

DevOps Enterprise Guidebook

White paper
May 2022

Authors: Eric Maxwell,
Sandeep Parikh,
Rob Edwards





Table of contents

02	What is this guide and why should you care?
03	What is DORA?
04	Why DORA?
09	The DORA DevOps Quick Check
12	The DORA Capabilities Assessment
13	How to use DORA to improve
16	Conclusion



What is this guide and why should you care?

Elite performers with Google Cloud are able to ship 6,570 times faster than their competitors.

The DevOps Research and Assessment (DORA) team is on a multi-year journey investigating the capabilities required for teams to improve software delivery and operations performance. This guide goes a step beyond identifying the capabilities and provides guidance on how to apply the findings. Google Cloud tools and technologies are used throughout the guide to solidify the recommendations.

We think Google Cloud has the best features to allow your organization to achieve elite performance on your way to shipping 6570 times faster than your competition. We believe Google Cloud excels at enabling companies to take full advantage of the five essential characteristics of cloud as defined by the National Institute of Standards and Technology ([NIST](#)). In this guide we take a very intentionally prescriptive approach that you can follow to uplevel your software delivery abilities.

This guide is not *just* for DevOps professionals. It is designed for anyone who is part of a team or leads a team that ships software, or anyone who is interested in better understanding modern practices for shipping software.


One of the beautiful things about the DORA research and the assessment is that it is tool- and platform-agnostic. DORA does not care if you're a Silicon Valley digital native startup shipping the next unicorn-driving code base, a specialty coffee multinational working on a remote plantation, or a government deploying to your own sovereign cloud—it simply meets you where you are and provides actionable insights and a way to measure improvements.

This guide, however, is a bit more opinionated.

What is DORA?

DORA, or the DevOps Research and Assessment, is the largest and longest-running scientific research project of its kind. It applies the scientific method to determine what enables companies to ship software quickly and reliably. The project has been running for seven years and studied tens of thousands of individuals around the world across all industries. The project has identified and validated a number of technical, process, measurement, and cultural capabilities that drive higher software delivery and organizational performance.

Recent publications from the DORA team include the updated [ROI of DevOps Transformation report](#), the [DevOps Quick Check](#), and the [2021 Accelerate State of DevOps Report \(SODR\)](#). If you're not familiar with the SODR, you'll get the most out of this guide if you take a few moments to read it first.



Why DORA?

DORA research is rigorous.

Research design

DORA research employs a cross-sectional, theory-based design. This theory-based design is known as **inferential predictive**, and is one of the most common types conducted in business and technology research today. Inferential design is used when purely experimental design is not possible and field experiments are preferred.

Target population and sampling

In 2021, the annual DORA State of DevOps Survey surveyed 1,200 individuals. Our target population was practitioners and leaders working in, or closely with, technology and transformations, especially those familiar with DevOps. We promoted the survey via email lists, online promotions, an online panel, and social media, and we asked people to share the survey with their networks (that is, snowball sampling).


Creating latent constructs

We formulated our hypotheses and constructs using previously validated constructs wherever possible. We developed new constructs based on theory, definitions, and expert input. We took additional steps to clarify intent to ensure that data collected from the survey had a high likelihood of being reliable and valid.¹



1. Churchill Jr, G. A. "A paradigm for developing better measures of marketing constructs," *Journal of Marketing Research* 16:1, (1979), 64–73

DORA uses rigorous research to identify the capabilities and practices that drive transformation and accelerate modernization.



Statistical analysis methods

Cluster analysis

We used cluster analysis to identify our software delivery performance profiles based on deployment frequency, lead time, time to restore service, and change failure rate. We used a latent class analysis² because we had no industry or theoretical reasons to have a predetermined number of clusters, and we used Bayesian information criterion (BIC) to determine the optimal number of clusters.

Measurement model

Prior to conducting analysis, we identified constructs using exploratory factor analysis with principal component analysis using varimax rotation³. We confirmed statistical tests for convergent and divergent validity and reliability using average variance extracted (AVE), correlation, Cronbach's alpha⁴, and composite reliability.

2. Hagenaars, J. A., & McCutcheon, A. L. (Eds.). (2002). *Applied latent class analysis*. Cambridge University Press.

3. Straub, D., Boudreau, M. C., & Gefen, D. (2004). Validation guidelines for IS positivist research. *Communications of the Association for Information systems*, 13(1), 24.

