pipefy

Automation with AI Agents:
the complete guide to transforming your processes

Table of Contents

Introduction

Why this guide matters

Artificial Intelligence (AI) is quickly becoming a key driver of digital transformation across businesses. Within this landscape, AI Agents are emerging as one of the most promising innovations—capable of acting autonomously, learning from data and decisions, and scaling processes with intelligence and efficiency.

In recent years, the use of these agents has grown exponentially. Enterprise platforms are adopting AI to automate repetitive tasks, enhance analysis accuracy, and support smarter decision-making. Still, many professionals remain unsure how these tools differ from traditional automation and how to implement them strategically.

That's exactly why this guide was created. More than just a theoretical introduction, it offers practical paths to understand, apply, and extract value from AI Agents in business operations—with a focus on intelligent automation, productivity gains, and competitive edge.

What are AI Agents?

AI Agents are computational systems that perceive their environment, make decisions based on data, and act to achieve specific goals. Unlike simple scripts or linear automations, agents operate with greater autonomy and adaptability.
They fall into two broad categories:

> **Reactive Agents:** respond to immediate stimuli, like chatbots that reply to specific commands.

Agentes Proativos: anticipate needs, learn from past interactions, and take initiative —like assistants that manage calendars, flag operational risks, or optimize workflows based on historical patterns.

The distinction between AI agents and traditional automation is critical. While automations follow fixed rules, AI agents learn, adapt, and act based on goals—making them more flexible and effective for dynamic or complex scenarios.

And in recent years, the use of these agents has grown exponentially. Enterprise platforms have been incorporating AI technologies to automate repetitive tasks, increase the accuracy of analyses, and improve decision-making. A report by the Boston Consulting Group indicates that, by 2025, one in three companies globally plans to spend more than US$25 million on solutions focused on artificial intelligence, which includes AI agents."

One in three companies across all markets are planning to spend $25milion+ on AI in 2025

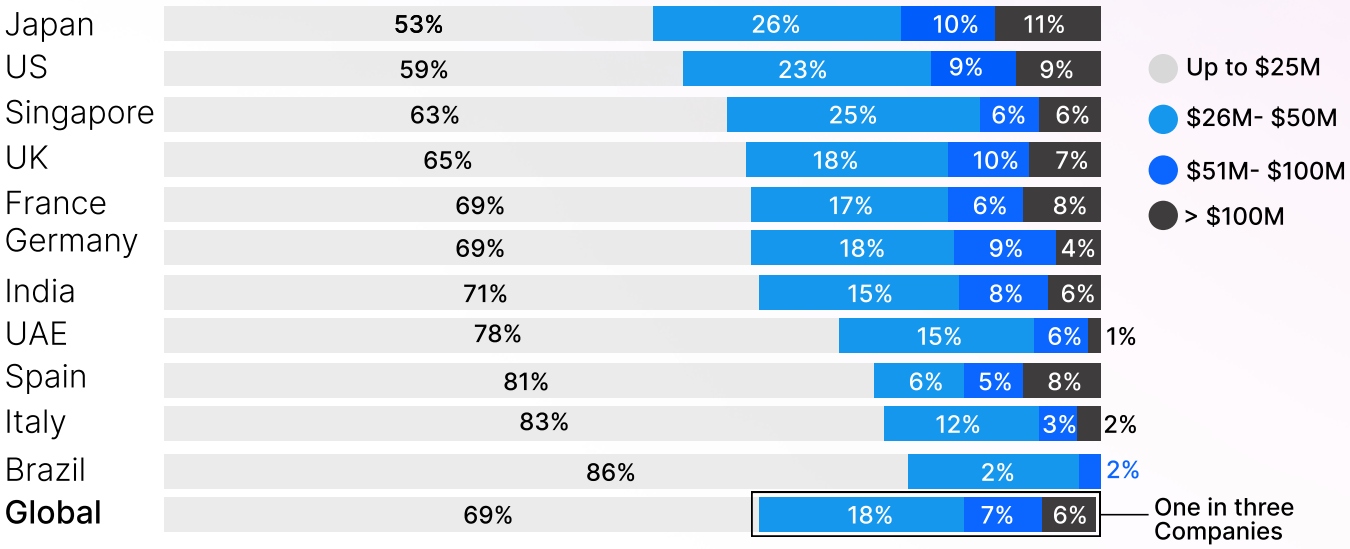

Source: BCG AI Radar Survey

Who should read this guide?

This material is designed for professionals leading innovation and operational efficiency within their organizations, including:

Operations professionals seeking to boost productivity and eliminate bottlenecks.

Team leaders interested in tools that scale results with less manual effort.

Innovation managers looking to understand how AI can drive meaningful change.

Product and tech teams exploring new ways to integrate AI into their solutions—from prototypes to final products.

