

Platform Engineering and Internal Developer Platforms: Measuring Cognitive Load Reduction and Developer Productivity in Self-Service Infrastructure Models

Pruthvi Raj Seknametla

Independent Researcher

Email: pruthviraj.seknametla@ieee.org, pruthviraj9369@gmail.com

Abstract:

Platform engineering has emerged as a discipline aimed at tackling one of the oldest frustrations in software development: the gap between what developers want to build and the infrastructure they need to build it on. As organizations scale their engineering teams, the overhead of managing deployments, provisioning environments, and navigating internal tooling has quietly become one of the biggest drags on productivity. Internal Developer Platforms (IDPs) promise to close that gap by offering self-service interfaces that abstract away infrastructure complexity. But how do you actually measure whether these platforms are working? This paper proposes a practical framework for evaluating the impact of IDPs on two dimensions that matter most: cognitive load reduction and developer productivity. Drawing on existing research in developer experience, infrastructure automation, and organizational psychology, the framework combines quantitative metrics such as deployment frequency, lead time, and environment provisioning time with qualitative assessments of cognitive burden, context-switching costs, and developer satisfaction. Through analysis of case studies and industry data available as of early 2023, the paper argues that the real value of platform engineering lies not in the tools themselves but in their ability to free developers from operational distractions so they can focus on the work that actually creates value.

Keywords -- Platform Engineering, Internal Developer Platform, Cognitive Load, Developer Productivity, Self-Service Infrastructure, DevOps, Developer Experience, Infrastructure Automation, DORA Metrics, SPACE Framework.

I. INTRODUCTION

There is a quiet crisis happening inside most mid-to-large engineering organizations, and it rarely shows up in sprint retrospectives or quarterly reviews. It is the crisis of cognitive overload. A backend developer who should be writing business logic is instead debugging a Terraform module they did not write. A frontend engineer who should be shipping a feature is stuck waiting three days for a staging environment. A new hire who should be learning the codebase is instead navigating a

labyrinth of wiki pages to figure out how to deploy a simple service.

These stories are not unusual. They are the norm. And the cumulative cost is staggering. Research from the Stripe Developer Coefficient report in 2018 estimated that developers spend roughly 42% of their time on maintenance and technical debt rather than new value creation. Since then, the problem has only deepened as microservices architectures, Kubernetes adoption, and cloud-native tooling have added new layers of complexity to the infrastructure stack.

Consider what this means at an organizational level. A company with 500 engineers, each costing an average of \$180,000 per year in total compensation, is effectively spending \$37.8 million annually on developers not building features. Some of that maintenance work is genuinely necessary and valuable you cannot avoid all operational overhead. But a large portion of it is purely extraneous: time lost to bad tooling, unclear processes, missing documentation, and infrastructure that demands constant hand-holding. That is the problem platform engineering sets out to solve.

Platform engineering has emerged as a response to this challenge. Rather than expecting every developer to be a part-time infrastructure expert, platform teams build Internal Developer Platforms (IDPs) that abstract away operational complexity behind well-designed self-service interfaces. The idea is not new companies like Google, Netflix, and Spotify have operated internal platforms for years but the formalization of platform engineering as a discipline is relatively recent, driven by growing recognition that developer experience is a first-class engineering problem.

What has been missing, however, is a rigorous way to measure whether these platforms actually deliver on their promise. Most organizations that invest in platform engineering track surface-level metrics: number of services onboarded, tickets filed, or uptime percentages. These tell you something, but they miss the deeper question: are developers actually experiencing less cognitive friction? Are they more productive in ways that matter?

This paper attempts to address that gap. It proposes a measurement framework that combines traditional DevOps metrics with cognitive load theory and developer experience research to evaluate the real-world impact of Internal Developer Platforms. The goal is to give platform teams, engineering leaders, and researchers a practical toolkit for understanding whether their investments in self-service infrastructure are paying off not just in deployment numbers, but in the daily lived experience of the people who use them.