4. Nunnally, J.C. *Psychometric Theory*. New York: McGraw-Hill, 1978

5. Hair Jr, J. F., Hult, G. T. M., Ringle, C. M., & Sarstedt, M. (2021). "A primer on partial least squares structural equation modeling (PLS-SEM)." Sage publications.

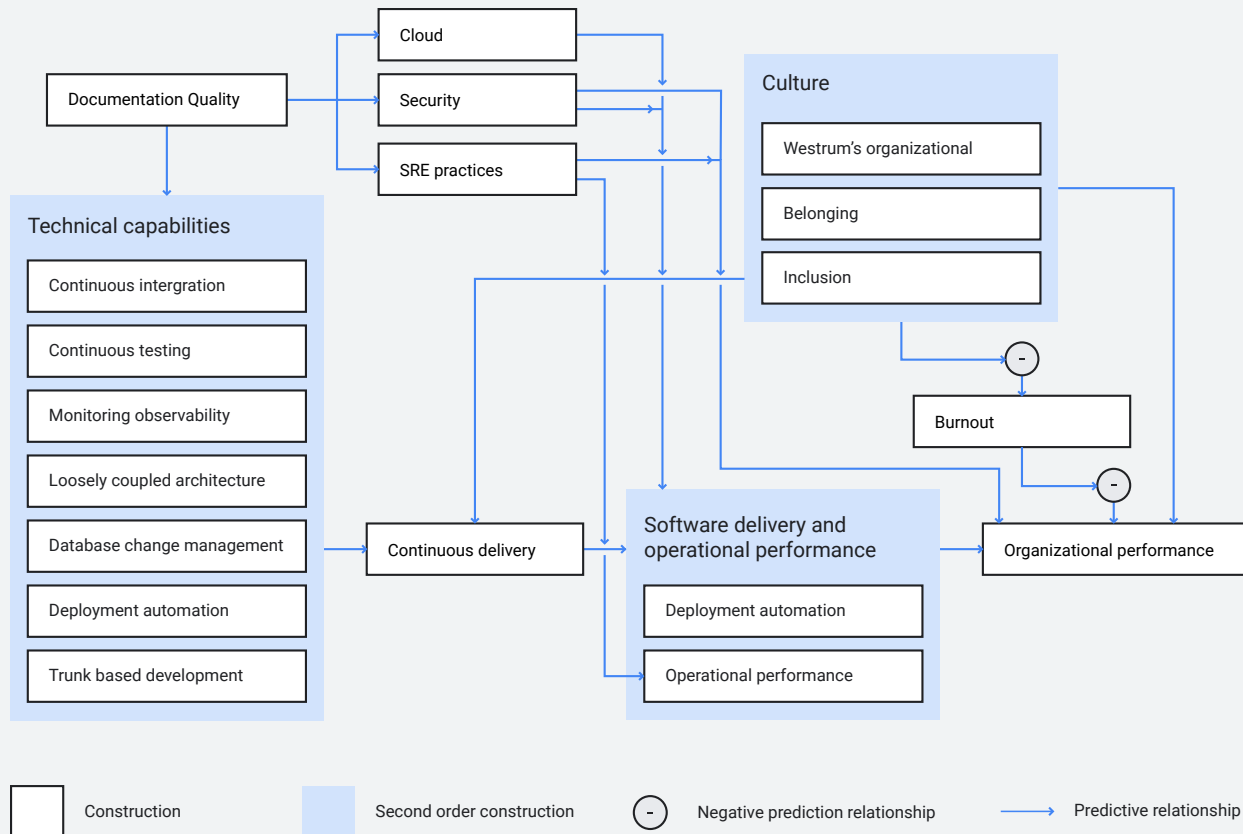
Structural equation modeling

From 2015 onwards, the DORA research program has used structural equation modeling (SEM), a predictive statistical modeling technique used to test relationships. We utilize PLS-PM (partial least squares path modeling) for our analysis, for several reasons: it does not require assumptions of normality in the data, it is well suited to exploratory and incremental research, and the analysis optimizes for prediction of the dependent variable (vs. testing for model fit of the data). All paths shown in the SEM figures are $p < .05$.

On page 6 you can find the SEM used in the 2021 analysis. Each box represents a construct we measured in our research, and each arrow represents relationships between the constructs. A larger box that contains boxes (constructs) is a second-order construct.

