spacelift

THE INFRASTRUCTURE AUTOMATION REPORT 2026:

The AI Readiness Gap

Based on a survey of 406 IT decision makers and platform engineering leads with responsibility for infrastructure decisions

PUBLISHED BY SPACELIFT / JUNE 2026

Contents

Introduction 03

Framing: From automation maturity to AI readiness 05

The AI-infrastructure gap 06

The AI-governance paradox 07

Vibe coding comes to infrastructure 08

The incidents are already happening 09

The agentic cliff 10

Platform engineering, the emerging answer 11

Metrics to meet the AI moment 12

Recommendations 13

Conclusion 14

Appendix 15

About Spacelift 16

It's 2026, and AI is top of mind for everyone in the tech world.

Regardless of your vertical, industry, or role, if you're working on software, AI is affecting everything from your strategic planning down to the tactical work on your to-do list. Nowhere is this more acute than in the development and infrastructure world. AI-assisted developers are moving faster than ever, going from idea to working code within hours, and it's putting unsustainable pressure on DevOps and platform teams.

What's worse is that these teams are still grappling with the effects of the last sea change in development, CI/CD. Implementing infrastructure as code (IaC) and adopting infrastructure automation was helping DevOps keep up before AI-assisted development exploded onto the scene (as we saw in our 2025 Infrastructure Automation Report). Now they are racing to keep pace and establish comprehensive frameworks for planning, implementing, and governing AI's integration with and application to their infrastructure.

In our 2026 Infrastructure Automation Report, we surveyed you, the DevOps and platform team leaders and practitioners, to get a real sense of how AI is affecting infrastructure automation, and how you're dealing with those effects day to day. What we've learned is interesting.

This year's findings led us to create the AI Maturity Index (AIMI) which assigns organizations to one of four segments based on their AI readiness: Exposed, Fragmented, Outpacing, or Pioneer.

Exposed organizations are using AI, but without the governance or frameworks to support it safely. What they are doing diverges significantly from what they have in place to manage it. They are operating without cover, and many of them do not know it. Twenty-four percent of organizations surveyed are categorized as Exposed.

AIMI Segments

Exposed

Using AI, but without governance or frameworks to support it safely

Fragmented

Using AI, but inconsistently and without standardization

Outpacing

Moving fast with AI, adopting aggressively, but governance hasn't kept up with that momentum

Pioneer

High AI adoption rates, but they are ahead because they built governance and automation before AI arrived

Fragmented organizations are somewhere in the middle. They use AI but not consistently or systematically. Some teams are doing things well, others are not, and governance is uneven.

They are aware of what good looks like but have not standardized it yet. Thirty-two percent of organizations surveyed fall into the Fragmented segment.

Outpacing organizations are moving fast with AI. They are experimenting heavily and adopting aggressively, but their governance has not caught up to their momentum. They are the ones using AI to generate software and IaC code most often but are also exposed to the risks of AI because speed without guardrails creates risk whether you see it or not. Twenty-five percent of organizations surveyed fall into the Outpacing segment.

Pioneer organizations built governance and automation before AI arrived, enabling them to use it effectively. They have formal policies, automated controls, and the structural discipline to handle whatever challenges AI-accelerated development throws at them. Nineteen percent of organizations surveyed fall into the Pioneer segment.

An organization's placement in these segments is behavioral rather than firmographic. A large enterprise with significant infrastructure spend can land in the Fragmented segment. A mid-size company can be a Pioneer. What separates them is not budget or headcount, but what they do with the tools they have.

What's interesting is that AI readiness and automation maturity of an infrastructure team are not always aligned. An organization can be highly automated but still be unprepared for AI. And the data from this year's report shows that many are.

