Teleport

The 2026

Infrastructure Identity Survey: State of AI Adoption

Contents

Teleport

PART 1

A New Actor: AI Identity in Infrastructure

How AI has moved from tool to teammate, widening the gap between adoption and security

A New Actor: AI Identity in Infrastructure

How AI moved from tool to teammate, widening the gap between adoption and security

Over the last year, enterprise infrastructure has been undergoing foundational change as the pace of AI adoption accelerates. AI in production infrastructure brings transformational promise: Autonomous agents can execute code, apply policies, retrieve sensitive information, and make decisions rapidly on their own — continuously and without checking in with human operators.

However, this enthusiasm around the promise of AI has resulted in innovation that is outpacing security. AI is now being deployed in production environments before needed security models, governance, and controls are fully in place. As a result, infrastructure and security leaders — already accountable for reliability and resilience — are finding themselves responsible for securing systems that are advancing faster than the protections designed to contain their risk.

To understand how this shift is affecting infrastructure leaders on the front lines, we commissioned Eleven Market Research in December 2025 to conduct an in-depth study. The research was based on structured telephone interviews with 205 infrastructure security leaders, including CISOs, VPs of Security, Security Architects, and Platform Engineers. Participants represented organizations ranging in size from 500 to more than 10,000 employees, across seven industries. These are leaders who define security policy, approve platforms, and ultimately, bear responsibility when systems fail.

The findings confirmed the growing unease around managing the balance between innovation and security. They also revealed something unexpected: a single factor that predicted security outcomes more accurately than industry, organizational maturity, or even confidence in AI deployment.

DEFINITION

AI in Infrastructure

AI-powered workloads. Agentic systems. Machine-to-machine communication. ChatOps. Compliance automation. Incident detection.

Not just tools. A new class of actor in your environment.

The State of Adoption

92% of companies surveyed are not simply experimenting; they've moved past the pilot phase into full-scale operational deployment. What is most relevant about this finding is that it is not a future goal or "aspirational" target but rather a current reality with a high degree of implementation across organizations.

KEY FINDING

92%

have near-term AI initiatives in infrastructure

Where Are You on the AI Adoption Curve?

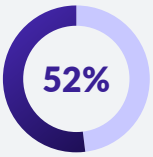

somewhat or extremely mature in AI adoption

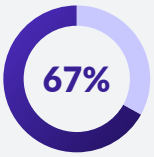

confident in their AI deployment

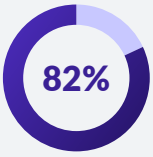

have deployed SaaS-managed AI

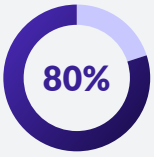

changed how they publish/ consume data

The use cases span productivity through security: from ChatOps interfaces that query internal knowledge bases to AI systems that detect incidents and automate compliance workflows.

Top AI Use Cases in Infrastructure

The Value Is Real

AI is delivering measurable results. Two-thirds report increased workplace efficiency, with specific gains that matter to infrastructure teams:

How Has AI Affected Performance?

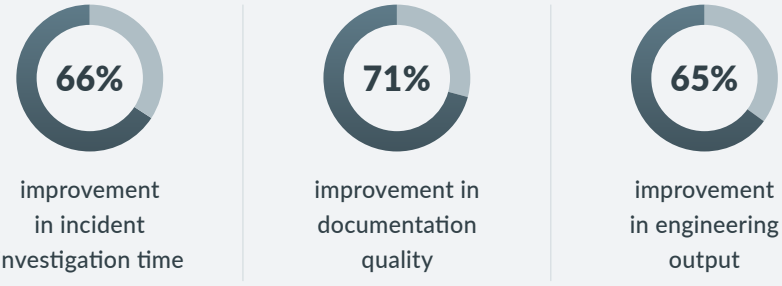

These efficiency gains extend beyond individual output to business impact. Faster incident investigation means shorter outages. Better documentation means fewer recurring support calls. Higher engineering output means more features shipped. The business case is solid, driving continuing adoption.

