STATE OF THE ROUTE TO LIVE 2026 REPORT

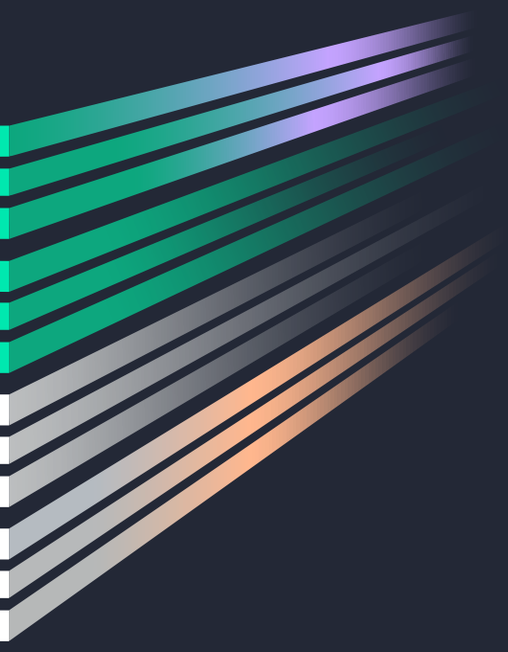

CLEARROUTE

STATE OF THE ROUTE TO LIVE 2026 REPORT

THE STATE OF THE ROUTE TO LIVE REPORT PROVIDES A DATA-DRIVEN EXAMINATION OF THE CURRENT STATE OF SOFTWARE DELIVERY.

TODAY'S ORGANISATIONS FACE UNPRECEDENTED PRESSURE TO DELIVER SOFTWARE FEATURES AND PRODUCTS AT AN EVER ACCELERATING PACE.

THE ABILITY TO EFFICIENTLY MOVE IDEAS FROM CONCEPTION TO PRODUCTION WHAT WE CALL THE "ROUTE TO LIVE" (RtL) HAS BECOME A CRITICAL COMPETITIVE DIFFERENTIATOR.

WHAT'S INSIDE

LIBERATE THE ROUTE TO LIVE

CLEARROUTE

What are the most common software delivery bottlenecks across the Route to Live? And how will AI accelerate governance in the enterprise?

The answer requires a fundamental reframing. Bolting AI onto the existing SDLC creates a chaotic amplifier. We make developers faster, but overwhelm the other 80% of the lifecycle. The real opportunity is to treat AI not as a tool invoked at discrete points by individual engineers, but as a governed participant embedded across the entire Route to Live, from requirements to production.

This requires a platform that manages AI with workload identity, policy-as-code, and a complete audit trail, answering the question every CISO will ask: "What did the agent do?," and "on whose behalf?" The enterprises that win will have treated this as a platform and governance problem before it becomes a production incident.

This State of the Route to Live report provides the data-driven case for this new model. It offers both strategic guidance for C-suite executives and actionable recommendations for engineering teams to build it.

Sarndeep Nijjar
Chief Technology Officer

STATE OF THE ROUTE TO LIVE 2026

THE GREAT AI AMPLIFIER

Benchmarks, bottlenecks, and breakthroughs: A data-driven roadmap for **liberating the route to live**, overcoming enterprise friction, and embedding AI at scale.

WHAT IS THE ROUTE TO LIVE?

The Route to Live (RtL) encompasses the end-to-end journey of getting ideas into customers' hands, including the processes, tools, and cultural practices that enable or impede delivery excellence.

Our analysis is built on four years of deep-dive Route-to-Live assessments, spanning an immense body of work across financial services, retail, healthcare, media, and tech.

Software shouldn't crawl to production. Yet for more enterprises, it does. The story has been remarkably consistent. Almost every organisation wrestles the same chokepoints: weeks of lead-time lost to manual gates[1]; brittle test suites masking defects[2]; and delivery teams cut off from the business they're meant to serve.

Last year, we talked about an industry on the precipice of generational change. By 2026, that AI inflection point has passed. What was once on the horizon is now a live reality, deeply embedded in the software lifecycle.

But it hasn't been the silver bullet many hoped for. Instead, it has become the great amplifier which is both good and bad. Our benchmarks show how AI magnifies the strengths of high-performing organisations and the dysfunctions of struggling ones.

AI MAGNIFIES THE STRENGTHS OF HIGH PERFORMING ORGANISATIONS AND THE DYSFUNCTIONS OF STRUGGLING ONES.

THE 80% PROBLEM

We're seeing that happen because of what we're calling **The 80% Problem**.

We've seen consistently how AI has made the 20% of the software lifecycle ***the act of writing code*** faster. But for many teams, the 80% ***the sprawling and complex system of testing, security, compliance and release orchestration*** is still broken. This is what we mean when we say delivery is trapped. Not by lack of ambition, but by a system that has accumulated weight to slow it down.

For teams battling mountains of toil and friction, this creates a new paradox: code gets written faster but time to market doesn't improve. Pressure on the existing bottlenecks just intensifies.

ANATOMY OF THE 80% PROBLEM

AI has successfully amplified the 20% of the software lifecycle: the act of writing code. But its impact stops abruptly when it collides with the other 80%: the sprawling system of manual process and organisational friction where real-world delivery speed is lost.

Our data reveals a consistent pattern of this friction across the industry. Across many large enterprises, we found:

1 PROCEDURAL GRIDLOCK
Change approval processes taking **four to eight weeks**.[3] This macro-level delay was compounded by micro-level friction, where even simple code changes languished for **two to four days** just waiting for a peer review.[4]

2 LEGACY QA BOTTLENECKS
Overwhelmingly manual testing processes. In one global financial institution, we found a critical system where **95% of testing was still manual**. This was a systemic issue where the CI pipeline provided no integration test coverage, forcing a manual, **three-week regression cycle** for every feature.[5]

3 INFRASTRUCTURE & DATA FRICTION
Convoluted environment management creating systemic friction and stalling developer velocity. One leading bank was a prime example, with a sprawling estate of **over 28 inactive dev/test environments** creating huge confusion and data bottlenecks. The process to provision test data into any of these environments was a manual service desk request that took **up to three days** to fulfill.

The upshot? Many organisations are investing heavily to implement AI into the RtL but not achieving anywhere near the ROI they were expecting.

Where organisations are taking a more holistic view moving towards Level 3 and beyond in our ***AI Maturity Framework in Section 4*** they're still held back from meaningful progress thanks to escalating agent sprawl. For these higher-maturity teams, we're seeing the exact same problem we highlighted in our 2025 report of "tool sprawl," only now the tools are autonomous AI agents.

Wherever you are today, this report will show you the explicit levers you can pull to move the needle, turning AI from a source of pressure into a genuine force multiplier.

A HOLISTIC APPROACH TO CHANGE

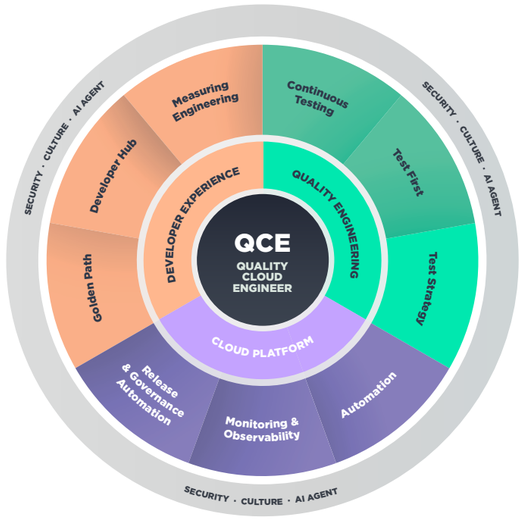

See full page diagram on page 39

Through the lens of the Quality Cloud Engineering framework

We've seen time and again that the true transformation happens when you look through one lens, to consider the Route to Live as a whole. That's just as true as organisations seek to integrate AI.

