

Agentic AI: The strategic imperative for modern enterprises

A whitepaper exploring intelligent automation for business success


Unlocking enterprise transformation

The enterprise technology landscape stands at an unprecedented inflection point. While artificial intelligence adoption has reached mainstream status—with 78 percent of organizations using AI in at least one business function—a paradox has emerged that challenges conventional wisdom about AI's transformative potential. Nearly eight in ten companies report using gen AI; yet just as many report no significant bottom-line impact, according to McKinsey's latest research.¹ The gap between AI investment and business results is driving a move toward autonomous, goal-driven agentic AI.

Unlike traditional AI systems that respond to specific queries or follow predetermined workflows, agentic AI represents a paradigm shift toward autonomous agents capable of perceiving context, making decisions, and executing complex tasks with minimal human intervention. According to Futurum Research, agent-based AI will drive up to \$6 trillion in economic value by 2028, accelerating AI's role in automating enterprise workflows.² This transformation is not merely incremental.

It represents a fundamental reimagining of how organizations operate, make decisions, and deliver value to stakeholders.

The momentum behind agentic AI is undeniable. Gartner estimates that by 2028, agentic AI will be integrated into 33 percent of enterprise software applications, automating 15 percent of day-to-day work decisions.³ A rapid adoption timeline underscores both the transformative potential and the need for organizations to develop implementation frameworks.



***By 2028, 15% of
day-to-day work
decisions will be
made autonomously.***

¹ McKinsey & Company

² The Rise of Agentic AI: The Leading Solutions Transforming Enterprise Workflows in 2025

³ Agentic AI – Ongoing coverage of its impact on the enterprise

According to IBM's research, industry leaders have identified a critical juncture where autonomous agents transition from experimental technology to operational necessity.

Microsoft envisions "an open agentic web, where AI agents make decisions and perform tasks on behalf of users or organizations," illustrating the

scope of this transformation. Organizations positioned strategically within this ecosystem will gain operational efficiencies and establish sustainable competitive advantages as the technology matures. The question facing enterprise leaders is not whether to embrace agentic AI, but how to architect the transformation journey, maximizing value while mitigating inherent risk.



Companies that position themselves strategically with **agentic AI** will not only capture immediate operational efficiencies but also establish competitive advantages.

Part 1

Executive summary



The AI revolution is here

Agentic AI represents the next evolutionary leap in artificial intelligence, moving beyond reactive systems to autonomous agents capable of independent reasoning, planning, and execution. Unlike traditional automation that follows predefined rules or conventional AI that responds to prompts, agentic AI systems demonstrate genuine autonomy, perceiving their environment, setting objectives, and adapting their strategies to achieve complex business goals without constant human oversight.

This technological shift arrives at a critical moment for enterprise leaders. While AI adoption has reached mainstream status across industries, the persistent gap between investment and measurable business impact has created urgency

for more sophisticated approaches. Agentic AI addresses this challenge by embedding intelligence directly into operational workflows, enabling autonomous decision-making that scales across enterprise functions from service delivery to strategic planning.

- **80% of common incidents could be resolved autonomously.**
- **MTTR reduction of 60-90% in service desk operations.**
- **One customer achieves a 65% deflection rate within 6 months of implementation.**

The business case for Agentic AI extends far beyond operational efficiency as organizations implementing autonomous agents report transformational outcomes. Up to 80 percent of common incidents could be resolved autonomously, with a reduction in time to resolution of 60 to 90 percent for service desk operations, while one customer achieves a 65 percent deflection rate within six months of implementation (and projections of 80 percent by the end of the year). These metrics represent not merely incremental improvements but fundamental shifts in how enterprises deliver value to customers and stakeholders.

Distinguishing agentic AI from traditional approaches like automation

The distinction between agentic AI and traditional automation extends beyond technical capabilities to encompass strategic implications for organizational design and competitive positioning. Traditional automation systems, while valuable for standardizing repetitive tasks, operate within fixed parameters and require explicit programming for each scenario. Though more sophisticated, conventional AI systems remain fundamentally reactive as they process inputs and generate outputs based on training data and human-defined objectives.

