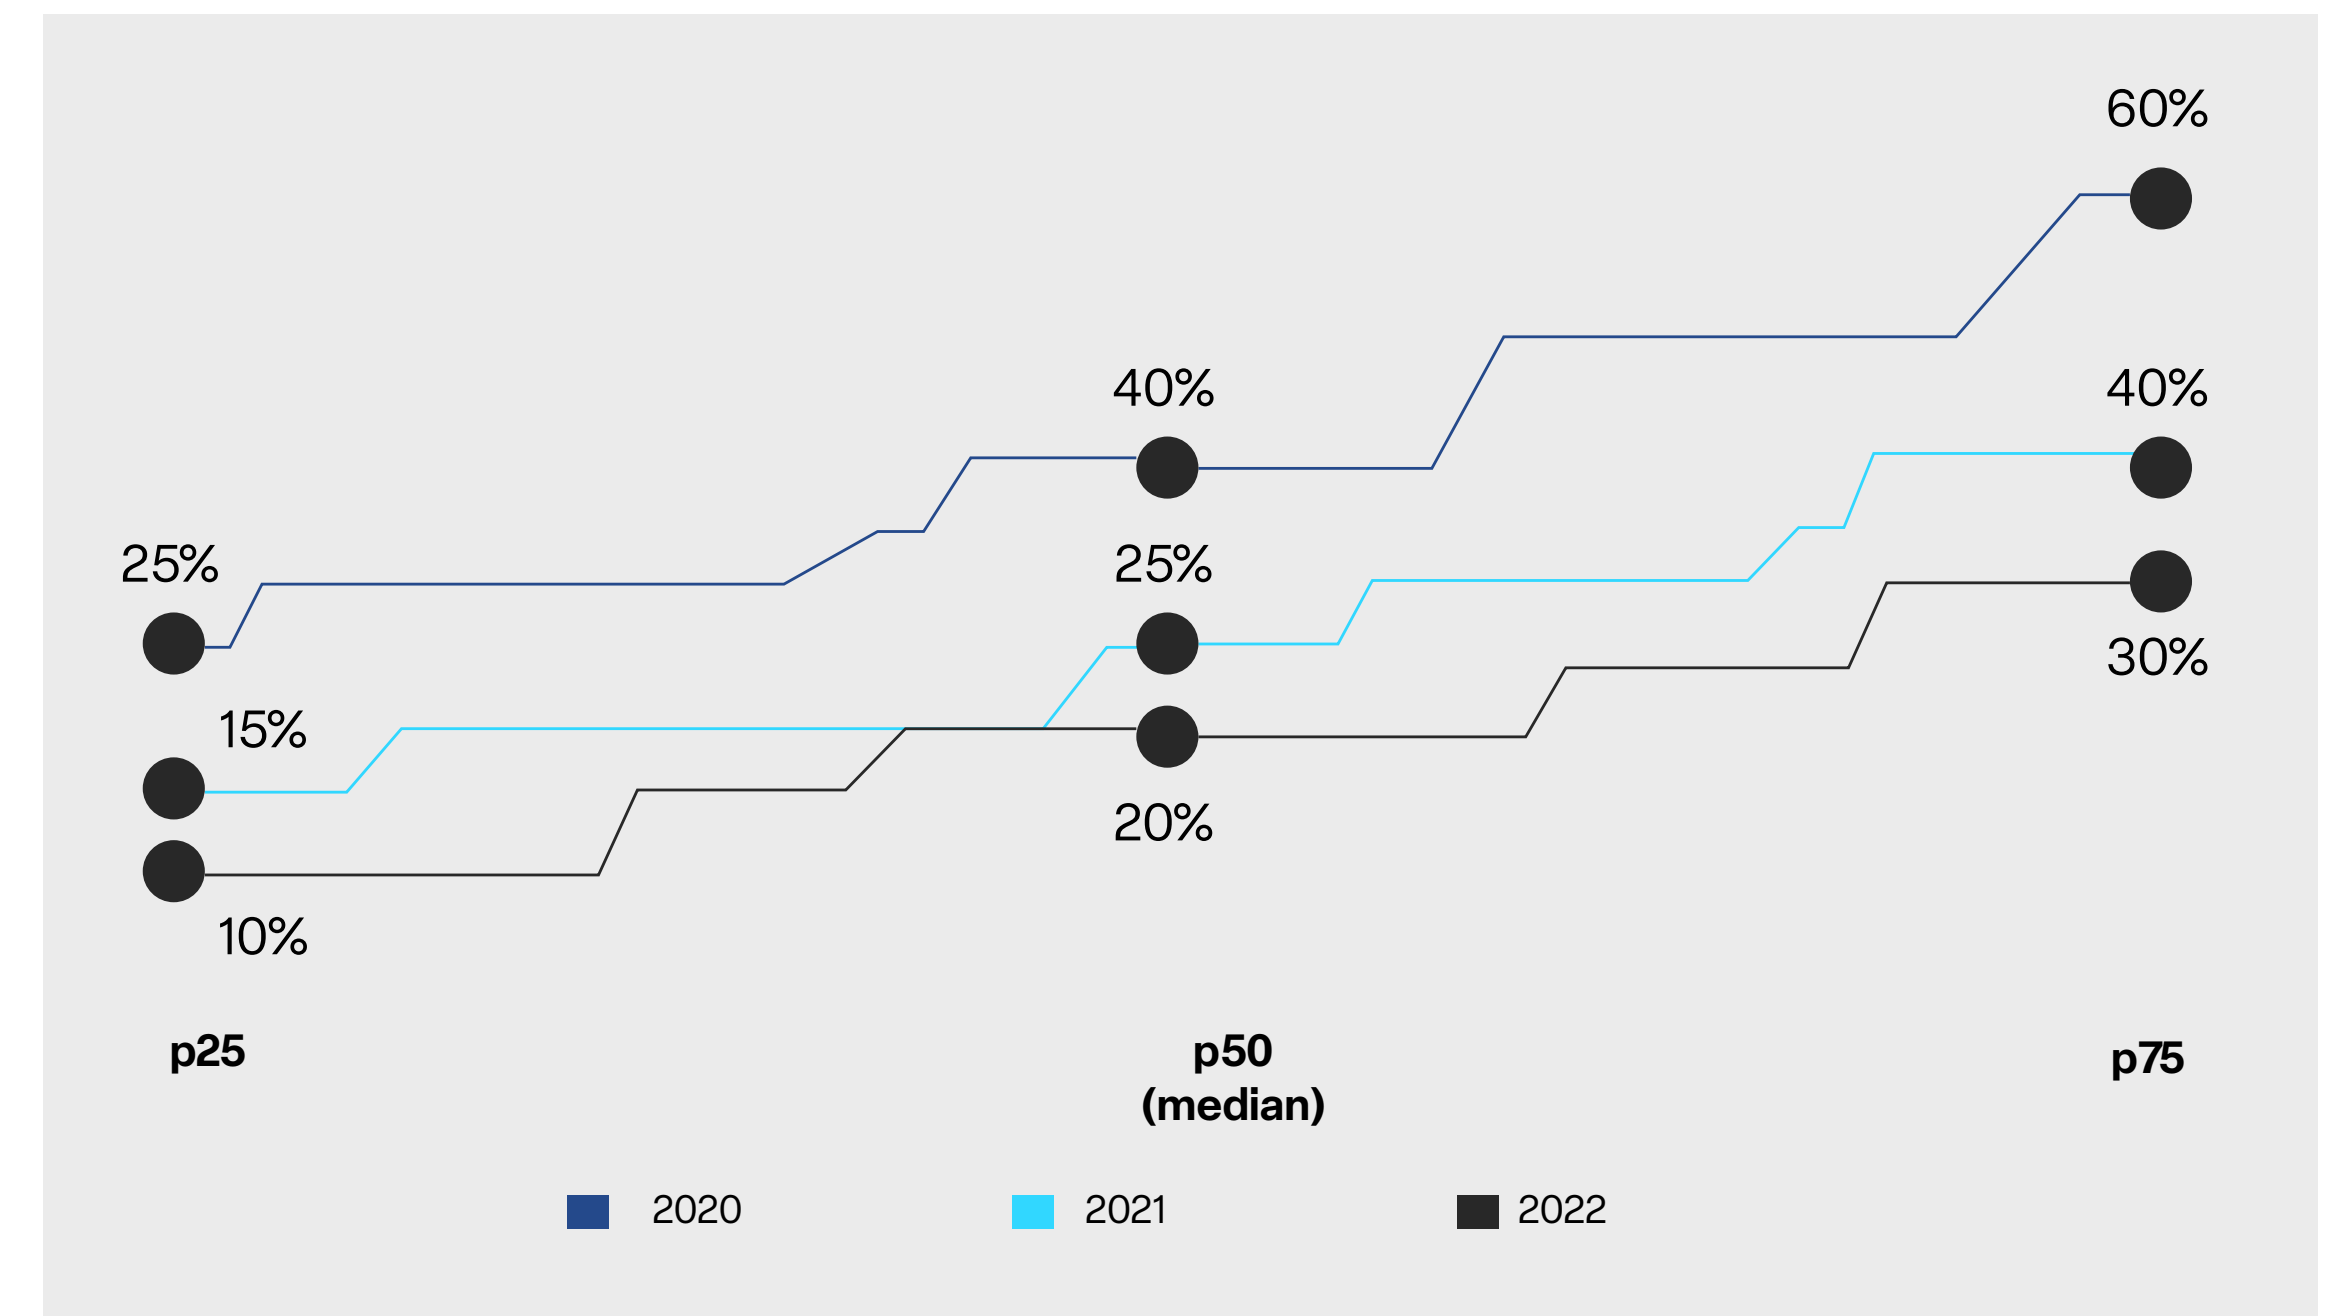
Median toil value has lowered by 5% compared with the previous report