The Quality Cloud Engineering framework, ***covering Cloud Platform, Developer Experience, and Quality Engineering and with a unified foundational layer of Security, Culture and AI***, gives teams the structure to systematically unblock complex, interacting barriers to production and integrate AI holistically at scale.

You'll see throughout this report how this philosophy comes to life, whatever teams' DORA profile and wherever you are on your SDLC and AI Maturity journey.

WHATEVER YOUR CURRENT STATE OF RTL MATURITY, AI ADOPTION, & AI PREPAREDNESS...

READ THIS REPORT TO:

- **Benchmark against a new, more nuanced model of team performance**
- **See how the five big myths of delivery have evolved in an AI-powered world**
- **Understand AI's real-world impact: the new challenges and opportunities**
- **Get a strategic roadmap for building a governable, AI-augmented RtL**

[1] Source: *State of the Route to Live 2025 Report*, which found a median lead time of 30–60 days for large enterprises.

[2] Source: *State of the Route to Live 2025 Report*, where we found end-to-end test success rates as low as 45% and manual testing cycles taking over 50 days in extreme cases.

[3] Source: *State of the Route to Live 2025 Report*, which found that a single feature change often required coordination across 4 to 10+ teams, with manual governance boards adding 4 to 8 weeks of delay.

[4] Source: *State of the Route to Live 2025 Report*. which found code review wait times averaged 4.5 days, validating the significant micro-friction that slows down development.

[5] Source: *State of the Route to Live 2025 Report*, which detailed CI pipeline failure rates of up to 98% in some organisations, explaining the reliance on lengthy manual regression cycles.

RUNNING THE NUMBERS

Running the numbers

IN 2026, ENGINEERING PERFORMANCE IS ABOUT HEALTH, STABILITY, AND SUSTAINABILITY. NOT JUST SPEED.

In 2025, our analysis focused on the hard numbers that define the Route to Live: lead time, change failure rates, and deployment frequency. Those metrics still matter.

But in 2026, the industry's understanding of performance has evolved. It's no longer just about the "what," how fast you ship, but also the "how," the health, stability, and sustainability of the teams doing the shipping.

This section presents the latest data through this new, more holistic lens. First, we'll establish a baseline with the familiar core metrics. Then we'll introduce the new, more nuanced model of team performance that's shaping today's industry conversation.

The scale of process friction is staggering.

Our assessments at one global financial institution revealed that a £2.76M annual engineering investment yielded only 5% of its value in delivered features. That's astronomically wasteful. And avoidable.

The inverse is equally powerful: one retail bank we partnered with unlocked over AUD 20M in annual engineering savings by building a unified DevEx platform to eliminate process friction.

[6] Source: *State of the Route to Live 2025 Report*, which identified a median lead time range of 30–60 days, indicating that while the top end has improved, the core challenge remains.

[7] Source: *State of the Route to Live 2025 Report*, showing elite teams achieve deployment frequencies up to 12x higher than their low-performing peers.

[8] Source: *State of the Route to Live 2025 Report*, which found change failure rates as high as 9.15%, aligning with the lower end of our 2026 observations for non-elite performers.

[9] Source: *State of the Route to Live 2025 Report*. Projections based on technical debt analysis showed a potential 15-25% reduction in rework from improved architecture, highlighting the high cost of existing instability.

[10] Source: *State of the Route to Live 2025 Report*, which cited that high-performing organisations typically recover from incidents in under one hour.

THE CORE METRICS AN INDUSTRY BASELINE

Our 2026 assessments reveal a consistent and troubling pattern: **while the best get better, many large enterprises remain stuck**. The gap between the top performers and the rest of the industry is widening, and AI is the primary accelerant.

Lead Time to Production

The chasm between the fastest and slowest is stark. We've observed lead times as low as **24 hours** for elite teams with fully automated pipelines, while many large organisations still measure lead time in **months, not days**. The median lead time for enterprises remains stubbornly high at **30-45 days**.[6]

Lead Time

24 hours ELITE vs **Months** LAGGING

Deployment Frequency

Top-quartile teams are now deploying multiple times per day using AI-assisted CI/CD to accelerate testing and validation.[7] But most enterprises are still deploying between 4 and 17 times per year, placing them firmly in what DORA would have previously classified as "Low" performers. This cadence is a major competitive disadvantage in a market that demands rapid iteration.

Deployment

Continuous ELITE vs **4x/year** LAGGING

In contrast: by implementing an agent-driven AI Workbench, we helped a UK-based national energy operator achieve a **3x uplift in development frequency**.

Change Failure Rate

Elite teams keep their Change Failure Rate below 5%, but we consistently observe rates between **10% and 20%**[8] in organisations with high degrees of manual testing and release management. The cost of this instability (in both rework and customer trust) is substantial.[9]

CFR

<5% ELITE vs **20%** LAGGING

Mean Time to Restore

High-performing teams can restore service **in under an hour**,[10] but for many enterprises MTTR is measured in **days**, hampered by complex dependencies, a lack of observability and manual rollback procedure.

MTTR

<1 hour ELITE vs **days** LAGGING

In contrast: a global travel and loyalty group battling a platform with critical P1 incidents every 36 hours achieved a **75% reduction in their change failure rate** by re-engineering their core platform and hardening their release process.

A NEW MODEL FOR PERFORMANCE
THE SEVEN TEAM PROFILES OF 2026

AI IS WIDENING THE PERFORMANCE DIVIDE. WE'RE SEEING THE BEST GET BETTER, WHILE THE REST REMAIN STUCK.

For years, the industry has measured performance using broad clusters like 'Elite', 'High', 'Medium', and 'Low'. In 2026, that model is no longer good enough. AI is creating new, more complex patterns of strengths and weaknesses that need new classification.

Reflecting this new reality, the latest industry research from DORA has deprecated the four-cluster model. In its place are seven distinct team profiles, giving a more nuanced way to understand performance that moves beyond simple speed and stability to look at the interplay between delivery, well-being and process maturity.

Your profile directly correlates to your ability to implement AI successfully, and the relationship is a symbiotic, virtuous circle. Higher performers have the right foundation to accelerate AI maturity, while AI maturity accelerates delivery performance. And the inverse is true too.

This is the AI amplifier effect.

THE SEVEN TEAM PROFILES

Wherever your organisation fits, it's clear enterprise engineering performance has evolved. In an AI-augmented world, excellence has become much more nuanced and complex, ***with a clear and direct correlation between AI maturity and SDLC maturity.***

1 FOUNDATIONAL CHALLENGERS

Struggling with the basics of performance, these teams have low stability and slow delivery. They are often mired in technical debt and manual processes, making any change difficult and risky.

THE AI AMPLIFICATION EFFECT

When Foundational Challengers adopt AI code-generation, the **Great Amplifier effect** is immediate and negative. The increased volume of code submissions collides with a development process already hampered by extreme friction. This results in an even higher number of failed builds, overwhelming CI systems that, according to our 2025 data, were already failing up to 98% of the time in similar environments.[5]

The promise of AI speed becomes a reality of increased frustration, as developers are forced to context-switch constantly to deal with a pipeline where end-to-end test success rates were already as low as 45%.[2] Ultimately, without a stable underlying platform, AI-driven development only creates more chaos.

2 THE LEGACY BOTTLENECK

Here teams are trapped by their own history, often supporting a critical legacy system that is difficult to change but essential to the business. They may have pockets of high capability but all innovation is choked by the constraints of the legacy core.

THE AI AMPLIFICATION EFFECT

For the Legacy Bottleneck, AI-driven development creates a dangerous pressure cooker. Faster feature creation in modern, peripheral systems only serves to highlight the core bottleneck, increasing business pressure to "just fix" the legacy component. Without a strategy to modernise, AI just accelerates the rate at which teams collide with the legacy monolith, amplifying frustration and reinforcing the belief that real change is impossible.