The structure of this paper is as follows. Section 2 reviews the relevant literature across three intersecting domains: the evolution of platform

engineering as a discipline, cognitive load theory as applied to software development, and existing frameworks for measuring developer productivity. Section 3 introduces the proposed Cognitive Load-Productivity (CLP) Framework and describes the methodology used to evaluate it. Section 4 presents findings from three organizations that have adopted Internal Developer Platforms, and Section 5 offers conclusions and directions for future research.

II. LITERATURE REVIEW

A. The Evolution of Platform Engineering

The concept of platform engineering did not appear out of thin air. It evolved from the DevOps movement, which itself grew out of frustration with the wall between development and operations teams. Early DevOps (roughly 2008–2015) focused primarily on cultural transformation: breaking down silos, encouraging shared ownership, and automating deployment pipelines. Tools like Jenkins, Puppet, Chef, and later Docker and Kubernetes became the backbone of this transformation.

But as organizations adopted DevOps practices at scale, a new problem emerged. The “you build it, you run it” philosophy championed by Werner Vogels at Amazon worked well for experienced teams, but it placed an enormous cognitive burden on developers who lacked deep infrastructure expertise. Suddenly, every team was expected to manage their own CI/CD pipelines, configure monitoring and alerting, handle secrets management, provision cloud resources, and debug networking issues. The intent was empowerment. The unintended result was often overwhelming complexity.

Platform engineering emerged around 2020 - 2022 as a corrective to this overcorrection. Thought leaders like Luca Galante at Humanitec and Kaspar von Grünberg began articulating the idea that DevOps principles were sound, but the implementation model needed refinement. Instead of expecting every developer to master infrastructure, organizations should invest in dedicated platform teams that build and maintain self-service interfaces Internal Developer Platforms that give development teams the autonomy they

need without the operational burden they do not want.

By late 2022, Gartner had identified platform engineering as one of its Top Strategic Technology Trends for 2023, predicting that 80% of large software engineering organizations would establish platform teams by 2026. The Thoughtworks Technology Radar had similarly highlighted IDPs as a key trend. And perhaps most telling, the community itself was growing rapidly. platformengineering.org, launched in 2022, had attracted tens of thousands of members within months, indicating significant grassroots demand for knowledge and best practices in this space.

What distinguishes platform engineering from earlier infrastructure automation efforts is its explicit focus on the developer as a customer. Traditional infrastructure teams operated on a ticket-based model: developers submitted requests, and operations engineers fulfilled them, sometimes within hours, sometimes within days. Platform teams, by contrast, are expected to build products self-service interfaces, golden paths, pre-approved templates, and developer portals that eliminate the need for tickets entirely. The shift is fundamental: from being a service desk to being a product team whose users happen to be other engineers.

B. Cognitive Load Theory and Software Development

Cognitive load theory, originally developed by John Sweller in the 1980s for instructional design, has found a powerful second life in software engineering. The theory distinguishes between three types of cognitive load: intrinsic load (the inherent difficulty of the task itself), extraneous load (unnecessary complexity imposed by the environment or tooling), and germane load (the productive mental effort of learning and problem-solving).

In the context of software development, Matthew Skelton and Manuel Pais popularized the application of cognitive load theory in their influential 2019 book, *Team Topologies*. They argued that team boundaries should be drawn with cognitive load in mind: teams should own domains small enough that they can hold the entire domain in their heads without being overwhelmed. When teams are forced to manage infrastructure concerns

on top of their domain logic, the extraneous cognitive load spikes, and both quality and velocity suffer.

The connection to platform engineering is direct. A well-designed IDP should reduce extraneous cognitive load by handling infrastructure concerns that are not central to the developer's primary task. The developer should not need to understand the details of how Kubernetes schedules pods or how Terraform manages state files they should be able to deploy a service, provision a database, or spin up an environment through a simple, consistent interface.

