



PagerDuty

# **Crossing the Operational Chasm:**

## **Moving From Complexity to Clarity**

March 2025

# Contents

From information overload to collaborative intelligence . . . . .	3
The evolution of AI in operations: Moving from noise reduction to strategic collaboration . . . . .	5
Agentic AI in practice: A four-pillar framework . . . . .	7
Elevating strategic work . . . . .	10
Agentic AI implementation roadmap . . . . .	11
The autonomous enterprise . . . . .	13

# From information overload to collaborative intelligence

Throughout history, groundbreaking innovations have reshaped industries. Electricity enabled the rise of factories, communication, and modern infrastructure. Assembly lines transformed car manufacturing. Computers ushered in the shift from analog to digital. These advancements didn't just improve efficiency—they redefined what was possible.

Fast forward to today. Cloud computing, subscription software, and artificial intelligence (AI) have once again reshaped how we work and innovate. These technologies promise smarter workflows, faster decisions, and seamless collaboration. Yet, for many organizations, the reality has been far less transformative. According to [McKinsey](#), while 90% of companies have initiated digital transformation efforts, they have only achieved one-third of the anticipated revenue gains.

We're doing so much to make technology work. Why doesn't it feel like it's working as hard as it should for us?

The truth is that today's business challenges are more complex, global, and real-time than ever before. In an effort to scale and innovate faster, organizations have added more tools—but this doesn't always lead to better outcomes. Instead, they generate more data points and signals, making it harder to cut through the noise.

Generative AI has incredible capabilities, but it has added to the noise as well. AI now generates content, recommendations, and notifications at machine speed—far outpacing humans' capacity to process information meaningfully. In 2020, [global data volume](#) stood at 64 zettabytes. By 2028, that number is projected to reach 394 zettabytes, with generative AI as a major contributor.

The result? More time spent triaging. Less time spent innovating.

## The operational chasm

This is the reality of modern digital operations. On one side: the promise of technology—faster innovation, seamless workflows, and AI-driven efficiency. On the other: the reality—siloed systems, overwhelming data, and teams stretched too thin to focus on what matters most.