Interestingly, the top use case for infrastructure leaders is in security and compliance.

#1 Goal for AI Adoption

Security and Compliance

Highest leadership engagement score: **3.92/5**

Developer output	**3.86**
Engineering velocity	**3.85**
Cost cutting	**3.77**
Business velocity	**3.76**

KEY INSIGHT

Security and compliance is the top AI use case being adopted by infrastructure leaders

Organizations understand that AI in their infrastructure can be applied to the hard problems of security and are focused there. Interestingly, this is also the central conflict in the concerns leaders have about adopting AI in the first place.

Security Lags Innovation

Despite the focus on security as a top use case, individuals closest to the systems are raising red flags.

When we asked what worried them most, the answers were illuminating, ranging from errors in configuration to secrets leakage, compliance gaps, and impact on team skills.

KEY FINDING

85%

concerned about AI risking infrastructure security

34% extremely concerned

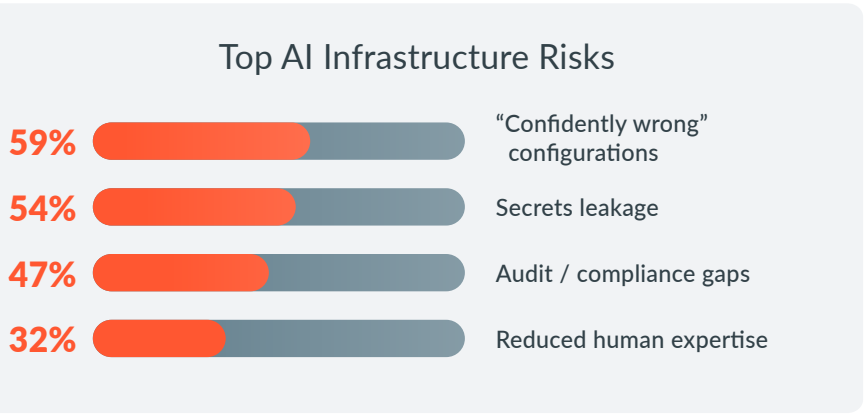

The top fear is AI being "confidently wrong." This refers to an AI system applying a configuration with absolute certainty, which may turn out to be incorrect. The traditional review processes were designed for humans who hedge, ask questions, and express uncertainty. AI doesn't do that. It proposes changes with the same confidence whether they're correct or catastrophically wrong. The top fear is that AI will push bad configurations through, bypassing review processes that would catch human mistakes.

Companies also suspect that they are already experiencing AI-led incidents:

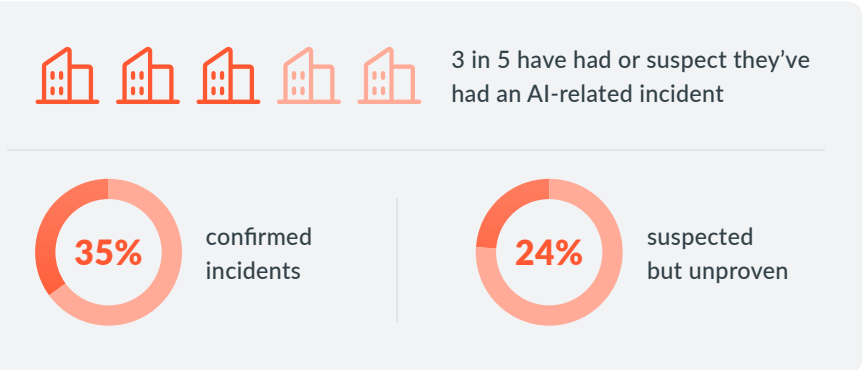

That 24% "suspected but unproven" category is particularly telling. These organizations think they experienced an incident due to AI, but can't definitively attribute it.