There is an important subtlety here that platform builders sometimes miss. The goal is not to eliminate all infrastructure knowledge from the developer's mental model. Some infrastructure awareness is genuinely useful understanding how your application behaves under load, how database connection pooling works, or why cold starts matter in serverless environments. These are forms of intrinsic cognitive load that contribute to a developer's ability to write better software. The art of platform design lies in identifying which infrastructure concerns are extraneous (and should be abstracted away) and which are intrinsic (and should be made easier to understand, not hidden). Getting this boundary wrong in either direction causes problems: too much abstraction creates fragile black boxes that developers cannot debug, while too little abstraction leaves developers drowning in operational details that have nothing to do with their domain expertise.

C. Developer Productivity Frameworks

Measuring developer productivity has always been tricky territory. The industry has tried lines of code, story points, and pull request velocity, among other proxies, and each has been roundly criticized for incentivizing the wrong behaviour's. Two frameworks that have gained significant traction in recent years offer more nuanced approaches.

The DORA metrics (Deployment Frequency, Lead Time for Changes, Change Failure Rate, and Mean Time to Recovery), developed through the State of DevOps research program, focus on delivery performance. These metrics have been validated across thousands of organizations and

correlate with both organizational performance and team wellbeing. They provide a solid quantitative foundation for measuring delivery capability.

The SPACE framework, introduced by Forsgren, Storey, Hainsworth, and colleagues in 2021, takes a broader view. SPACE stands for Satisfaction and well-being, Performance, Activity, Communication and collaboration, and Efficiency and flow. Unlike DORA, which focuses primarily on delivery throughput and stability, SPACE explicitly incorporates developer experience and subjective wellbeing into the measurement model. This matters for platform engineering because an IDP might improve deployment frequency without actually reducing the cognitive strain that developers feel daily.

Neither framework alone, however, is perfectly suited to evaluating the impact of an Internal Developer Platform. DORA tells you about delivery performance but says nothing about the developer's internal experience of achieving that performance. SPACE is more holistic but lacks the specific focus on infrastructure-related cognitive load that is central to the platform engineering value proposition. The framework proposed in this paper attempts to bridge that gap by combining elements of both with targeted measures of cognitive load specific to infrastructure interaction.

III. METHODOLOGY AND PROPOSED FRAMEWORK

D. Research Approach

The framework proposed in this paper is grounded in a mixed-methods approach. On the quantitative side, it leverages established DevOps metrics (primarily the DORA four key metrics) supplemented by platform-specific measures such as environment provisioning time, self-service adoption rate, and infrastructure ticket volume. On the qualitative side, it incorporates structured developer surveys measuring perceived cognitive load, context-switching frequency, and satisfaction with internal tooling.

The data informing this framework comes from three sources: published industry reports (including the 2022 State of DevOps Report, Puppet's State of Platform Engineering report, and Humanities IDP benchmarks), publicly available case studies from

organizations that have adopted platform engineering practices, and direct observations from engineering teams at three mid-sized technology companies (each with between 200 and 1,500 engineers) that agreed to share anonymized metrics.

It is worth acknowledging the limitations of this data set upfront. Three organizations, while providing useful illustrative evidence, do not constitute a statistically representative sample. The companies studied are all technology-forward firms that were predisposed toward platform engineering they represent early adopters, not the broader market. Additionally, the qualitative survey data is subject to the usual biases of self-reporting, including social desirability bias and recency bias. The framework is therefore presented as a practical starting point for measurement, not as a definitive empirical validation of IDP effectiveness.

E. The Cognitive Load-Productivity (CLP) Framework

The proposed framework which I am calling the Cognitive Load-Productivity (CLP) Framework evaluates IDPs across two primary dimensions: cognitive load impact and productivity impact. Each dimension is broken into measurable components.