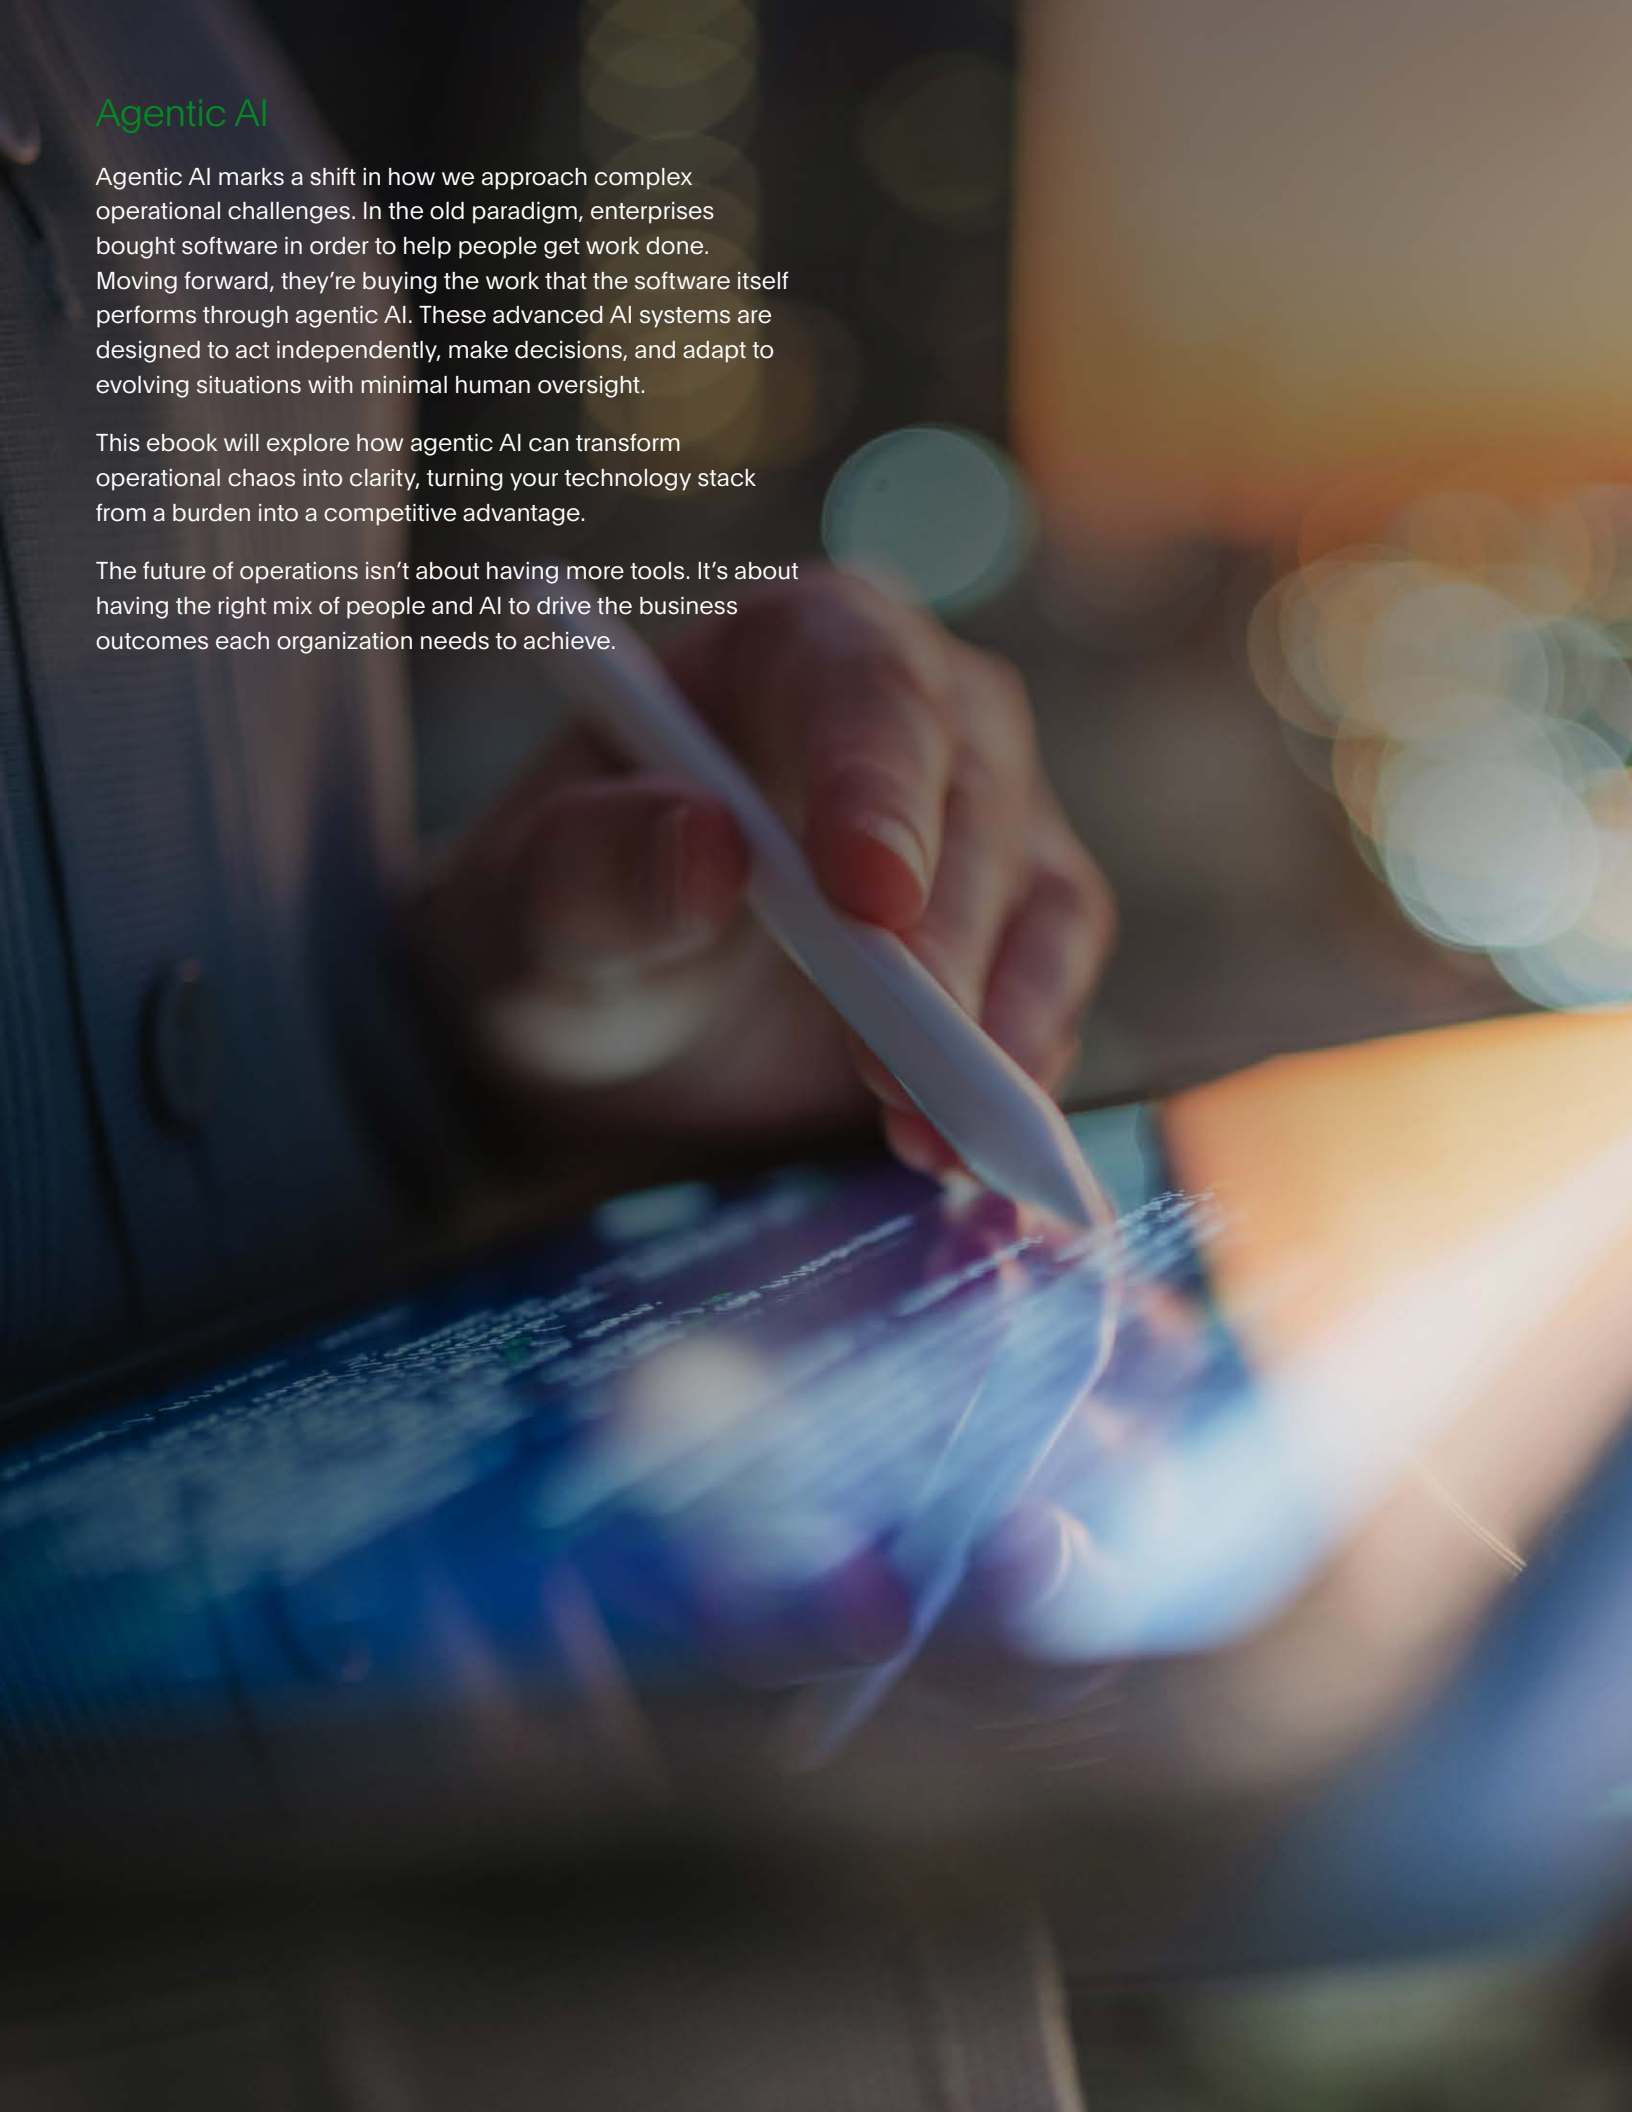
Bridging this gap requires more than just another tool. It requires a fundamental shift in how organizations approach technology. A move toward AI that acts as a true strategic partner rather than just another source of data.

