PagerDuty

AI reality check:

Strategic truths for CIOs and CTOs on derisking AI transformation

May 2025

Contents

AI is the great accelerant of the modern era

With each industrial revolution came a slew of new technologies that brought in the new era. In this modern era, AI is causing seismic shifts in how work gets done. Now, an organization's competitive advantage doesn't distill down to simple product market fit or supply and demand. It's about harnessing AI to deliver more results with less cost, while keeping resilience and customer experience at the forefront of everything you do.

AI has the power to transform how work gets done—freeing up teams for creativity, speeding up decision-making, and delivering value at scale. But without the right guardrails, that same power becomes a risk multiplier. We've already seen the headlines: AI hallucinations, rogue bots, hidden model drift, and million-dollar outages that no one saw coming. This isn't hype—it's happening.

For CIOs, CTOs, and operations leaders, the stakes couldn't be higher. You're being asked to move fast, show impact, and keep everything running—all while figuring out how to govern something that's still evolving in real time. AI for the sake of AI doesn't solve anything. In fact, it often makes things worse. Without a clear plan and operating model, AI initiatives can quickly become expensive science projects—or worse, brand-damaging liabilities.

So, what does it actually take to build a resilient AI strategy? One that moves fast but stays grounded? That empowers teams instead of replacing them? That scales across the org without breaking everything in its path?

In this guide, we'll walk through the truths of modern AI transformation—what's working, what's failing, and how to lead your organization into an AI-powered future with confidence, not chaos.

Because when things go sideways (and they will), you don't want to be asking "what now?". You want to already know the answer.

Truth 1

Your AI will fail

(it's how you recover that matters)

Myth

"If we just build the model right, we won't have problems."

Truth

Even the best models can fail in unpredictable ways. That's why smart teams design for resilience–so issues are caught early, the right response kicks in fast, and learning loops close the gap before it happens again.

Let's get to something straight up front: your AI will fail.

Maybe not today. Maybe not next quarter. But eventually, something will go sideways–an integration will break, a mode will drift, an agent will act unpredictably, or a key dependency will go dark without warning. And when it does, the stakes won't just be technical. They'll be business-critical.

What most teams overlook is that the most costly AI failures don't start in code, they start in communication. Failures escalate when no one knows who owns what, when teams aren't aligned, or when engineers don't trust what the models are doing. If your AI stack moves faster than your people can respond, you're not accelerating but compounding risk.

Truth 1

Why traditional safeguards fall short

Legacy IT safeguards weren't built for this. Static runbooks, siloed monitoring tools, and manual escalation paths can't keep pace with the real-time complexity of AI-driven systems. And traditional incident management assumes the failure is known. AI failures are often silent, subtle, or unpredictable until damage is already done.

Despite the glossy keynote demos, there's no such thing as a perfect AI. Models are probabilistic. Systems are complex. Dependencies shift constantly. AI doesn't live in a vacuum–it's woven into your infrastructure, data flows, and human decision-making. And when one piece falters, the ripple effects can hit fast and hard.

Warning sign

Your incident response strategy doesn't account for AI-specific failure modes. Or worse, no one knows who's responsible when something breaks.

Building your recovery muscle

AI is everywhere and always evolving. You don't just need AI models. You need a new kind of operational muscle–one that's built for AI. That means:

- **Real-time detection** that spots anomalies across AI models, data inputs, and agent actions
- **Automated response workflows** that bring the right people (and bots) into the loop instantly
- **Continuous learning** from every AI-related incident, feeding insights back into the systems, processes, and teams.

You can't prevent eavery AI failure, but you can reduce the blast radius, recover faster, and come out stronger on the other side.

Action item

Audit your AI stack for points of failure–across infrastructure, integrations, and model outputs. If an agent goes rogue or a model silently degrades, what's your response plan?

PagerDuty POV

PagerDuty has spent over 15 years helping teams prepare for the unknown. Our platform was built to detect the earliest signals of disruption, coordinate a fast and intelligent response, and turn every incident into a moment of learning. That's not just table stakes–it's how you build confidence in a future where AI isn't just a tool, but a critical part of how your business runs.