The core problem remains the high cost of change and risk of failure, which our 2025 data shows can be as high as 9.15% even for planned changes in such systems.[8] This makes the case for a platform approach that isolates the legacy system and provides safe, stable integration points for new, AI-driven services.

3 CONSTRAINED BY PROCESS

These teams may have good technical skills, but they are suffocated by bureaucratic, manual, and risk-averse organisational processes. Every step requires a ticket, a manual approval, and a lengthy wait in a queue.

THE AI AMPLIFICATION EFFECT

AI amplifies the pain of manual governance. These teams can now generate code and even test plans in hours, only to see them sit idle for weeks awaiting approval from a manual CAB or security review.[3]

The speed mismatch turns AI from a productivity tool into a spotlight on organisational dysfunction. Developers become demotivated as their AI-accelerated work simply piles up in a longer queue, with the 2025 data on a 35% time-spend on non-productive toil feeling like an underestimate.[12]

There's a clear case here for a platform strategy that automates these manual gates into intelligent, software-defined guardrails.

4 HIGH IMPACT, LOW CADENCE

Often found in highly regulated industries, these teams deliver significant, valuable, and stable software, but do so infrequently. Releases are large, high-stakes events, often planned months in advance.

THE AI AMPLIFICATION EFFECT

AI-powered code generation provides little value here. The bottleneck is not the speed of coding but the lengthy, high-ceremony release process. AI may help generate documentation or test cases, but it cannot accelerate the business cycles of validation, legal review, and market timing that dictate their low cadence.

This leads to a perception that AI is "not for us". That's not the case. What these teams need is a platform that provides provably compliant, automated release pathways to de-risk frequent deployments.

5 STABLE & METHODICAL

These teams are the silent majority in many enterprises. They aren't elite but they aren't failing either. They deliver with a predictable, although not remarkable, cadence and maintain a reasonable level of stability. They often follow established processes without question.

THE AI AMPLIFICATION EFFECT

Stable and Methodical teams see modest productivity gains from AI in code generation but they don't challenge their underlying processes and so miss the bigger opportunity for transformative change.

AI helps them do the same things, slightly faster. They risk falling into a "local maximum", where AI adoption solidifies existing practices rather than inspiring the leap to a truly elite performance model enabled by a cohesive platform. The impetus to change can feel gentle here because AI isn't causing disastrous issues. But there's an enormous missed opportunity cost that will become existential over time.

6 PRAGMATIC PERFORMERS

These are capable teams who have carved out an effective bubble within a wider, more dysfunctional organisation. They have automated their pipelines, adopted modern practices, and deliver quickly and reliably, but often in defiance of official processes.

THE AI AMPLIFICATION EFFECT

AI amplifies Pragmatic Performers' bubble, *and its limitations.* These teams can use AI to go even faster but their impact remains constrained by their interfaces with the rest of the slow-moving organisation. The friction at their boundaries, where they must interact with manual governance or legacy systems, becomes even more pronounced and frustrating.

Here, a platform strategy becomes essential to help them scale their bubble of good practice to other teams and safely expose their services across the enterprise.

7 HARMONIOUS HIGH-ACHIEVER

The ideal state. These teams are fast, stable, and aligned with the business. They have a mature and automated Route to Live, a culture of continuous improvement, and the psychological safety to experiment and learn.

THE AI AMPLIFICATION EFFECT

For these teams, AI is a true force multiplier. With a stable, automated platform already in place, they can safely experiment with and deploy AI-driven capabilities across the entire lifecycle.

AI doesn't just help them write code faster; it helps them test smarter, release more safely, and gain deeper insights from operational data. Their platform allows them to accelerate their already virtuous cycle of improvement, widening the gap with all other profiles.

They embody the 2025 finding that high-performing organisations achieve 12x higher deployment frequency and suffer 75% fewer production defects.[11]

ON AVERAGE, INDIVIDUALS SPEND

4.5 DAYS PER FEATURE CHANGE, WAITING FOR CODE REVIEWS

35% OF TIME ON TOIL

23 TOOLS PER DAY NAVIGATED

IN SEVERAL CASES, MORE THAN HALF OF DELIVERY TIME WAS SPENT ON MANUAL ACTIVITIES.

[11] Source: *State of the Route to Live 2025 Report*, benchmark data for high-performing organisations.

[12] Source: *State of the Route to Live 2025 Report*, where developer surveys revealed that over 35% of a typical developer's week was spent on manual toil and non-productive tasks.

INSIGHTS FROM THE EXPERTS

Insights from the experts

THE EVOLVING MYTHS OF AN AI-POWERED WORLD

In 2025, we debunked five pervasive myths that impact how engineering happens and consistently hurt organisations' Route to Live.

But in 2026, these myths haven't disappeared. With AI acting as the great amplifier, they've evolved into more subtle, more insidious forms. And challenging them is more critical to performance than ever.

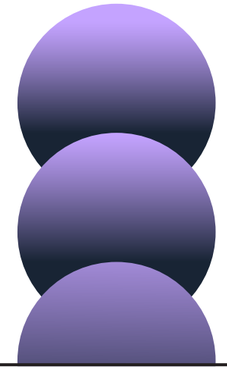

MYTH 1
MORE PROCESS = MORE SAFETY

THE OLD MYTH

For most enterprises, governance has typically been a manual, paper-based process. Teams believed the mountain of documents, checkpoints, and serial sign-offs was a necessary evil for safety and stability.

THE NEW MYTH

AI-GENERATED CODE IS *INHERENTLY SAFE*.

The old myth of manual governance is being replaced by blind trust in AI's output. The speed of generation creates an illusion of reliability, and a dangerous assumption that the code produced is secure and correct by default.

At one firm we assessed, the average lead time for a business-critical feature, or "Epic," was 266 days. Even more alarming, the average time to fix a critical production bug was 158 days. These metrics expose the truth: without a holistic view of the Route to Live, sprint-based work offers only an illusion of speed.

Automated processes, however, fundamentally change this equation. At a global mobility company, we saw how agentic workflows could turn this on its head. By creating a closed loop between production defects and test generation and ensuring the test suite was continuously learning from real-world failures, the enterprise reduced production defects by over 50%.

NEW REALITY

GOVERNANCE *MUST BE AUTOMATED*, NOT ABANDONED.

Research and real-world data tell a different story. While AI often boosts speed, it often introduces subtle security flaws, bugs, and insecure dependencies.

The adoption of AI has intensified the need for governance, but not the old manual kind. The consensus is that AI's speed must be balanced by a robust set of automated guardrails. We're seeing leading organisations shift focus to **automated governance**, where:

- All AI output is treated as an untrusted draft.
- Automated tests are required for any AI-generated code *before* it's merged.
- Policy-as-code programmatically enforces security and compliance standards.

AI often boosts speed but at the cost of introducing subtle flaws, bugs and dependencies. Robust automated governance guardrails are non-negotiable.

MYTH 2
MORE QA = HIGHER QUALITY

THE OLD MYTH

Bloated QA functions felt like an insurance policy. Enterprises kept hiring testers or throwing resources at test automation, believing volume equated to quality.

THE NEW MYTH

AI TESTING REPLACES THE NEED FOR QA PROFESSIONALS.

A new misconception is emerging, positing that AI-powered testing tools, with their ability to generate tests and self-heal, will make the QA role redundant.

A sophisticated AI strategy can rebalance the test pyramid. At a capital markets technology firm, we observed how a reliance on slow, end-to-end tests created a classic "ice cream cone" anti-pattern, with core C++ modules having zero unit test coverage. By using code-exploring AI agents to generate targeted unit tests, they achieved 80% coverage on these critical modules and improved build times by 65%.