## Agentic AI

Agentic AI marks a shift in how we approach complex operational challenges. In the old paradigm, enterprises bought software in order to help people get work done. Moving forward, they're buying the work that the software itself performs through agentic AI. These advanced AI systems are designed to act independently, make decisions, and adapt to evolving situations with minimal human oversight.

This ebook will explore how agentic AI can transform operational chaos into clarity, turning your technology stack from a burden into a competitive advantage.

The future of operations isn't about having more tools. It's about having the right mix of people and AI to drive the business outcomes each organization needs to achieve.



# The evolution of AI in operations: Moving from noise reduction to strategic collaboration

Imagine your company experiences a critical cloud infrastructure failure that threatens to disrupt operations for thousands of customers. But you have a system that has already detected anomalous patterns in your cloud infrastructure, identified a previously unknown misconfiguration, implemented a temporary workaround, and notified the operations team all before the morning stand-up.

This isn't science fiction. It's the emerging reality of operations management for enterprises that have embraced the shift from **reactive to proactive AI operations**. And it highlights a fundamental transformation occurring across industries: the evolution from noise-filtering AIOps to truly collaborative agentic AI systems.


## The limitations of first-generation AIOps

Traditional **AIOps** platforms emerged as an answer to a pressing problem: the overwhelming volume of notifications and monitoring data produced by modern digital systems. As enterprise IT environments grew more complex, operations teams found themselves drowning in signals—what has been termed “**alert fatigue**.”

Companies built increasingly sophisticated monitoring systems over the last decade. But that sophistication created its own problem—more data than any human team could meaningfully process. We have better visibility into our systems than ever before, but paradoxically, this makes it harder to identify what actually matters.

First-generation AIOps tools addressed this challenge through correlation and noise reduction. They aggregated similar notifications and helped prioritize issues. While valuable, these systems ultimately functioned as sophisticated filters rather than autonomous collaborators.

This reactive paradigm left organizations in a perpetual state of response rather than prevention.



## Next generation of AIOps: Agentic AI

Agentic AI represents a fundamental shift in this operational environment. Rather than simply processing and presenting information, these systems possess three critical capabilities that their predecessors lacked:



### Autonomous decision-making

The ability to evaluate options and select appropriate actions without human guidance.



### End-to-end execution

The capacity to implement decisions across multiple systems and workflows.



### Continuous learning

The capability to improve decision quality based on outcomes and changing conditions.

