CloudDirect®

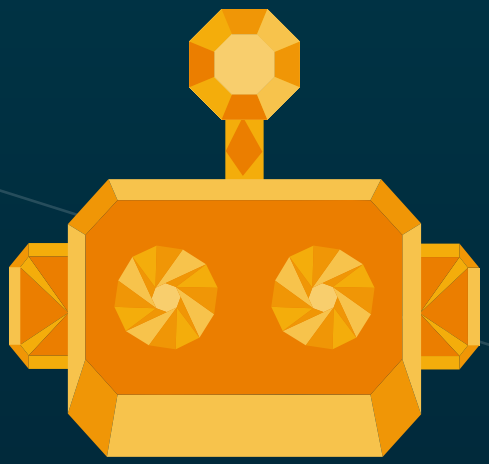

Managing Tomorrow:

A practical guide to working with AI Agents

2026

Contents

What to expect

The workforce is changings. Teams are now working alongside AI, rather than with it. This guide is designed to help you move beyond agent experimentation and into operational impact.

You'll learn:

- What agentic AI is and how it differs from existing AI tools
- Where AI agents are delivering measurable value today, across different business units
- Why so many AI agent initiative fail to scale
- What the Microsoft AI agent ecosystem looks like, and how it works
- How to successfully design, deploy, and work alongside AI agents

Most importantly, it will provide a **practical, actionable framework** that enables your organisation to adopt AI agents with confidence.

Jump to the framework

Managing Tomorrow:
The rise of AI agents

88%

of organisations are using AI in at least one business function.

McKinsey

82%

are planning to integrate AI agents by 2029.

Index

Artificial intelligence is no longer a future consideration.

It's reshaping how organisations work right now, and over the past two years tools like Copilot and other generative AIs have driven meaningful gains in individual productivity and enabled employees to move faster, reduce manual effort, and improve the quality of their work.

A more fundamental shift is now underway. AI is moving beyond assistance and into action, transitioning from tools that respond to systems that can plan, execute, and optimise independently. This evolution introduces AI agents: intelligent systems capable of operating across workflows, systems, and data to deliver outcomes with minimal human intervention.

While 88 per cent of organisations are now using AI, widespread adoption has not yet translated into widespread impact. Many organisations are still applying AI at the task level – improving isolated activities rather than rethinking how work gets done end to end. As a result, the true value of AI remains largely untapped.

What is emerging is a growing divide between organisations that are experimenting with AI, and those that are operationalising it.

AI agents represent the next stage of maturity. They enable organisations to move beyond productivity gains and towards genuine transformation—automating complex processes, coordinating actions across systems, and delivering consistent, data-driven outcomes at scale.

And the direction of travel is clear, with 82 per cent of organisations already planning to integrate AI agents within the next three years.

The organisations that succeed will not simply be those that adopt AI quickly, but those that successfully embed it into the fabric of their operations.

What Is Agentic AI?

From assistance to autonomy

Artificial intelligence is evolving rapidly – but not all AI is equal. While many organisations have already adopted AI tools to improve productivity, fewer fully understand what comes next.

The shift to agentic AI represents a step change in how technology is applied within the business.

Until now, AI has largely been used to support individual tasks—generating content, analysing data, or summarising information on demand. These capabilities have delivered clear value, but they remain inherently reactive, relying on human input to initiate and guide every action.

AI agents change this dynamic.

They move beyond responding to instructions and begin executing work independently – interpreting goals, making decisions, and operating across systems to deliver outcomes. This is what defines agentic AI, and why it is being positioned as the next phase of enterprise transformation.

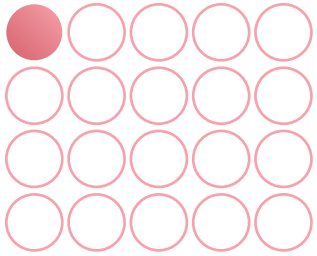

5% of enterprise applications included AI agents in 2025.

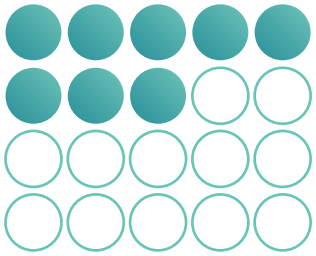

By the end of 2026, that is expected to increase to **40%**.

Understanding the shift

AI tools respond.

Traditional AI systems respond to prompts, generating outputs based on user input. They are effective, but entirely dependent on human direction.

Copilots assist.

Copilots work alongside users, streamlining tasks and improving efficiency. They reduce effort but still rely on user initiation and oversight.

Agents act.

AI agents introduce a new level of capability by interpreting objectives, planning actions, and execut multi-step processes across systems.

What makes an AI agent different?

AI agents are not a single product or tool—they are systems that combine multiple capabilities to operate effectively within a business environment.

At a high level, an AI agent is designed to:

- **Understand context** from data, systems, and business rules
- **Plan actions** based on defined goals
- **Execute tasks** across applications and workflows
- **Adapt over time**, improving performance through feedback

This enables agents to move beyond isolated tasks and operate across entire processes—something traditional AI tools cannot achieve.

Where AI agents create value

AI agents are already being applied across functions – not as isolated tools, but as embedded components within workflows.

It stands to reason that **64%** of AI agent use case focus on business process automation.