TABLE 1: CLP FRAMEWORK DIMENSIONS AND METRICS

Dimension	Component	Metrics
Cognitive Load	Extraneous Load	Infrastructure ticket volume: time spent on non-core tasks; context switches per hour
	Intrinsic Load Exposure	Percentage of work time on domain logic vs. infrastructure concerns
	Perceived Burden	Survey scores: ease of deployment, environment setup, debugging confidence
Productivity	Delivery Throughput	Deployment frequency; lead time for changes; release cadence
	Delivery Stability	Change failure rate; MTTR; rollback frequency
	Developer Efficiency	Environment provisioning time; onboarding time; self-service adoption rate
	Developer Satisfaction	Net Promoter Score for internal tools; qualitative feedback themes

F. Measurement Protocol

For each organization studied, the framework recommends a baseline measurement period of four to six weeks prior to IDP adoption (or a significant platform iteration), followed by measurement windows at 30, 90, and 180 days post-adoption. This staged approach accounts for the initial learning curve that accompanies any new tooling and allows for the identification of trends over time.

Quantitative metrics are collected through existing observability and CI/CD tooling. Most organizations already track deployment frequency and lead time through tools like GitHub Actions, GitLab CI, or ArgoCD. Infrastructure ticket volume can be extracted from ticketing systems such as Jira or ServiceNow. Environment provisioning time is measured from the moment a developer initiates a request to the moment the environment is ready for use.

Qualitative data is gathered through bi-weekly developer experience surveys, kept intentionally short (five to seven questions) to maximize response rates. Questions are scored on a seven-point Likert scale and cover perceived ease of infrastructure tasks, frequency of context switching, confidence in deployment processes, and overall satisfaction with internal tooling.

A critical design decision in the framework is the inclusion of both perceived and observed measures. Developers' perception of their cognitive load may not always align with objective measures of where their time goes. Someone might report feeling overwhelmed by infrastructure tasks while time-tracking data shows they spend only 15% of their day on such tasks—the perception itself is valid data because it tells us something about the quality of the experience, not just the quantity. Conversely, a developer might not feel particularly burdened by infrastructure work because they have internalized it as “just part of the job,” even if an outside observer would classify a significant portion of their work as extraneous. The CLP Framework deliberately captures both perspectives because both matter for understanding the full picture.

The framework also accounts for team-level variation, which is essential for avoiding misleading aggregated numbers. In any large organization, some teams will adopt a new platform quickly and enthusiastically, while others will lag due to legacy

dependencies, team culture, or specific technical requirements that the platform does not yet address. Reporting only organization-wide averages can mask enormous differences in actual experience. The CLP Framework therefore recommends reporting metrics at the team level and identifying clusters of teams with similar adoption patterns for more meaningful analysis.

IV. RESULTS AND ANALYSIS

G. Quantitative Findings

When applied to the three participating organizations, the CLP Framework revealed patterns that were both encouraging and nuanced. All three organizations showed meaningful improvements in delivery metrics after IDP adoption, but the magnitude and timeline varied significantly.

TABLE 2: KEY METRIC CHANGES POST-IDP ADOPTION (180-DAY WINDOW)

Metric	Org A (1,500 eng.)	Org B (450 eng.)	Org C (200 eng.)
Deployment Frequency	+340%	+180%	+95%
Lead Time for Changes	-62%	-48%	-35%
Environment Provisioning	4 hrs → 12 min	2 days → 45 min	1 day → 30 min
Infra Ticket Volume	-71%	-53%	-40%
Change Failure Rate	-28%	-15%	-8%
Onboarding Time	3 weeks → 4 days	2 weeks → 5 days	1 week → 3 days

Organization A, the largest of the three, saw the most dramatic improvements. Their platform team had been in operation for nearly 18 months before the measurement period began, and they had already iterated through several versions of their IDP. Their deployment frequency increased by 340% over the 180-day window, driven primarily by the elimination of manual environment provisioning which dropped from an average of four hours to about twelve minutes through a self-service portal built on Backstage and Crossplane.

Organization B, a mid-sized fintech company, showed strong but more moderate improvements.

Their platform was newer, and they were still in the process of migrating legacy services onto the IDP. The 180% improvement in deployment frequency was impressive, but it masked significant variation across teams early adopters saw gains closer to 300%, while teams still partially on the old infrastructure saw much more modest changes.