While traditional AIOps help humans make better decisions, agentic AI makes and executes decisions independently—consulting humans only when necessary. For example, an agentic AI system monitoring cloud infrastructure doesn't just detect a potential database performance degradation; it automatically adjusts resource allocation, modifies query parameters, and optimizes performance before users experience any impact.

**Gartner predicts** that 33% of enterprise software applications will incorporate agentic AI by 2028 (up from less than 1% in 2024). As this evolution continues, the fundamental question for business leaders isn't whether to embrace agentic AI in their operations but how quickly they can make the transition. In an age where operational excellence increasingly determines market leadership, the shift from reactive to proactive AI operations may be the most consequential transformation of the digital era.





# Agentic AI in practice: A four-pillar framework

Adopting agentic AI isn't just about integrating new technology—it requires a fundamental shift in how organizations structure data, processes, and decision-making. To move beyond basic automation and realize AI-driven autonomy, enterprises must lay a strategic foundation that ensures reliability, adaptability, and efficiency.

Based on emerging best practices from early adopters, we propose a four-pillar framework for effectively implementing agentic AI.

## PILLAR 1

### Unify your data— The foundation for agentic action

At the heart of any successful agentic AI implementation is a unified data ecosystem. This foundation allows AI agents to interact seamlessly with various tools and environments, ensuring they can execute tasks and receive information effectively.

Many organizations struggle with fragmented data silos—customer data lives in one platform, operational data lives in another, and historical records are scattered across various repositories. Others have created what amount to “**islands of automation**”—well-meaning efforts that may have some short-term impacts but do not drive long-term efficiency improvements. These structures create blind spots that limit AI's ability to execute decisions autonomously. Without seamless access to structured, real-time data, AI agents cannot perform complex reasoning, automate workflows, or collaborate with human teams effectively.

#### Key actions for leaders:

- Build a federated data architecture that allows AI agents to pull in the relevant information in real time.
- Standardize APIs and data exchange protocols to ensure AI can interact across ERP, CRM, and cloud environments.
- Adopt real-time data streaming to improve situational awareness and reduce latency in AI-driven decision-making.

## PILLAR 2

### Context-aware diagnosis— LLMs as the organizational subject matter expert

Agentic AI systems must understand not just what is happening but why it matters. Unlike traditional automation, which follows predefined rules, agentic systems leverage foundation models trained on vast, unstructured datasets, enabling them to adapt to novel scenarios.

A prime example is financial fraud detection. Traditional rule-based systems rely on static parameters, such as flagging transactions above a certain amount. But fraudsters constantly evolve their tactics, rendering static rules ineffective. An AI-powered fraud detection system equipped with context-aware reasoning can analyze historical fraud cases, detect emerging patterns, and dynamically adjust risk thresholds.

To make this possible, organizations must first transform their operational knowledge into a living history. This involves capturing organizational processes and subject matter expertise that serve as the foundation for AI learning. This gives AI the required context to dynamically contribute to tasks.

## PILLAR 3

### Autonomous or assisted resolution— AI-powered workflows

The true power of agentic AI lies in execution. While first-generation AIOps surfaced insights, agentic AI actively resolves issues, managing workflows through reasoning, past experience, and tool integration.

Consider a cybersecurity operations center dealing with real-time threats. When an anomaly is detected, traditional systems generate notifications, requiring security analysts to investigate and respond manually, often leading to notification fatigue. An agentic AI system, however, can analyze historical attack patterns, correlate threats across multiple systems, and autonomously deploy countermeasures before a breach escalates.

By filtering out false positives and prioritizing only the most critical threats, AI agents reduce notification fatigue and enable operators to focus their expertise on the most complex incidents—where human insight is most valuable.

#### Key actions for leaders:

- Train AI models on company-specific historical incidents, decision logs, and best practices.
- Develop context-aware AI agents that can correlate new challenges with past resolutions.
- Leverage natural language interfaces to allow teams to query AI for insights.

Positioning AI as an organizational subject matter expert (SME) empowers agents to diagnose and respond to issues with deep institutional knowledge—reducing redundant errors and accelerating resolution times.