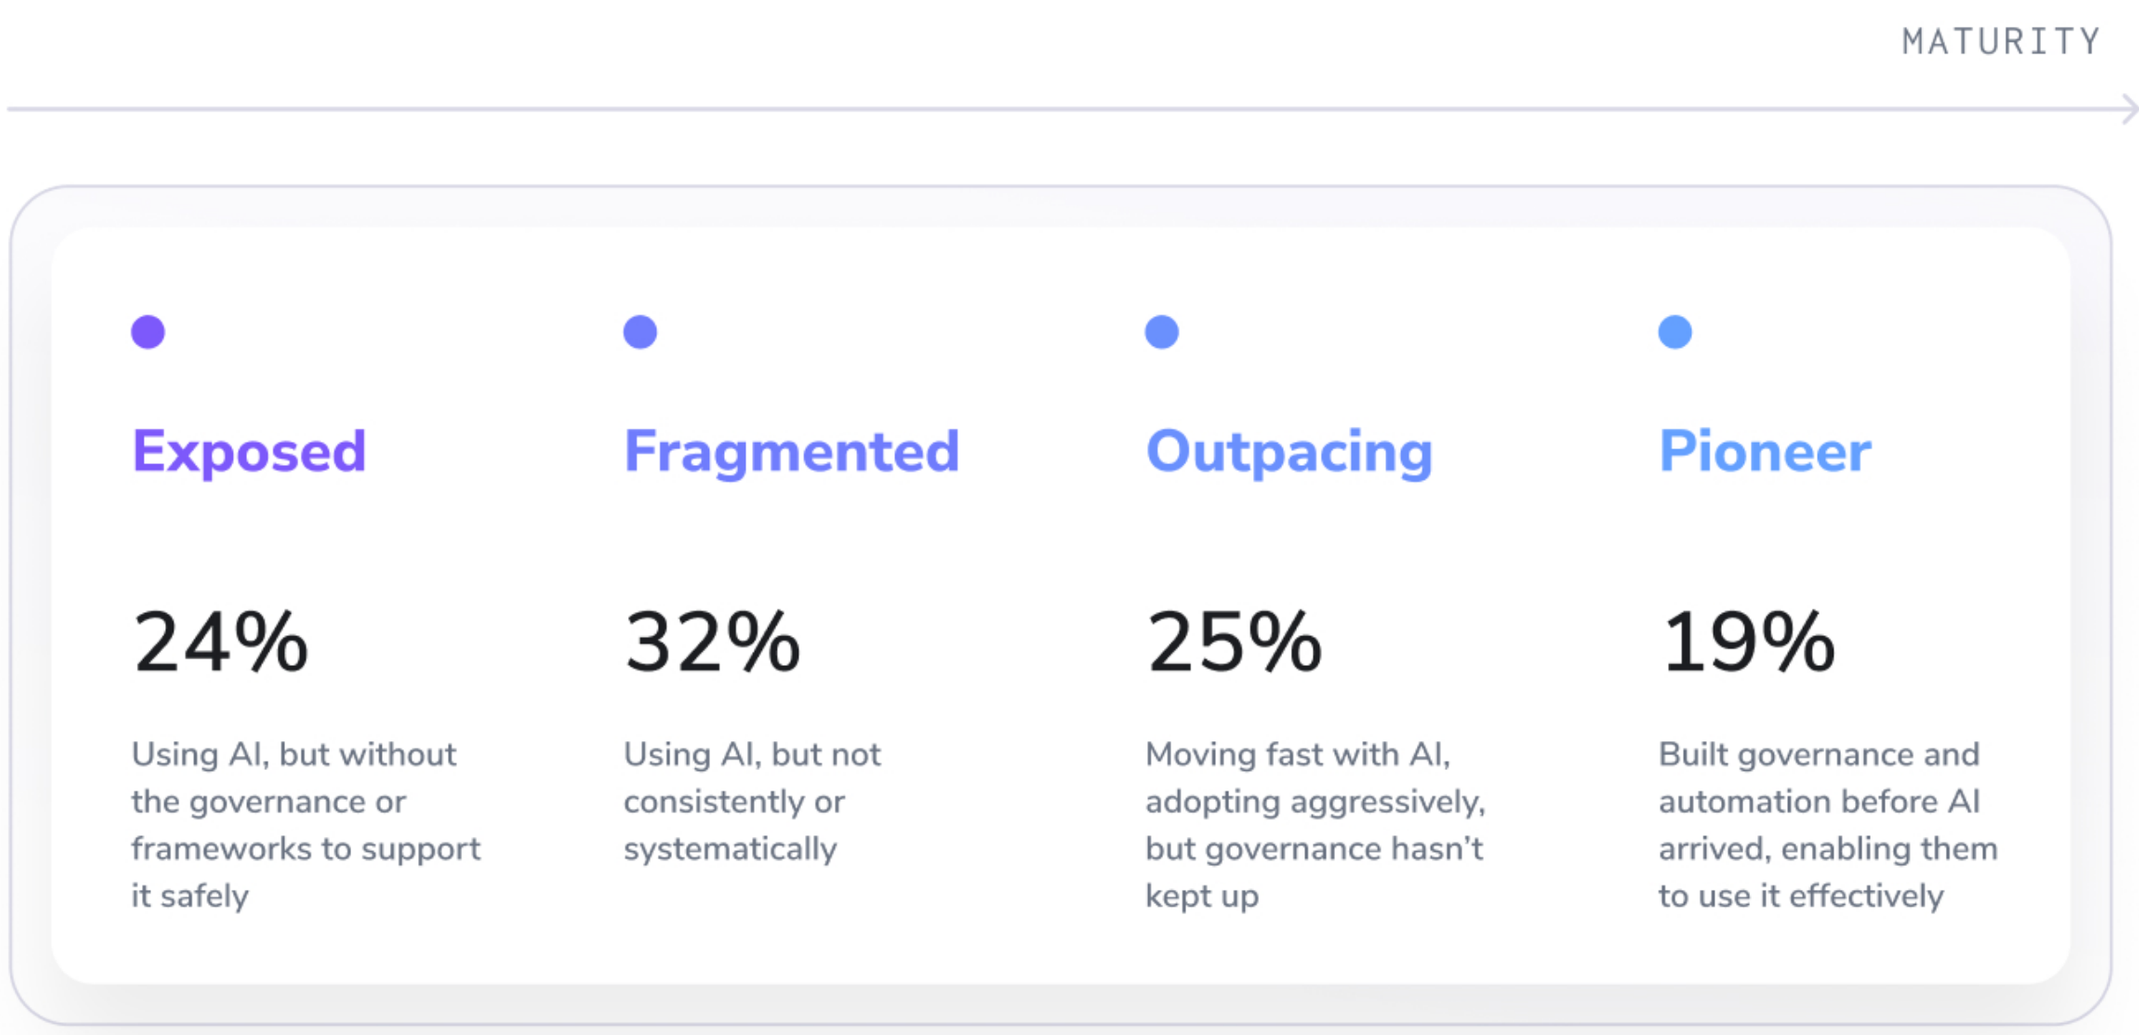

[1] The report is based on a survey conducted by Panterra Group in April 2026 among n=406 IT decision makers and platform engineering leads with responsibility for infrastructure decisions. All respondents were based in North America.
[2] The AIMI is explained in greater detail in the Appendix.