To interpret the model, all arrows can be read using the words predicts, affects, drives, or impacts.





Tool- and platform-agnostic

DORA's research does not mention or measure usage of any specific technologies, platforms, or products, but instead focuses on the capabilities that drive improvements. That is, rather than measuring the impact of adopting a specific continuous integration (CI) tool (e.g., Jenkins, Circle CI, or Google Cloud Build), DORA's research measures whether teams that practice continuous integration realize improvements in their software delivery performance. For example, when your team commits code, does that kick off a series of automated tests? Including automated tests in your CI pipeline is predictive of improved software delivery performance, but simply having an instance of Jenkins running doesn't ensure that your pipeline implements CI capabilities.

Why does this matter?

Arguably, one of the most compelling findings to come from previous State of DevOps Reports (SODRs), detailed in Accelerate (Forsgren, et. al., 2018), is the direct link between software delivery performance and organizational performance. Said another way, companies that can ship software at an elite level—on demand, multiple times a day—are over twice as likely to meet or exceed their business goals. So software delivery has a direct impact on profitability, productivity, market share, the number of customers, operational efficiency, customer satisfaction, and much more.

In 2021 we have even more exciting data to share in this area. We can now say that there is also a connection between having a secure supply chain and organizational performance, and that there is a direct connection between Site Reliability Engineering (SRE) and organizational performance. Companies that “shift left” on security—meaning they focus on security sooner in the development process—are [1.8 times more likely to meet or exceed their business goals, and companies that implement SRE practices are 1.6 times](#) more likely to do the same.

Companies are

1.8x

more likely to meet or exceed their business goals when they ‘shift left’ on security

Companies are

1.6x

more likely to do the same when they implement SRE practices



In our 2018 report we stated that “high-performing teams are twice as likely to be developing and delivering software in a single, cross-functional team.” A [generative](#) organizational culture that optimizes for information flow, trust, innovation, and risk-sharing is predictive of high performance.

We expanded our research in 2021 and found that high-performing organizations are more likely to have a culture that encourages employees to take calculated and moderate risks without fear of negative consequences. We also looked at the impact of COVID-19 with regard to culture and burnout. We found that teams with a generative team culture, composed of people who felt included and felt like they belonged on their team, were half as likely to experience burnout during the pandemic.



Acquiring talent is expensive and often a difficult task, especially in today's disconnected and competitive market. Additionally, talent [attrition](#) can be disruptive to product development and to team morale. Maintaining a culture that thrives in information sharing should be a key strategy for all organizations.

To calculate the return on investment (ROI) of software delivery we looked at two categories, cost and value. You can reduce costs by cutting the time it takes to resolve outages and avoiding downtime as much as possible. On the value side, we looked at enhanced efficiency through the reduction of unnecessary rework, a reduction in toil, and the potential revenue gained by reinvesting the time saved in new offer capabilities. If your engineers are spending 60% of their time trying to release new code instead of innovating, you're losing 60% right off the top. Imagine if instead of spending 40% of their time innovating they could spend 90% of their time building new features and adding value to the business! We have published a [white paper](#) to help you quantify the ROI of DevOps as it pertains to your specific organization. We also have an internal value calculation tool we would be happy to run with you. For information, reach out to your Google Cloud Representative.

The DORA DevOps Quick Check

The [DORA DevOps Quick Check](#) equips you to measure your software delivery performance against~your peers—and helps you prioritize possible improvements. The Quick Check is a five-question multiple-choice survey that focuses on DORA metrics along with your organization's principal industry. Based on the DORA research, four key metrics provide the strongest indicators of software delivery performance: lead time for changes, deployment frequency, time to restore, and change fail percentage.

[Take the 5 question Quick Check](#)



The first two Quick Check questions cover the “speed” metrics—lead time for changes and deployment frequency:

- For the primary application or service you work on, what is your lead time for changes (that is, how long does it take to go from code committed to code successfully running in production)?
- For the primary application or service you work on, how often does your organization deploy code to production or release it to end users?



The next two Quick Check questions cover the “stability” metrics—time to restore and change fail percentage. When evaluating time to restore, it's important to consider incidents or defects (not planned outages or maintenance) that impact end users, and to determine how long it takes to bring services back to standard functionality. With change fail percentage, the survey measures the rate of changes you ship or release that lead to service outages or service degradations.