#### How it works:

1. The AI agent processes the issue and formulates a resolution workflow.
2. A manager agent breaks down the task and assigns it to specialized subagents.
3. Subagents collaborate using domain expertise and organizational data, ensuring a precise and coordinated response.
4. AI dynamically requests user input where needed and refines its actions based on feedback.

#### Key actions for leaders:

- Implement AI-driven runbooks that enable agents to resolve routine issues autonomously.
- Define human-in-the-loop criteria for handling novel or high-risk scenarios.



#### PILLAR 4

## Continuous learning— Intelligent operations evolution

Agentic AI is built to continuously learn. As agents interact with the real world, they refine their models, workflows, and decision protocols—ensuring ongoing relevance and resilience.

This self-learning capability ensures organizations stay ahead of emerging threats rather than constantly playing catch-up.

#### Key actions for leaders:

- Implement real-time AI performance monitoring to detect accuracy drift. As data and threats evolve, AI models can become less reliable. Continuous monitoring ensures models stay accurate and effective over time.
- Establish AI-to-human feedback loops for validation and refinement. Human oversight helps AI learn from real-world decisions, improving accuracy and preventing errors from compounding.
- Apply bias-correction mechanisms to prevent unintended overcorrections. AI models can unintentionally amplify biases as they learn. Proactive bias detection ensures fair, consistent, and reliable decision-making.

Agentic AI is still in its early stages, and the journey toward full autonomy is complex. However, by anchoring implementation in these four foundational pillars, organizations can build resilient, adaptive, and intelligent AI ecosystems.



# Elevating strategic work

To move from reactive to proactive AIOps, rethinking your operations toolset is critical. PagerDuty Operations Cloud is designed to help enterprises solve their most high-impact, mission-critical issues using AI and automation.

## The leading operations ecosystem for agentic AI in operations

PagerDuty's comprehensive ecosystem unifies disparate tools, reducing "tech sprawl" and creating an environment for agentic AI to operate effectively. This integration of 700+ tools allows for real-time connection of information and workflows needed for AI agents to analyze patterns, automate repetitive tasks, and make informed decisions in real time. By eliminating data silos and ensuring smooth interoperability, PagerDuty creates a foundation where agentic AI can deliver faster incident resolution, smarter operational insights, and optimized resource allocation—all of which drive operational efficiency and business value.

## PagerDuty AI agents in the operations cloud

In the face of mounting technical debt and an overwhelming deluge of data from tech tools, the new PagerDuty AI agents play a crucial role in enhancing operational efficiency. These AI agents, introduced as part of PagerDuty's agentic AI functionality, enable organizations to automate repetitive tasks, resolve incidents more swiftly, and ensure seamless customer experiences.

1

### Agentic site reliability engineer (SRE)

This agent autonomously identifies, classifies, and contextualizes operational issues, offering actionable recommendations that shorten resolution times and mitigate business risks.

For example, in the event of a network outage, the agent can correlate notifications from multiple monitoring systems, determine the root cause, and provide the on-call engineer with the exact configuration changes needed, reducing time to resolve from 45 minutes to 5 minutes.

By handling routine incidents, agentic SRE frees up human engineers to tackle more complex operational challenges.

2

### Agentic operations analyst

Moving beyond immediate incident resolution, this agent does cross-tool data analysis to transform operational data into insights that drive process improvements, enabling teams to optimize performance and align operations with broader business objectives.

Picture an AI agent correlating customer support tickets with system performance metrics. It might reveal that even minor latency spikes (50-100ms) in the checkout API trigger a 15% rise in abandoned carts, pinpointing a critical area for performance optimization. This empowers teams to prioritize efforts with the greatest business impact.

3

### Agentic scheduler

Optimizing on-call scheduling is critical to minimize downtime and ensure continuous coverage. This agent dynamically manages staffing rotations, reducing costs while enhancing customer impact by ensuring that the right resources are in place at the right time.

For example, if you're managing a global team across multiple time zones, this agent can arrange rotations so no engineer faces on-call duty during their nighttime hours more than once a month while still maintaining 24/7 coverage and evenly distributing workload across the team.