Google defines toil as “the kind of work tied to running a production service tending to be manual, repetitive, automatable, tactical, devoid of enduring value, and that scales linearly as a service grows.” Hence, the need to “automate all the things” is a key SRE tenet.

	p25	p50 (median)	p75
2022	10%	20%	30%
2021	15%	25%	40%
2020	25%	40%	60%

This year, the median toil value dropped to 20% compared to 25% in 2021 and 40% the year before (pre-pandemic). Last year, we theorized this drop was temporary due to the widespread imposition of work-from-home policies and this would rise again as offices reopened. However, despite some of the loudest voices in the business world continuing to rally to get their workers back in offices, the large-scale predicted return to the office has not materialized. The hybrid workplace, it appears, is here to stay - for now.

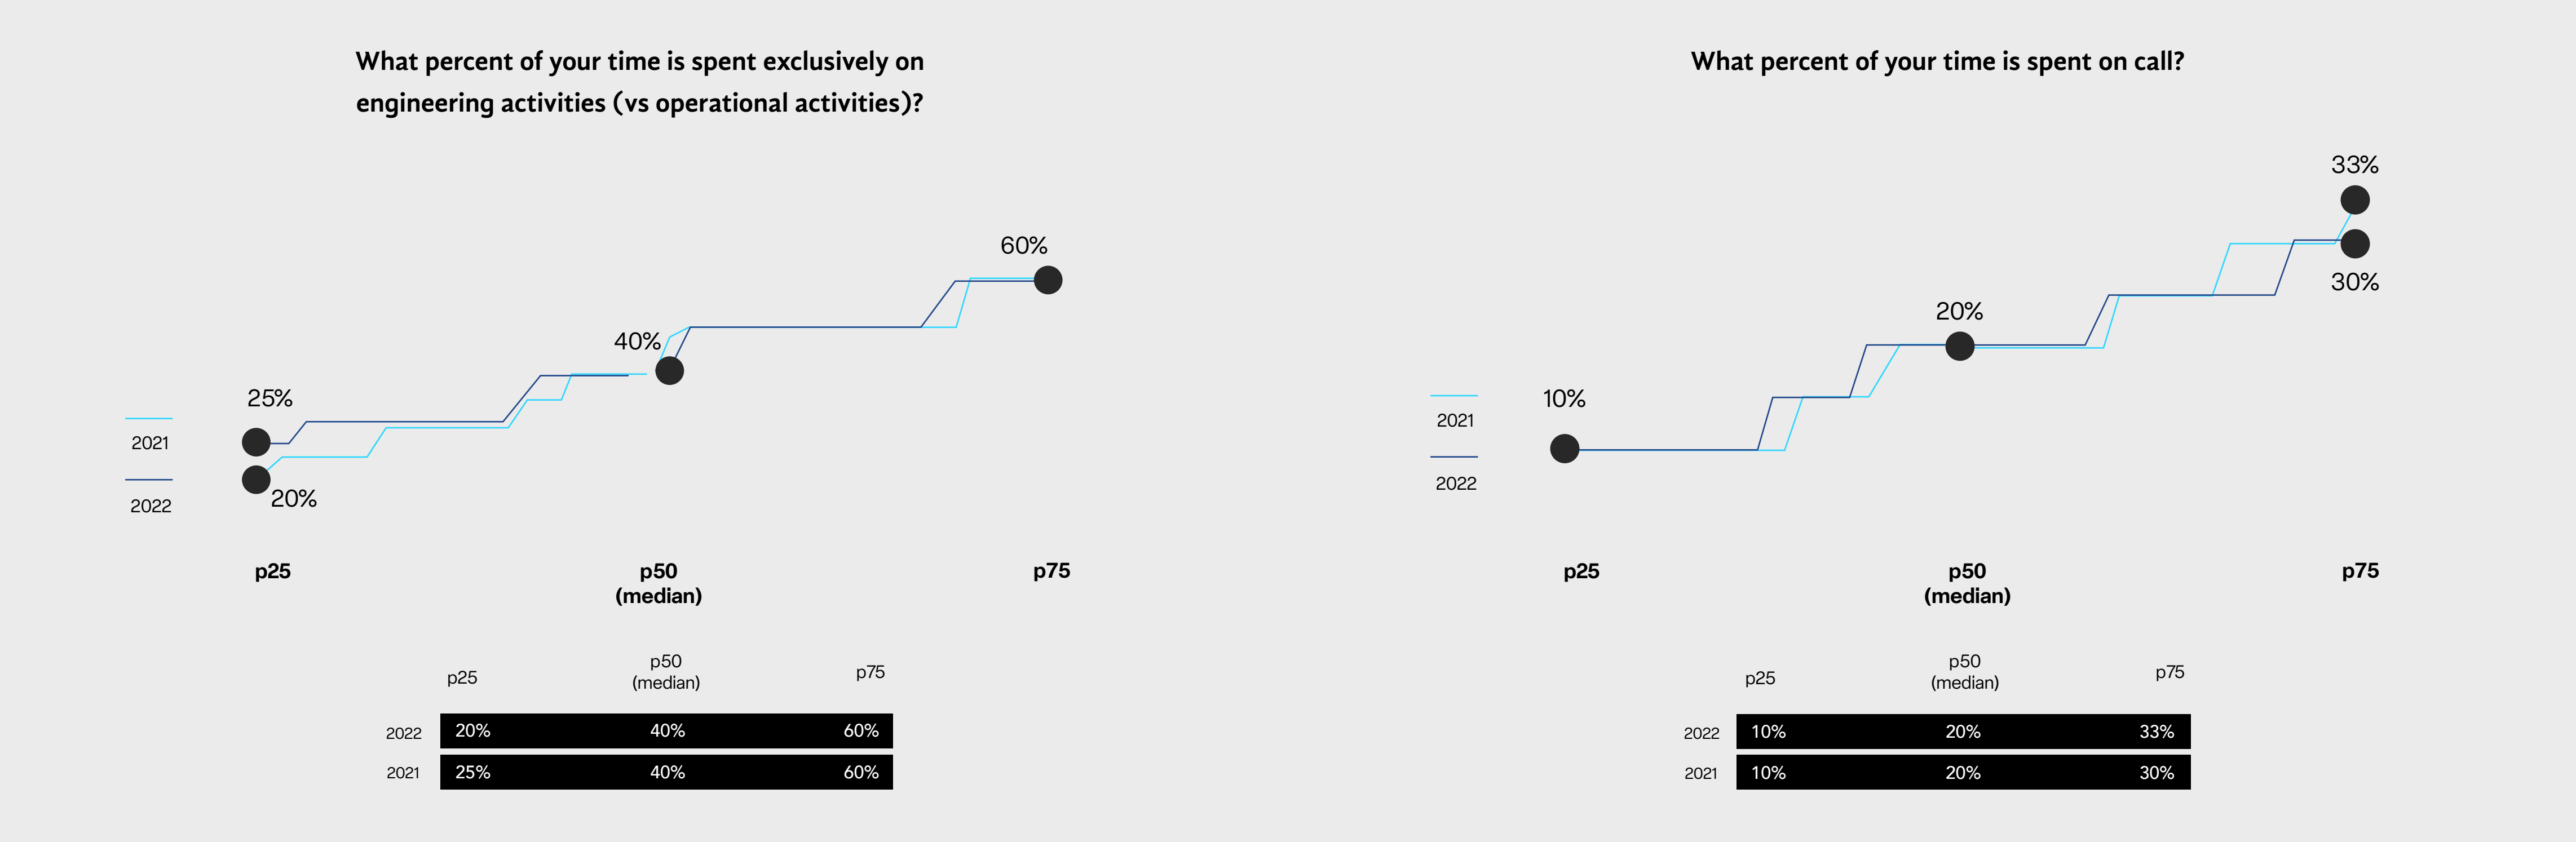
What percent of your work, on average, is toil?