If you're involved in strategic decision-making, process management, or tech development, this guide is for you.

Enjoy the read!

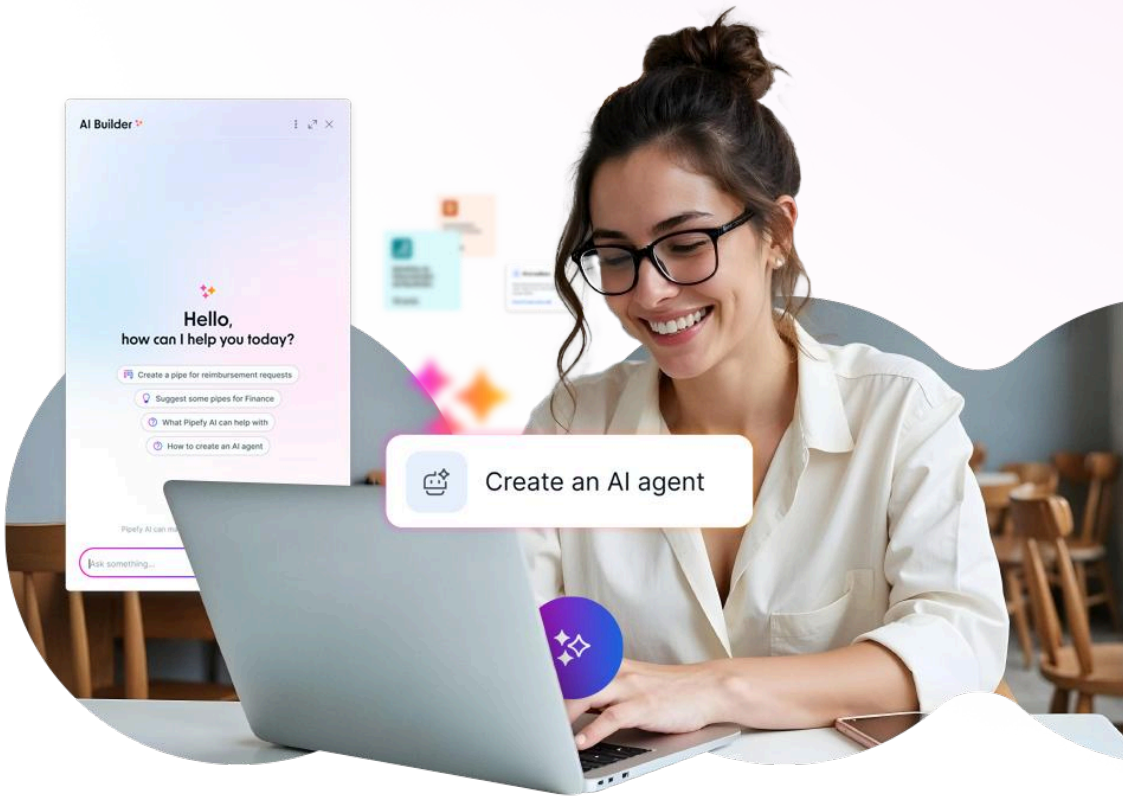

pipefy

Chapter 1

Foundations of AI Agents

Key Concepts

Differences between Generative AI, Rule-Based Automation, and AI Agents
The advancement of AI has brought a range of approaches, each with different use cases and complexity levels. To understand the role of AI agents, it's essential to distinguish three key concepts:

- **Generative AI:** Based on models like LLMs (Large Language Models), generative AI can create original content—text, images, code, even music—based on human prompts. It excels at creative tasks, natural language analysis, and automated content generation.

- **Rule-Based Automation:** These systems use fixed logic to execute specific tasks. Common in RPA (Robotic Process Automation), they handle repetitive, predictable activities like filling forms or generating reports. But they don't learn or adapt to context changes.

- **AI Agents:** These stand out for their ability to perceive, reason, and act. They are autonomous systems that make decisions based on goals and contextual data, capable of integrating with multiple systems, adapting behavior, and acting proactively. They combine the best of generative AI with adaptive decision-making.

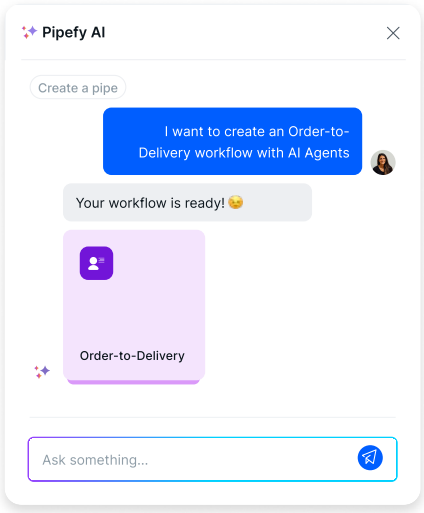

The Building Blocks of an AI Agent: Perception, Reasoning, and Action

To understand how AI agents work, it helps to think of their behavior as built from three core components:

1. **Perception:** The ability to gather and interpret environmental information—from corporate systems, user responses, documents, or operational logs.
2. **Reasoning:** The agent's "brain." Based on perception, the agent evaluates context, applies logic, leverages AI models, and determines the best course of action. This is where techniques like machine learning, generative models, and decision engines come into play.
3. **Action:** Once a decision is made, the agent executes a task—whether it's sending an email, updating a system, alerting a team, or triggering automated workflows.