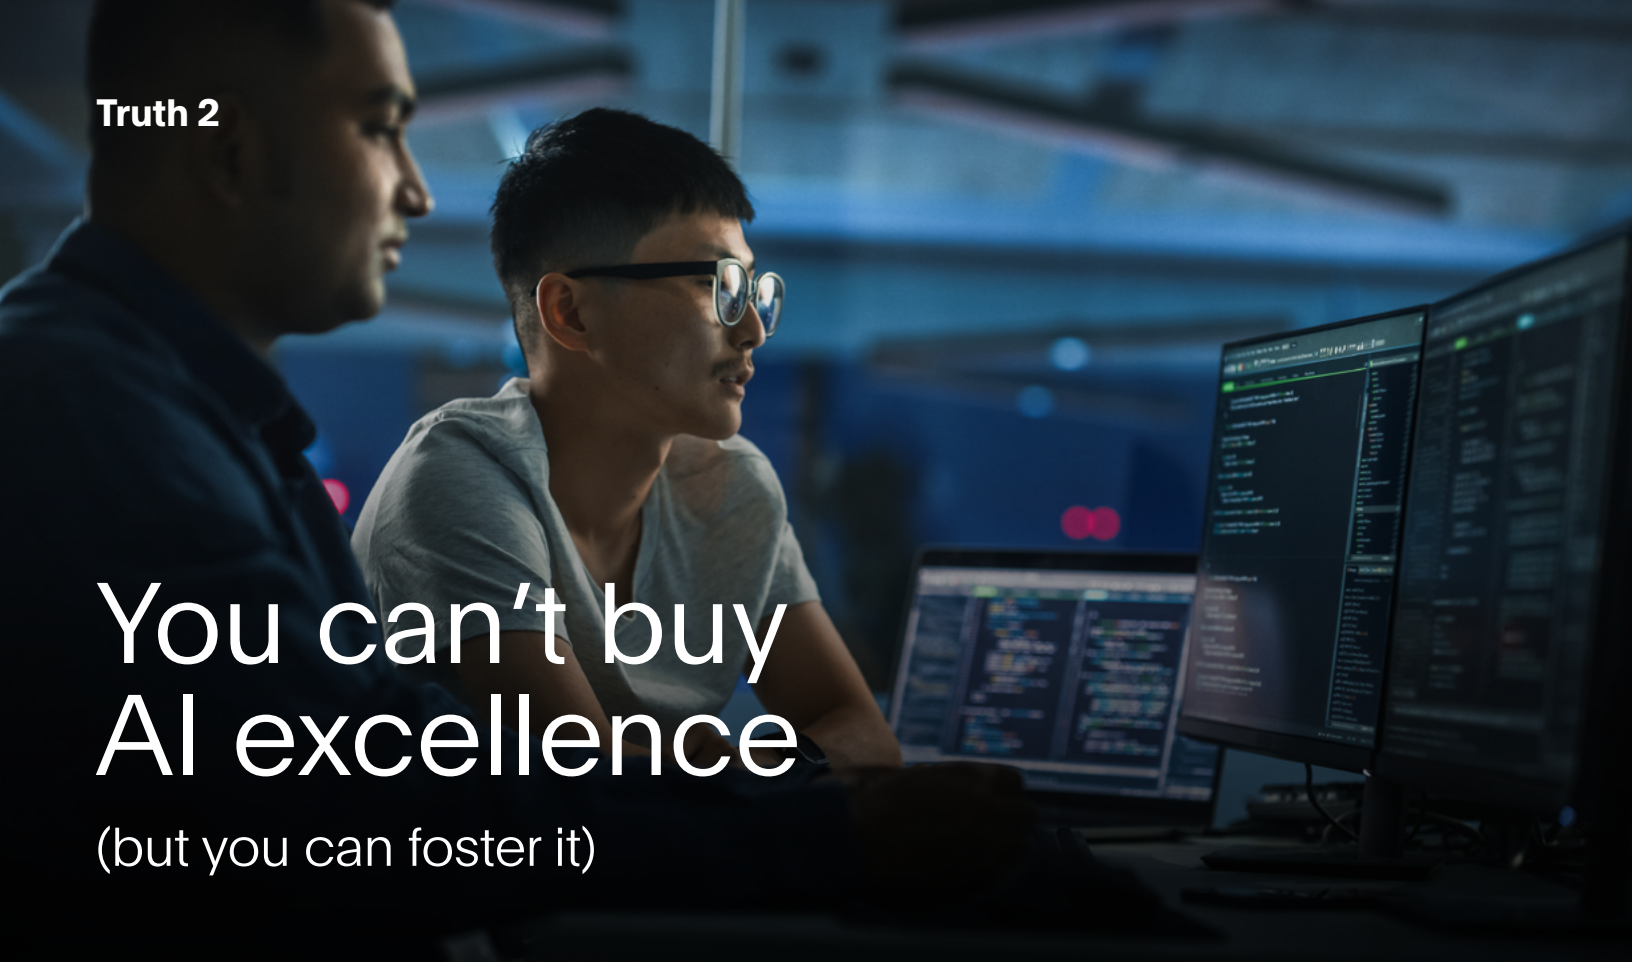

Truth 2

You can't buy AI excellence

(but you can foster it)

Myth

"We'll just buy a few tools and plug them in."

Truth

AI isn't a product—it's a practice. Long-term success requires the right skills, systems, and structure to support it.

The belief that you can buy your way to AI excellence is one of the biggest risks organizations face right now. It leads to bloated tool stacks, shallow integrations, and massive technical debt. What looks like a shortcut is often just a detour to even more complexity.

"Bolt-on" AI fails because it's reactive, not foundational. It doesn't adapt to your workflows, people, or data. It creates silos instead of solving them. And when it breaks, it leaves your teams scrambling to fix systems they never truly understood.

Truth 2

Buying AI is easy. Making it work is the hard part

AI that's slapped on instead of built in leads to:

- Shadow operations and inconsistent workflows
- Fragmented ownership and unclear accountability
- Surprise outages and untraceable errors
- Ongoing maintenance of hard-to-manage systems that increase risk

You can buy AI, but that doesn't mean it fits into how your teams actually work or delivers results you business can use. Tools alone won't drive impact if they don't integrate with your systems, support your people, or align to your goals.

Warning sign

You've bought multiple AI-powered tools, but your team still struggles to show measurable outcomes or collaborate effectively across teams.

The real advantage: Building your AI foundation

Leading organizations don't just deploy AI and cherry-pick a feature or two to use–they build fluency across teams and learn from every success and failure. They invest in:

- **Internal competencies:** AI Ops practices, incident response integration, governance models
- **AI literacy:** Upskilling teams to understand, monitor, and improve AI workflows
- **Maturity models:** Continuous assessing how ready they are to scale AI responsibility

It's not about doing it all at once–it's about maturing over time, with intention. The difference between "AI project" and "AI transformation" is depth. Real success comes from embedding AI into your operational fabric, using it to make decisions faster, resolve issues proactively, and free up your people for higher-value work.

Action item

Take inventory of your AI efforts. Are they helping your teams move faster–or creating more work and confusion? Where can you invest in internal capability-building instead of chasing features?

PagerDuty POV

We've seen what happens when AI is just another bolt-on. And we've seen what happens when it's treated like a core capability. PagerDuty helps organizations operationalize AI from the ground up–through embedded AI across the incident lifecycle, generative features that support human decision-making, and agents that deploy automation and adapt to your dynamic ecosystem.

Truth 3

Your governance is probably killing innovation

Myth

"Tighter controls mean less risk."

Truth

Inflexible governance slows innovation and increases the risk of shadow AI. Modern AI needs adaptive guardrails that scale with speed and complexity.

Governance is supposed to keep your organization safe. But for many companies trying to scale AI, it's doing the opposite–stalling progress, blocking experimentation, and creating a culture of fear around failure. The problem isn't that governance exists. The problem is that old governance frameworks weren't built for AI.