DevOps split and time on call

Google places a 50% cap “on the aggregate ‘ops’ work for all SREs — tickets, on-call, manual tasks, etc.” In theory, this would enable an SRE team to spend the rest of their time working on development activities. Our data shows that organizations are spending more time on operations than development. At 40%, the median value of time spent exclusively on engineering activities has not changed year over year.

There has been no median change in time spent on call. The Google guideline is that SREs spend no more than 25% of their time on call. With a median value of 20%, most of our respondents are close to that sweet spot.




Did you get to read this entire report without interruption?

In a typical week (when you are not on call) what percentage of your time is consumed responding to interrupts?

20%

Half of this year's survey respondents said they spent up to 20% of their non-on call time responding to interrupts

This was a new question this year. Our respondents reported that when they're not on call, they spend up to 20% of their working week being interrupted. Add on the overhead for task switching and that number could easily double. Apparently, over the course of a year, employees can spend up to 5 working weeks, or 9% of their annual time at work [toggling between applications](#). It's important, however, to note that there's a distinction between useful - or necessary - interruptions versus wasteful ones. How does your company compare to this baseline?



Conclusions

Steve McGhee

Reliability Advocate, SRE, Google Cloud

Let's be honest, nobody loves surveys. Ok, well I sure don't. But surveys satisfy a huge need in our demand for insights into complex human-computer, *sociotechnical systems*. It turns out that we've been measuring the computer part *pretty well*, but the humans – not as easy to keep track of. When Google SRE first defined **toil** as a metric we wanted to reduce, we spent far too long trying to quantify it numerically based on tooling and insights from computer systems. It turned out to be easy: just ask the humans. Never stop, just keep asking them. We never found a better measure for toil, and I don't expect we will.

So, a survey is a powerful tool, but it takes work. It takes unbiased, structured questions, *lots* of respondents who actually take it seriously, and lots of analysis at the end. Without all of these critical elements, surveys are too often a waste of time that end up regurgitating existing biases and forgone conclusions. And those are easy to spot. I was excited to help with the creation and analysis of this particular survey. I felt that the questions were well-considered and the analysis was thoughtful. I wish there had been more respondents, but alas, there's always next year.

My takeaway from the conclusions this year is the theme of **SRE empowerment**. SREs in my experience thrive the most when they feel truly empowered: when their organizations trust them to do the right thing and they're given the resources and freedom they need. This means leaders must listen to their needs and support them, without inserting preconceived notions or interpretations. SRE is a very young field. There is a lot of interpretation at play here.