Industry Adoption Signals

Some industries, such as financial services, are known for being innovation leaders and have more complex infrastructure. The data pattern in reported incident rates by industry suggest that some may be the tip of the spear in highlighting the gap between innovation / complexity and security for agentic workflow adopters.

KEY INSIGHT

Financial services leaders reported the highest rate of AI-related security incidents

Who's Getting Hit?

Incident rates by industry (had or suspected AI-related incident)

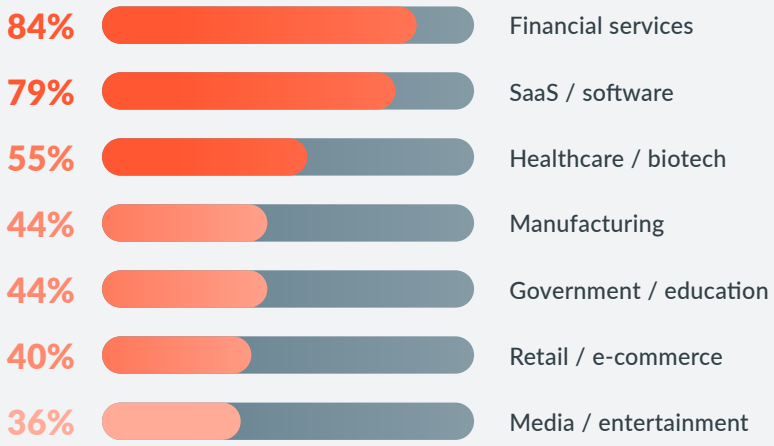

Financial services led at 84%, SaaS at 79%. To be clear, that doesn't mean these industries are inherently less secure. It is more likely that they are further along the innovation curve, previewing what others will face as their own agentic deployments ramp up.

The autonomous nature of agent workflows compounds this. We asked how often AI systems make infrastructure changes without human oversight:

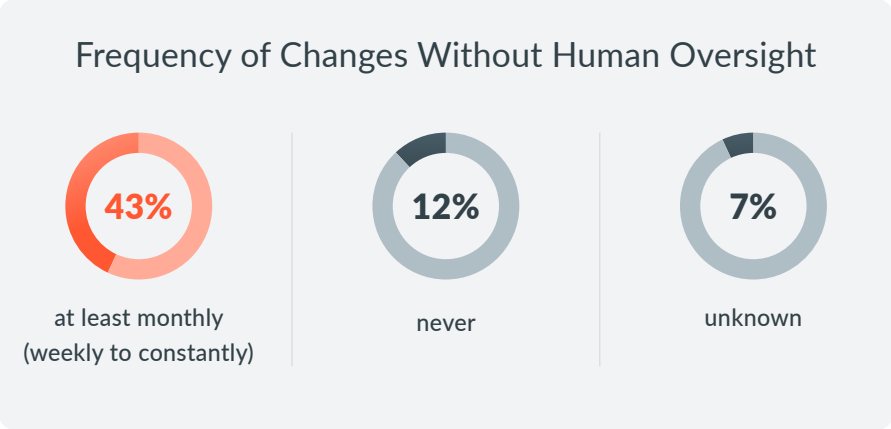

43% report AI affecting infrastructure at least monthly without human review. But 7% can't even answer the question. They don't know how often their AI systems are making autonomous changes. That visibility gap increases the risk of security incidents — you can't govern what you can't see or investigate.

The Confidence Paradox

We also explored whether confidence played a role in incident outcomes. Ostensibly, one might expect that less confident organizations are earlier in their journey, less mature, or making more mistakes. The data showed the opposite.