This triad allows agents to go beyond task execution and behave in dynamic, context-aware ways—much like a human decision-maker.

Common Use Cases

AI Agents are already being deployed across various sectors to boost efficiency and productivity. Here are some real-world examples:

- **Customer Support:** Agents that combine generative AI with internal data to deliver contextual answers, automate support, and scale service 24/7.
- **Internal Operations:** Agents that monitor processes and workflows, anticipate bottlenecks, and handle tasks like ticket creation or SLA alerts.
- **Sales:** Agents that track leads, suggest data-driven outreach strategies, and automate follow-ups.

HR: Agents that guide new hires through onboarding, answer FAQs, and support tasks like feedback, time-off, and recruitment.

Finance: Agents that handle reconciliations, forecast cash flow, generate financial reports, and run audits based on anomaly detection.

Analyzing the expected effects of generative AI deployment by business function, a survey by McKinsey consulting firm found that respondents generally expect a reduction in headcount in service operations, such as customer service and field service, as well as in supply chain and inventory management. In IT and product development, however, respondents are more likely to expect an increase than a reduction in headcount.

These examples show how agents can serve as “digital teammates,” freeing up human bandwidth and boosting process accuracy.

Expected change in business function's of employees as a result of gen AI use, next 3 years.

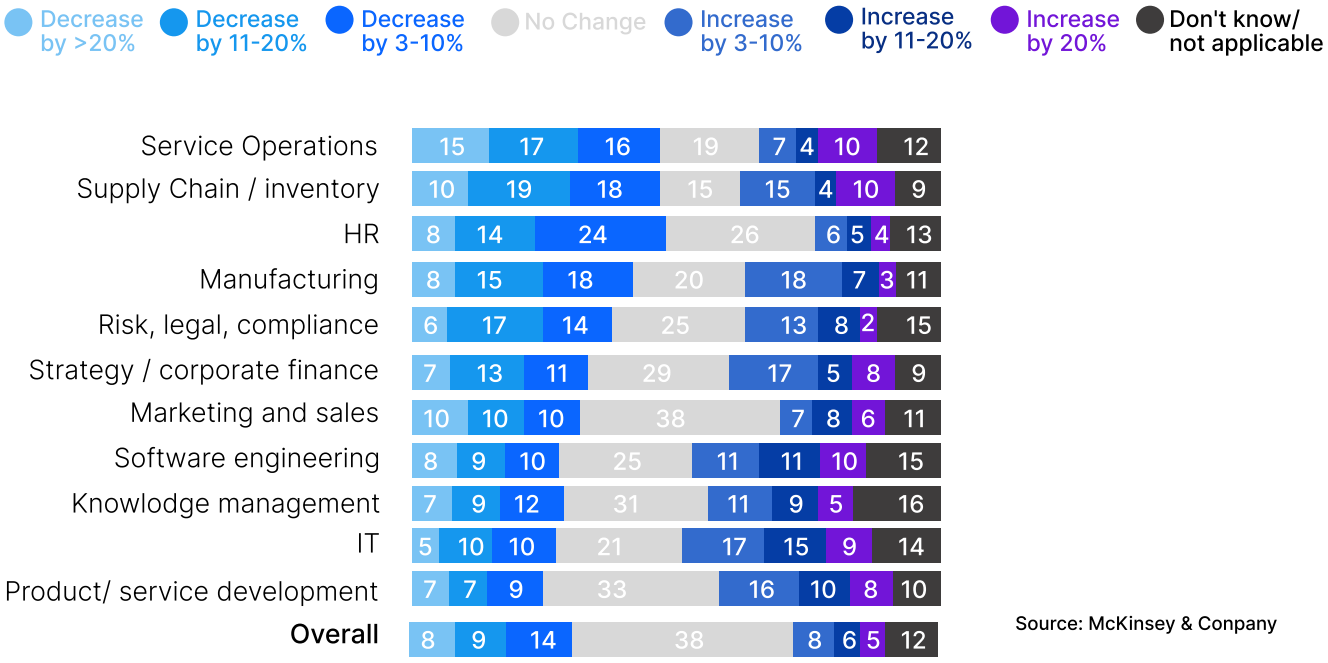

Source: McKinsey & Conpany

Why Now?