Truth 3

Over-restriction is its own kind of risk

Traditional IT governance models are designed for predictable systems, linear processes, and change that happens in quarterly release cycles–not in real time. But AI doesn't wait for executive sign-off. It evolves dynamically. It learns from new data. It behaves differently in production than it did in testing.

- When you try to apply static rules to a dynamic system, two things happen:
- Teams slow down or stop innovating entirely
- Risk doesn't go away–it just hides in shadow AI projects no one wants to talk about

The outcome is a false sense of security at the cost of actual progress.

Ironically, some of the most risk-averse orgs are the ones introducing the most risk. When guardrails become roadblocks, teams go around them–launching shadow AI projects with zero visibility or oversight.

Warning sign

Experimentation is happening outside sanctioned tools–or not happening at all. Teams are unclear on what's allowed or who approves it.

Build guardrails, not walls

The goal isn't to remove governance–it's to evolve it. Resilient AI-first organizations create governance models that are:

- **Adaptive:** Able to respond to change without starting from scratch
- **Risk-aware:** Focused on impact, not fear
- **Enabling:** Designed to empower experimentation with safety nets, not roadblocks

This means monitoring for anomalies and giving teams room to explore. It means tracking accountability without requiring five signatures to ship a workflow. And it means building policy into the platform so it scales with you–not against you.

Action item

Review your current AI governance practices. Are they enabling innovation–or adding friction? Look for places where automation, monitoring, and policy-as-code can lighten the load.

PagerDuty POV

We've helped some of the world's most complex organizations shift from rigid controls to flexible risk-aware frameworks. As part of PagerDuty's continuing commitment to our customers and continually working to earn our customer's trust in us as a vendor, we have developed our key current guidelines for the safe and secure use of generative AI throughout the product life cycle, from design to development to test and ultimately to providing to our customers. We want our customers to understand how we approach the use of generative Artificial Intelligence technologies and how we are working to advance the responsible use of AI and generative AI within PagerDuty and within PagerDuty's products and services. PagerDuty's guidelines for the safe and secure use of generative AI are based on our core values, including our commitments to championing the customer, and our robust Information Security Policies, Procedures, and Standards and are intended to ensure the secure, safe, and ethical use of AI within PagerDuty and within PagerDuty products.

Truth 4

Your AI is only as strong as your weakest integration

Myth

"Once the model works, we're done."

Truth

The model is just one part of the equation. Resilient AI depends on strong, observable, and fault-tolerant integrations.

AI doesn't operate in isolation. It relies on inputs from dozens of systems—data pipelines, APIs, observability tools, third-party models, workflow automation engines, and more. Every one of those touchpoints is a potential failure point. And if even one of them breaks, your AI breaks with it.

Truth 4

The integration imperative

AI systems are uniquely sensitive to disruptions. Slight changes in input data, service delays, or dropped handoffs can trigger bad decisions, missed escalations, or complete system outages. The more integrated your AI stack becomes, the more fragile it gets—unless you're actively designing for resilience. And the truth is that most teams don't realize there's a problem until after it causes damage.

As AI systems evolve, so does the web of dependencies supporting them. A GenAI assistant might depend on a real-time data stream, which depends on a third-party connector, which depends on cloud infrastructure that's been updated without notice. This isn't just complexity—it's cascading risk.

Common weak links:

- Data pipelines with poor monitoring or no fallback
- Overreliance on fragile APIs or vendor endpoints
- Lack of observability between services
- Missed signals due to siloed incident management tooling

Your AI is only as strong as your most brittle dependency—and often, that dependency is invisible until it breaks.

Warning sign

You're still reacting to integration failures manually and only after they've caused customer-facing issues.