NEW REALITY

QUALITY IS BECOMING AN ENGINEERING DISCIPLINE.

The focus is shifting from the quantity of tests to the intelligence of the overall quality process. Rather than finding individual bugs, sophisticated teams are leveraging AI to design strategic systems that prevent defects, not just catch them.

For example, a UK energy operator saw a 96% reduction in test automation runtime after building an AI workbench, proving that smarter automation leads to faster, not just more, testing.

AI is transforming quality assurance into a true engineering discipline. The most effective teams are evolving their quality capabilities, with an emphasis on:

- Curating and validating AI-driven test strategies.
- Analysing results and identifying patterns that AI might miss.
- Engineering a holistic, automated quality *system*.

AI is shifting focus from running more tests to building better systems.

MYTH 3
MICROSERVICES ALWAYS IMPROVE DELIVERY

THE OLD MYTH

The love affair with microservices sold a simple story: smaller pieces mean faster change. So many enterprises adopted them, expecting an automatic boost in velocity.

THE NEW MYTH

AI WILL DESIGN OUR MICROSERVICES FOR US.

With AI's ability to rapidly generate boilerplate code for new services, a new myth is taking hold: that AI can handle the complex architectural thinking required for a successful microservices strategy.

At one technology firm, the landscape was littered with high-friction, home-grown tooling. A legacy test framework took hours to run basic checks while another could take days for a full regression. Their primary backend monorepo had swelled to over 20GB, creating a complex thicket of dependencies that actively worked against developer productivity.

NEW REALITY

ARCHITECTURAL DISCIPLINE IS *MORE CRITICAL*, NOT LESS.

AI is an amplifier. If you have a solid, well-defined architecture, AI can help you implement it faster. But if your architectural discipline is weak, AI will help you create a distributed monolith at an astonishing speed.

The speed of AI code generation makes strong architectural oversight more critical, not less. AI is a powerful tool for implementing a well-defined architecture, but it can't (yet) replace the critical thinking, business context and trade-off analysis of a skilled architect.

If your architectural discipline is weak, AI will help you create a distributed monolith at astonishing speed.

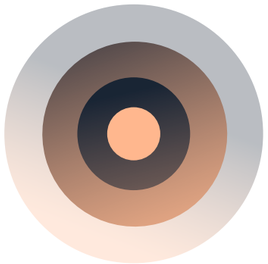

MYTH 4
ENGINEERS DEVELOP WHAT THE BUSINESS NEEDS

THE OLD MYTH

There's a common assumption that because engineers are close to the product, they're well-aligned to business context. So they'll naturally optimise for value.

THE NEW MYTH

AI-DRIVEN VELOCITY EQUALS BUSINESS VALUE.

The speed of AI-powered code generation has turned velocity into a dangerous vanity metric. The new myth is that if teams are using AI to ship code faster, they must be delivering more value.

We helped a global loyalty company shift time-to-market for new customer propositions from years down to just a few weeks through focus on strategic alignment, showing the clear link between engineering improvement and business agility.

NEW REALITY

SPEED WITHOUT DIRECTION IS WORTHLESS.

AI is a powerful implementation engine but it (currently) lacks genuine business context. The most successful organisations are making business alignment a core part of their AI strategy, ensuring that AI's incredible speed is pointed squarely at measurable business outcomes.[14]

With AI handling more of the "how", the value of engineering and product leaders is shifting emphasis to the strategic "what" and "why."

AI can build the feature right, but it can't tell you if you're building the right feature.

MYTH 5
DEVELOPER EXPERIENCE DOESN'T MATTER

A world-class DevEx platform acts as an incredible force multiplier. For example, at a major Australian retail bank facing intense regulatory pressure, a new unified developer platform reduced a critical regression suite from **8 hours to just 2 minutes** and cut the time to stand up a new service from over a month to under 5 minutes. They also achieved 99% uptime on their pre-production environments, giving developers a reliable foundation to build upon. This stability is a prerequisite for rolling out AI effectively.

THE OLD MYTH

Leadership has historically treated DevEx as a perk or an afterthought: nice-to-have but not essential.

The cost of a poor developer experience is often clearest during onboarding. At a major financial services institution, we found it took new engineers between **three and six months** to make their first contribution.

This wasn't a talent gap; it was a systemic failure caused by a high-friction environment. That's an enormous speed restriction, *and an expensive one*.

THE NEW MYTH

AI REDUCES DEVELOPER TOIL.

The initial narrative was that AI assistants would eliminate the tedious, frustrating, repetitive parts of a developer's day, reducing toil and improving the developer experience by default.

NEW REALITY

AI OFTEN *SHIFTS* TOIL, FROM TYPING TO VALIDATING.[12]

The reality is more nuanced. AI often replaces the toil of writing code with the high-stakes cognitive load of validating, debugging, and refactoring plausible-but-wrong suggestions. One 2025 study found experienced developers could be 19% slower when using AI because of this extra toil.

In the age of AI, an excellent developer experience is no longer a perk; it's a prerequisite.[13] A core goal of a good developer platform is creating stability, giving developers a reliable foundation to build on. To harness AI successfully, you must reduce the cognitive load across the entire system. This means providing clean APIs, clear documentation, stable environments, and a low-friction "golden path" to production. Without this foundation, AI tools just add more noise, and amplify the frustration of a broken system.

To harness AI successfully, you first need to reduce cognitive load across the entire system. Or AI tools will just add more noise and amplify frustration.

AI IS NOT A FIX-ALL FOR EVERY DELIVERY ISSUE YOU'VE EVER HAD

All these myths are the perfect example of the 80% problem.

- Governance bottlenecks? **In the 80%.**
- Quantity-led QA? **In the 80%.**
- Architecture friction? **In the 80%.**
- Value alignment? **In the 80%.**
- Developer cognitive load? **In the 80%.**

We're seeing teams take the 20%, ***AI coding tools***, and assume that's the whole pie. Because, yes, AI genuinely accelerates the act of writing code. But this selective evidence creates only the illusion of progress.

The reality is that these AI tools do not approve change requests. They don't schedule penetration tests. They don't manage environment queues, compliance evidence, or release trains.

This creates a paradox: we've given developers a faster pen, but engineering stays just as slow and manual. If not more so, as faster code generation increases pressure on downstream bottlenecks.

Because AI accelerates the 20% but does nothing to fix the other 80%.

And the problem persists because gains are concentrated while the pain is distributed. The engineers who use AI tools daily experience the 20% acceleration viscerally. But the resulting bottlenecks are the governance team's problem; the compliance team's problem; the infrastructure team's problem. Someone else's problem.

A strategic organisational problem.

We're seeing time and again how the organisations positioned to capture the most value, ***truly transformative value***, from AI are those that take a holistic, strategic approach, addressing the whole infrastructure surrounding their pipelines.

AI IS NOT A CURATIVE FIX-ALL FOR EVERY DELIVERY PROBLEM YOU'VE EVER HAD. FOR ORGANISATIONS STRUGGLING WITH FOUNDATIONAL WEAKNESSES, ADDING AI TO BROKEN PROCESSES PRODUCES MORE BROKEN OUTPUT, FASTER.

[13] Source: *State of the Route to Live 2025 Report*, which found that developers in high-performing organisations spend 81% of their time on value-adding activities compared to just 51% in low-performing ones.

[14] Source: *State of the Route to Live 2025 Report*, which found that organisations with clear, communicated objectives achieved 47% higher feature adoption rates and teams with business-aligned metrics delivered 3.2x more value-generating features.

THE IMPACT OF AI

The Impact of AI

AI MOVING FAST
FROM AMPLIFIER TO ORCHESTRATOR

As teams build their AI maturity, the AI conversation is moving beyond simple code completion to how AI agents can and should interact with the entire delivery ecosystem. This section looks beyond the hype to look at the practical, real-world impact of AI on the Route to Live.