PagerDuty's AI functionality is embedded across the PagerDuty Operations Cloud, ensuring that all of our customers can take advantage of agentic AI.

# Agentic AI implementation roadmap

Transitioning from legacy AIOps to autonomous AI agents requires a structured and phased approach. This roadmap provides a clear, actionable framework to help organizations integrate agentic AI effectively.

## Phase 1

### Ecosystem unification

Weeks 1-8

#### Focus

Establish a unified data ecosystem by integrating existing systems, ensuring AI agents have real-time access to the information necessary for effective decision-making. This phase sets the stage for automation by breaking down silos and establishing a single source of truth.

#### Actionable steps

- **Optimize data pipelines for AI readiness** – Streamline data ingestion, storage, and retrieval processes to ensure AI agents can efficiently access and analyze relevant operational data in real time.
- **Leverage PagerDuty's integration hub** – Utilize 700+ pre-built connectors to integrate disparate systems rapidly, creating a seamless flow of operational data across tools.
- **Establish a performance baseline** – Measure current mean time to repair (MTTR) and service level agreement (SLA) performance to establish a clear benchmark. These metrics will serve as the yardstick for evaluating agentic AI's impact post-transition, providing concrete evidence of improvement and justifying further investment.

## Phase 2

### AI agent deployment

Weeks 9-16

#### Focus

Introduce AI agents to automate critical operational workflows, enhance incident response, and drive data-driven decision-making.

#### Actionable steps

- **Implement business impact scoring** – Configure AI models to assess and prioritize incidents based on their potential business impact, ensuring that critical issues receive immediate attention.
- **Roll out process automation for L1 incidents:** Deploy AI agents to handle Level 1 incident resolution, such as basic alerts or routine resets, demonstrating immediate value by cutting response times and freeing human engineers from repetitive tasks. This "quick win" showcases agentic AI's potential and boosts adoption.
- **Monitor AI performance and fine-tune models** – Continuously track AI-driven incident resolution times, accuracy, and overall impact on key operational metrics, refining models as needed to improve efficiency.

## Continuous learning and optimization

Weeks 17+

### Focus

Develop ongoing learning mechanisms to refine AI models, expand automation capabilities, and drive continuous operational improvements.

### Actionable steps

- **Enhance AI feedback mechanisms** – Establish real-time feedback loops where human operators can validate or adjust AI-driven recommendations, improving AI learning and adaptability.
- **Enable auto-generated postmortems** – Automate incident analysis and knowledge capture, allowing AI to learn from past events and refine future responses. This reduces manual effort in root-cause analysis and accelerates incident resolution.
- **Expand AI agent responsibilities** – Gradually introduce AI into higher-level decision-making processes, such as predictive risk mitigation, capacity planning, and workflow optimization.
- **Continuously evolve AI models** – Use real-world performance data, incident trends, and business priorities to update AI models, ensuring they remain effective and aligned with organizational needs.

From week 17 onward, the system becomes self-improving, with AI agents refining their performance through data-driven feedback loops and teams uncovering new ways to harness their capabilities.




# The autonomous enterprise

In an era where organizations face constant pressure to deliver seamless customer experiences, maintain uptime, and adapt to rapidly evolving demands, operational improvement and resilience are no longer optional—they are critical imperatives for success. For decades, organizations have added tools to address persistent operational challenges, but these efforts have often fallen short of delivering true agility and efficiency.

Agentic AI represents a fundamentally different approach. By transitioning from information to action with virtual coworkers capable of completing complex workflows, this technology paves the way for a new era of efficiency and innovation. Through intelligent automation and seamless collaboration, organizations can shift from a reactive, firefighting mode to a proactive, strategic stance.

Organizations that embrace this vision today will be leaders tomorrow. PagerDuty is here to guide you every step of the way. Together, we can redefine what's possible in operations, transforming challenges into opportunities and paving the way for a future of unparalleled agility, resilience, and innovation.

The autonomous enterprise is within reach. Are you ready to lead the way?



## 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](https://www.pagerduty.com).