FRAMING

From automation maturity to AI readiness

Last year's report measured where organizations stood on the infrastructure automation maturity curve. This year, we needed a different lens. Automation maturity tells you how much of your infrastructure lifecycle is codified and automated. AI readiness tells you whether your infrastructure can absorb AI-generated code safely, govern it consistently, and scale it without the wheels coming off.

You can have a fully automated infrastructure pipeline and still be completely unprepared for what happens when AI starts writing the Terraform, policies, and compliance checks that flow through it. Automation is a prerequisite for AI readiness, but it is not a substitute for it.

The AIMI measures readiness across five dimensions: how deeply AI is integrated into your workflows, how mature your governance controls are, how automated your infrastructure processes are, how much risk exposure you are carrying from AI-generated changes, and how ready your infrastructure automation layer is to enforce controls at the speed AI demands. Together, those five dimensions determine which segment an organization falls into.

What makes this segmentation useful is that it does not care about the size of your company or the size of your infrastructure budget. It cares about behavior: how much of your infrastructure is codified, how consistently governance is applied, whether your automation is encouraging collaboration across teams or fracturing it. Those are the things that predict outcomes and separate Pioneer organizations from Exposed ones.

> "I expect the role of AI to increase significantly over the next two to three years because AI is the future and early adoption means safety and continuity for your business infrastructure."
>
> HEAD OF IT INFRASTRUCTURE

[1] https://spacelift.io/infrastructure-automation-survey

The AI-infrastructure gap

Here is the core tension of this report: AI is making developers faster, but all of that AI-generated code has to land somewhere, and it's landing on infrastructure teams.

Eighty-nine percent of organizations say developer velocity has increased because of AI, and 82% report that between 25% and 74% of their code is now AI-assisted. AI-assisted development is now a mainstream practice, and 25% of organizations call it a top strategic priority. The problem is that all of this code has to be deployed, governed, and maintained by DevOps and platform teams that aren't reaping the same acceleration benefits of AI.

Given that 62% of respondents report that development is ahead of infrastructure in terms of AI adoption, it is hardly surprising that 86% say AI has increased demands on infrastructure teams.

The downstream effects are already evident: 40% say security vulnerabilities are appearing faster, 40% say governance is getting harder, 37% report higher change volume, 35% report increased pipeline strain, and 35% report growing infrastructure drift. These are no longer forward-looking concerns but issues infrastructure teams are dealing with today as a direct consequence of AI-accelerated development. This is the beginning of the AI-infrastructure gap.