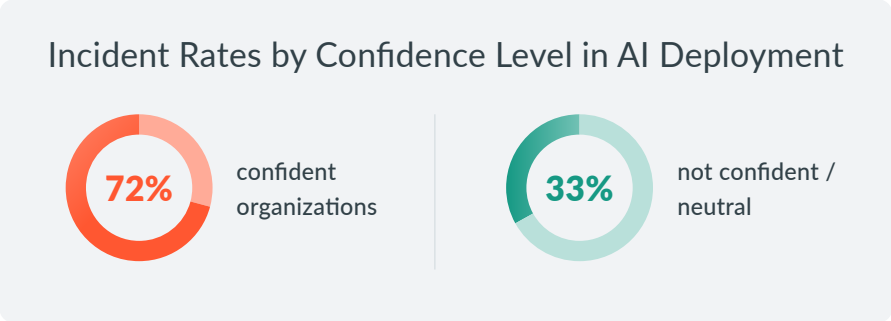

KEY FINDING

2.2x

the incident rate for confident organizations

Confidence does not equal safety

Organizations with stated confidence in their deployments have an incident rate more than twice that of those without confidence.

Regardless of the causal mechanism, the data suggests that feeling secure about AI deployment is not the same as being secure. Confidence doesn't predict safety.

Teleport

PART 2

The Identity Problem

Why 69% of security leaders believe identity management must fundamentally change and why the data proves they're right

A Consensus Emerges

We asked a simpler question: What do security leaders think needs to change? The consensus was striking.

More than two-thirds of security leaders believe the way we handle identity must fundamentally shift to effectively secure AI. Only 2% disagree. This confirms the view that traditional human-centered methods of managing identity and access, such as credentials, monitoring, and manual approvals are insufficient for agent-centered workflows.

However, an additional insight can be found in how companies have implemented access controls for AI agents. We asked organizations how the privileges they are granting AI systems compare to those that a human performing the same task(s) would receive:

KEY FINDING

69%

agree AI adoption will require significant changes to identity management

Only 2% disagree

Seven in 10 organizations have given their AI systems greater access than they would give a human employee doing the same job. Almost one in five have given their AI system dramatically greater access.

Does the scope of access that an organization grants its AI system affect security? We cross-referenced access configurations with incident history. The results were clear.

The Access Equation

We divided organizations into two groups: those whose AI systems have more access than needed (over-privileged) and those whose AI systems have access appropriate to the task they are performing. Then we looked at their incident rates.

This is the most important finding in the survey. Organizations with over-privileged AI have a 76% incident rate. Those with least privileged controls: 17%. This is the single most predictive factor for AI-related incidents that we found — more predictive than the industry, maturity level, or stated confidence.

> “It’s not the AI that’s unsafe. It’s the access we’re giving it.”
>
> – **Ev Kontsevoy,** CEO, Teleport

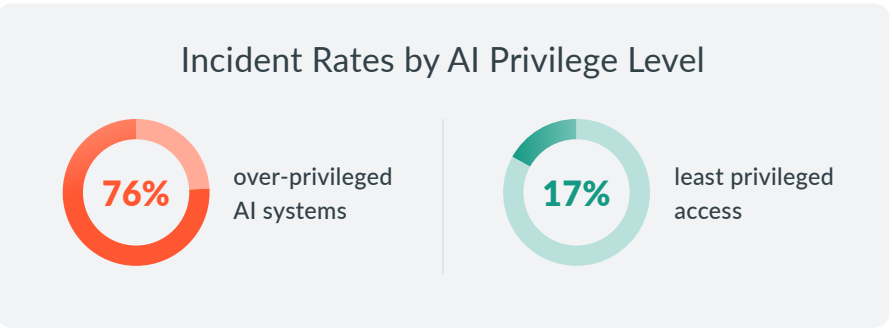

The implications are profound. Organizations that apply the principle of least privilege to their AI systems have dramatically fewer incidents. Those that don’t are nearly 4.5 times more likely to experience security problems.

THE CORE FINDING

Organizations that scope AI access properly have 4.5x fewer incidents than those that don’t. The question isn’t whether AI is dangerous — it’s whether you’ve given it more access than it needs.

The AI-Accelerated Static Credential Risk

What drives AI over-privileging in the first place? The data points to a familiar culprit: static credentials.