Build resilient AI ecosystems

Operational excellence in AI doesn't mean zero incidents—it means zero avoidable surprises. If you knew it was coming, you should've been able to prevent it. That's why the goal isn't perfection—it's managing risk with the right safeguards in place. To get there, organizations need:

- **End-to-end observability** across every AI integration point
- **Automated response workflows** when key services fail
- **Proactive monitoring** that detects anomalies at the edges of your stack
- **System-aware AI agents** that act with context, not blind automation

Resilient AI isn't just about the model—it's about the mesh of systems that surround it.

Action item

Audit your AI pipeline for weak links. What happens if a data feed fails? An endpoint times out? Do you have detection, fallback, and response in place?

PagerDuty POV

PagerDuty's platform is built to surface weak signals before they become loud failures. With deep integrations across your stack, automated runbooks, and intelligent routing, we help you identify and resolve integration issues in real time—long before your customers feel the impact. Because if your AI can't recover from a broken link, it's not ready for production.

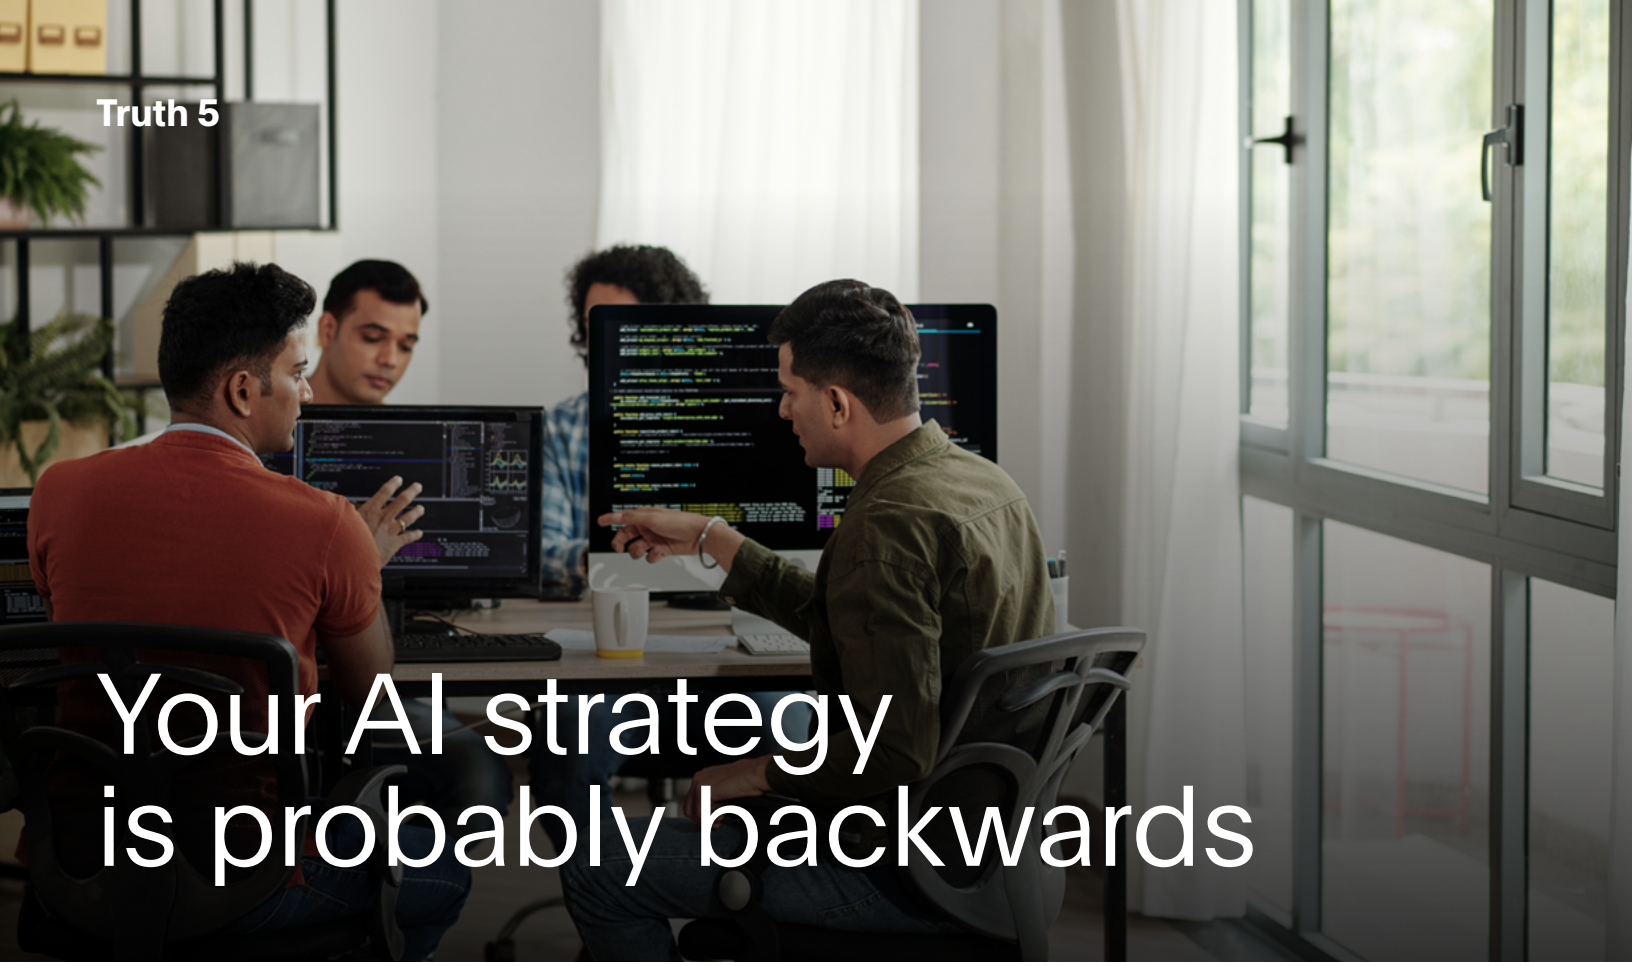

Truth 5

Your AI strategy is probably backwards

Myth

"We'll adopt AI, then figure out how to use it."

Truth

You don't retrofit strategy after implementation. AI excellence starts with clarity, not curiosity.

If your AI strategy starts with a tool and ends with "we'll figure it out," you're not alone. Too many organizations rush to implement AI because they feel behind. This results in a patchwork of disconnected tools, unclear ownership, and no real value to show for it. It's a strategy built on FOMO, not foresight.

Truth 5

The cost of building from hype