What happens when AI-generated code lands on infrastructure teams

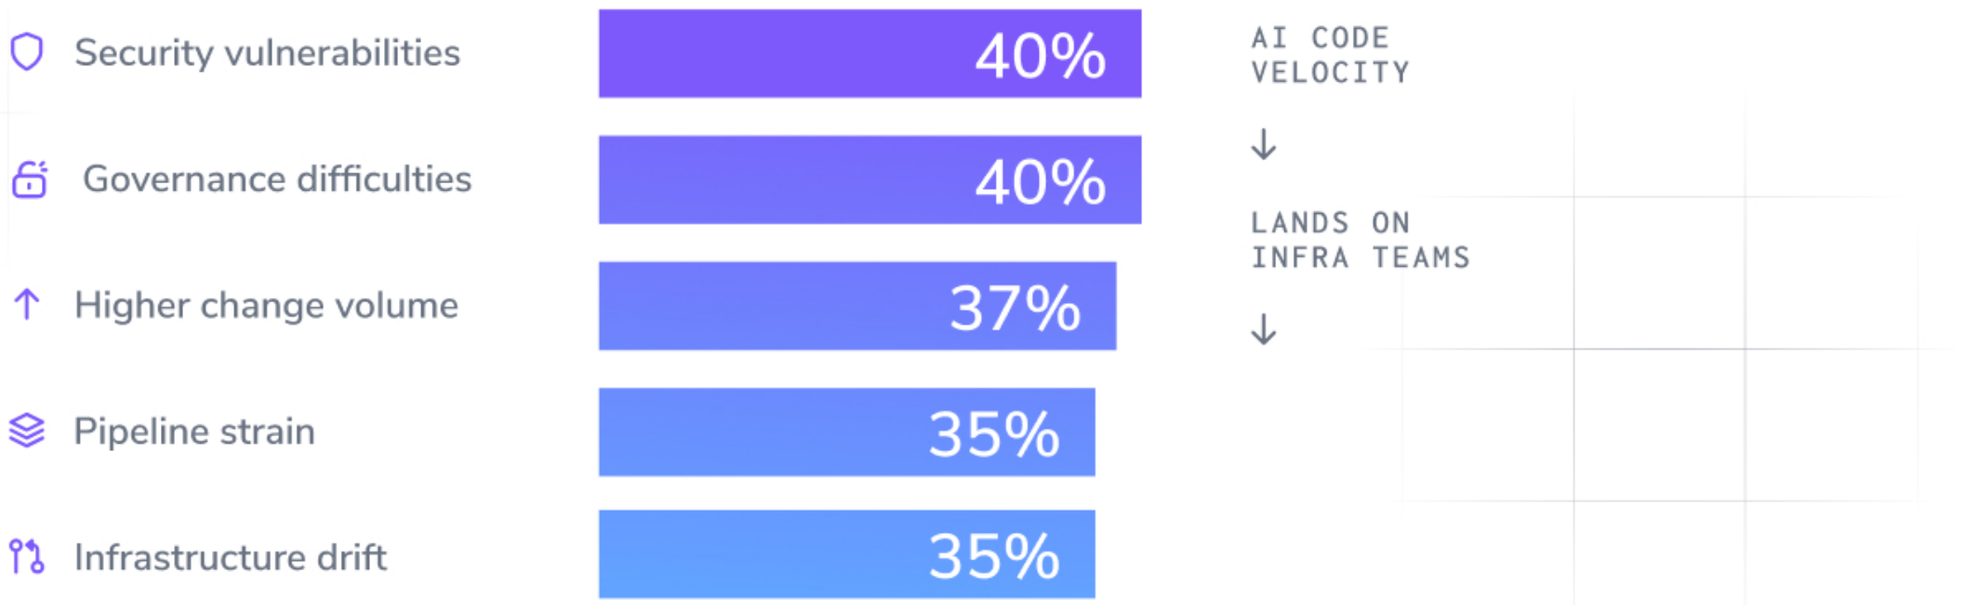

Closing the gap requires two things. The first is automation. Organizations that have deeply automated their infrastructure processes are better able to absorb the increased volume. Organizations that have not are drowning. Among Pioneer organizations, 92% describe their processes as fully or mostly automated, compared to just 14% of Exposed organizations.

The second is governance. Even well-automated infrastructure cannot safely absorb AI-generated code without formal policies, approval workflows, and automated controls. Automation without governance just creates a faster path for deploying risk into your production environment.

Pioneer organizations are closing the AI Infrastructure gap because they have focused on automation and governance equally. They've built automation deep enough to handle the volume and governance strong enough to validate what flows through it. Fragmented and Exposed organizations are falling behind because they're missing one or both, and every month that AI-generated code volume increases, the gap gets wider.

However, the survey responses reveal a unique cognitive dissonance around governance and AI that is having a drastic effect on organizations' AI readiness.

The AI-governance paradox

When we asked infrastructure leaders whether they are confident in their organization's ability to govern AI, 86% said yes. When we asked whether they have a formal AI governance policy in place, only 30% do. There is a gulf between confidence and action that doesn't bode well for AI governance.

This is the governance paradox: Many teams believe they are governing AI effectively, but they lack the supporting strategy, policies, and controls.

However, Dominant organizations have avoided the governance paradox. Seventy-one percent have a formal governance policy that they actively enforce, and when you ask them about their AI governance concerns, 24% say they have none because risks become manageable when controls are in place. The risks do not disappear; they just become something they are aware of and can work with. Among Exposed organizations, only 4% have a formal governance policy, yet 70% are confident in their governance capabilities.

> "I intend to automate governance and compliance audits in multi-cloud systems using AI."
>
> VP / DIRECTOR / HEAD OF PLATFORM ENGINEERING

Vibe coding comes to infrastructure

Considering the AI infrastructure gap and the paradox of AI governance, you'd think vibe coding would be limited to the development team. Should this thing we don't completely trust and don't have policies for be trusted to generate the Terraform code that runs our infrastructure? Shouldn't it at least be heavily reviewed? Unfortunately, the data says otherwise.

When we asked whether teams use AI to generate developer code, infrastructure as code, and policy as code without thorough review, the rates were nearly identical: 79% for developer code, 78% for IaC and HCL, and 78% for policy as code. Vibe coding has fully penetrated the infrastructure and governance layers.