Conclusions

AIOps sounds amazing, heck even the name sounds cool. But listen to the practitioners who are actually trying it out, not the sales pitch. What is it actually doing today? Does it actually solve a problem you have right now? If not, move on. Don't be lured by the siren song of all-seeing, all-dancing AI. Don't make a decision for the SREs, **empower them** to choose (or not choose) tools based on their understanding of the current system and their own needs for operating it for the *immediate* future. Remember, the root of all evil is [premature optimization](#).

Tool sprawl sounds scary. Too many tools? Sounds expensive! I know when I go to a mechanic or a woodworking shop, I look for the place with the fewest tools on the walls and workbench. Wait, that's not right. When it comes to skilled labor, or "operations" perhaps, you want teams to be able to reach for the right tool at the right time, not to be impeded by earlier decisions about what they think they might need in the future. Also, what counts as a tool? If I combine two unix commands in a script, does that make a third? Why are we even stressing about this? Cost is the boogeyman here. Teams either have a culture of IT as a cost-center, which must be reduced over time, or they've been bitten by runaway costs of Cloud. APIs are powerful! Especially when you're not watching your billing statement. Instead of forcing SREs to rationalize every tool and prevent every possible overlap of functionality, **empower the SREs**. Give them transparency into cost, let them assess the value judgment as a group, inform them of contract details like renewal dates, let them propose alternatives.

Blamelessness is working. What better example of seeing the benefit of a psychologically safe environment? This is another form of empowerment. Knowing that they can be trusted with a complex system, despite fallible humanity (we all make mistakes!) **empowers an SRE** and results in a stable, sustainable system. This is a great datapoint to see reflected in the survey.

Why do ICs and Execs disagree so broadly? Why aren't they aligned? One interpretation is that Execs are looking at the bigger picture and ICs are focusing on a smaller portion, missing the context. However, that's not the only way to work. That's certainly the traditional (Taylorist) model that is employed at many Enterprises today, but **we can do better**. By providing transparency, context, and rationale around budgets, revenue and loss, teams can better understand tradeoffs made "above them" instead of simply throwing POs up to management to see what sticks. **SREs fight for the user**. Don't tie their hands, instead empower them to provide a big-picture solution. They can do it if you let them.

The last few years have been a heckuva ride. WFH is here to stay, remote work is only growing, even the allure of the 4-day work week is approaching, depending on who you talk to. Is this possible? Is this **great**? Is this scary? How about all of the above. I don't think this bell can be un-rung, nor should we want it to be. Being knowledge workers in the age of the Cloud means you don't have to be datacenter-adjacent, or even in-office. SRE is about creating higher levels of abstraction by which to control the systems that society ever more depends on. Let this happen, don't tie it down with old models of working, lest they come back to haunt you by way of attrition, burnout, and checked-out, uninspired employees. **Trust your SREs, empower them to defend the user (within clear expenditure limits), give them time and resources to be creative, push and reward sustainable behavior**.

Demographics

The SRE Survey was open for the month of June and closed on July 15th, 2022. The survey received 559 responses from all across the world, and from all types of reliability roles. Catchpoint and Blameless made donations in the amount of \$5,590 to the [International Red Cross](#) and [Girls Who Code](#).

How many employees does your company have?

One to 100	16.1%
101 - 1,000	33.1%
1,001 - 10,000	34.9%
10,001 - 100,000	13.6%
More than 100,000	2.3%

How many reliability engineers are in your company?

Zero to Ten	67.8%
11-100	23.3%
101-1,000	7.2%
More than 1,000	1.8%

Where are you personally located?

North America	78.0%
Europe	10.2%
Asia	7.5%
Australia/Oceania	1.8%
South America	1.6%
Africa	0.9%

Which most closely describes your role?

Individual practitioner/subject matter expert	28.8%
Team lead/supervisor	17.2%
Manager	13.8%
Senior manager (director/vice president),	11.6%
Architect,	10.9%
Project/Program manager	6.1%
C-Suite executive	3.8%
External consultant/contractor/coach	1.1%
Student	0.4%
Other	6.4%



Appendix I: What is the number one challenge hindering successful reliability implementations?