TO HARNESS AI'S TRUE POWER, ORGANISATIONS NEED A CLEAR ROADMAP TO NAVIGATE THE JOURNEY FROM SIMPLE AMPLIFIER TO FULL-SCALE ORCHESTRATOR.

Our AI Maturity Framework maps five distinct levels of AI maturity:

1 **PREDICTIVE ASSISTANCE**
Foundational AI offers pattern-based suggestions like code autocompletion, with minimal deep organisational context. Human oversight is paramount.

2 **CONTEXT-AWARE GENERATION**
AI leverages organisational context (via RAG) from codebases and documentation to generate more relevant outputs, like tailored user stories or initial test cases.

3 **ENVIRONMENT-CONNECTED INTELLIGENCE**
A critical shift where AI gains bidirectional integration with enterprise tools (Jira, GitHub, CI/CD), enabling it to retrieve data and initiate actions (like updating Jira epics or triggering test suites).

4 **AUTONOMOUS REASONING LOOPS**
Advanced AI agents reason, decompose tasks, and interact with environments to craft solutions. At this collaborative stage, AI augments human expertise (like in feasibility analyses or feature development) and humans provide crucial oversight, guidance, and validation.

5 **MULTI-AGENT ORCHESTRATION**
At the pinnacle of AI maturity, collaborative AI ecosystems of specialised agents achieve complex, end-to-end RTL objectives (such as requirements negotiation, feature implementation, and deployment) with minimal human intervention.

ROUTE TO LIVE
AI MATURITY FRAMEWORK

LEVELS OVERVIEW

LOW MATURITY

1 Predictive Assistance
Basic generative AI providing autocompletion and suggestions based on patterns, with minimal context awareness.

2 Context-Aware Generation
Generative AI with organisational context through RAG, understanding codebases, documentation, and organisational guardrails.

3 Environment-Connected Intelligence
AI systems with bidirectional integration to organisational tools (Jira, GitHub, CI/CD), able to both retrieve real-time data and make updates to these systems.

4 Autonomous Reasoning Loops
AI agents that can reason through complex problems, break down tasks, and interact with development environments to implement, test, and refine solutions.

5 Multi-Agent Orchestration
Collaborative AI systems where specialised agents work together under orchestration, delegating tasks among themselves to achieve complex objectives with minimal human intervention.

HIGH MATURITY

CROSSING THE CHASM
PASSIVE AI TO ACTIVE INTELLIGENCE

Over the past year, we've seen the industry make major maturity leaps, with many teams successfully crossing the chasm from passive AI assistance to active, environment-connected intelligence.

Levels 1 and 2, represented by tools like GitHub Copilot, are now table stakes. Code completion and generation are ubiquitous.

Mainstream adoption is now happening at **Level 3**, defined by AI that can not only generate content but also interact with and act within the development environment.

We are seeing this in practice with the rise of:

- **Self-Correcting CI/CD**: AI agents that automatically detect and fix dependency issues in pull requests.
- **Agentic Code Reviews**: AI agents that participate directly in the code review process, going beyond static analysis to provide contextual feedback.
- **Autonomous Testing**: AI agents that execute tests, analyse failures, and perform self-healing actions on test scripts directly within the pipeline.

Level 4 is the new frontier for advanced teams while **Level 5**, Multi-Agent Orchestration, remains the long-term vision. We're working with many global enterprises on landmark enterprise AI deployments that push into this bracket, laying the foundations in real-time for AI-first operations at scale.

The new chasm between **Level 3** and **Level 4** is not a gap in AI model capability: ***it is a gap in platform maturity.***

We'll explore on the next pages but an organisation's ability to answer core questions of Identity, Access, and Discovery for its non-human workers is the single deciding factor in its ability to move beyond simple co-pilots to true, orchestrated AI at scale.

WE'RE WORKING WITH ONE OF THE WORLD'S LARGEST BANKS TO UNLOCK AI-DRIVEN DATA MIGRATION, RESOLVING A £19M MIGRATION PROBLEM AND BUILDING A REPEATABLE ENGINE FOR CHANGE.

THE NEW CHALLENGES
OF THIS AI-MATURE WORLD

As AI becomes increasingly integrated into the delivery lifecycle, it's creating a new set of challenges and imperatives that leaders must address.

The Delivery Orchestration Imperative

For years, platform engineering has focused on providing infrastructure and tools: the "what." In the age of AI, this is no longer enough. To solve the paradox raised by the 80% problem, platform engineering must evolve to own the end-to-end journey: the "how."

This is the delivery orchestration imperative. It's a shift in accountability, where platform engineers are no longer just providers of a platform, but orchestrators of the entire delivery process. Their role expands from infrastructure provisioning to owning the automated, secure, and compliant path from committed code to production value.

This is how organisations begin to fix the 80%.

PLATFORM ENGINEERS ARE NO LONGER JUST PROVIDERS OF A PLATFORM, BUT ORCHESTRATORS OF THE ENTIRE DELIVERY PROCESS.

From Agent Sprawl to Governed Scale

As organisations embrace Level 3 AI maturity and beyond, the evolution from "tool sprawl" to "agent sprawl" is the primary barrier to scaling AI.

The industry is making the same mistake it made a decade ago, before the rise of the holistic platform engineering approach: buying best-in-class tools without a central plane of control to pull everything together. And the result is leaders left with an ecosystem that's not just inefficient but fundamentally ungovernable.

To move beyond single-use-case AI tools and truly solve the 80% problem, enterprises must stop building individual AI solutions and start building an agentic platform: a foundational control plane for all non-human work.

When Identity, Access, and Discovery are treated as first-class platform capabilities, the benefits are transformative.

Every new AI use case no longer needs to suffer through a full, heavyweight enterprise SDLC. And the agents themselves don't have to build or test for these critical concerns because the platform provides it by default. This is how you achieve scale, minimise risk, and create a truly governable, AI-augmented delivery ecosystem at scale.

This is happening in highly regulated environments today. At a national energy system operator, we built and deployed a multi-agent AI Workbench to over 2,000 employees in six months. By treating governance, including model observability, deterministic replay, and managed identities, as a first-class citizen, they were able to ship production agents that answered complex policy queries in seconds instead of hours.

Designing an Agentic Platform

Just as a Developer Experience Platform removes toil for human developers, an agentic platform removes the security and governance burden from every AI agent.

This is achieved by building three core capabilities directly into the platform itself:

- **Identity (The Who)**: How do you prove an agent is what it claims to be? The platform must provide a mechanism to issue, validate, and manage a verifiable identity for every non-human worker.
- **Access (The What)**: What systems, tools, and credentials can that identity use? The platform must enforce access control, ensuring agents are granted only the permissions they need, for only the time they need them. Leaving access management to individual agents or hardcoding credentials creates an unmanageable security risk.
- **Discovery (The How)**: How does an agent dynamically and securely find the resources it's allowed to access? The platform must provide a managed discovery service so that agents can interact with the correct systems without being hardcoded to specific environments.

Shadow AI and the Citizen Developer

A final, powerful effect of AI's amplification is the rise of the "citizen developer", seeing non-technical teams in marketing, finance, and operations use AI assistance to build and deploy production-facing tools and applications.

This creates a governance crisis, as the application security risk surface expands much faster than traditional, centralised security teams can manage. It's a huge shadow AI problem: ungoverned applications are built and deployed outside of established engineering and security oversight, creating a new vector for risk.

The response can't be to lock down these tools. It's to extend automated governance and golden paths, to make it easy for everyone to build safely.

THE APPLICATION SECURITY RISK SURFACE IS EXPANDING EXPONENTIALLY FASTER THAN THE SECURITY TEAMS DESIGNED TO MANAGE IT.