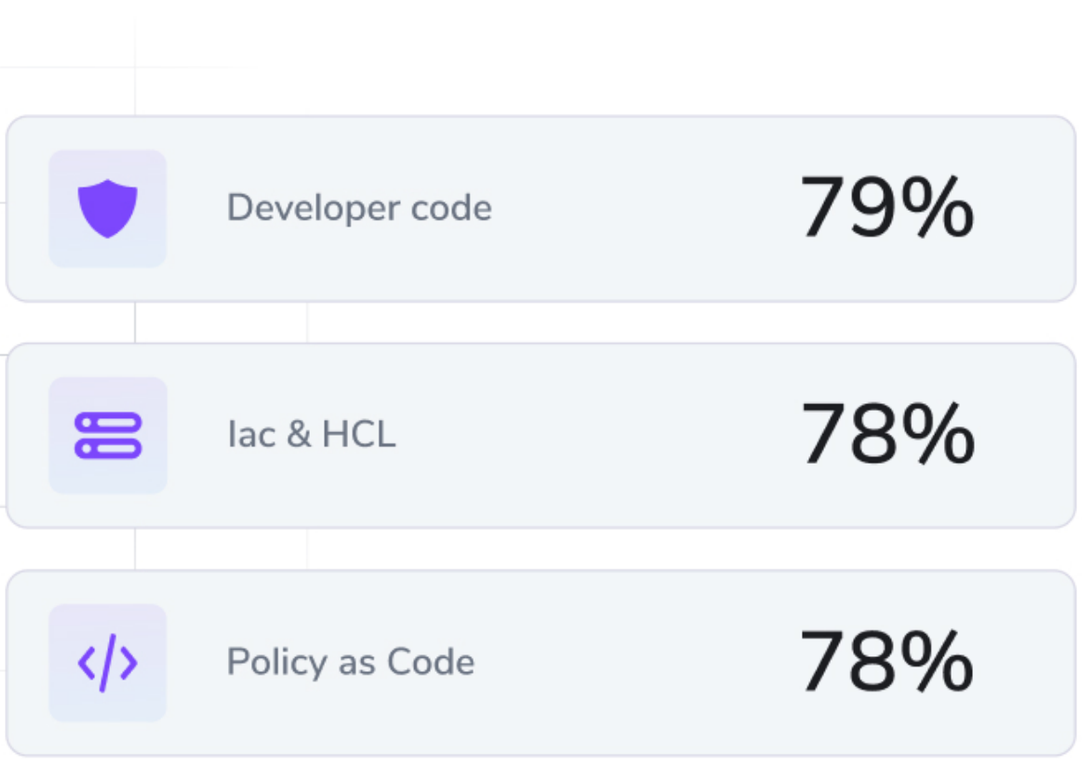

Last year, we found a gap between organizations' self-perceived infrastructure automation maturity and their actual execution. This year, we're seeing a similar gap with governance and trust. They're saying that infrastructure needs to meet a higher standard of governance, and they're acutely aware of the AI-caused issues, but they're letting AI generate Terraform code and shipping it. In our survey, 33% of infrastructure teams say they would apply AI-generated HCL directly to production without any review at all. A further 43% say they would apply it with only minimal review. Together, that is 76% who would ship AI-generated infrastructure code with little or no scrutiny.

This matters because infrastructure code is not application code. When a vibe-coded application function misbehaves, you get a bug ticket. When a vibe-coded infrastructure policy misbehaves, you get a resource misconfiguration that reaches production.

This means an unplanned incident response, a scramble to identify what changed and why the policy didn't catch it, and a remediation cycle that pulls your team off the work they were supposed to be doing. Multiply that by the volume of AI-generated IaC flowing through your pipelines every week, and you are looking at a team that spends more time firefighting AI-generated problems than building the tooling for which they are responsible. The stakes are categorically different.

That said, ignoring AI is not the answer. Pioneer organizations vibe-code IaC at a higher rate than Exposed ones: 86% versus 69%. The difference is that they do it inside governed pipelines. AI-generated code plus automated governance is what makes the practice sustainable and increases deployment velocity to keep pace with AI-assisted development.

The incidents are already happening

We're already seeing the results of letting AI-generated Terraform go to production unchecked. Ninety-three percent of organizations have experienced at least one AI-caused infrastructure incident. Over the past 12 months, DevOps and platform leaders report rework from AI-generated changes (37%), security misconfigurations that reached production (36%), compliance violations (36%), infrastructure drift attributable to AI changes (35%), and incidents caused by agentic systems (33%).

Incidents types reported — last 12 months

Rework from AI-generated changes	37%
Security misconfigurations that reached production	36%
Compliance violations	36%
Infrastructure drift attributable to AI changes	35%
Incidents caused by agentic systems	33%

Who experiences them

EXPOSED ORGANIZATIONS

97%

report at least one AI-caused infrastructure incident

PIONEER ORGANIZATIONS

17%

report no AI-caused infrastructure incidents

Exposed organizations are the most prone to these issues: Only 3% report having experienced none of these incident types. Contrast this with Pioneer organizations, 17% of which report none. The difference between the two segments is governance. The organizations with formal controls experience fewer incidents because failures and misconfigurations are caught before they reach production. The reason is straightforward: automated validation. Manual review might catch issues, but it is slower, less consistent, and does not scale with the volume that AI-generated code produces.

It is worth understanding the compounding effect here: Without governance, manual review processes accumulate. More manual work requires more engineering capacity. Teams start to feel like the automation they've put in place is not saving them time, and in a narrow sense, they are right, because the ungoverned workflows are generating overhead. These infrastructure team workload issues ultimately become incidents that require remediation.

The agentic cliff

It's no surprise that most of the infrastructure community is trying to figure out how to harness the power of AI. However, several concerning factors in our findings suggest a chaotic future: the gap between AI-accelerated development and platform teams' ability to manage it, the governance paradox, the willingness to deploy unreviewed infrastructure code for the sake of speed, and the resulting incidents when this fails. Our survey data shows that 89% of organizations plan to adopt agentic AI for infrastructure. Twenty-four percent plan to do so within six months.

Agentic AI is faster, more complex, and harder to govern than AI-assisted development. When a developer uses AI to generate a Terraform module and then reviews it before applying, there is a human in the loop. When an agentic system makes infrastructure decisions autonomously, there is no platform engineer to serve as a checkpoint. Instead, the controls have to be embedded in the workflow itself because there is no manual review stage to catch what the system misses.