Agentic AI transcends these limitations through autonomous goal setting, dynamic planning, and adaptive execution. These systems continuously assess the environment, identify optimization opportunities, and modify their approaches

based on real-time feedback. Where traditional automation might handle routine password resets, an agentic AI system proactively monitors system performance, predicts potential issues, and implements preventive measures while optimizing resource allocation across the entire IT infrastructure.

This autonomous capability fundamentally alters the value proposition for enterprise technology investments. Rather than replacing human workers with automated processes, Agentic AI creates intelligent partnerships where autonomous agents handle complex operational tasks while human professionals focus on strategic initiatives, creative problem-solving, and relationship management. This shift represents an opportunity for organizations to improve operational efficiency and enhance their human capital's strategic contribution.

The implications extend to organizational agility and competitive responsiveness. Traditional systems require significant lead times to adapt to market changes or operational requirements. Agentic AI systems continuously learn and adapt, enabling real-time optimization and rapid response to changing business conditions. This capability becomes particularly valuable in dynamic markets where competitive advantage increasingly depends on operational agility and customer responsiveness.

The table below illustrates the key differences between agentic AI and traditional AI.

Feature	Traditional Automation	Conventional AI	Agentic AI
Task Execution	Rule-based, static	Predictive, data-driven	Autonomous, adaptive
Human Involvement	High	Moderate	Minimal to collaborative
Learning Capability	None	Supervised/unsupervised	Continual, contextual
Decision-Making	Predefined	Assisted	Independent, goal-oriented
Integration	Siloed	Augmented	Embedded across workflows

Part 2

Deep dive into real-world scenarios



Agentic AI isn't a futuristic concept; it is already reshaping core enterprise functions. Here are real-world use cases where autonomous agents are delivering a tangible business impact:

Service Desk

The service desk is often a bottleneck for productivity, with a high volume of repetitive Tier-1 requests. An agentic AI can serve as a "Tier-0" resolution engine.

- **Current state:** A user submits a ticket for a password reset. A chatbot provides a link to an instruction page, but the user is still stuck. A human agent then takes over and manually resets the password.
- **Agentic future:** An agentic AI receives the same request. It identifies the user and the system, accesses the necessary tool with an API call, and autonomously initiates the password reset. If the action fails, the agent reflects on the error (e.g., "account is locked"), communicates the reason to the user, and escalates the ticket to the appropriate team with the necessary context. This reduces the Mean Time to Resolution (MTTR) and frees human agents to focus on complex, high-value issues.

SOC (Security Operations Center)

Security analysts are overwhelmed by alert fatigue, with thousands of daily alerts, many of which are false positives.

- **Current state:** An analyst sees an alert for unusual login activity. They manually investigate by cross-referencing log files, checking network traffic, and looking up threat intelligence databases. This process is slow, reactive, and prone to human error.
- **Agentic future:** An autonomous security agent ingests the alert. It then proactively begins an investigation. It pulls user login history from the identity provider, checks for concurrent logins

from other locations, queries threat intelligence for the IP address's reputation, and isolates the affected system if a threat is confirmed. It then generates a comprehensive report, summarizes the findings in plain language, and even suggests or initiates a remediation plan, such as blocking the IP or forcing a password change.

NOC (Network Operations Center)

NOC teams face constant pressure to maintain network uptime and resolve performance issues quickly.

- **Current state:** A network monitoring tool flags a performance degradation on a server. A human NOC engineer must log in, run diagnostics, and manually attempt to resolve the issue. If the problem is complex, they may need to involve other teams, causing delays.
- **Agentic future:** An agentic AI, after detecting the same performance issue, initiates an autonomous troubleshooting workflow. It checks CPU utilization, memory, and disk space. If it identifies a resource bottleneck, it can autonomously scale up the resources or clear unnecessary caches. If a root cause is found (e.g., a misconfigured service), the agent can attempt to revert the change or apply a known fix. Only if the agent's pre-defined resolution paths fail does it escalate the ticket, providing the human engineer with a complete diagnostic log and a summary of the attempted resolutions, drastically reducing MTTR.

Field Services

Field service operations are a complex mix of scheduling, logistics, and on-site problem-solving. An agentic AI can transform this function from a reactive system to a proactive, self-optimizing one.