Three major factors explain why this is the right moment to embrace AI agents:

1. **Tech Advancements:** The rise of LLMs and access to advanced AI architectures now make it possible to build robust, intelligent agents capable of learning, adapting, and using natural language.

2. **Tool Accessibility:** Platforms like LangChain, AutoGen, CrewAI, and others have democratized agent development, letting businesses build custom solutions with fewer technical barriers.

3. **Market Maturity:** There are now real-world success stories, active communities, and established best practices—reducing risk, accelerating implementation, and building trust in the tech.

Today's landscape doesn't just allow for experimentation—it's primed for strategic adoption of AI agents as core assets in modern operations.
Now that you understand what AI agents are and why they represent a major leap from traditional automation, it's time to look inward. Which of your processes could benefit from agents? Where are the bottlenecks or scaling opportunities?

In the next chapter, we'll walk through how to identify these critical points, map internal workflows, and prepare the ground for a successful agent implementation—because every transformation starts with a clear diagnosis.

Chapter 2

Assessing Your Company's Processes

Successfully implementing AI agents doesn't start with technology—it starts with clarity. Before building or integrating any smart solution, you need to know where it will deliver the most value. This chapter helps you analyze your company's internal processes with a strategic lens—pinpointing real opportunities for intelligent automation.

Mapping Existing Processes

The first step in bringing AI agents into your operations is mapping out your current processes. This means documenting workflows, inputs and outputs, responsibilities, and critical steps in each activity.

How to Spot High-Potential Opportunities:

Repetitive tasks: Activities performed frequently in the same way (e.g., sending emails, system updates).

Operational bottlenecks: Steps that involve delays, rework, or heavy manual validation.

Clear decision points: Process stages where logic or inference can guide decision-making.

Multi-system interactions: Situations requiring data movement across platforms—a sweet spot for agents that integrate and automate.

Use simple flowcharts or BPM (Business Process Management) tools to clearly visualize inefficiencies and areas with improvement potential.

Criteria for Agent-Based Automation

Not every process is suited for AI agents. Some are better handled by simpler automation; others still need human judgment. To allocate resources wisely, use the following criteria:

Criterion	What to Observe
Volume	Does the process happen often enough to justify the effort?
Frequency	Does it occur daily, weekly, or sporadically?
Impact	Could automation reduce cost, speed things up, or cut errors?
Decision Complexity	Are decisions based on clear logic or nuanced judgment?
Rule Clarity	Are decisions rule-based or dependent on ambiguous context?

Use agents when:

- The process involves high contextual complexity but follows detectable patterns.
- Tasks have multiple input variables (data, language, context).
- Proactivity and adaptive learning make a strategic difference.

Diagnostic Tools and Frameworks

To support your analysis, consider structured visual tools to help prioritize what to automate first.

1. Process Evaluation Checklist

- Does the process involve repetition?
- Large data volumes or multiple steps?
- Rework or heavy people involvement?
- Direct impact on customers or key business goals?
- Multiple systems or data sources?

The more “yes” answers, the higher the automation priority for AI agents.

2. Impact vs. Effort Matrix

	Low Effort	High Effort
High Impact	Priority 1 (Go all-in with agents)	Priority 2 (Evaluate carefully)
Low Impact	Priority 3 (Use simple automation)	Avoid for now

This framework helps visualize where AI agents will deliver the greatest strategic return.

3. Practical Example

Take a reimbursement process that involves:

- Document collection → repetitive
- Criteria validation → rule-based
- Status updates in two systems → integration
- Communication with requestor → natural language

This is a prime candidate for an AI agent: it combines repetition, rules, volume, and the need for multichannel automation.
Now that you know how to spot the right processes, it's time for the next step: choosing the right type of agent and making sure it's equipped to perform effectively.