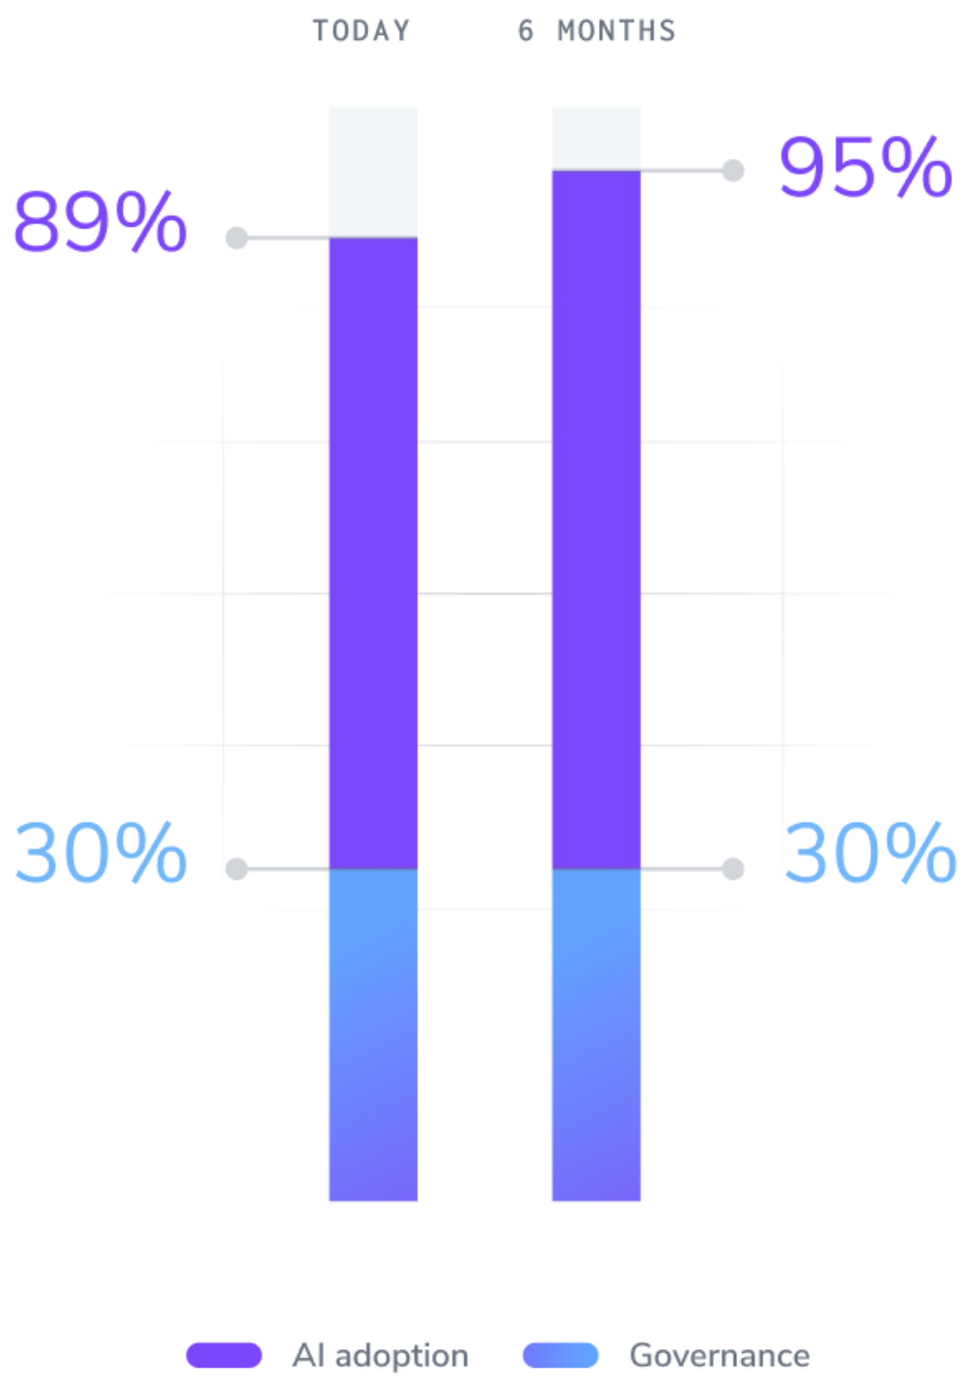

The agentic incident rate among early adopters is already running at 33%.

The organizations that build governance infrastructure before expanding into agentic use cases are in a fundamentally better position than those that try to bolt it on after their first agentic incident. Retrofitting governance is always harder and more expensive than building it in advance. The Pioneer organizations already know this. That is why they are Pioneer organizations.

"AI is going to be at the front lines of detection and security alerts so humans can have the time and ability to creatively solve the problems."

DIRECTOR OF DEVOPS

Platform engineering, the emerging answer

The reason governance and platform engineering keep showing up in the same conversations is that they are converging.

Sixty percent of infrastructure leaders say their infrastructure automation tool or vendor determines their ability to govern AI-generated infrastructure. The infrastructure automation layer is not just a delivery mechanism anymore. It has become the governance layer where policies get enforced, approval workflows run, drift gets detected, and the blast radius gets controlled. If your infrastructure automation pipelines don't have those built in, you are building governance by hand, which does not scale.

Eighty-three percent of organizations say they are considering a shift to platform engineering. Among Pioneer organizations, 25% have already made that shift. Among Exposed organizations, only 2% have. The gap between intending to adopt platform engineering and actually doing it is one of the defining patterns in this data, and it aligns with the gap between organizations that are governing AI effectively and those that are not (71% for Pioneer versus 4% for Exposed).