67% of organizations reported high utilization of static credentials, such as passwords, API keys, and long-lived tokens, and that reliance has a direct correlation to the number of reported AI incidents.

KEY FINDING

20%

point increase in incidents where static credentials are prevalent

Incident Rates by Static Credential Reliance

What About Governance?

We also explored whether organizations have implemented governance controls.

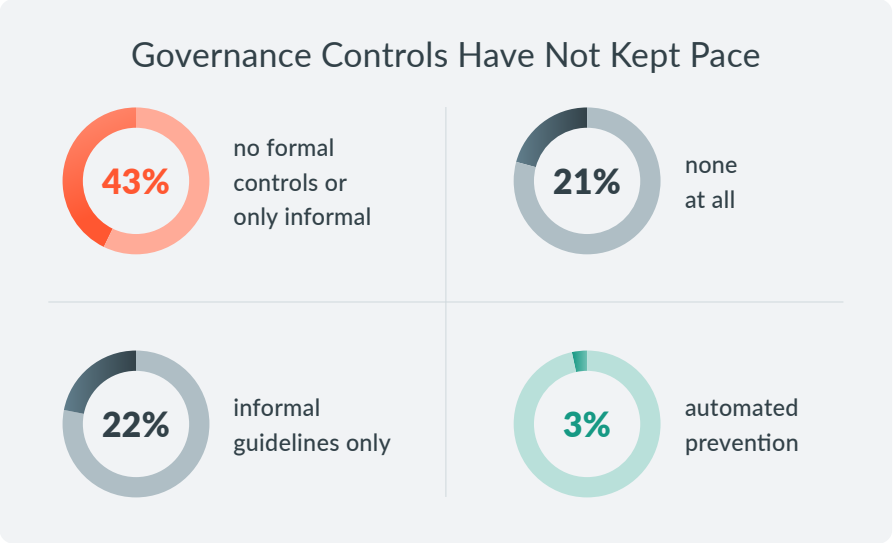

43% have no formal governance controls or only informal guidelines in place to guide their AI use. Only 3% of organizations currently use controls or automated prevention to monitor and govern AI systems as they run.

The Preparedness Gap

The governance challenge becomes acute when we look at what's coming next. Agentic AI — systems that can plan, execute multi-step tasks, and operate with minimal human oversight — is no longer just a theory.

KEY INSIGHT

AI adoption is outpacing organizational preparedness

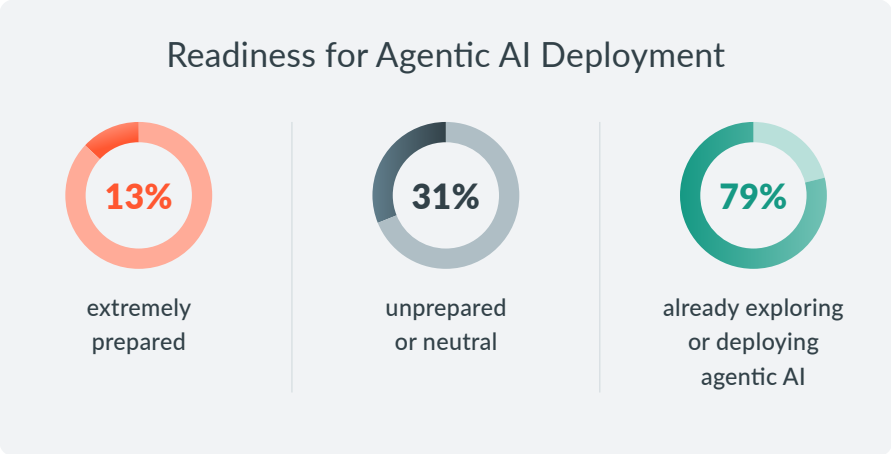

79% of organizations are actively evaluating or deploying agentic AI. Only 13% feel extremely prepared for it. 31% are unprepared or neutral. This gap between adoption and preparedness is the security challenge that needs to be addressed.