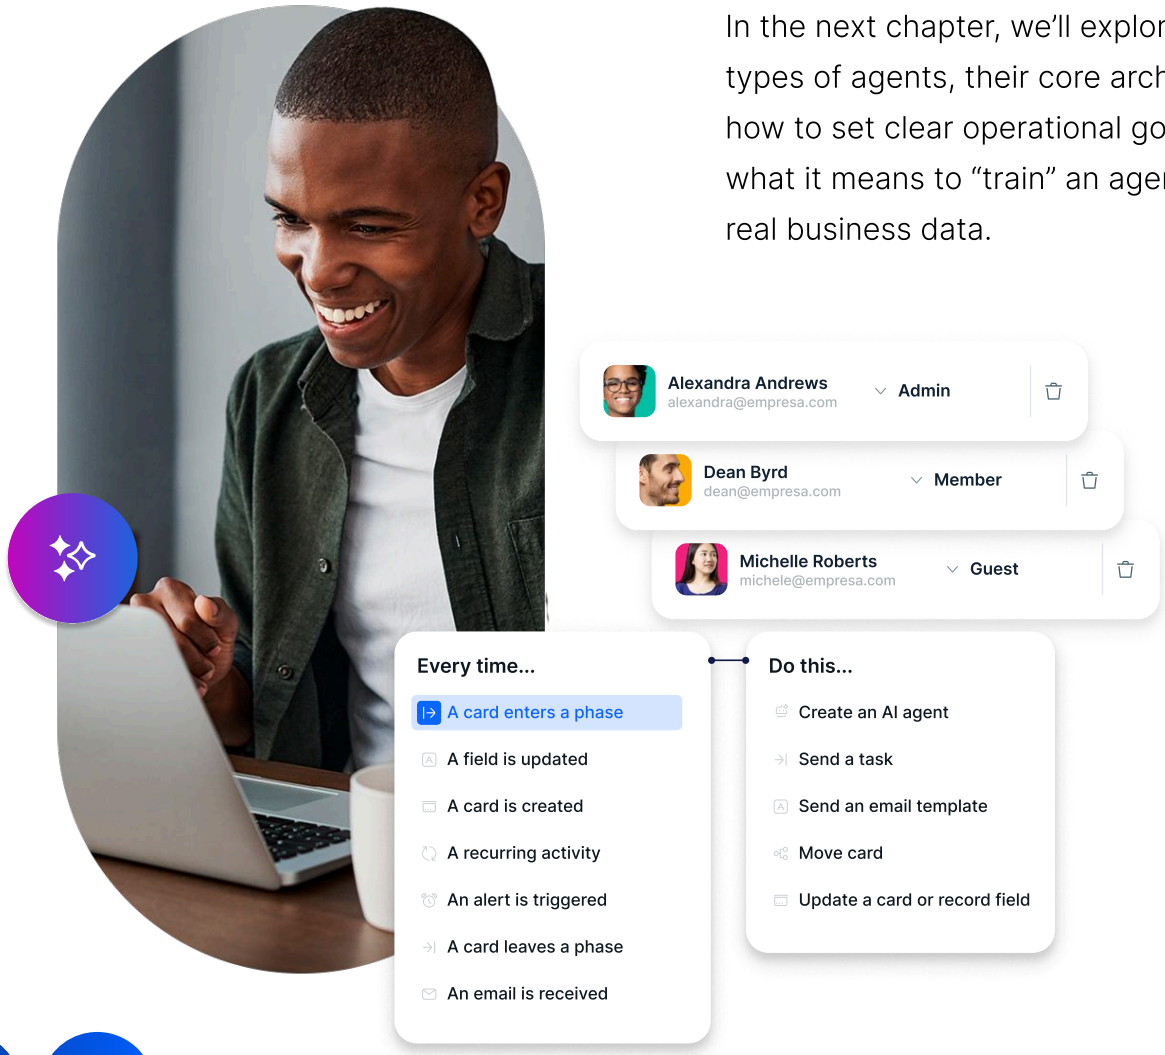

In the next chapter, we'll explore the types of agents, their core architectures, how to set clear operational goals, and what it means to "train" an agent using real business data.

Chapter 3

Choosing and Training Your Agents

With your company's processes mapped and prioritized, it's time to make more technical and strategic decisions: what type of AI agent should you use and how do you prepare it to operate effectively? This chapter introduces the main categories of available agents, the most relevant tools on the market, and best practices for instructing them successfully.

Types of Agents

AI agents are not all the same—they differ in autonomy, reasoning ability, and their role within a corporate process. Understanding these differences is essential for making the right choices.

1. Autonomous AI Agents

These act independently to achieve a goal. Once given a task, they can break it into subtasks, query data sources, make decisions, and act with little or no human supervision.

> Best use cases: Research projects, system testing, complex automation, and decisions based on multiple data sources.

2. Copilots (or Intelligent Assistants)

These work alongside humans, offering suggestions, summaries, partial automations, and decision-making support. They rely on commands or clear context and don't act fully autonomously.

> Best use cases: Highly variable processes requiring constant adjustment, like ticket classification, task allocation, or personalized experiences.

3. How to Train and Instruct an Agent

These agents learn or adjust based on results. After each action cycle, they evaluate the outcome and refine future behaviors. While "learning" here doesn't always imply machine learning, it refers to an architecture with greater adaptive intelligence.

> Best use cases: Highly variable processes requiring constant adjustment, like ticket classification, task allocation, or personalized experiences.

How to Train and Instruct an Agent

Having a powerful tool is just the beginning. The true differentiator lies in how you train and guide the agent to understand your context, interpret commands accurately, and act safely.