This is not a coincidence. Platform engineering is the structural answer to the governance problem. When developers have a governed, self-service path to production that is faster than working around it, they use it. When the governed path is slower or harder than the ungoverned one, they go around it. Platform teams exist to make governance the path of least resistance, and organizations that have figured this out are the ones showing up in the Pioneer segment.

Cross-team collaboration data backs this up. Sixty-eight percent of Pioneer organizations say automation built by the platform engineering team supports collaboration among engineering, platform, and security teams effectively. For Exposed organizations, collaboration is rated as "somewhat effective" or neutral. When platform engineering builds governance at scale it doesn't just give organizations strong controls, it creates a collaboration structure.

Metrics to meet the AI moment

Most organizations are measuring infrastructure the way they did before AI arrived: team productivity (26%), deployment frequency (23%), and security incidents (23%). These are reasonable baselines, but they're not sufficient.

The problem is not that organizations are measuring the wrong things, it's that they are not measuring the new things. When AI starts generating infrastructure code, you need to know how that code is performing relative to human-authored code. You need to know the error rate of AI-generated changes specifically. You need to know how much AI-generated IaC is moving through your pipelines. You need to know whether drift is increasing because of AI-generated changes.

Without those signals, you cannot tell whether your governance is working. You can only observe the consequences after the fact.

Very few organizations track these AI-specific metrics today. The volume of AI-generated IaC moving through pipelines is tracked by only 15%. Error rates of AI-generated changes are tracked by only 20%. These are the signals that would tell you whether governance is doing its job, and almost nobody is collecting them. You optimize what you measure. If you are not measuring AI-specific signals, you are optimizing in the dark. Governance without measurement is hope.

WHAT MOST ORGANIZATIONS MEASURE TODAY		AI-SPECIFIC SIGNALS ALMOST NOBODY TRACKS	
Team productivity	26%	Volume of AI-generated IaC in pipelines	15% track
Deployment frequency	23%	Error rates of AI-generated changes	15% track
Security incidents	23%	Drift attributable to AI changes	—

> "AI in infrastructure is expected to move from assisting with individual tasks to more integrated pipeline automation - provisioning, compliance checks, and remediation tasks."
>
> HEAD OF IT OPERATIONS

RECOMMENDATIONS

Five things to do now

Adopting platform engineering is the key strategic change required for any IT organization seeking to achieve AI readiness. Combine this with the following five tactical steps to reach the Pioneer segment of the AI Maturity index:

1

Make IaC coverage your first priority.

You cannot govern what is not codified. Pioneer organizations have 75%+ IaC coverage. Exposed organizations are at 15%. Every other recommendation that follows depends on closing this gap first. If your infrastructure is not in code, none of the governance controls in the world can help you because there is nothing for them to attach to.

2

Build governance before you need it.

The policy gap between Pioneer and Exposed organizations is the single most important structural difference in this report: 71% versus 4%. Remember, 38% of organizations using governance-capable infrastructure automation tooling never turned governance on. They had the capability, but they did not adopt it. Use it.

3

Treat shadow IaC as a governance emergency.

At 90% of organizations, AI-generated infrastructure code is being produced outside governed IaC orchestration workflows, bypassing the controls that platform teams have put in place. Rather than mandating compliance with a slower process, the answer is creating governed paths that are faster and easier than the ungoverned alternatives. If the governed path is harder, people will go around it. Platform engineering exists to make the right thing the easy thing.

4

Measure AI outputs, not just infrastructure outputs.

Error rates of AI-generated changes, drift attributable to AI, and volume of AI-generated IaC moving through your pipelines are the signals that indicate whether governance is working. The highest-performing organizations track them, whereas the lowest-performing organizations do not. Start tracking them.

5

Plan for agentic governance now.

The controls that are adequate for AI-assisted development, where a human reviews every change before it is applied, are not sufficient for autonomous systems. Policy enforcement is the mechanism that catches misconfigurations before they become incidents, and it is the only mechanism for agentic workflows. Forty-three percent of Pioneer organizations planning agentic adoption are doing so with governance infrastructure already in place. If you are planning agentic adoption without that foundation, you need to build it before the first agentic workflow goes live, not as a remediation step after the first agentic incident.

CONCLUSION

The window is closing

The performance gap highlighted in this report opened years before AI arrived, in decisions about automation investment, governance discipline, and platform engineering adoption. **AI accelerated the consequences of those decisions.**

The organizations outperforming on every metrtic are the ones that built a governance framework and actually use it. The most cautionary pattern in this data is the organization that is highly active, running infrastructure at scale, deploying frequently, and doing it all with almost no governance. At 90% of organizations, AI-generated infrastructure code is being produced outside governed workflows. Ninety-three percent have experienced at least one AI-caused incident. Using AI heavily to generate application and infrastructure code is not the same as governing AI well, and the pace of AI-accelerated development is making the difference between them impossible to ignore.

Agentic AI will raise the stakes further. Infrastructure decisions made autonomously by AI systems eliminate the human review step that currently catches governance gaps. You're no longer able to catch issues and remediate on the fly, you're finding out about issues after they've caused a problem in production. The blast radius of an ungoverned agentic action is larger than the blast radius of an ungoverned human one.

The infrastructure automation layer is now inseparable from governance capability. Approval workflows, automated validation, drift detection calibrated to AI-generated change velocity, blast-radius controls, and audit trails are table stakes for operating safely in this environment. The organizations that treat them as such will be positioned to move quickly when agentic AI matures. The rest will be managing incidents.