- For the primary application or service you work on, how long does it generally take to restore service when a service incident or a defect that impacts users occurs (for example, 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 (for example, lead to service impairment or service outage) and subsequently require remediation (for example, require a hotfix, rollback, fix forward, patch)?

The final Quick Check question asks you to select your organization's principal industry—so we can show your survey results as compared to your industry peers.

We aggregated and analyzed the responses from the 2021 DORA State of DevOps Survey and, through cluster analysis, identified and labeled the groups based on the four key metrics. With statistically significant distinctions between clusters, we were able to categorize the groups into distinct software delivery performance profiles (low, medium, high, and elite). Using those performance profiles and your responses to the Quick Check, we provide an assessment of your software delivery performance and rank it against

other organizations in your industry. The Quick Check results are grouped into three categories: areas for improvement, industry performance comparison, and common performance characteristics.

Based on your performance profile, the first thing you'll see are three DevOps capabilities you could adopt that will have the most positive impact on your software delivery performance. Next you can examine the performance comparison and see how your team compares against those in your specific industry, or across all industries. Finally, you can look through the performance profile breakdown and dive into the differences at each level between low, medium, high, and elite.

How to use the Quick Check

Now that you have the results of the Quick Check, how do you interpret and use them to start driving change? Starting with the performance profile breakdown, you can take a look at individual aspects of software delivery performance and see how they compare across low, medium, high, and elite tiers. The comparisons demonstrate the differences between tiers and do not necessarily reflect the choices you made when you completed the Quick Check.

The industry performance comparison provides a view of how your software delivery performance compares across all industries and your specific industry. Software delivery performance and DevOps capabilities are noted in the State of DevOps Report as strategic differentiators across industries, and this comparison highlights where you stand amongst your peers and competitors.

The three improvement areas noted provide a starting point when considering how to enhance software delivery performance. To identify the most critical improvement, the Quick Check provides a secondary set of questions to help determine how you perform and which area to focus on first. After you provide answers to those questions, the three capabilities are ordered from most impactful to least impactful. Finally, each DevOps capability includes links to expanded explanations and examples, walking you through the specifics of how to implement, common pitfalls, and how to measure success.



Limitations of the Quick Check

The purpose of the Quick Check is to provide a baseline level of guidance and understanding around your team's software delivery performance, and to identify DevOps capabilities that can offer the most improvement. However the Quick Check only scratches the surface; it is impossible for these questions to fully capture your organization's internal initiatives or desired outcomes. Instead, consider the highlighted DevOps capabilities as attributes of a successful DevOps outcome. For example, deployment automation is often cited as a critical DevOps capability, but cannot be delivered in a vacuum—it should be delivered as part of an organizational investment in Cloud adoption or Infrastructure as Code initiatives.

As you review the DevOps capability recommendations, be sure to consider how they should be incorporated into a larger business or technical organizational outcome. Finally, it's important to remember that the Quick Check is limited to a single viewpoint and does not aggregate responses across a team or an organization. The results you see may vary from one individual to the next, depending on how they view your organization's software delivery challenges. Instead, think of the Quick Check as a guide, helping you start a discussion with your teams and begin the journey towards improving your software delivery performance.



The DORA Capabilities Assessment


We've discussed the Accelerate State of DevOps Report global research program and the DORA Quick Check, which is driven by the data collected in the SODR surveys, but what if you want more than just a quick online check? Enter the [DORA Capabilities Assessment](#). The DORA Capabilities Assessment is an assessment we deploy in your organization that gives us a robust measurement of your organization's capabilities as they pertain to software delivery.

The assessment looks at four distinct dimensions: technology, process, measurement, and culture. The capabilities assessment allows us to get a detailed look at your line of business (LoB) and the teams that compose that LoB. We compare your abilities to others in your industry, against high performers in your industry, and run a comparison across all industries. The detailed report that is generated creates a roadmap for your improvement journey and allows you to objectively determine which areas to focus your improvement efforts on first to maximize your ROI. The assessment and report help you prioritize your journey.

“That sounds great, how do I get that?”