AI access control will be much more difficult as agentic systems become increasingly complex. Agentic systems are not like other systems that respond in deterministic ways to queries; instead, they operate autonomously and with non-deterministic outcomes. Today's identity and access problems will increase exponentially as agentic systems become more embedded.

PART 3

The Organizational Shift

Who governs AI security? The answer reveals a fundamental restructuring of how security decisions get made.

Power Has Shifted

Who is responsible for AI security decisions? The data reveals a fundamental restructuring.

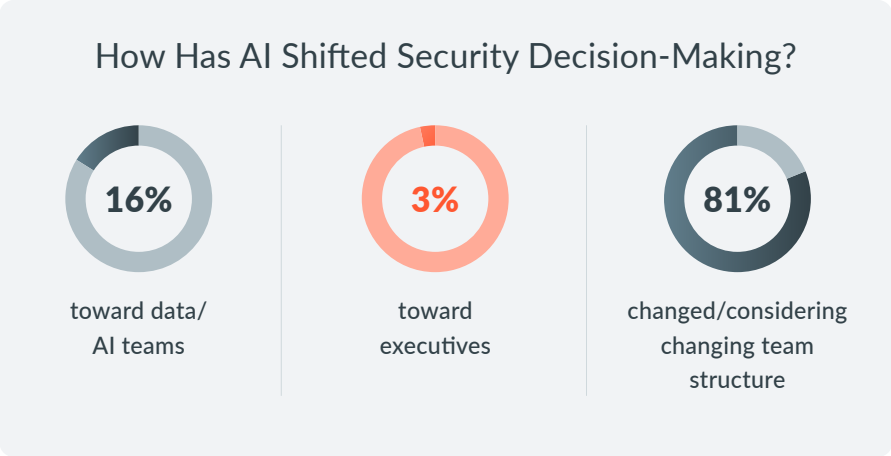

KEY FINDING

52%

report security decision-making power is shifting to platform / infrastructure teams

More than half of organizations report that decision-making power for AI security has shifted toward platform and infrastructure teams. Only 3% say it's moved toward executives. The complexity of AI infrastructure requires that people closest to the systems drive strategy. Correspondingly, AI is emerging as a platform team responsibility and occupying leadership mindshare. Platform teams now have a seat at the table, given their understanding of the technical complexity.

PART 4

The Way Forward

The organizations that build proper identity foundations now will not just be more secure. They'll be able to adopt AI faster because their governance scales with their capability.

The Three-Year Horizon

We also asked about the anticipated horizon for change.

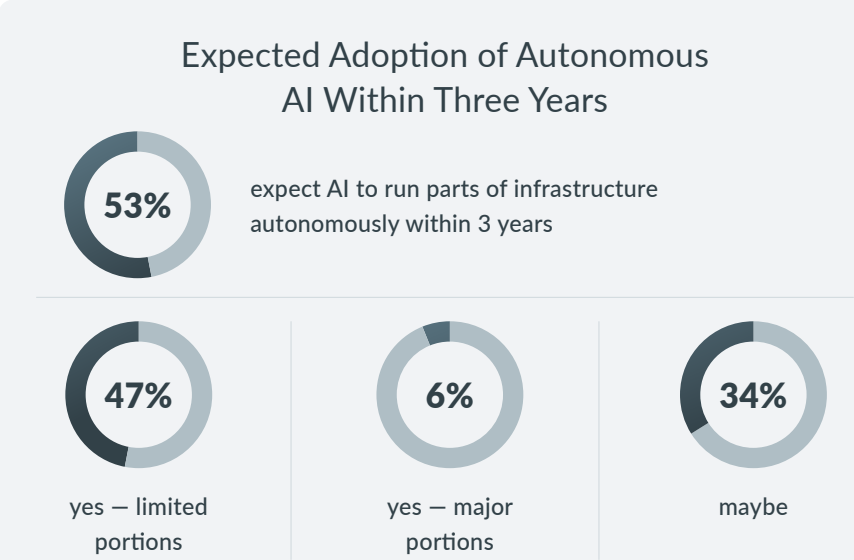

More than half of respondents believe AI will run parts of infrastructure autonomously within the next three years. The lack of access control reported by organizations today, resulting from over-privileging, static credentials, and governance gaps, will potentially result in disastrous outcomes where AI systems make decisions without human-in-the-loop checks.