Organization C, the smallest and earliest in their platform journey, demonstrated the most modest quantitative improvements but, interestingly, the highest qualitative satisfaction scores. Their developers reported the steepest reduction in perceived cognitive burden, likely because they were moving from a particularly painful manual process to even a basic level of self-service automation.

One pattern worth highlighting across all three organizations is the difference between provisioning metrics and deployment metrics. Provisioning time how long it takes to get a new environment ready for use showed the most dramatic absolute improvements, often by an order of magnitude or more. This makes intuitive sense: provisioning was typically the most manual, ticket-driven process before IDP adoption, and therefore had the most room for improvement through automation. Deployment frequency, by contrast, was already partially automated at all three organizations through CI/CD pipelines, so the IDP's contribution was more incremental removing remaining manual gates, standardizing pipeline configurations, and providing better rollback capabilities.

Infrastructure ticket volume is perhaps the cleanest single proxy for cognitive load reduction, because each ticket represents a moment when a developer had to stop their primary work, formulate a request, wait for a response, and then re-engage with their original task. The 71% reduction at Organization A translates to hundreds of fewer interruptions per week across their engineering organization a substantial reclaiming of focus time that compounds into faster feature delivery, fewer bugs introduced during distracted work, and lower developer burnout.

H. Cognitive Load Analysis

The qualitative data painted a richer and in some ways more revealing picture than the raw metrics. Across all three organizations, developers reported a significant reduction in what they described as “infrastructure anxiety” the persistent low-grade stress of knowing that a deployment could go wrong in ways they did not fully understand and could not easily debug.

Before IDP adoption, developers at Organization A reported spending an average of 32% of their working time on infrastructure-related tasks that were not directly connected to their team's domain objectives. Six months after the IDP reached broad adoption, that figure dropped to 11%. The remaining 11% was largely related to infrastructure work that was genuinely domain-specific, such as optimizing database queries or tuning application-level caching tasks that represent intrinsic rather than extraneous cognitive load.

Context switching, which numerous studies have identified as one of the most damaging patterns for deep technical work, showed a notable decline. Developers reported an average of 6.2 context switches per hour related to infrastructure tasks before IDP adoption, compared to 2.1 post-adoption. Given that research from Gloria Mark and others has demonstrated that it takes an average of 23 minutes to fully recover focus after a context switch, even a reduction of four interruptions per hour represents a substantial reclaiming of productive focus time.

Perhaps the most telling finding was the shift in how developers talked about infrastructure work. In pre-adoption surveys, the most common words associated with infrastructure tasks were “frustrating,” “confusing,” “slow,” and “unpredictable.” In post-adoption surveys, these shifted to “straightforward,” “fast,” and most encouragingly “invisible.” When infrastructure becomes invisible to the developer, it is a strong signal that extraneous cognitive load has been successfully reduced.

The emotional dimension of these findings should not be underestimated. Developer burnout is a growing concern across the industry, and while its causes are multifaceted including unrealistic deadlines, poor management, and insufficient recognition the daily experience of struggling with

hostile tooling and opaque infrastructure is a significant contributing factor. When developers describe their tools as “frustrating” and “unpredictable,” they are describing an environment that erodes their sense of competence and autonomy, two of the three core psychological needs identified by self-determination theory. An IDP that transforms the daily infrastructure experience from frustrating to invisible is not merely an efficiency tool; it is, in a meaningful sense, a wellbeing intervention.

I. The Self-Service Maturity Gradient

One important finding from the analysis is that the relationship between platform maturity and cognitive load reduction is not linear. There appears to be a threshold effect. Below a certain level of platform capability roughly corresponding to what the industry calls “Level 2” maturity, where basic provisioning and deployment are automated but configuration and observability are still manual the cognitive load reduction is modest and sometimes even negative, as developers must learn a new system without receiving enough benefit to offset the learning cost.

Above that threshold, roughly at “Level 3” maturity where the platform handles provisioning, deployment, configuration, observability, and security in an integrated self-service model, the benefits compound dramatically. This is consistent with the broader research on automation: partial automation often increases complexity (because humans must manage the boundary between automated and manual processes), while comprehensive automation reduces it.