1. Effective Prompt Design

The foundation of communication with generative agents is the prompt—the text that tells the AI what's expected. A good prompt is:

- Clear and concise
- Contextualized with the agent's role
- Specific about the desired response format

Weak example: "Analyze this email."
Good example: "Review this sales email to make it more persuasive and suggest improvements to the CTA."

2. Defining Objectives

Before deploying the agent, define:

- The ultimate goal of its role
- The tasks it should or should not perform
- Its decision-making boundaries (e.g., human approval required, when to notify)

3.Context Injection

To help agents make smart decisions, you need to provide:

- Historical data or prior examples
- Internal documentation, policies, or manuals
- Access to relevant APIs or systems

Context transforms a generic executor into a true digital teammate with business understanding.

4. Action Control

Especially in enterprise environments, it's critical to define:

- Activity logs
- Execution permissions (e.g., read vs. write access)
- Review and approval mechanisms before critical actions

These controls ensure security, traceability, and confidence in agent operations.

Now that you know how to choose and train your agents, it's time to put that into practice. But where should you begin? Which areas will benefit most? What kind of agent fits each scenario?

In the next chapter, we'll explore concrete use cases across departments—like customer service, HR, operations, legal, and more—with actionable ideas you can tailor to your company.

Chapter 4

Practical Applications by Department

The real power of AI agents lies in their ability to act contextually across different parts of a company, adapting to each department's specific challenges and workflows. In this chapter, we highlight how agents can be applied across corporate functions with real-world, easy-to-adapt examples.

Finance

The finance department handles large volumes of data, complex analysis, and a high demand for precision and compliance. AI agents can revolutionize this space by:

- Automating bank reconciliations and statement comparisons
- Performing predictive cash flow analysis based on historical data
- Generating management reports and dashboards automatically
- Classifying and processing expenses and reimbursements by matching receipts with internal policies
- Detecting accounting anomalies or potential fraud based on irregular patterns

Human Resources

HR is one of the most promising areas for intelligent automation, especially for administrative tasks and employee communication. AI agents can:

- Automate responses to FAQs about payroll, benefits, and internal policies
- Handle digital onboarding, delivering materials, guides, and personalized training
- Track hiring process status, send updates to candidates, and schedule interviews
- Support HR teams by analyzing organizational climate and interpreting qualitative feedback
- Monitor performance review deadlines, notifying managers and consolidating updates

Marketing

With multiple channels, tight deadlines, and constant content demands, AI in marketing boosts both scale and creativity. Agents can:

- Generate campaign content variations (emails, social media, landing pages) based on simple prompts
- Analyze campaign performance and recommend data-based optimizations
- Summarize market research from multiple external sources
- Automate social media or CRM responses in brand-consistent language
- Assist with content calendar planning by prioritizing high-engagement topics

Sales

In sales, speed, personalization, and pipeline management are major competitive advantages. AI agents deliver value through:

- Copilots for SDRs and reps, suggesting outreach strategies, lead prioritization, and email drafts
- Summarizing client meetings and updating CRM automatically
- Automating prospecting using ICP (ideal customer profile) and external data
- Drafting commercial proposals and contracts using smart templates
- Forecasting deal closure probabilities based on historical interactions

Legal

Even in a highly regulated area like legal, well-trained and controlled agents bring efficiency:

- Reviewing and suggesting contract clauses based on standard templates
- Summarizing contracts and flagging risks or inconsistencies
- Automating answers to recurring internal questions, like privacy or compliance policies
- Organizing and categorizing large volumes of legal documents
- Tracking legal deadlines and alerting teams to renewals or expirations

Customer Support

In many companies, agents already lead the way in scaling customer interactions without sacrificing quality:

- Intelligent chatbots powered by generative AI and internal knowledge bases
- Agents that classify and route tickets by priority and department
- Personalized response generation based on customer history and brand tone
- Collecting and analyzing post-service feedback to suggest improvements
- Monitoring SLA performance and customer satisfaction levels automatically

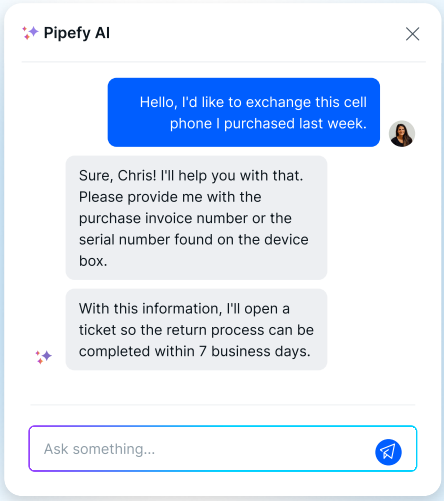

IT & Operations

In teams that support company infrastructure and efficiency, agents can be integrated to:

- Monitor systems and issue real-time alerts on failures or anomalies
- Automate responses to IT service requests (e.g., password resets, access provisioning)
- Execute system maintenance and updates on a scheduled basis
- Orchestrate operational workflows across teams and systems
- Audit logs and processes, flagging potential risks

With so many opportunities at hand, the challenge becomes: how do you ensure your agents operate efficiently, securely, and in alignment with your business goals?
In the next chapter, we'll cover key practices to make your AI agent operations scalable, controlled, and sustainable—covering security, scope control, human review, performance metrics, and usage policies, everything your company needs to mitigate risks and maximize value.