The lack of access control reported by organizations today, resulting from over-privileging, static credentials, and governance gaps, will potentially result in disastrous outcomes where AI systems make decisions without human-in-the-loop checks.

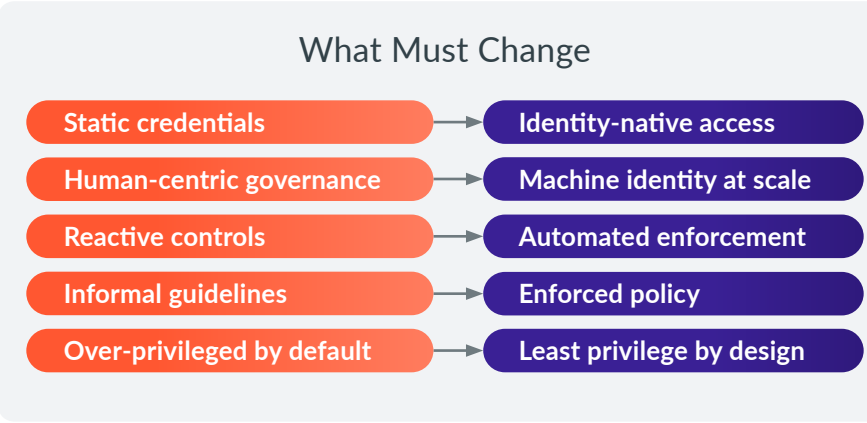

The good news: The survey data reveals a clear target for improvement. Organizations that solve for least privileged access and reduce their reliance on static credentials will have a stronger foundation for autonomous AI. Those that do not will be compounding risk.

Recommendations

As companies consider how to close the gap between AI innovation and effective security, three key ideas emerge from the survey data:

Reduce incident risk by improving identity and access controls.

Implement least privileged access controls. The 76% versus 17% finding represents your highest-ROI intervention. Identify which AI systems have more access to your systems and data than is required, and restrict them accordingly. Implementing these restrictions will be a stronger indicator of your security posture than anything else you do.

Reduce reliance on static credentials.

Static API keys, passwords, service accounts, and other forms of static credentials are associated with persistent, over-privileged access. Eliminate your reliance on static credentials so that identities cannot be stolen, shared, or lost.

Reshape identity management teams.

Platform and engineering teams are playing a larger role in identity management and security, as AI moves into production infrastructure. Unify fragmented identity management silos and resource these teams appropriately.

Methodology	
Sample	205 infrastructure security leaders (CISOs, VPs, Architects, Platform Engineers)
Method	In-depth telephone interviews conducted by Eleven Market Research
Field dates	December 2025
Organization size	500 to 10,000+ employees across seven industries
Qualification	Respondents define, approve, or own infrastructure security platforms

Teleport

Teleport, the AI Infrastructure Identity Company, prepares organizations for an AI future by establishing a unified identity layer for infrastructure, with humans, machines, workloads, and AI agents secured cryptographically with a hardware root of trust rather than vulnerable credentials. By replacing fragmented identity and access management systems with Infrastructure Identity, Teleport scales zero trust across cloud and on-prem environments, eliminating the complexity and risk created by identity fragmentation and credential sprawl. Teleport protects infrastructure from identity attacks, accelerates engineering by reducing infrastructure complexity, and secures non-deterministic agentic workflows.

For more information, visit **www.goteleport.com** or follow **@goteleport**.

Follow us on: