EverOps

June 2026

EVEROPS RESEARCH

The Infrastructure Imperative

Why enterprise AI success depends on platform foundations

Insights for technology leaders in SaaS, fintech, and high-scale consumer technology

Build Faster. Spend Smarter. Sleep Better.

everops.com

ABOUT THIS REPORT

This report synthesizes publicly available research from leading analyst firms and industry organizations to examine the relationship between infrastructure maturity and AI outcomes. All statistics are sourced and verified through primary publications. No proprietary client data is referenced.

Sources include Deloitte, McKinsey, Gartner, IDC, the Cloud Native Computing Foundation (CNCF), Google's DORA research program, the FinOps Foundation, NVIDIA, Salesforce/MuleSoft, and other credible research organizations. Publication dates range from late 2025 through early 2026, reflecting the most current data available at time of writing.

While the findings apply broadly to any enterprise investing in AI, this report gives particular attention to industries where real-time performance, massive concurrency, and rapid iteration cycles define competitive advantage, including SaaS platforms, fintech, consumer technology, and e-commerce. These sectors operate some of the most complex distributed infrastructure in the world, and the relationship between platform foundations and AI readiness is especially pronounced.

Contents

THE STATE OF PLAY

1 Executive Summary

2 The State of Enterprise AI Adoption (2025–2026)

3 The Infrastructure Readiness Gap

THE FOUNDATIONS

4 Kubernetes and Container Orchestration as AI Foundation

5 Observability as the Data Plane for AI

6 CI/CD and Software Delivery in the AI Era

7 FinOps and AI-Driven Cost Optimization

8 Developer Experience and Internal Developer Platforms

9 Security, Governance, and AI Readiness

WHAT IT UNLOCKS

10 What Infrastructure Standardization Unlocks

11 Addressing the Counter-Arguments

12 Two Trajectories: What the Research Suggests

13 Conclusion

REFERENCE

Appendix: AI Infrastructure Readiness Self-Assessment

Sources and References

01

Executive Summary

The gap between AI ambition and AI results is not a strategy problem. It is an infrastructure problem.

Enterprise investment in artificial intelligence has reached unprecedented levels. Global AI spending is projected to exceed $300 billion in 2026, according to IDC's Worldwide AI Spending Guide, while 86% of organizations surveyed by NVIDIA report plans to increase their AI budgets. Adoption is near-universal: McKinsey's 2025 Global AI Survey found that 88% of organizations now use AI in at least one business function, up from 78% the prior year.

Yet behind this enthusiasm lies a stark reality. Nearly two-thirds of organizations remain stuck in experimentation or piloting and have not begun scaling AI across the enterprise (McKinsey, 2025). Gartner predicts that through 2026, organizations will abandon 60% of AI projects unsupported by AI-ready data. Only about 6% of companies qualify as high performers generating meaningful EBIT impact from AI.

This report examines the foundational investments enterprises must make before AI can function at scale. Drawing on research from Deloitte, McKinsey, Gartner, CNCF, Google's DORA research program, and others, it argues that infrastructure standardization, including consolidated Kubernetes environments, unified observability, modern CI/CD pipelines, and disciplined cloud financial management, is not a separate initiative from AI strategy. It is the AI strategy's foundation.

The same investment that drives operational efficiency today creates the substrate on which all future AI capabilities will be built. Organizations that treat infrastructure modernization as merely a cost optimization exercise are undervaluing the investment by an order of magnitude.

The research points to two diverging trajectories. Organizations that build the foundation first, meaning consolidated platforms, unified observability, and governed pipelines, see AI capabilities compound on top of it. Those that buy AI tools first and defer the foundation watch individual productivity gains stall at the team level while costs and fragmentation grow. The gap between these paths is already widening, and it is set by infrastructure, not by model choice.

02

The State of Enterprise AI Adoption

2.1 Spending and Growth Trajectories

Enterprise AI spending has entered a new phase of scale. IDC projects global spending of $301 billion in 2026, up from $223 billion in 2025. Gartner projects AI software alone will account for $157 billion. The investment is broad-based: according to EY, 88% of mid-to-large organizations now spend more than 5% of their total IT budget on AI, with many targeting 25% in the near term.

Menlo Ventures reports that departmental AI spending hit $7.3 billion in 2025, a 4.1x year-over-year increase, with coding tools representing $4.0 billion of that total, making software engineering the single largest AI spending category. IT operations tools reached $700 million as teams automated incident response and infrastructure management.

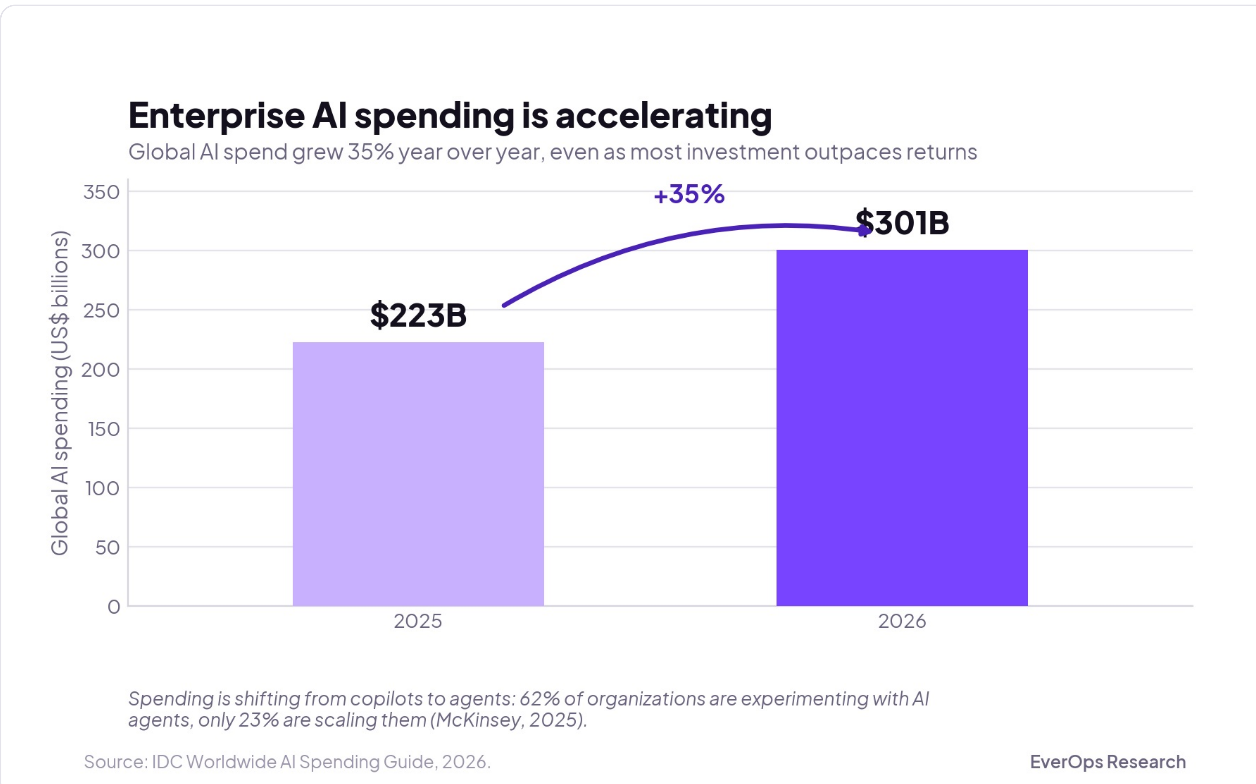

Figure 1. Global enterprise AI spending, 2025–2026.

$301 billion in projected global AI spending in 2026, up 35% from 2025.

IDC Worldwide AI Spending Guide

2.2 The Pilot-to-Production Gap

Despite massive investment, the transition from experimentation to enterprise-scale deployment remains the central challenge. McKinsey's 2025 survey of nearly 2,000 respondents across 105 countries found that approximately one-third of organizations have begun scaling AI, while roughly two-thirds remain in experimenting or piloting stages. ISG's 2025 report found that only 31% of AI use cases studied had reached full production.

The blockers are consistent: data quality and architecture limitations, integration challenges with legacy systems, workflow rigidity, unclear ROI measurement, and insufficient AI-specific talent. McKinsey found that 39% of organizations report insufficient AI skills. Gartner's survey of 506 CIOs revealed that 72% report their organizations are breaking even or losing money on AI investments.

> **Roughly 66% of organizations have not yet begun scaling AI across the enterprise.**
>
> McKinsey Global AI Survey, 2025

2.3 The Shift from Copilots to Agentic AI

The most consequential shift is the transition from copilot-style tools to agentic AI systems capable of autonomous planning and execution. McKinsey found that 62% of organizations are experimenting with AI agents, though only 23% report scaling agentic systems. Gartner predicts that by 2029, 70% of enterprises will deploy agentic AI in IT infrastructure operations, up from less than 5% in 2025.

This transition fundamentally changes infrastructure requirements. Copilots augment human decisions within existing workflows. Agents act autonomously, requiring standardized APIs, consistent execution environments, governance guardrails, RBAC policies, and audit trails that most enterprise infrastructure cannot provide. The gap between what agentic AI requires and what organizations have deployed is the central tension driving the infrastructure imperative.

03

The Infrastructure Readiness Gap

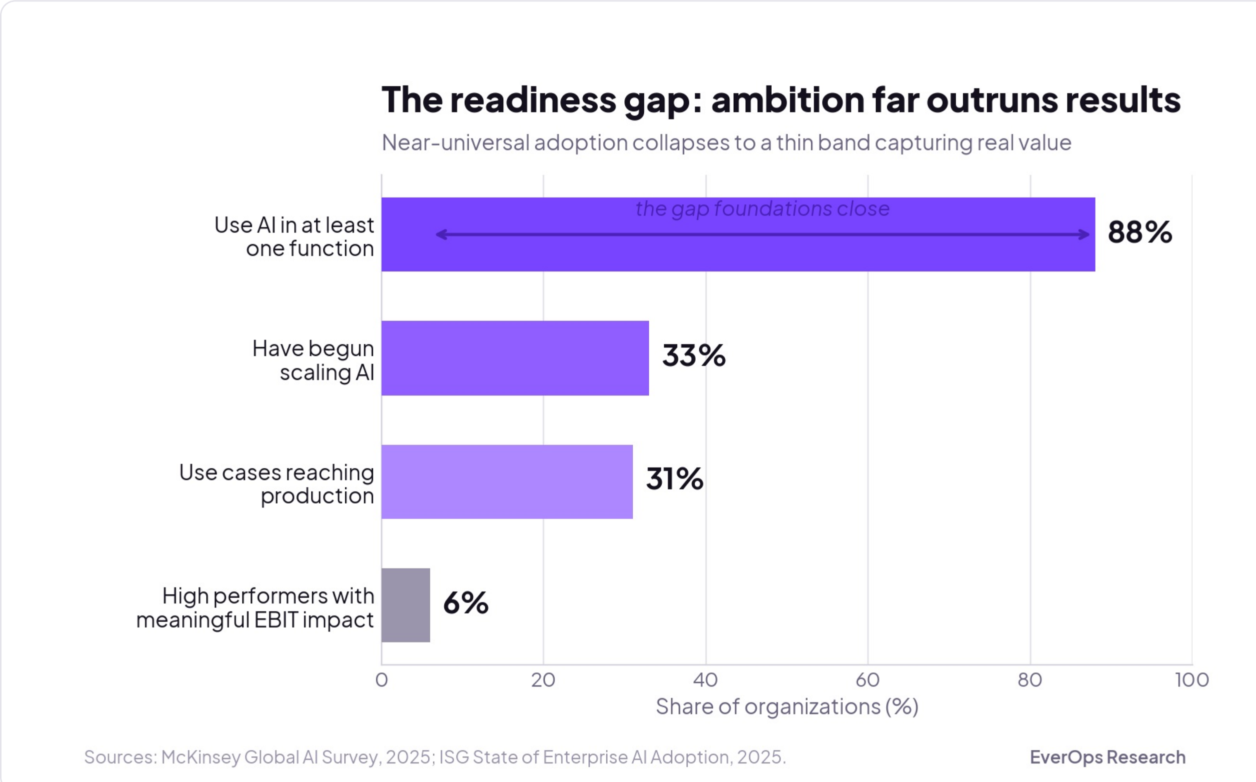

Figure 2. From near-universal adoption to a thin band of value capture.

3.1 Preparedness Is Declining as Ambitions Grow

Perhaps the most alarming finding across recent surveys is that infrastructure preparedness is declining even as AI ambitions accelerate. Deloitte's 2026 State of AI report, surveying 3,235 leaders globally, found that while 42% believe their strategy is highly prepared for AI adoption, they simultaneously report feeling less prepared in terms of infrastructure, data, risk, and talent. The gap between strategic confidence and operational readiness is widening.

Gartner predicts that through 2026, 60% of AI projects will be abandoned due to insufficient AI-ready data. Their survey of data management leaders found that 63% of organizations either do not have or are unsure whether they have the right data management practices for AI. Traditional data management, Gartner notes, is too slow, too structured, and too rigid for AI teams.

> **60% of AI projects will be abandoned due to insufficient AI-ready data through 2026.**
>
> Gartner, February 2025

3.2 Why Legacy Architectures Cannot Support AI at Scale

AI systems require unified data planes where agents can access consistent, real-time information across services. They require standardized APIs that enable autonomous tools to interact with infrastructure programmatically. They require consistent deployment patterns for safe validation. And they require real-time telemetry for monitoring, diagnosis, and remediation.

Most enterprise environments are characterized by the opposite: fragmented infrastructure spread across teams and business units, inconsistent tagging, siloed observability tools, and bespoke deployment processes. This is especially acute in organizations that have grown through acquisition or operate multiple semi-autonomous product lines. Salesforce's 2025 MuleSoft Connectivity Benchmark Report, surveying more than 1,050 enterprise IT leaders, found that 90% of organizations say data silos are creating active business challenges, and that organizations now manage an average of 897 applications, with only a fraction integrated. In these environments, AI tools cannot reason across the stack and cannot be trusted with autonomous action.

> **90% of enterprise IT leaders say data silos are creating active business challenges; organizations manage 897 applications on average.**
>
> Salesforce MuleSoft Connectivity Benchmark Report, 2025

3.3 What Leading Analysts Say

Gartner's December 2025 report on I&O trends for 2026 identified agentic AI and hybrid computing as the two most impactful trends, noting enterprises must adopt composable architectures. McKinsey's high-performer analysis found that the 6% of organizations generating meaningful AI returns are distinguished not by model choice but by their investment in workflow redesign, data foundations, and platform infrastructure.

The 2025 DORA report from Google Cloud stated the finding most directly: AI does not create organizational excellence; it amplifies what already exists. For high-performing organizations with solid foundations, AI becomes a powerful accelerator. For those with fragmented systems, it magnifies existing dysfunction.

> **"AI doesn't fix a team; it amplifies what's already there."**
>
> 2025 DORA Report, Google Cloud

04

Kubernetes as AI Foundation

4.1 The New Operating System for AI

Kubernetes has evolved from a container orchestration tool into what the CNCF now describes as the operating system for artificial intelligence. The CNCF's 2025 Annual Cloud Native Survey found that 82% of container users run Kubernetes in production, up from 66% in 2023. 98% of surveyed organizations reported adopting cloud native techniques, confirming Kubernetes has moved decisively beyond the early-adopter phase.

The CNCF found that Kubernetes is increasingly the platform of choice for AI workloads, though only 7% of organizations deploy AI models daily. Most prioritize reliably and cost-effectively operating pre-trained models, a use case for which standardized Kubernetes is well suited. As the report noted, success requires treating AI and ML as a first-class infrastructure challenge, not just an algorithmic one.

> **82% of container users run Kubernetes in production in 2025, up from 66% in 2023.**
>
> CNCF Annual Cloud Native Survey, January 2026

4.2 Why Consolidation and Standardization Matter

For AI tools to operate effectively, they require consistent, predictable environments. Fragmented Kubernetes clusters with varying configurations, policies, and tooling prevent AI systems from generalizing across the stack. An AI agent trained to diagnose issues in one cluster configuration cannot reliably transfer that knowledge to a differently configured environment. This problem scales with organizational complexity: multi-product SaaS platforms, fintech firms operating across multiple business lines, and consumer technology companies managing global infrastructure all face the same challenge of fragmented environments preventing AI from operating at the organizational level.

The CNCF survey identified a clear link between maturity and standardization. 0% of organizations classified as “explorers” report using GitOps extensively, compared to 58% of “innovators.” This gap represents the difference between organizations where AI tools can reason about infrastructure and those where they cannot.

4.3 Platform Engineering as the Delivery Mechanism

The platform engineering movement has emerged as the organizational response to cloud native complexity. The 2025 DORA report found that 90% of organizations have adopted at least one internal platform and established a direct correlation between platform quality and an organization's ability to unlock AI value. Platform engineers now represent 37% of Argo CD users (CNCF, 2025).

For AI capabilities to reach engineering teams at scale, they must be delivered through platforms that abstract complexity and provide standardized interfaces. This includes consistent resource quotas, RBAC policies, governance frameworks, and deployment pipelines that AI agents can interact with safely.

05

Observability as the Data Plane for AI

5.1 The Single Most Critical Prerequisite

If AI-powered operations have one non-negotiable prerequisite, it is unified observability. AI SRE tools, anomaly detection systems, and autonomous remediation agents all depend on consistent, correlated telemetry across logs, metrics, traces, and events. When observability is siloed, AI agents cannot perform root cause analysis, cannot correlate symptoms across services, and cannot build the situational awareness required for autonomous action.

The AIOps market is valued at approximately $18.95 billion in 2026 and projected to reach $37.79 billion by 2031 (Mordor Intelligence). Dynatrace's 2025 State of Observability report found that 70% of respondents saw increased observability budgets and 75% expected further increases.

> **$18.95 billion AIOps market value in 2026, growing at 14.8% CAGR to $37.79 billion by 2031.**
>
> Mordor Intelligence, January 2026

5.2 OpenTelemetry and the Standardization Movement

OpenTelemetry has become the second-highest-velocity CNCF project, with more than 24,000 contributors. An EMA survey of 400 IT professionals found that 48.5% of organizations are already using OpenTelemetry in some capacity, with 25.3% planning implementation. Among adopters, 46.4% report greater than 20% ROI, and 84% of those who reduced costs achieved at least a 10% decrease. Production adoption jumped from 6% to 11% between 2025 and 2026 (Dimensional Research/Elastic).

For organizations operating at massive scale, including high-growth SaaS platforms scaling toward millions of users, fintech systems handling rapidly rising transaction volumes, and consumer apps managing fast-growing global traffic, the ability to correlate telemetry across the entire stack in real time is not a luxury. It is the minimum viable foundation for AI-assisted operations. Organizations using AI in IT operations already report 31% fewer critical incidents and 28% faster mean time to resolution, but these gains are only possible with unified, standardized telemetry.

5.3 From Reactive Monitoring to Predictive Reliability

The shift from reactive monitoring to predictive, AI-driven reliability requires real-time anomaly detection, automated root cause analysis, predictive capacity planning, and ultimately self-healing capabilities. IBM's 2026 observability trends analysis notes that agentic AI integration with observability data is the most significant emerging trend, as the intelligence required to maintain AI-powered systems demands more innovative observability capabilities.

06

CI/CD and Software Delivery in the AI Era

6.1 AI Is Transforming How Code Is Written and Deployed

The 2025 DORA report confirmed that AI adoption among software professionals has surged to 90%, with a median of two hours daily spent working with AI tools. More than 80% indicate AI has enhanced their productivity, and 59% report positive impacts on code quality.

Coding has become the dominant AI spend category. Menlo Ventures reports coding represents $4.0 billion of $7.3 billion in departmental AI spending, with 50% of developers using AI tools daily. Teams report 15%+ velocity gains across the development lifecycle.

> **90% of software professionals now use AI assistance at work; 80%+ report increased productivity.**
>
> DORA 2025 State of AI-Assisted Software Development

6.2 The Governance Challenge: AI as Mirror and Multiplier

The DORA report's most important finding is that AI does not automatically improve delivery performance. It acts as a multiplier of existing engineering conditions. AI adoption is now linked to higher throughput but continues to have a negative relationship with delivery stability. Without robust testing, mature version control, and fast feedback loops, increased velocity leads to instability.

Faros AI's telemetry analysis of 10,000+ developers found the “AI Productivity Paradox”: AI assistants boost individual output (21% more tasks completed, 98% more pull requests merged) but organizational delivery metrics remain flat. Without standardized pipelines that validate and safely deploy AI-generated artifacts, individual gains are absorbed by downstream friction.

6.3 Standardized Pipelines as the Safety Net

As AI-generated code volumes increase, the CI/CD pipeline becomes the critical governance layer, enforcing consistent quality gates regardless of whether code was human-written or AI-generated. CNCF data shows 60% of organizations have adopted CI/CD platforms for most cloud native applications, a 31% increase year-over-year. 77% of organizations have adopted GitOps to some degree.

For any organization where a failed deployment carries significant revenue or reputational consequences, including high-traffic SaaS releases, fintech platform updates, payment processing changes, and consumer app launches, standardized deployment pipelines are not optional. They are the safety net that makes AI-accelerated development viable.

07

FinOps and AI-Driven Cost Optimization

7.1 The Scale of Cloud Waste

Cloud spending continues to grow rapidly, with Forrester projecting public cloud spending of approximately $1.03 trillion in 2026. Yet waste remains high. Analyses consistently find that 27–32% of spending is wasted, a figure stable since 2019 (Flexera, Harness, Datadog). Kubernetes clusters average just 10% CPU and 20% memory utilization (CNCF FinOps Survey). At current spend levels, this represents more than $270 billion annually in unnecessary expenditure.

> **27–32% of cloud spending is wasted: more than $270 billion annually at current spend levels.**
>
> Flexera, Harness, Datadog (2019–2025)

7.2 AI-Powered Optimization Requires Unified Foundations

AI-driven cost optimization requires unified cost visibility and consistent resource tagging. Only 43% of organizations track costs at the unit level (Gartner, May 2025). The FinOps Foundation's 2026 report found that AI spend management surged from 31% of respondents in 2024 to 98% in 2026. AI is simultaneously the biggest new cost category and the most powerful optimization tool, but only for organizations with the data foundations to leverage it.

7.3 Standardized Infrastructure Simplifies Cost Management

Fragmented, per-team infrastructure is inherently difficult to optimize. When every business unit runs its own configurations, tagging conventions, and provisioning patterns, cost visibility becomes impossible. Standardized platforms reduce this complexity by enforcing uniform tagging, centralized provisioning, and automated lifecycle management. Organizations implementing structured FinOps programs report 25–30% reductions in monthly cloud spend.

08

Developer Experience and Internal Platforms

8.1 Platforms as the AI Delivery Mechanism

The DORA report established a direct correlation between platform quality and AI value capture. With 90% of organizations now using at least one platform, the question is how to build platforms that serve as the distribution layer for AI capabilities. Internal developer platforms abstract infrastructure complexity, providing standardized interfaces through which developers interact with compute, storage, networking, and AI services.

8.2 Productivity Gains and the Infrastructure That Enables Them

Menlo Ventures found that teams report 15%+ velocity gains from AI across the development lifecycle. GitHub Copilot research indicates 40–55% more code per week from AI-assisted developers. These gains are contingent on infrastructure quality. The DORA research makes clear that without mature version control, strong automated testing, and high-quality platforms, AI-driven productivity does not translate to organizational performance.

In the $465 billion global SaaS market, where engineering velocity, reliability, and cost efficiency directly determine competitive position (Precedence Research, 2026), the compound effect of developer productivity is not abstract. An engineering organization gaining 15% velocity from AI-assisted tooling has the effective output of a team 15% larger, but only if the surrounding infrastructure can absorb and deploy that increased output safely. For fintech and consumer tech companies, where a single deployment failure can disrupt a large share of the user base and, in regulated segments, trigger compliance scrutiny, the quality of the platform is not an engineering concern alone. It is a business continuity and trust concern.

> **$465 billion projected global SaaS market in 2026; engineering velocity and reliability are the primary competitive levers.**
>
> Precedence Research, 2026

8.3 The Rise of AI-Focused Platform Engineering

A new specialization is emerging: platform teams focused on building infrastructure that AI tools and agents operate on. This includes AI model serving, standardized agent configurations, governance frameworks for autonomous actions, and telemetry optimization for AI consumption. The CNCF found 15.6 million developers now use cloud native technologies, with platform engineering growing as a distinct discipline.

09

Security, Governance, and AI Readiness

9.1 Shadow AI as a Governance Failure

Shadow AI, the use of AI tools outside organizational governance, has emerged as a critical enterprise risk. More than 80% of employees use unapproved AI tools (UpGuard, 2025). IBM's 2025 Cost of Data Breach Report found shadow AI incidents account for 20% of all breaches with a $670,000 cost premium per incident. Yet only 37% of organizations have policies to detect shadow AI.

Gartner predicts that by 2030, more than 40% of enterprises will experience security or compliance incidents from unauthorized AI. Fragmented infrastructure makes this problem exponentially harder: when environments are distributed with inconsistent tooling and access controls, detecting unauthorized AI usage requires searching across dozens of uncoordinated systems. Standardized platforms provide centralized visibility.

> **80%+ of employees use unapproved AI tools; only 37% of organizations have governance policies.**
>
> IBM, UpGuard, Menlo Security (2025)

9.2 Governance Requirements for Autonomous Systems

As AI transitions from copilots to agents, governance intensifies. Autonomous systems require audit trails, reasoning traces, decision artifacts, and rollback mechanisms. Gartner's December 2025 report emphasized that as autonomy increases, organizations require stronger operational control. Standardized platforms provide the enforcement point: consistent RBAC, centralized audit logging, automated compliance, and policy-as-code applied uniformly across all deployments.

10

What Infrastructure Standardization Unlocks

The research reviewed in this report consistently shows that infrastructure standardization is the prerequisite for enterprise AI at scale. This section maps specific AI capabilities to their infrastructure prerequisites.

10.1 AI-Powered Incident Detection and Remediation

Organizations using AI in IT operations report 31% fewer critical incidents and 28% faster mean time to resolution. These capabilities, including anomaly detection, automated root cause analysis, and autonomous remediation, require unified observability as a prerequisite. Without correlated telemetry across the entire stack, AI tools have no data to reason about.

Infrastructure prerequisite: Unified observability with standardized telemetry, OpenTelemetry adoption, and correlated logs, metrics, and traces across all environments.

10.2 AI-Assisted Code Migration and Modernization

McKinsey's research shows that AI high performers are 2.8x more likely to have fundamentally redesigned workflows. One of the most promising applications is automated codebase analysis: AI systems that compare existing service architectures to target platform standards, identify deltas, and generate migration plans. This capability requires standardized target architectures for AI to evaluate against.

Infrastructure prerequisite: Consolidated, standardized Kubernetes platform with documented target-state architecture and consistent service patterns.

10.3 Predictive Capacity Scaling

AI-driven capacity forecasting uses historical usage patterns, event calendars, and real-time demand signals to predict scaling needs before they materialize. For industries with spike-driven usage patterns, predictive scaling eliminates both the performance risk of under-provisioning and the cost waste of over-provisioning.

Infrastructure prerequisite: Standardized infrastructure with consistent resource management, unified metrics, and clean historical usage data.

10.4 AI-Driven Cost Optimization

FinOps Foundation data shows that organizations implementing structured optimization recover 25–30% of cloud spend. AI dramatically improves this through automated right-sizing, idle resource detection, and commitment optimization. With 27–32% of cloud spend wasted annually, AI-powered FinOps represents one of the clearest ROI cases for infrastructure standardization.

Infrastructure prerequisite: Consistent resource tagging, centralized cost visibility, standardized provisioning patterns, and unit-level cost tracking.

10.5 Automated Deployment Validation

The DORA report found that AI increases delivery throughput but decreases stability when governance is weak. AI-powered deployment validation, including automated code review, security scanning, and regression testing of AI-generated code, addresses this directly. As AI-generated code volumes increase, the pipeline becomes the critical quality gate.

Infrastructure prerequisite: Standardized CI/CD pipelines with consistent quality gates, automated testing, and governed deployment processes.

10.6 Intelligent Observability and Predictive Reliability

IBM's 2026 analysis identifies agentic AI integration with observability as the top emerging trend. This includes AI agents that ingest telemetry data, identify patterns invisible to human operators, and proactively address issues before they impact users. The AIOps market's growth to $37.79 billion by 2031 reflects the scale of this opportunity.

Infrastructure prerequisite: Unified observability with OpenTelemetry instrumentation and real-time data accessibility for AI agents.

11

Addressing the Counter-Arguments

An honest assessment of the infrastructure-AI relationship must address the strongest objections to the thesis presented in this report.

"Most organizations are losing money on AI. Why invest more?"

Gartner's 2025 survey found that 72% of CIOs report their organizations are breaking even or losing money on AI investments. This is a real finding. But the conclusion is not that AI investment should be reduced. The conclusion is that most organizations are investing in AI incorrectly, funding models, tools, and experiments while neglecting the infrastructure that would allow those investments to generate returns. The answer is not less AI investment; it is differently allocated AI investment.

"We should wait for AI to mature before investing in infrastructure."

This argument assumes that deferring infrastructure investment is cost-neutral. The evidence suggests otherwise. McKinsey's high-performer analysis reveals a compounding advantage: organizations with strong foundations deploy AI more successfully, generating returns that fund further investment. Organizations without foundations spend on experiments that fail, consuming budget that could have built the infrastructure for future success. Every month of delay adds new fragmentation: new services deploy with new patterns, new teams adopt new tools, and data accumulates in new silos.

"Infrastructure standardization is just a cost play. We already budget for that."

This is the most consequential misconception. If infrastructure modernization is framed as a cost optimization exercise, it will be evaluated against the annual savings it produces. That is a legitimate but incomplete valuation. The same investment also creates the platform on which AI-powered operations, AI-driven cost optimization, automated deployment validation, predictive reliability, and intelligent capacity management will all be built. Organizations that frame infrastructure modernization as cost optimization will fund it at cost-optimization levels. Organizations that recognize it as AI-enabling infrastructure will fund it at the level the opportunity warrants.

12

Two Trajectories

The data reviewed in this report points to two divergent trajectories for organizations at similar stages of AI adoption. These are not predictions; they are patterns observed in the research.

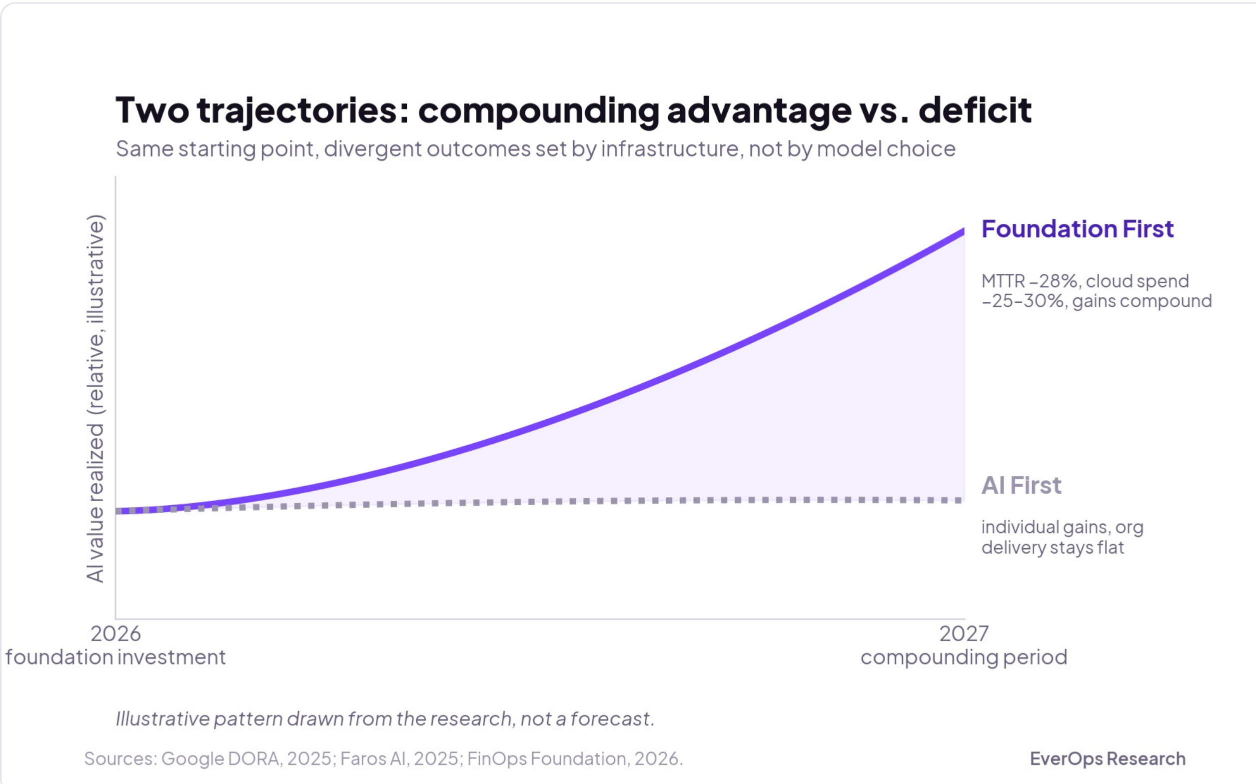

Figure 3. Two trajectories from the same starting point. Illustrative pattern, not a forecast.

Trajectory A — Foundation First

Organization A invests in platform standardization in 2026: consolidated Kubernetes, unified observability, standardized CI/CD, consistent tagging and governance. By 2027, AI capabilities begin compounding on top of that foundation. AI-assisted incident response reduces MTTR by 28%. AI-driven cost optimization recovers 25–30% of cloud spend. AI-assisted deployment validation absorbs increased code velocity without sacrificing stability. Each AI capability builds on the same infrastructure, with marginal cost decreasing as the foundation matures.

Trajectory B — AI First, Foundation Later

Organization B invests directly in AI tools and experiments without addressing infrastructure fragmentation. Individual developers see productivity gains, but organizational delivery metrics stay flat (Faros AI, 2025). AI-powered operations tools are purchased but cannot operate effectively across inconsistent environments. Shadow AI proliferates as approved tools fail to meet needs. Cloud costs grow faster than value. By 2027, the organization recognizes it needs the infrastructure work it deferred, now at higher cost, with more technical debt, and with competitors on Trajectory A already compounding AI capabilities.

McKinsey's data shows these trajectories are already diverging. AI high performers are 3x more likely to pursue transformative change and 2.8x more likely to have redesigned workflows. The 6% at the top are pulling away from the 94% still in pilot purgatory. The differentiator is not which AI tools they chose. It is the infrastructure and organizational foundations they built first.

13

Conclusion

Infrastructure standardization is not a separate initiative from AI strategy. It is the AI strategy's foundation.

The evidence across every major analyst firm, industry survey, and research program converges on the same finding: the barrier to enterprise AI success is not the availability of models, tools, or talent. It is the readiness of the infrastructure on which AI must operate.

The same investment that drives operational efficiency, reduces cloud waste, improves developer productivity, and strengthens security posture also creates the platform on which all future AI capabilities will be built. The 2025 DORA report's central finding should serve as the organizing principle for enterprise technology strategy: AI amplifies whatever an organization already is. Building strong foundations creates a compounding advantage. Deferring that investment creates a compounding deficit.

For technology leaders in SaaS, fintech, and consumer technology — high-growth organizations operating some of the most demanding, fastest-scaling platforms in their markets — the calculus is particularly clear. These organizations already operate at the edge of infrastructure complexity, managing rising concurrency, real-time performance requirements, regulatory exposure, and increasingly global distribution. The same platform investments that solve today's operational challenges also build the foundation for every AI capability the organization will need over the next five years.

The convergence of mature open standards, proven platform patterns, and favorable investment conditions creates a window for organizations that are prepared to act. The research suggests that the organizations that use this window to build AI-ready infrastructure will define the competitive landscape of the next decade.

Build Faster. Spend Smarter. Sleep Better.

everops.com

A

AI Infrastructure Readiness Self-Assessment

The following questions are derived from the research findings in this report. Each maps to a capability the research identifies as a prerequisite for successful AI deployment at scale.

Are all production workloads running on a standardized Kubernetes platform with consistent configurations across teams and business units?	CNCF: innovators are defined by standardized Kubernetes plus GitOps (Section 4)
Do you have correlated telemetry (logs, metrics, traces) accessible across all services and environments through a single platform or standard?	Observability is the #1 prerequisite for AI-powered operations (Section 5)
Have you adopted or planned adoption of OpenTelemetry for vendor-neutral instrumentation?	48.5% of organizations already use OpenTelemetry; 2nd highest-velocity CNCF project (Section 5)
Do all teams deploy through standardized pipelines with consistent quality gates, automated testing, and governed processes?	AI increases throughput but decreases stability without governed pipelines (Section 6)
Can you attribute cloud costs to specific services, teams, and business units with 90%+ accuracy?	Only 43% track costs at unit level; prerequisite for AI-driven optimization (Section 7)
Do you operate an internal platform that provides self-service infrastructure to development teams with consistent governance?	DORA: direct correlation between platform quality and AI value capture (Section 8)
Do you have documented policies governing which AI tools are approved, what data can be shared, and how AI outputs are validated?	Only 37% of organizations have AI governance policies (Section 9)
Do you use declarative, version-controlled infrastructure management (GitOps) across your environments?	0% of explorers vs. 58% of innovators use GitOps extensively (Section 4)

How to read your score

Organizations answering “yes” to six or more of these questions have the infrastructure foundations in place to pursue AI at scale. Those answering “yes” to fewer than four should consider infrastructure standardization as a prerequisite investment before scaling AI initiatives.

Sources and References

All sources verified through primary web research, March–April 2026.

Analyst Reports and Surveys

Deloitte. “The State of AI in the Enterprise, 2026.” January 2026. 3,235 leaders surveyed.

McKinsey / QuantumBlack. “The State of AI in 2025.” November 2025. 1,993 participants, 105 countries.

Gartner. “Lack of AI-Ready Data Puts AI Projects at Risk.” February 2025.

Gartner. “Predicts 2026: AI Agents Will Transform IT Infrastructure and Operations.” December 2025.

Gartner. “Top Trends Impacting Infrastructure and Operations for 2026.” December 2025.

Gartner. “IT Symposium/Xpo Survey of 506 CIOs.” October 2025.

IDC. “Worldwide Artificial Intelligence Spending Guide.” 2026.

NVIDIA. “State of AI Report 2026.” March 2026.

ISG. “State of Enterprise AI Adoption Report 2025.” September 2025.

EY. “Enterprise AI Adoption Survey.” 2025.

Cloud Native and Platform Engineering

CNCF. “Annual Cloud Native Survey 2025.” January 2026.

CNCF / SlashData. “State of Cloud Native Development.” November 2025.

CNCF. “Argo CD End User Survey 2025.” July 2025.

Software Delivery and Developer Productivity

Google Cloud / DORA. “2025 State of AI-Assisted Software Development.” September 2025.

Menlo Ventures. “2025: The State of Generative AI in the Enterprise.” December 2025.

Faros AI. “AI Productivity Paradox.” July 2025.

Integration, FinOps, Observability, Security

Salesforce / MuleSoft. “2025 Connectivity Benchmark Report.” July 2025.

FinOps Foundation. “The State of FinOps 2025/2026.”

Flexera. “State of the Cloud Report.” 2023–2025. Forrester. “Public Cloud Market Outlook.” 2026.

Mordor Intelligence. “AIOps Market Report 2031.” January 2026. Dynatrace. “The State of Observability 2025.”

Dimensional Research / Elastic. “The Landscape of Observability in 2026.” EMA. “OpenTelemetry Adoption Survey.” 2025.

IBM. “Observability Trends 2026” (March 2026); “2025 Cost of Data Breach Report.”

Precedence Research. “Global SaaS Market Size Report.” 2026. UpGuard. “State of Shadow AI Report.” 2025.