Chapter 5

Best Practices and Governance

Once AI agents are integrated into business processes, the priority becomes ensuring that they operate securely, efficiently, and in alignment with organizational guidelines. Like any strategic initiative, deploying agents requires clear governance, ongoing monitoring, and continuous adjustments.
This chapter outlines best practices to ensure agents operate ethically, are auditable, and deliver measurable returns for the organization.

These good governance practices are essential to reducing incidents involving the use of generative AI solutions. A survey conducted by the consulting firm McKinsey and published by Stanford University indicates that 59% of organizations reported between three and ten incidents involving AI programs in 2024.

Number of artificial intelligente (AI) incidents reported by organizations globally in 2024

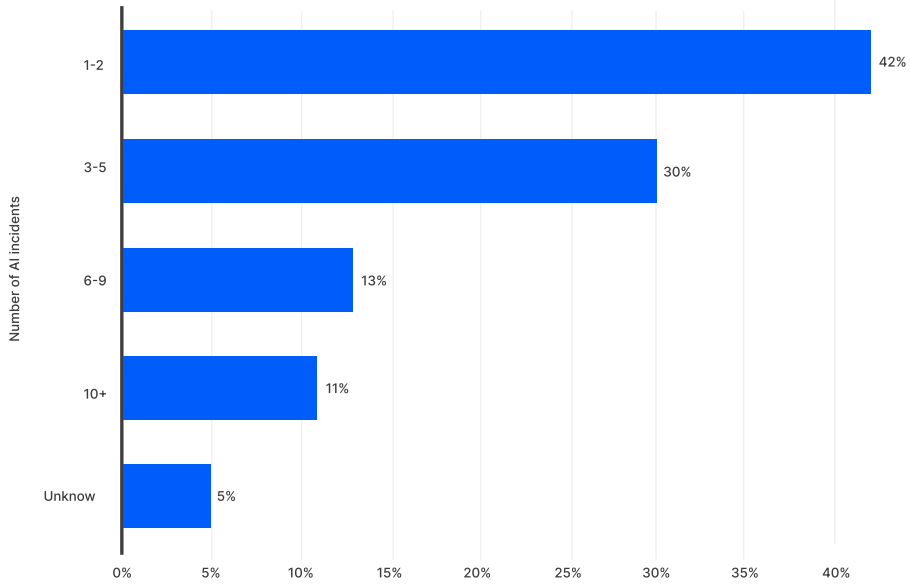

Sources
Stanford University,
McKinsey & Company

Best practices for security and compliance:

> **Scope and permission control:** Ensure agents only have access to what they need—no more. Never grant unnecessary permissions.
> **Data anonymization:** Remove or obscure personal information whenever possible when using data to train or run agents.
> **Data retention policy:** Define how long data processed by agents will be stored and under what conditions it will be deleted.
> **Regular audits:** Periodically review agent logs, decisions, and access, especially in regulated environments.
> **Informed consent:** Always inform users or customers when automated systems are being used during interactions.

Responsible AI adoption means ethics and technology must go hand in hand.

Human Oversight and Control

Even advanced agents should operate under human supervision, especially in critical processes or those that directly affect the business or customers. Governance isn't about restricting agents - it's about building a reliable monitoring and control system.

Essential mechanisms include:

> **Fallback to humans:** When agents encounter uncertainty, ambiguity, or tasks outside their scope, they should escalate the issue to a human for review.
> **Critical action review:** Create checkpoints requiring human validation before agents execute sensitive tasks—like sending contracts, legal decisions, or financial transfers.
> **Activity logs:** All agents should log their actions, decisions, inputs, and outputs. This enables audits, learning, and transparency.
> **Smart alerts:** Set up alerts for unusual behavior, outlier patterns, or operational failures.
> **Ongoing team training:** Humans need to understand how agents work so they can supervise effectively and interpret outputs accurately.

Measuring Impact

AI agent implementation must be assessed with objective metrics. Measuring results is essential for justifying investments, guiding improvements, and aligning operations with strategic goals.

Recommended KPIs:

> **Time saved:** Compare task durations before and after automation.
> **Volume of tasks automated:** Track how many tasks agents complete over time.
> **Accuracy:** Measure how often agent decisions are correct or acceptable against a defined standard.
> **Internal NPS:** Gauge satisfaction among employees who interact with or benefit from the agents.
> **Escalation rate:** How often the agent needs to call for human help—and whether this signals a need for reconfiguration.
> **Cost per task:** Determine the operational cost of running the agent vs. the previous method.