Good question! We are currently deploying surveys to select customers that are interested in enhancing their software delivery capabilities on Google Cloud. The assessment takes a couple hours to set up and each employee's survey lasts between 20–40 minutes. The quality of the recommendations and the report we generate depends heavily on the amount of data we can collect. Because of this fact, we need to send the survey out to at least 100 individuals, and we must have a minimum of 50 responders to generate a report. Remember, this is the bare minimum. After we collect the data and generate a report, we go over the results with you and work with you to create a plan for improvement.

To learn more, reach out to your Google Cloud Account Team or visit our [Cloud Application Modernization Program \(CAMP\)](#) site.



For a deeper look into your organization, consider the DORA capabilities assessment.



How to use DORA to improve

Now that you have an understanding of what DORA is and the various ways we can use the research to analyze your own abilities, let's take a look at how we approach improvement.


Find your true north

Your “true north” is an aspirational, system-level business goal set by leadership. It could be an ideal that can't be achieved, such as zero change failures. Or it could be a tough goal that is one to three years out, such as a tenfold increase in productivity. We recommend aiming high and thinking 10x. Challenge yourself or challenge your team to set goals that seem out of reach. Missing a goal is not failure if the objective is continuous growth. It's much more satisfying to hit 50% of a seemingly impossible goal than it is to hit 100% of an easy benchmark. Remember, this journey is one of continuous improvement. There is no end and that's the point. In fact, as we've seen year over year with the SODR, if you are not continuously learning and improving, you are actually falling behind.

Measure your current performance

The next step is to understand your condition as it exists today. You can't set targets and goals until you establish a baseline. The DORA Quick Check and the DORA Capabilities Assessment can help you understand how you're doing in terms of your software development capabilities and outcomes.

Missing a goal is not failure if the objective is continuous growth.





Set a target

After you have a firm understanding of your capabilities as they exist today, it's time to set a target. These targets could be described using a format like OKRs (Objectives and Key Results), which begin with a qualitative objective, and then specify measurable key results (target conditions). For example, HSBC's CIO for Global Banking and Markets set every team's goal "to double, half and quarter every year: double the frequency of releases, half the number of low-impact incidents, and quarter the number of high-impact incidents."

A large global telco defined their two-year targets as "accelerating their release cycle from 3 months to under 2 weeks, increasing their product launch velocity from 7 months to under 2 months, and decreasing their mean-time to restore from over 2 hours to under 30 minutes".

What are your targets?

Select an area of improvement and improve

Regardless of your chosen assessment path, the Quick Check or the full DORA improvements. Both assessments will suggest capabilities that you can focus on to start making improvements. It is up to you to prioritize the capabilities based on your specific situation and goals, but you will be pointed in a general direction.

Dive into our [DevOps Capabilities website](#) to get a deep understanding of the different capabilities including:



How to implement them



How to improve them



How to monitor them



How to avoid common pitfalls associated with these capabilities

This site allows you to start up-leveling your own capabilities and shows you how to measure them.

Thinking of these capabilities as items that directly correlate to attaining your targets, you or your teams should begin experimenting with ways to achieve your goals. Foster an environment of exploration and tinkering. Teams should be running experiments daily to try to move towards the target conditions or key results.

Everybody on the team should ask themselves the following five questions every day:

01

What is the target condition?

02

What is the current condition?

03

What obstacles do you think are preventing you from reaching the target condition and which one are you addressing now?

04

What is your next step and what outcome do you expect?

05

When can the results be evaluated to see what can be learned from taking that step?

What happens once you've reached your target?

You start over and repeat the process.

Designate Fridays as “No meeting Friday”—and stick to it. No exceptions. Use this time to really go deep, to try radical things, to go far outside the box. Make sure to always take risks and not be afraid to fail—this is how you learn. This might sound like a far-off dream, but the more you relentlessly focus on improving your DevOps capabilities, the more this will become a reality and the more comfortable you will feel.

Repeat

We often say DevOps is a “continuous journey” or a “journey with no end” and it's true. What do you do now that you have reached your target? You start over and repeat the process.

Conclusion

As you study the capabilities, we also highly recommend you read the article [How to Transform](#).

We know that a DevOps transformation is about much more than just tooling, but that doesn't make the tools any less important. We have a variety of tools designed to maximize your DevOps capabilities and to allow you to get the most out of your cloud computing environment.

This is the first of many forthcoming chapters in the DevOps Guide. We will be digging much deeper into *how* to make these improvements and showing you how Google Cloud tools and technology can help you reach your goals. We realize this is a big journey and we are here to support you. Until next time, be well.