ORGANISATIONS MUST LEARN TO EMPOWER THE CITIZEN DEVELOPER TO BUILD AND DEPLOY SAFELY, OR RISK SNOWBALLING SHADOW AI.

ROADMAP FOR THE NEXT 12–18 MONTHS

Roadmap for the next 12–18 months

Surviving and thriving in this new landscape demands a deliberate strategy. Based on the trends and challenges we surfaced in our research, here are five high-level, actionable themes for your 12–18 month roadmap.

These five themes represent the critical strategic levers for success. While all are important, the right starting point and sequence will depend on your organisation's unique context, challenges and goals.

We encourage you to use this as a framework to build your own tailored, high-impact roadmap for the months ahead. (And as always, our team is available for a conversation if you need support).

GOVERNANCE

THEME 1

Formalise AI Strategy & Governance

Move from human-enforced policies to a platform that enforces AI governance through its core technical capabilities.

KEY GOVERNANCE FEATURES

1 **CENTRAL MODEL GATEWAY**

All calls to large language models (both internal and third-party) must be routed through a single, centrally-managed gateway. This provides a critical control point for logging all prompts and responses, enforcing content policies, managing costs, and preventing the leakage of sensitive data and intellectual property. Without a gateway, there is no observability and no control.

2 **ZERO-HOP POLICY ENFORCEMENT**

Governance decisions must be enforced at the source, not after the fact. The platform must provide a "zero-hop" policy-as-code engine that can make real-time authorisation decisions for agents. For example, before an agent is allowed to access a production API, the policy engine should be able to instantly verify its identity and credentials without relying on slow, out-of-band approval workflows.

3 **DYNAMIC TOOL DISCOVERY**

Agents should never have hardcoded access to tools or credentials. The platform must provide a dynamic, secure discovery service. When an agent needs a tool, whether it's a test environment, a deployment pipeline, or a database, it should query the discovery service, which then authenticates the agent and provides it with ephemeral, least-privilege access, strictly for the duration of the task.

WHY THIS MATTERS

The rush to adopt AI without guardrails creates new avenues for data leakage. For example, a post-breach analysis at one technology firm traced an incident to an employee's use of a non-approved AI summarisation tool, which resulted in sensitive intellectual property being sent directly to an unauthorised third party.

DEVELOPER EXPERIENCE

THEME 2

Invest in DevEx as an AI Force Multiplier

Optimise the entire development environment to reduce cognitive load and maximise the effectiveness of AI tools.

ACTIONS

1 **MEASURE AND REDUCE COGNITIVE LOAD**

Survey your developers to find the biggest points of friction in their day. Go beyond simple output metrics to understand the time they spend waiting for builds, debugging flaky tests, or navigating complex environments.

2 **BUILD A WORLD-CLASS INTERNAL DEVELOPER PLATFORM (IDP)**

Provide a seamless, self-service platform that abstracts away infrastructure complexity. A great IDP integrates AI tools in a consistent and low-friction way, making the golden path the easiest path.

3 **TREAT DEVEX AS A PRODUCT**

Appoint a product manager for your developer experience. Give them a dedicated roadmap and the authority to make improvements. Your developers are your internal customers; treat them as such.

THE CORE INSIGHT

The most effective way to counteract AI's toil-shift is to invest in world-class Developer Experience. AI amplifies everything, *including friction*. A high-friction environment with AI tools produces high-friction output faster.

QUALITY ENGINEERING

THEME 3

Evolve QA into an Engineering Discipline

Evolve quality assurance from a manual, end-of-stage gate into a continuous, automated quality engineering discipline that's owned by the entire delivery team.

ACTIONS

1 **EMPOWER ENGINEERS WITH AI-DRIVEN TESTING TOOLS**

Allocate budget for modern tools that provide autonomous test generation and self-healing capabilities and place them directly in the hands of developers.

2 **DEVELOP FULL-LIFECYCLE QUALITY CAPABILITIES**

Upskill engineers to become curators of AI-driven testing strategies. The goal is to build team-wide skills in analysing complex failure patterns and improving the quality of the system, not just finding individual bugs.

3 **MASTER NON-DETERMINISTIC VALIDATION**

Develop new skills across the team for testing and validating systems that rely on non-deterministic AI components. This is a new and critical engineering expertise.

THE CORE INSIGHT

The old model of QA as a safety net is fundamentally incompatible with the speed of AI-driven development. The focus must shift from finding bugs at the end to engineering quality in from the start.

ARCHITECTURE

THEME 4

Re-evaluate
Architectural Principles

Make conscious, deliberate architectural decisions that leverage the speed of AI without accumulating massive architectural debt.

ACTIONS

1 **STRENGTHEN ARCHITECTURAL OVERSIGHT**

Ensure experienced architects are guiding the use of AI for code generation. Do not allow AI tools to make unsupervised architectural decisions. Their role is to provide leverage, not replace judgment.

2 **EXPLORE HYBRID ARCHITECTURES**

Re-evaluate your defaults. Patterns like the "modular monolith" may be more appropriate for stable, core functionality, while microservices remain the right choice for highly scalable or experimental components.

3 **PROTOTYPE AND TEST FOR PERFORMANCE**

For services that will use AI-intensive workloads, prototype and test different architectural patterns early to understand their real-world impact on latency, cost, and performance.

THE CORE INSIGHT

AI gives you the power to build faster than ever before. This is a double-edged sword. Without strong architectural oversight, you can create a distributed monolith at an alarming rate. Re-examine your architecture strategy for an AI-native world.

BUSINESS ALIGNMENT

THEME 5

Operationalise
Business Alignment

Create the systems and incentives to ensure AI-accelerated development delivers measurable business value.

ACTIONS

1 **REDEFINE SENIOR ENGINEERING AND PRODUCT ROLES**

Explicitly shift the focus of your senior talent from tactical execution to strategic direction. Their job is to define the "why" and ensure that AI is being used to achieve it.

2 **ALIGN INCENTIVES WITH OUTCOMES**

Stop rewarding teams for story points or lines of code. Link performance evaluations to business outcomes like feature adoption, customer satisfaction, or revenue impact.

3 **IMPLEMENT VALUE STREAM MANAGEMENT**

Use modern tools to gain visibility into the entire value stream, from idea to production. Use AI to help identify bottlenecks and forecast the value of work in progress, making it easier to make data-driven decisions about what to build next.

THE CORE INSIGHT

Velocity without value is a liability. In an AI-augmented world, the most important role for human leaders is to ensure that AI's power is focused on the right business problems.

ROUTE TO LIVE
INITIATIVE PRIORITISATION MATRIX

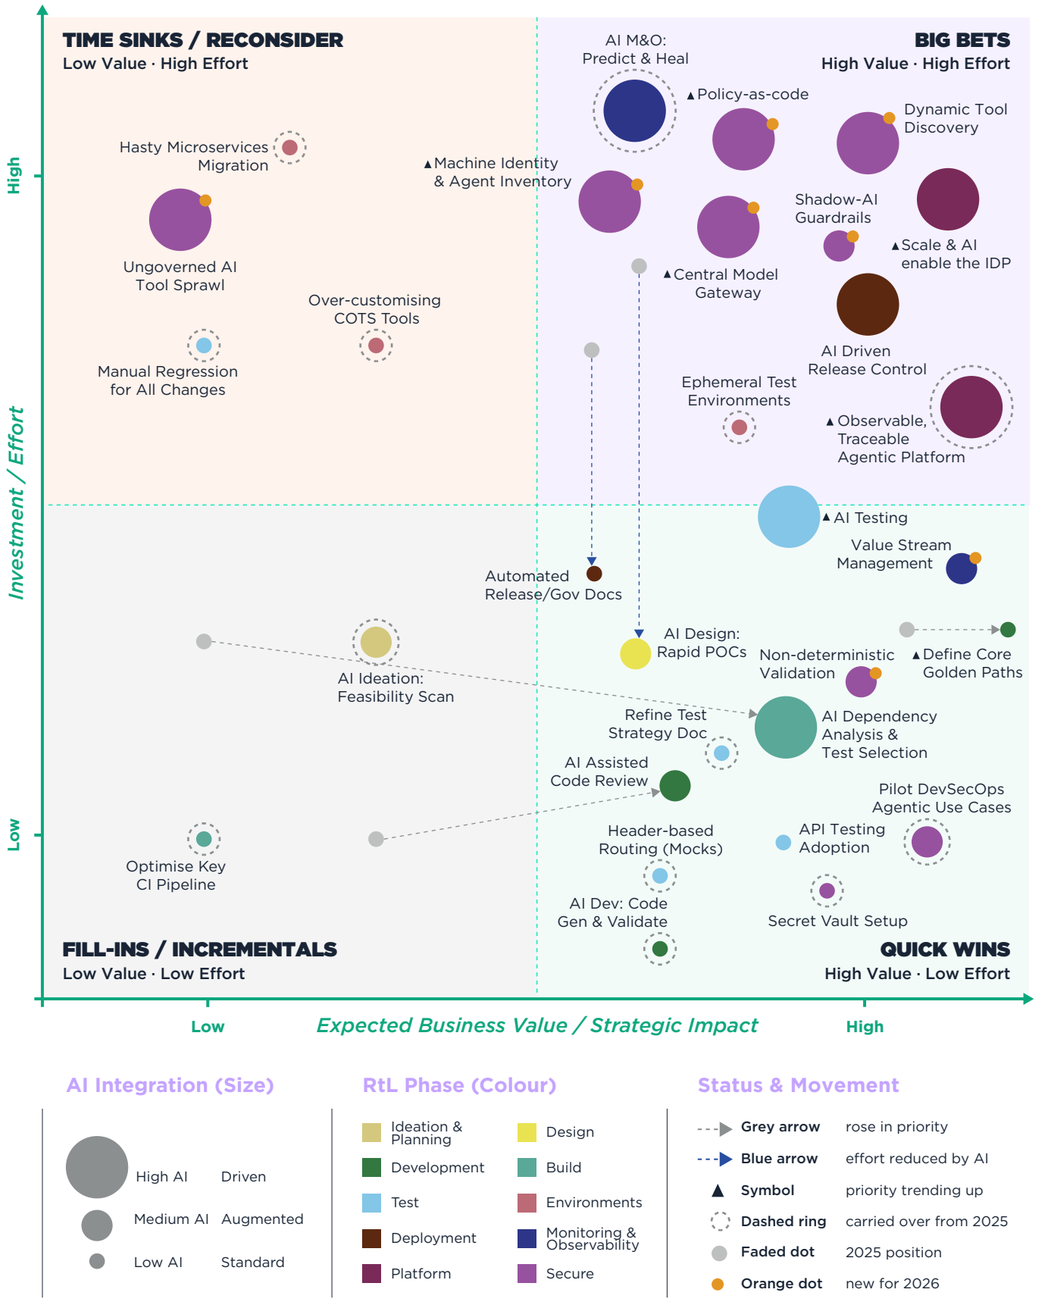

AI Integration (Size)

High AI	Driven
Medium AI	Augmented
Low AI	Standard

RtL Phase (Colour)

Ideation & Planning	Design
Development	Build
Test	Environments
Deployment	Monitoring & Observability
Platform	Secure

Status & Movement

Grey arrow	rose in priority
Blue arrow	effort reduced by AI
Symbol	priority trending up
Dashed ring	carried over from 2025
Faded dot	2025 position
Orange dot	new for 2026

WHAT WILL THE GREAT AMPLIFIER AMPLIFY?

THE CENTRAL QUESTION FOR EVERY LEADER IN 2026 IS, WHAT AI WILL AMPLIFY: THE FRICTION OF THE 80% PROBLEM, OR TRUE VALUE DELIVERY?

THE DATA PROVES THE DIFFERENCE IS NOT THE QUALITY OF THE AI TOOLS, BUT THE MATURITY OF THE DELIVERY ECOSYSTEM THEY PLUG INTO. IN ENVIRONMENTS RIDDLED WITH MANUAL PROCESSES AND LEGACY CONSTRAINTS, AI-DRIVEN SPEED SIMPLY POURS FUEL ON THE FIRE.

The Escape Route from the Paradox

FOR ELITE PERFORMERS.

For elite performers who understand this, the Harmonious High Achievers, the path to Level 4 and 5 AI maturity is clear. Agentic code reviews and self-healing pipelines are not science fiction; they are the next logical steps on a journey built on a solid foundation.

This is the essence of platform engineering: architecting a control plane of central gateways, real-time policy engines, and dynamic discovery services that solves both the 80% delivery problem and the looming agent sprawl crisis.

FOR EVERYONE ELSE.

Attempts to scale advanced AI without first solving the underlying delivery problems will only make things worse.

It's how you ensure that when your developers feel faster, your business actually becomes faster, and safer.

A generational opportunity now exists to get and stay in the fast lane. For those that hesitate, the performance divide will soon become insurmountable. But for those that act now, treating the 80% as the real battlefield and building a governable, AI-native future, the rewards will be decisive.

This is the delivery revolution. Set delivery free. Liberate the Route to Live.

AI becomes a true force multiplier, *accelerating every stage of a mature, automated pipeline*

Agentic code reviews and self-healing pipelines are the next logical step on a solid foundation

Governance, identity, and discovery built into the platform, not bolted on after the fact

The performance divide widens further, *and becomes increasingly insurmountable for those behind*