Talent (hiring, retention, assimilation)	7.9%	Complexity of architecture	7.5%	Business value is hard to realize	6.7%
Lack of End-to-end visibility	6.3%	Alignment or prioritization	4.2%	Time management	3.8%
Communication or Collaboration	3.8%	Knowledge, training, or education	3.3%	Lack/misuse of resources	2.9%
Cost or Budget	2.5%	Perpetual evolution or change	2.5%	Balance - velocity versus reliability	2.5%
Lack of buy-in	2.5%	Sprawl - tools	2.1%	Culture	2.1%
Silos	1.7%	Lack of standards	1.7%	Lack of automation (too much toil)	1.7%
Interruptions (Incidents, complaints)	1.3%	Lack of vision or direction	1.3%	Lack of testing	1.3%
Lack of data-driven decisions	1.3%	Lack of reward or recognition	1.3%	Lack of customer centricity	1.3%
Entrenched modalities	0.8%	Sprawl - systems	0.8%	Laziness, shortcuts, or apathy	0.8%
Money	0.8%	Compensation	0.8%	Morale or attitude	0.8%
Technical debt	0.8%	Human error	0.8%	Lack of trust	0.8%
Unnecessary complexity	0.8%	Silent Success (until shit hits the fan)	0.8%	Not treating reliability as a feature	0.8%
Adapting to cloud technology	0.4%	Lack of growth/progression	0.4%	Price increases due to COVID-19	0.4%
Lack of hardware integrations	0.4%	Lack of clarity (too much ambiguity)	0.4%	Too much red tape	0.4%
Alert fatigue	0.4%	Weediness (can't see forest for trees)	0.4%	Poor architecture	0.4%

Appendix I: What is the number one challenge hindering successful reliability implementations?


#allthethings	0.4%	Rigidity	0.4%	Lack of ownership / accountability	0.4%
Sprawl - purchasing	0.4%	Lack of project management	0.4%	Lack of comments / readability	0.4%
Lack of resiliency	0.4%	Lack of documentation	0.4%	Lack of security	0.4%
Code instrumented incorrectly	0.4%	Poor or missing procedures	0.4%	Administrative overhead	0.4%
Prioritizing profit above everything else	0.4%	Business telemetry is hard to push	0.4%	ROI	0.4%
Hypergrowth	0.4%	Delivering reliable services	0.4%	Ignoring social aspects	0.4%
Lack of chaos engineering	0.4%	Inability to measure reliability	0.4%	Complexity of tools	0.4%
Difficulty measuring user impact	0.4%	Lead time (short runways)	0.4%	Time to ROI	0.4%
Lip service	0.4%	Management roadblocks	0.4%	Lack of integrity	0.4%
Word salad	0.4%	Lack of investment	0.4%	Lack of managed services	0.4%
Engineering teams operating beyond saturation point making goal sacrifices	0.4%				
Lack of attention to mental health and proper work-life balance	0.4%				
Lack of social contracts across the entire engineering org	0.4%				
Remote work impact on knowledge exchange	0.4%				

Appendix II: How do reliability engineering practices add business value?