We've seen this before. In 2022, the GenAI boom triggered a wave of rushed deployments, "AI-first" initiatives, and experimentation without guardrails. Everyone wanted in—fast. But many organizations fell into the same trap: chasing AI without a strategy. The result: tools that didn't talk to each other, models no one trusted, and roadmaps that looked impressive on paper but delivered little actual value.

Chasing AI for the sake of AI leads to:

- Wasteful spending on tools that don't integrate or scale
- Confusion around priorities, ownership, and success metrics
- Tech debt that's hard to unwind once business-critical systems depend on it

Without a clear outcome in mind, every implementation feels important—but none of it feels strategic.

Warning sign

You're measuring AI success by the number of models deployed not by the impact on business outcomes or operations.

Flip the paradigm

The most successful organizations don't start with capabilities. They start with asking tough questions like:

- What problems are we trying to solve?
- Where will AI create leverage?
- How do we balance speed with risk, innovation with governance?

That's how you build sustainable AI, by aligning strategy, resourcing, and operations from the start.A strong AI strategy is intentional. It ties directly to business outcomes, and it evolves as those outcomes shift. That means measuring:

- Time-to-value from AI deployments
- Operational maturity (not just model accuracy)
- Risk-adjusted velocity (how fast you move without breaking things)
- Team fluency and readiness—not just tech adoption

AI isn't a side hustle. It's a transformation. And transformations need roadmaps.

Action item

Pressure-test your AI strategy. Is it aligned to outcomes? Do you have the right team, tech, and timeline to get there—without burning out or introducing unnecessary risk?