- **Current state:** A field technician is dispatched based on a static schedule. Upon arrival, they may find the required part is missing or the problem is more complex than anticipated. They must then make calls, consult manuals, and manually order parts, leading to delays and repeat visits.
- **Agentic future:** An autonomous agent receives a service ticket. It proactively assesses the problem, identifies the needed parts, and checks the technician's van inventory or the nearest warehouse in real time. It then optimizes the technician's route to include a stop at the warehouse if necessary. During the service call, the agent acts as a digital copilot, accessing the full knowledge base of technical manuals and past repair logs to provide real-time guidance to the technician. If an unexpected issue arises, the agent can autonomously order the correct part for next-day delivery and reschedule the follow-up visit, all without human intervention. This leads to a higher first-time fix rate and a more efficient use of human resources.



Cloud Ops

Managing complex cloud environments is a constant battle against performance bottlenecks, cost overruns, and security vulnerabilities. Cloud Ops teams are often overwhelmed by the sheer volume of data and alerts.

- **Current state:** A cloud engineer gets an alert that a specific application server is experiencing high latency. They must manually log into the dashboard, check CPU and memory usage, review log files for errors, and determine if they need to scale the server. This is a time-consuming, reactive process.
- **Agentic future:** An agentic cloud ops platform receives the same alert. An autonomous agent, designed for performance management, immediately begins an investigation. It queries the application's performance metrics, correlates them with recent code deployments or traffic spikes, and identifies the root cause. The agent can then autonomously execute a remediation plan, such as scaling up the virtual machine, optimizing a database query, or deploying a patch. The system works with other agents, such as a cost-optimization agent, to ensure the proposed solution is cost-effective. The result is a self-healing infrastructure where issues are resolved before they can impact end-users, freeing engineers to focus on strategic cloud architecture.

Security Ops

In security operations, the volume of threats and the speed of attacks outpace human capabilities. A Security Operations Center (SOC) is often a reactive function, with analysts triaging an endless stream of alerts.

- **Current state:** A security analyst receives an alert about a potential malware file on a corporate endpoint. They must manually pull the file, detonate it in a sandbox, analyze the network traffic it generates, and then write a report. This can take hours, during which the threat may spread.
- **Agentic future:** An autonomous security agent receives the same alert. It immediately initiates a multi-step investigation: it isolates the endpoint from the network to prevent further spread. Then, it extracts and analyzes the file in a secure virtual environment and queries threat intelligence databases for known indicators of compromise. The agent then correlates its findings with other data sources, such as user behavior analytics and network logs, to build a comprehensive picture of the attack. It can then autonomously block the malicious IP address at the firewall and roll back the system to a clean state. The human analyst is only notified when the issue is fully resolved, or if the agent requires a human decision on a particularly complex or novel threat. This proactive approach reduces the mean time to detect and respond to threats from hours to minutes.



Finance

In finance, autonomous agents can streamline complex, multi-stage processes that require data validation and cross-system communication.

- **Current state:** An invoice arrives. A human must manually enter the data, cross-reference it against a purchase order, get approval, and then process the payment in a separate system.
- **Agentic future:** A financial agent ingests the invoice document, extracts key data points, and automatically validates the information against the purchase order and vendor master data in the ERP system. It then initiates the appropriate approval workflow based on the amount and vendor and schedules the payment upon approval. The agent can also continuously monitor for fraud by flagging unusual invoice amounts or new vendor accounts for human review.

Human Resources

HR teams are burdened with routine administrative tasks, from onboarding to benefits questions.

- **Current state:** A new employee joins the company. An HR representative must manually send out welcome emails, create accounts in different systems, and enroll the employee in training programs.
- **Agentic future:** An HR agent can autonomously orchestrate the entire onboarding process. Upon a hiring decision, it can trigger the creation of a new employee profile, set up accounts in various IT systems, schedule introductory meetings, and automatically enroll the new hire in mandatory training modules. The agent can also serve as an “HR assistant” for employees, autonomously answering policy questions or helping them navigate benefits enrollment, only escalating to a human HR representative when the query is nuanced or sensitive.