Quality Cloud Engineering (QCE) Framework

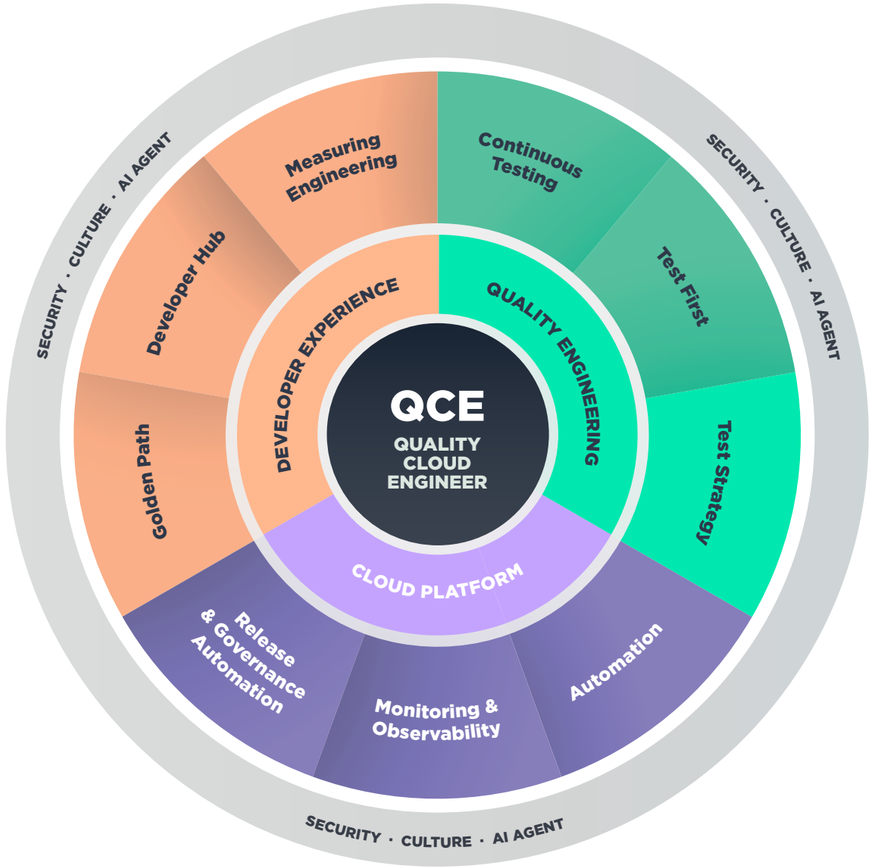

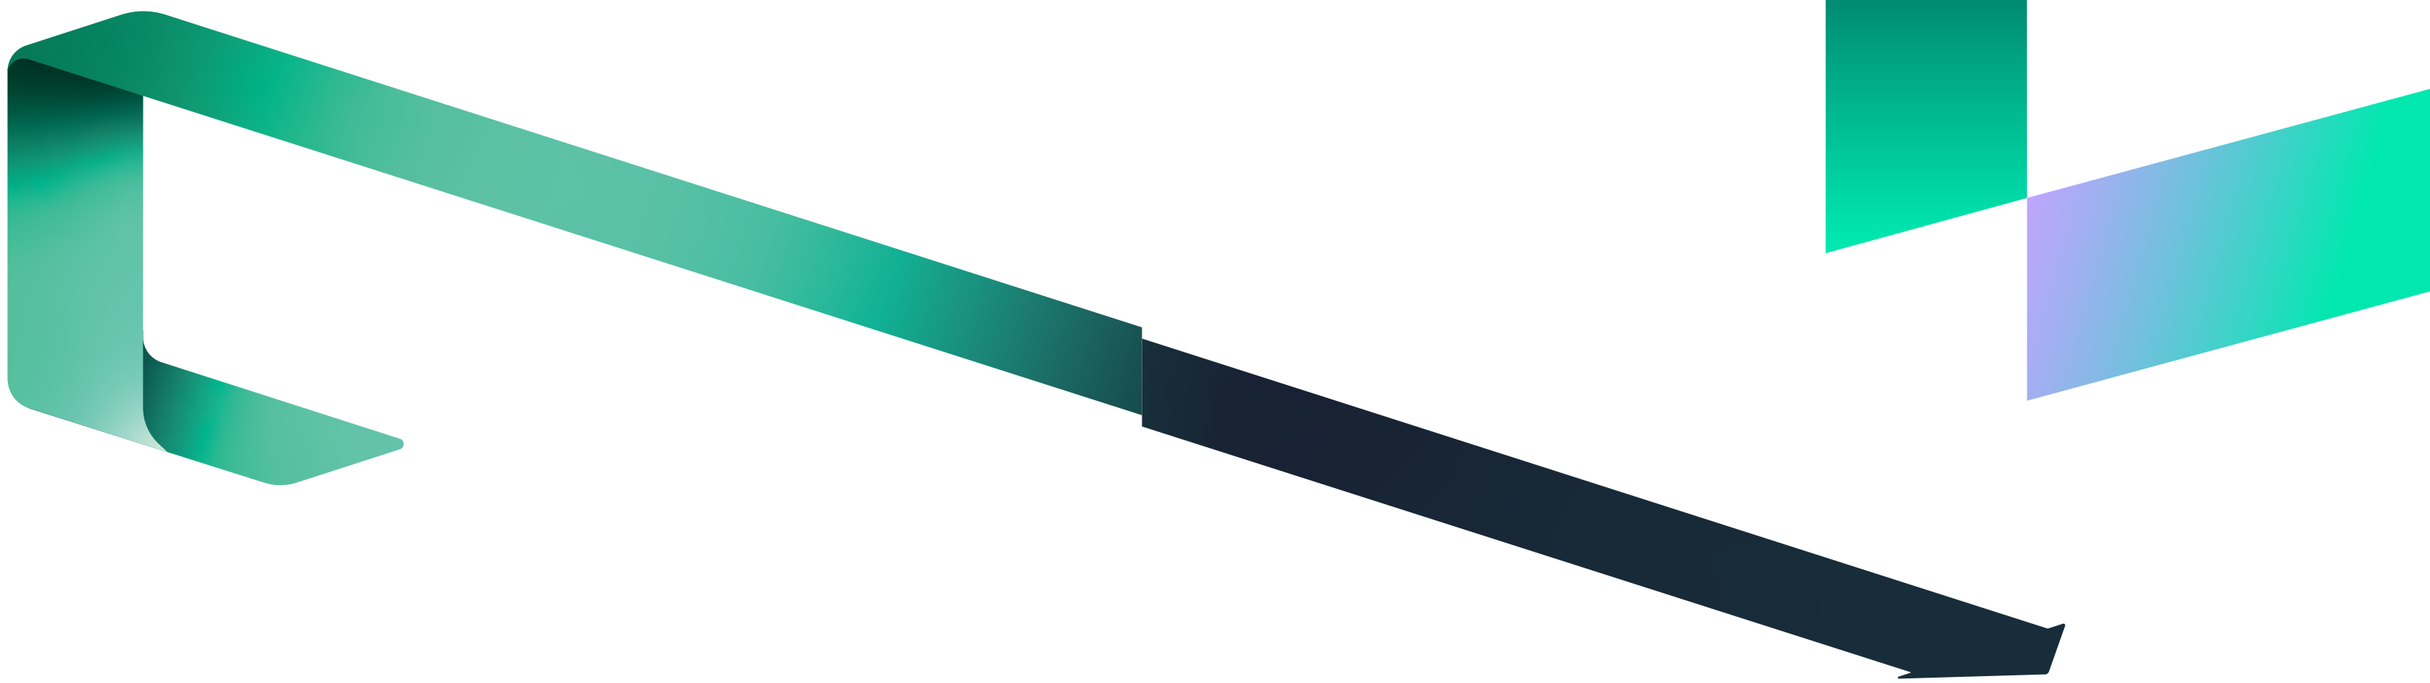

DELIVERY SHOULDN'T BE TRAPPED. NOT IN BOTTLENECKS. NOT IN SILOES. NOT IN ALMOST. WE SAY POWER TO THE ENGINEERS. POWER TO THE TEAMS WHO SHIP. WE BREAK THE BLOCKERS. YOU BUILD THE FUTURE. JOIN THE DELIVERY REVOLUTION.

LIBERATE THE ROUTE TO LIVE.

STATE OF THE ROUTE TO LIVE 2026 REPORT

ABOUT CLEARROUTE

CLEARROUTE IS A GLOBAL AI-NATIVE PLATFORM ENGINEERING CONSULTANCY THAT HELPS ENTERPRISES RELIABLY BRING HIGH-IMPACT DIGITAL PRODUCTS TO MARKET FASTER, CHEAPER, AND SAFER.

Report Authors

Sarndeep Nijjar
Chief Technology Officer
sarndeep.nijjar@clearroute.co.uk

Michael Powers
Director of Presales
michael.powers@clearroute.io

Justin Wilkin
Global Head of Engineering & AI
justin.wilkin@clearroute.io

Contributors

James Jarvis
Chief Executive Officer
james.jarvis@clearroute.io

Kouros Aliabadi
Regional Managing Director, ClearRoute UK
kouros.aliabadi@clearroute.io

Mike Hopewell
Regional Managing Director, ClearRoute US
mike.hopewell@clearroute.io

Glyn Darkin
Global Head of Delivery
glyn.darkin@clearroute.io

LONDON, UK
WeWork Farringdon
26 Hatton Garden
London EC1N 8BR

AUSTIN, USA
3300 N Interstate 35
Suite 700
TX 787705

SOFIA, BULGARIA
myFlex Megapark
115G Tsarigradsko Shose Blvd, fl 9
1784 Sofia

SYDNEY, AUSTRALIA
WeWork ANZ Tower
161 Castlereagh Street, Sydney
NSW 2000

PUNE, INDIA
WeWork Raheja Woods
8, Central Avenue, Kalyani Nagar, Pune,
MH 411006

www.clearroute.io