PagerDuty POV

PagerDuty helps organizations make the shift from reactive AI experiments to forward-looking operational strategy. We provide the governance frameworks, accountability, and automation to help you align investments with outcomes—so AI drives your business, not the other way around. Because if your AI strategy doesn't start with the end in mind, it's not a strategy—it's a gamble.

Truth 6

The fastest learners will win

Myth

"We'll scale AI once we've figured everything out."

Truth

The real leaders are scaling while learning—and pulling ahead because of it.

If your AI strategy starts with a tool and ends with "we'll figure it out," you're not alone.Too many organizations rush to implement AI because they feel behind. This results in a patchwork of disconnected tools, unclear ownea In an AI-powered business, success isn't just about how fast you move—it's about how fast you learn. The companies that are winning right now aren't the ones with the flashiest models or biggest budgets. They're the ones that turn every failure into insight, every incident into improvement, and every deployment into a more resilient system.

The future of AI belongs to the operationally self-aware.

Truth 6

Resilience is a mindset—and a moat

We've seen what happens when organizations chase AI without strategy. The GenAI boom led to rushed rollouts, ungoverned agents, and tools that looked impressive—until they broke. The companies pulling ahead now are doing something different:

- They're using every incident to make their AI stack smarter.
- They're embedding resilience into people, processes, and platforms.
- They're learning faster—and making it harder for others to catch up.

Every incident holds a clue—whether it's a hallucinated output, a missed escalation, or a misfired automation. The key is recognizing the pattern, not just resolving the symptom.

This kind of learning isn't just technical—it's cultural. The strongest orgs build teams, workflows, and tooling that make continuous improvement a shared habit, not an afterthought.

When incidents become data points, and data points become insight, you shift from firefighting to foresight. And that's where competitive advantage compounds.

Warning sign

Your team treats AI incidents as one-offs—with no clear owner, no feedback loop, and no shared learnings to build from.

From incident to improvement

Resilient organizations:

- Detect AI issues early
- Respond faster with confidence and coordination
- Learn from failure without blame
- Scale AI safely across teams and use cases
- Keep improving—because they've built feedback loops that actually work

If you're still in "wait and see" mode, you're already falling behind. The gap between experimenters and operators is getting wider by the day. But it's not too late—if you start now.

Action item

Start building resilience now. Use PagerDuty to detect early signals, assign clear ownership, and create learning loops that strengthen your operations over time.

PagerDuty POV

PagerDuty empowers organizations to build AI-powered operations with resilience at the core. From real-time detection to automated response and post-incident learning, we help teams turn friction into fuel. With over a billion real-world incidents informing our platform, we give you the operational intelligence to adapt fast—and lead with confidence.

Because AI that doesn't teach you anything is just a cost center. AI that helps you learn? That's a competitive moat.

How PagerDuty powers resilient AI operations

he PagerDuty Operations Platform is AI-powered—built on a foundation of operational intelligence from billions of real-world incidents. When your business is on the line, you can't afford to gamble on unproven solutions. We offer a proven path to AI-powered operations that help you scale with confidence and control.

PagerDuty's AI-first operations helps you:

- **Accelerate velocity** by automating routine tasks and freeing up teams for higher-value work
- **Reduce cost** by streamlining processes and scaling more efficiently
- **Mitigate risk** by detecting and resolving issues before they escalate

What sets PagerDuty apart:

In a market flooded with one-size-fits-all solutions, this is what sets PagerDuty apart with four key advantages:

- **Deep data and domain expertise:** AI is only as good as the data it's built on. Ours is trained on years of real-world incident data—giving us the operational context others can't match.
- **Enterprise-grade guardrails:** Built-in governance controls minimize hallucinations and reduce risk—so you can deploy AI confidently without compromising compliance or control.
- **Fast time to value:** Our AI works out of the box. No infrastructure changes, no vendor lock-in—just instant impact using the best model for each use case.
- **Ecosystem-first design:** With 750+ integrations and secure agent interoperability, PagerDuty fits into your stack without disruption. No rip-and-replace required.