Lower cost	12.5%	Focus	0.9%	Reduce incident severity	0.5%
Customer experience, sentiment, or satisfaction	12.5%	Balance - Innovation versus Reliability	0.9%	Work-life balance	0.5%
Maintain reliability, performance, or uptime	11.1%	They keep us out of trouble	0.5%	Reducing opportunity cost	0.5%
Retain existing customers	6.5%	Happy users happy teams	0.5%	Correlate IT metrics with business KPIs	0.5%
Avoid SLA penalties	6.0%	Reduce human error	0.5%	Improved consistency / repeatability	0.5%
Increase operational efficiency	5.6%	Improved morale	0.5%	Is this for real?	0.5%
Increase new logo count or revenue	4.6%	I am not sure to be honest	0.5%	Secure data	0.5%
Talent attraction/retention	3.7%	Improved project management	0.5%	Less burnout	0.5%
Self-evident business value	2.8%	Promote services as utilities	0.5%	Stay current	0.5%
Improve time to repair	2.8%	Creates a successful product	0.5%	Unclear, we're not exposed to the business side	0.5%
Trust or integrity	2.3%	Reduce time spent on maintenance	0.5%	improved documentation	0.5%
Increase or preserve brand integrity or reputation	1.9%	Increase conversions	0.5%	Improved communication/collaboration	0.5%
Increase innovation velocity	1.4%	Betterment of entire socio-technical system	0.5%	This is completely unknown	0.5%
increase business competitiveness	1.4%	Balance - everything between cost and quality	0.5%	Modernize company	0.5%
Remove technical debt	1.4%	Synergy	0.5%	Improved knowledge	0.5%
Increase productivity	1.4%	Dedicate resources	0.5%	No surprises	0.5%
Improve quality	0.9%	Improve morale	0.5%	Better code	0.5%
Improve engineering experience or confidence	0.9%	Allows you to follow your dreams and goals	0.5%	Increased customer centricity	0.5%
Resiliency	0.9%	Reduce bugs	0.5%		
Culture	0.9%	Improved accountability	0.5%		

Appendix III: DORA IV

For the primary application or service you work on, how long does it take to go from committed code to successfully-running-in-production code?




For the primary application or service you work on, how often does your organization deploy code to production or release it to end users?




Appendix III: DORA IV

For the primary application or service you work on, how long does it take to restore service when a user-impacting incident or defect occurs (e.g., unplanned outage, service impairment)?



For the primary application or service you work on, what percentage of changes to production or releases to users result in degraded service (eg., lead to service impairment or service outage) and subsequently require remediation (eg., hotfix, rollback)?



Acknowledgements

Catchpoint and Blameless wish to thank the SRE community for providing the survey data which was used to create The SRE Report. We also want to thank the entire report production team for their hard work, dedication, and grit. Without the community and production team, this report would not have been possible.

Survey Authors and Report Analysts

[Kurt Andersen](#) and [Leo Vasiliou](#)

Report Writers

[Denton Chikura](#), [Kurt Andersen](#), and [Leo Vasiliou](#)

Report Editors

[Anna Jones](#) and [Leo Vasiliou](#)

Report Graphic Designer

[Michael Dittner](#) ([dittco](#))

Report Production Team

[Anna Jones](#), [Deirdre Mahon](#), [Denton Chikura](#), [Gleb Kapusto](#), [Jinyu Qian](#), [Kurt Andersen](#), [Leo Vasiliou](#), [Mark Towler](#), [Sergey Katsev](#), [Tony Ferrelli](#), the [Blameless](#) Marketing Team, and the [Catchpoint](#) Marketing Team.

Report Contributors

[Adrian Cockcroft](#), [Adriano Velasco Nunes](#), [J. Paul Reed](#), [Keri Melich](#), [Laura Nolan](#), [Niall Murphy](#), [Premkumar C. Ingersoll](#), [Simone Cibba](#), [Steve McGhee](#), and [Tamara Miner](#)

Executive Sponsors

[Drit Suljoti](#), [Gerardo Dada](#), [Howard Beader](#), and [Mehdi Daoudi](#)



About Catchpoint

Catchpoint is the Internet Resilience Company™. The leading global brands rely on Catchpoint and its deep observability across thousands of global vantage points to increase their internet resilience by catching any issues before they impact their customers, workforce, networks, website performance, applications, and APIs. The Catchpoint platform combines synthetics, RUM, performance optimization, high fidelity data and flexible visualizations with advanced analytics to provide unparalleled observability into the internet stack.

Learn more: www.catchpoint.com



About Blameless

Blameless drives resiliency across the software lifecycle by operationalizing Reliability Engineering practices. Teams share a unified context during incidents, efficiently communicate, and resolve quickly. With detailed retrospectives, teams up-level and continuously improve end-user experiences. Leading brands such as Procore, Under Armour, Citrix, Mercari, Fox, and Home Depot, rely on Blameless to embrace a culture of continuous improvement, do more with less, and protect their customers. Blameless is backed by Lightspeed Venture Partners, Accel, Third Point Ventures, and Decibel.

Learn more: www.blameless.com