Part 3

Adoption and change management



Strategic implications for organizational structure and workflow design

Implementing agentic AI systems profoundly impacts organizational structure, workflow design, and talent management strategies. Organizations must reconsider traditional hierarchical models designed for human decision-making when autonomous agents can independently execute complex operational tasks.

The integration of autonomous agents requires new organizational models that optimize human-agent collaboration while maintaining strategic oversight and accountability. Successful implementations establish clear governance frameworks defining agent autonomy boundaries, escalation procedures, and performance measurement criteria. These frameworks ensure autonomous agents enhance rather than replace human expertise while maintaining organizational control over strategic objectives.

Workflow redesign becomes essential as organizations integrate agentic AI capabilities. Traditional linear processes must evolve into dynamic and adaptive workflows where autonomous agents continuously optimize execution paths based on real-time conditions. This transformation requires a comprehensive process analysis to identify opportunities for autonomous execution while preserving critical human touchpoints for strategic decision-making and stakeholder relationship management.

The talent implications are equally significant. Organizations implementing agentic AI must invest in workforce development programs that prepare employees for strategic roles requiring creativity, emotional intelligence, and complex problem-solving. This shift represents an opportunity to enhance job satisfaction and employee value by focusing human capabilities on activities that leverage uniquely human strengths.

Part 4

A forward-looking perspective



The future enterprise operating model

The trajectory of agentic AI development points toward enterprise operating models that fundamentally differ from current organizational structures. Future enterprises will operate as ecosystems of autonomous agents and human professionals collaborating to achieve strategic objectives through adaptive, self-optimizing systems.

At its zenith, it could become a fully autonomous digital workforce capable of understanding and executing complex business objectives. This vision extends beyond operational automation to encompass strategic planning, market analysis, and innovation management through autonomous agents capable of complex reasoning and decision-making.

The competitive implications are profound. Organizations implementing comprehensive agentic AI strategies will achieve sustainable advantages through superior operational efficiency, enhanced customer experiences, and accelerated innovation cycles. These advantages compound over time as autonomous agents continuously learn and optimize, increasing returns on AI investments.

However, this future requires strategic planning and investment in foundational capabilities. Organizations must develop robust data architectures, governance frameworks, and talent strategies to support autonomous agent deployment while maintaining strategic control and accountability. The window for establishing a competitive position in the agentic AI era is narrowing as technology capabilities accelerate and early adopters demonstrate measurable business advantages.

Putting it all together: Pomeroy's vision for enterprise transformation

The evidence is unequivocal: agentic AI represents the most significant transformation in enterprise technology since the advent of cloud computing. In addition to improving operational efficiency, organizations that embrace this shift will fundamentally redefine their competitive status in an increasingly automated global economy.

Pomeroy's leadership in this transformation stems from our deep understanding of enterprise operations combined with advanced AI capabilities. We recognize that successful agentic AI implementation requires more than technological deployment; it demands comprehensive organizational transformation that aligns autonomous capabilities with strategic business objectives.

Our approach integrates proven enterprise expertise with cutting-edge AI innovation to deliver autonomous solutions that enhance rather than replace human capabilities. We partner with organizations to design, implement, and optimize Agentic AI systems that transform operational challenges into competitive advantages while preserving the strategic value of human expertise.

The opportunity before us is unprecedented. Capable of perceiving context, agents have the potential to deliver tangible and measurable impact by boosting productivity, accelerating innovation, and creating new pathways to growth. Organizations that act decisively to implement comprehensive agentic AI strategies will establish sustainable competitive advantages that compound over time through continuous learning and optimization.

However, the complexity of this transformation demands strategic partnership with proven experts who understand both the technological possibilities and the organizational realities of enterprise change management. Pomeroy combines deep enterprise experience with advanced AI capabilities to guide organizations through successful agentic AI transformations that deliver measurable business value while building foundations for future innovation.

The future belongs to organizations that successfully integrate autonomous intelligence with human creativity and strategic thinking. Pomeroy is committed to leading this transformation, helping our partners navigate the complexities of agent-based AI implementation while capturing the unprecedented opportunities this technology creates. The question is not whether your organization will embrace agentic AI, but whether you will lead or follow in this fundamental transformation of enterprise operations.

The time for strategic action is now. [Contact Pomeroy](#) to begin your agentic AI transformation journey.