Feedback Loop:

1. Collect usage and performance data
2. Identify patterns of success and failure
3. Refine prompts, scopes, and access
4. Re-deploy and track improvement

Measuring and adjusting is a fundamental part of an AI agent's lifecycle in the enterprise. Until now, we've covered how to identify processes, choose the right agents, train them, and operate them responsibly. But what's next?

In the next chapter, we explore the trends shaping the future of AI agents—including multi-agent systems, IoT integration, fusion with RPA, and AI's growing role in strategic decision-making.
Get ready to glimpse what's ahead—and how your company can stay ahead of the curve.

Chapter 6

The Future of AI Agents

The adoption of AI agents marks just the first wave of a deeper transformation in how businesses operate. The future will bring even more sophisticated agents capable of working in networks, learning continuously, and integrating seamlessly with enterprise tech ecosystems.
This chapter highlights the trends redefining the role of agents and how your organization can prepare for this new era.

Emerging Trends

1. Multi-Agent Systems

Imagine not just one, but several agents collaborating—each with a specific role in a process. One gathers data, another analyzes it, another suggests actions, and another executes. This team-based model, inspired by human workflows, increases complexity and automation autonomy.
Example: A virtual team made up of a financial analysis agent, a legal review agent, and a communication agent working together to automatically approve contracts.

2. Agents with Memory and Continuous Learning

Agents that don't just react—but remember. Memory allows them to consider previous interactions, preferences, and decisions. Combined with ongoing learning, this results in more personalized, efficient, and adaptable agents.
Example: An HR agent that learns how each manager gives feedback and adjusts future communication accordingly.

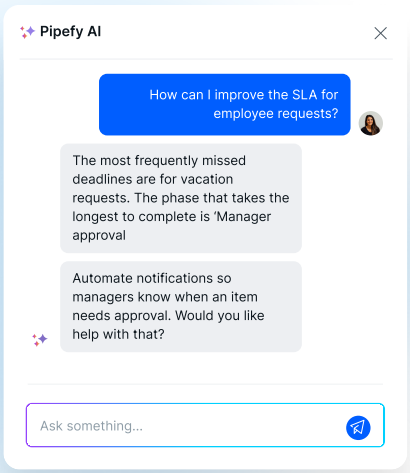

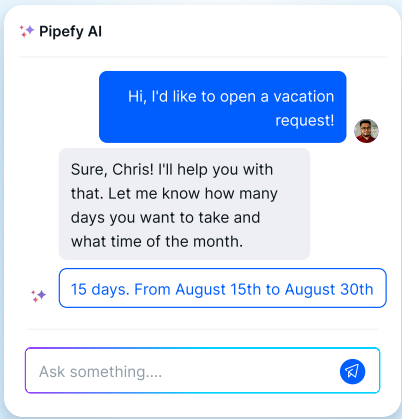

3. Integration with APIs and External Ecosystems

Future agents will become smart hubs—integrating ERPs, CRMs, legacy systems, data platforms, productivity tools, and third-party APIs. This will enable cross-system, cross-department automations.
Example: A sales agent that integrates CRM data, website behavior, and email history to automatically prioritize leads.

4. Domain-Specific AI

We'll see the rise of highly specialized agents by function (e.g., legal, finance, compliance), trained on specific datasets and regulatory frameworks. These "digital experts" will perform tasks with quality comparable to human consultants.

Preparing Your Team for the Future

The technology is ready—but is your company? To unlock the benefits of the next generation of AI agents, you'll need more than tools. You'll need cultural transformation, ongoing training, and new ways of working.

1. Culture of Experimentation

Companies that encourage smart testing and iteration will lead the pack. Create innovation squads, host internal hackathons, and encourage teams to test agents in small processes before scaling. Practical innovation starts with freedom to experiment.

2. AI Training

AI fluency will be a baseline skill in many roles. Invest in training programs for both technical and non-technical teams, focused on:

- AI and machine learning fundamentals
- Prompt design
- Interpreting automated decisions
- Ethics and responsible AI use

3. Redesigning Roles and Processes

As agents take on operational tasks, new human responsibilities emerge: curation, supervision, strategic analysis, and training the agents themselves. Prepare your teams to transition from executors to AI operators and strategists.

pipefy

Take the next step: See Pipefy's AI Agents in action

You already understand the transformative power of AI agents in business operations—now it's time to see how they work in practice at your company.
Book a personalized demo with our experts and discover how Pipefy's AI agents can automate tasks, eliminate bottlenecks, and deliver results in as little as 15 days.

Click here to schedule your free demo

Don't miss the opportunity to put innovation to work for your operations.