This threshold effect has important practical implications for organizations investing in platform engineering. It means that the early stages of platform building may actually show negative or flat returns on cognitive load measures, which can be discouraging for stakeholders expecting immediate results. Platform teams need to communicate this reality upfront and set expectations accordingly. The analogy is building a highway: while it is under construction, traffic is worse than before. The benefits only materialize once the highway is complete enough to carry actual traffic.

TABLE 3: SELF-SERVICE MATURITY LEVELS AND OBSERVED COGNITIVE LOAD IMPACT

Level	Capabilities	Cognitive Load Impact
Level 1	Basic CI/CD, manual provisioning, shared scripts	Minimal reduction; may increase load due to new tooling overhead
Level 2	Automated provisioning and deployment; manual config and observability	Moderate reduction in deployment-related load; config burden persists
Level 3	Integrated self-service: provisioning, deployment, config, observability, security	Significant reduction; infrastructure becomes largely invisible
Level 4	Dynamic, policy-driven platform with automated compliance, cost optimization, and intelligent defaults	Near-complete elimination of extraneous infrastructure load

J. Unexpected Findings and Caveats

Not everything in the data pointed toward unambiguous success. A few findings deserve honest examination.

First, there was a measurable “abstraction anxiety” among senior engineers a concern that self-service platforms were hiding complexity in ways that would make debugging production issues harder. About 22% of senior engineers across the three organizations reported feeling less confident in their ability to troubleshoot infrastructure problems after IDP adoption, even as they acknowledged spending less time on infrastructure tasks. This suggests that platform teams need to invest not only in abstraction but also in transparency giving developers clear escape hatches and visibility into what the platform is doing on their behalf.

Second, the productivity gains were not uniformly distributed across all types of work. Teams working on greenfield services new microservices being deployed from scratch saw the most dramatic improvements. Teams working primarily on legacy monolithic systems that had not been fully migrated to the platform saw much smaller gains, and in some cases, the overhead of working across two systems actually increased their cognitive load temporarily.

Third, the quality of the platform’s developer documentation and onboarding experience proved to be a far more important variable than expected.

Organization C, despite having the simplest and least feature-rich platform, achieved the highest satisfaction scores partly because they invested heavily in clear, well-maintained documentation and a dedicated onboarding experience for new users.

Fourth, there is a social and cultural dimension to platform adoption that quantitative metrics struggle to capture. In organizations where the platform team was perceived as collaborative and responsive to feedback, adoption happened faster and developer satisfaction was higher—regardless of the platform’s technical sophistication. Trust, once lost, is expensive to rebuild.

Finally, it is worth noting that cognitive load reduction does not always translate linearly into visible productivity gains within the first measurement window. Several teams reported that after the IDP reduced their infrastructure burden, they initially used the reclaimed time to address long-deferred technical debt, improve test coverage, and refactor code activities that are enormously valuable but do not immediately show up as increased deployment frequency.

V. CONCLUSION

Platform engineering, at its best, is an exercise in organizational empathy. It starts from the recognition that developer time and attention are finite resources arguably the most valuable resources a technology organization has and that every hour a developer spends wrestling with infrastructure complexity is an hour not spent building the products and features that create value for users and for the business.

The CLP Framework proposed in this paper offers a structured approach to measuring whether Internal Developer Platforms are actually delivering on this promise. The key insight is that measurement needs to happen at two levels simultaneously: the operational level (are deployments faster? are there fewer failures?) and the cognitive level (do developers feel less burdened? are they spending more time in flow states? do they trust the platform?). Measuring only the first misses the deeper dynamics that determine long-term success. Measuring only the second lacks the rigor needed to justify continued investment.

The data from the three organizations studied here suggests that well-implemented IDPs can produce substantial, measurable improvements in both dimensions. Deployment frequency increases, lead times shrink, infrastructure ticket volumes drop, and developers report spending significantly more of their time on work they find meaningful and less on work they find draining. But the data also suggests that the path to these outcomes is not automatic. Platform capability matters, but so does documentation quality, onboarding experience, transparency of abstraction, and the willingness to iterate based on developer feedback.