We know AI transformation isn't about chasing trends—it's a strategic decision to position your organization for the future while delivering value today. PagerDuty brings both the technology and operational experience to help you succeed.

PagerDuty isn't just built for AI—it's built for the operational realities of modern enterprise. We know critical work, and we know what it takes to modernize with confidence. When AI moves fast, we help you move smart. Learn more about PagerDuty AI here.

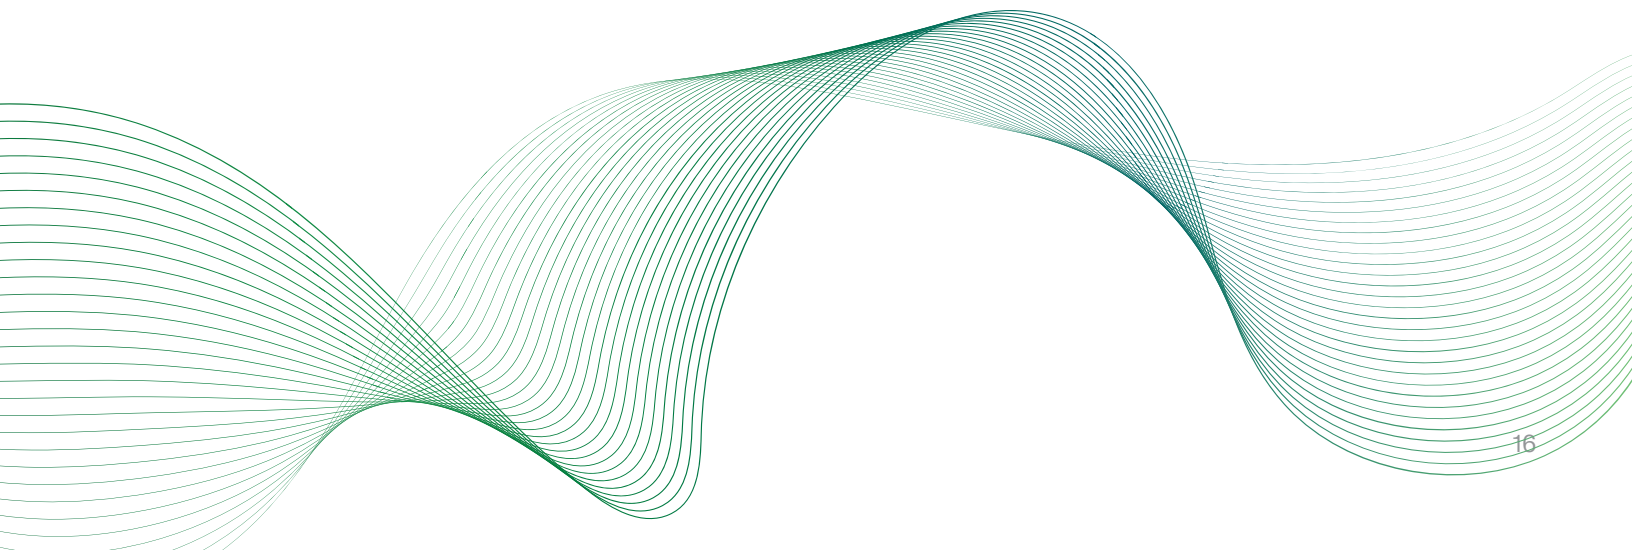

About PagerDuty

PagerDuty, Inc. (NYSE:PD) is a global leader in digital operations management, enabling customers to achieve operational efficiency at scale with the PagerDuty Operations Cloud. The PagerDuty Operations Cloud combines AIOps, Automation, Customer Service Operations and Incident Management with a powerful generative AI assistant to create a flexible, resilient, and scalable platform to increase innovation velocity, grow revenue, reduce cost, and mitigate the risk of operational failure. Half of the Fortune 500 and nearly 70% of the Fortune 100 rely on PagerDuty as essential infrastructure for the modern enterprise. To learn more and try PagerDuty for free, visit www.pagerduty.com.

PagerDuty