> “I expect AI to move beyond simple automation. I anticipate AI agentic systems will be able to manage selfhealing and optimized resource allocation.”
>
> INFRASTRUCTURE AUTOMATION DIRECTOR

APPENDIX

RESEARCH METHODOLOGY

The survey included 406 respondents sourced from a leading global online panel provider. They were selected from the panel based on geographic quotas, as well as screening questions based on role in economic or financial research, company size, and how long they have been in their analyst role. All participants were IT decision-makers, platform engineering leaders, and DevOps professionals responsible for infrastructure decisions. Selected respondents were further screened based on their self-reported knowledge of AI and infrastructure, as well as their attentiveness to survey questions.

ROLE QUOTAS and AIMI SEGMENTATION

Respondents were asked to select the role from a list of 27 options that most closely reflected their primary responsibility, even if none were an exact match or if they held multiple roles. Based on their stage of AI adoption, respondents were classified into four segments: Pioneer (19%), Outpacing (25%), Fragmented (32%), and Exposed (24%). The sample was intentionally designed to capture a range of perspectives .

GEOGRAPHIC QUOTAS

This survey was limited to respondents from North America, and responses from other regions were excluded during screening.

RESPONDENT SCREENS

Role: All respondents were required to indicate that they were responsible for or had influence in infrastructure and automation strategy.

Company size: All respondents must self-report that their companies have a minimum of 250 employees. All potential respondents from smaller companies were excluded. In total, the survey includes 24% of respondents from companies with 250 to 499 employees, 30% from companies with 500 to 999 employees, 32% from companies with 1,000 to 4,999 employees, 14% from companies with 5,000 or more employees.

Information level: In our experience, it is possible to have “qualifying respondents” who nevertheless prove to have too little information or knowledge about the space to provide useful data from which to draw insights. We therefore apply an “information” screen to respondents as well. Specifically, we ask whether or not respondents could explain certain terms to their colleagues if asked to do so. In order to qualify for this survey, a respondent must indicate they can explain the term "Infrastructure as Code (IaC)".

“Attention” level: It is easy for respondents to speed through surveys or not pay enough attention to provide useful data. We make an effort to exclude these respondents as well, as they provide generally less useful data. In this survey, respondents were screened out for “attention” reasons if they said they could explain the made-up term “ Greenfield as a Service (GaaS) ” to a colleague in the same question used for the Information Screen noted above.

RESPONDENT SCREENS

It is technically impossible and improper to list a margin of error for a survey of this type. The respondents for this sample were drawn from an online panel with an unknown relationship to the total universe, about which we also do not know the true demographics. As such, the exact representativeness of this, or any similarly produced sample, is unknown.

AIMI SEGMENTATION

The AI Maturity Index (AIMI) places organizations into one of four segments based on scored responses across five dimensions. Segment placement is behavioral, not demographic. Company size, revenue, and infrastructure spend do not predict which segment an organization falls into. What matters is how teams operate.

AI integration depth measures how broadly and deeply AI is embedded across development and infrastructure workflows, as in whether AI is used experimentally by a few individuals, standardized across teams, or fully integrated into the deployment pipeline. Organizations scoring high on this dimension are using AI not just to generate code but to make infrastructure decisions at speed and scale.

Governance maturity measures whether formal controls are in place to govern AI-generated outputs before they reach production. This includes the presence of a formal AI governance policy, the use of approval workflows for AI-generated infrastructure changes, automated testing and validation of AI-generated IaC, and human-in-the-loop controls for critical changes. Organizations scoring low on this dimension may be actively using AI without any formal mechanism to validate what it produces.

Infrastructure automation maturity measures the breadth and sophistication of automated infrastructure processes, such as how much of the infrastructure lifecycle is managed through automated workflows rather than manual intervention. This dimension captures not just whether automation exists but whether it is deep enough and consistent enough to absorb the volume that AI-accelerated development produces.

Risk exposure measures the incidence and severity of AI-caused negative outcomes. This includes security misconfigurations that reached production, compliance violations, infrastructure drift attributable to AI-generated changes, rework required from poor-quality AI outputs, and incidents caused by agentic systems. Organizations scoring high on this dimension have experienced fewer of these outcomes, either because their governance controls prevented them or because their measurement infrastructure detected and resolved them early.

Platform readiness measures the organizational structure and tooling that govern how infrastructure changes move from development to production. This includes whether the organization has adopted a platform engineering model, whether developers have access to governed self-service pathways, and whether the IaC management platform in use natively supports the controls needed for AI-governed infrastructure operations.

Each dimension is scored independently based on survey responses. Together, the five scores determine segment placement. The four resulting segments and their share of the survey population are described in the Framing section of this report.

ABOUT SPACELIFT

Spacelift is the infrastructure orchestration platform built for the AI-accelerated software era. It manages the full lifecycle of both traditional IaC and AI-provisioned infrastructure, integrating with tools like Terraform, OpenTofu, CloudFormation, Pulumi, and Ansible in a single governed workflow.

Spacelift Intelligence adds an AI-powered layer for natural language provisioning, diagnostics, and operational insight across traditional and AI-driven workflows, helping platform teams move faster while keeping infrastructure secure and compliant at scale.

Learn more about the Spacelift platform and how it can help you overcome your infrastructure challenges at Spacelift.io.


spacelift