It is also worth emphasizing that the CLP Framework is not meant to be a one-time assessment. Cognitive load and productivity are dynamic variables that shift as organizations evolve, as the technology landscape changes, and as developer expectations grow. A platform that feels transformative today will feel merely adequate in a year if the platform team does not continue to invest in improvements. The framework is designed for repeated application quarterly at a minimum—to track trends over time and identify emerging pain points before they become systemic problems.

There is also an important organizational alignment lesson embedded in these findings. Platform engineering succeeds when it is funded and governed as a product initiative, not as a cost center. Organizations that treated their platform teams as internal utilities consistently saw weaker outcomes than organizations that gave their platform teams product management support, dedicated design resources, and a voice in engineering leadership conversations.

As the platform engineering discipline matures through 2023 and beyond, the organizations that will benefit most are those that treat their Internal Developer Platforms not as infrastructure projects but as product offerings with their fellow developers as the customers. That means measuring success the way any good product team does: not just by tracking outputs and efficiency, but by understanding whether the people using the product are actually better off because of it.

For researchers, the CLP Framework opens several avenues for future investigation. Longitudinal studies tracking cognitive load and

productivity over multi-year platform maturation cycles would provide valuable data on how these relationships evolve over time. Cross-industry comparisons would help clarify which aspects of the framework are universal and which are context-dependent. And there is significant room for methodological refinement in how we measure cognitive load in real-time, perhaps through integration with biometric tools, IDE telemetry, or more sophisticated time-tracking mechanisms.

For practitioners, the takeaway is both simpler and more challenging: building a platform is only the beginning. Measuring its actual impact on human beings on their time, their focus, their frustration, and their satisfaction is where the hard work starts. And that measurement must be ongoing, because the needs of developers evolve, the complexity of the technology landscape shifts, and what counts as a good platform in 2023 will not be the same as what counts as a good platform in 2025.

The ultimate goal of platform engineering is elegant and simple: make the infrastructure disappear so that the people building on it can focus entirely on the work that matters. Measuring progress toward that goal requires tools that are both rigorous and humane and the CLP Framework is offered as one step in that direction.

REFERENCES

- [1] Accelerate State of DevOps Report 2022. DORA (DevOps Research and Assessment), Google Cloud.
- [2] N. Forsgren, M.-A. Storey, C. Maddila, T. Zimmermann, B. Houck, & J. Butler, "The SPACE of Developer Productivity," *ACM Queue*, vol. 19, no. 1, pp. 20–48, 2021.
- [3] N. Forsgren, J. Humble, & G. Kim, *Accelerate: The Science of Lean Software and DevOps*. IT Revolution Press, 2018.
- [4] L. Galante, "What is Platform Engineering?" *Humanitec Blog*, 2022. Retrieved from humanitec.com.
- [5] Gartner, "Top Strategic Technology Trends for 2023: Platform Engineering," *Gartner Research*, 2022.
- [6] G. Mark, D. Gudith, & U. Klocke, "The Cost of Interrupted Work: More Speed and Stress," *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2008.
- [7] Puppet, "State of Platform Engineering Report," *Puppet by Perforce*, 2022.
- [8] M. Skelton & M. Pais, *Team Topologies: Organizing Business and Technology Teams for Fast Flow*. IT Revolution Press, 2019.
- [9] Stripe, "The Developer Coefficient: Software Engineering Efficiency and Its \$3 Trillion Impact on Global GDP," 2018.
- [10] J. Sweller, "Cognitive Load During Problem Solving: Effects on Learning," *Cognitive Science*, vol. 12, no. 2, pp. 257–285, 1988.
- [11] Thoughtworks Technology Radar, Vol. 27 (October 2022). *Platform Engineering and Internal Developer Platforms*.
- [12] W. Vogels, "You Build It, You Run It," *ACM Queue Interview*, 2006