They are commonly used to:

- **Automate** customer and employee interactions
- **Orchestrate** workflows across multiple systems
- **Manage** approvals and internal processes
- **Support** data-driven decision-making in real time

Unlike traditional automation, which is rule-based and rigid, AI agents introduce flexibility and reasoning—allowing them to adapt to changing inputs and scenarios.

Microsoft's agentic ecosystem

Agentic AI is not a standalone capability – it is part of a connected platform designed for scale, governance, and integration.

Its key components include:

Copilot Studio
Build and manage AI agents

M365 Copilot
Embed agents into everyday work

Power Platform
Automate workflows and integrate systems

Azure AI Foundry
Develop and scale AI capabilities

Microsoft Fabric
Ground agents in governed, trusted data

Agent 365
Govern and manage AI agents at scale

Together, these Microsoft technologies provide the foundation for enterprise-ready AI agents that can operate securely and effectively across the business.

From productivity to transformation

The first wave of AI focused on productivity, on helping individuals work faster and more efficiently. While valuable, these gains are incremental.

AI agents shift the focus to something more significant: process transformation.

By enabling end-to-end automation and decision-making, they allow organisations to:

- **Eliminate** manual handoffs
- **Reduce** operational function
- **Deliver** more consistent outcomes at scale

But what does that look like in reality?

In the next section of this guide, we'll look at real, business unit-focused examples of how organisations are making the most of their agents.

Productivity gains of up to

40%

are being reported in certain business operations.

BCG

Agents In Action

Where impact is already being made

AI agents are no longer theoretical. Across industries, organisations are deploying them into real workflows – automating processes, reducing manual effort, and delivering measurable improvements in performance.

The shift is subtle, but significant. Rather than supporting isolated tasks, AI agents are beginning to operate within processes, coordinating activity across systems and teams to deliver outcomes with greater speed and consistency.

What this looks like in practicevaries by function, but the pattern is consistent: identify a repeatable process, embed intelligence, and remove friction.

Sales

Accelerating pipeline and qualification

Sales teams are increasingly using AI agents to manage early-stage engagement and qualification, allowing human teams to focus on high-value interactions.

What agents are doing:

- **Responding** to inbound enquiries in real time
- **Qualifying** leads based on predefined criteria
- **Scheduling** meetings and follow-ups
- **Logging** activity directly into CRM systems

The outcomes:

- **Faster** response times
- **More consistent** qualification
- **Increased** conversion rates

Marketing

Orchestrating campaigns at scale

Marketing teams are moving beyond content generation and using AI agents to coordinate full campaign execution.

What agents are doing:

- **Managing** campaign workflows across channels
- **Personalising** messaging based on customer data
- **Monitoring** performance and adjusting activity in real time
- **Automating** reporting and insight generation

The outcomes:

- **Reduced** time to launch
- **Improved** campaign performance
- **Greater** consistency across channels

Legal

Accelerating document review and risk analysis

Legal teams are using AI agents to reduce the time and effort required to review, analyse, and manage large volumes of documents.

What agents are doing:

- **Reviewing** contracts and extracting key clauses
- **Identifying** risks, obligations, and deviations from standard terms
- **Comparing** documents against policies and templates
- **Supporting** due diligence across large document sets

The outcomes:

- **Faster** document review and turnaround times
- **More consistent** identification of risks and issues
- **Reduced** manual workload for legal teams

Finance

Reducing manual processing and risk

Finance functions are leveraging AI agents to automate highly repetitive, process-driven tasks while improving accuracy and oversight.

What agents are doing:

- **Processing** invoices and reconciling transactions
- **Identifying** anomalies or inconsistencies
- **Managing** approval workflows
- **Producing** reports and forecasts

The outcomes:

- **Reduced** manual processing times
- **Improved** data accuracy
- **Greater** visibility and control

HR

Streamlining employee experiences

HR teams are using AI agents to improve internal processes and deliver more consistent employee experiences.

What agents are doing:

- **Managing** onboarding workflows
- **Responding** to employee queries
- **Coordinating** training and compliance activities
- **Supporting** documentation and policy access

The outcomes:

- **Faster** onboarding
- **Reduced** administrative burden
- **Improved** employee satisfaction

IT and Operations

Automating systems and support

Within IT, AI agents are playing a growing role in both service delivery and operational efficiency.

What agents are doing:

- **Monitoring** systems and identifying issues
- **Automating** incident response and remediation
- **Managing** service desk requests
- **Coordinating** workflows across platforms

The outcomes:

- **Reduced** downtime
- **Faster** isuse resolution
- **Lower** operational overheads

What these examples reveal

While use cases vary, with each and every business unit having its unique challenges to overcome, they all follow a clear pattern.

Focus on repeatable, high-volume processes.

Combine automation with decision-making capability.

Measure outcomes based on time, cost, and consistency improvement

Embed agents within workflows, not alongside them.

Individually these use cases deliver incremental improvements, but together they represent something more significant – collectively, we're building a digital workforce.

This is where the real opportunity lies.

But success isn't guaranteed. To move to individual use cases to scalable impact, organisations need a clear, structured approach – one that guides how AI agents are identified, built, deployed, and managed across the business.

Organisations adopting AI agents report up to

20%

cost reductions in IT and engineering functions.

McKinsey

From Idea To Impact

How to successfully adopt AI agents

The difference between organisations that experiment with AI and those that realise its full value is simple: **execution**.

AI agents are not difficult to build. But building an agent is not the objective. The objective is to design something that delivers measurable, repeatable business outcomes.

This requires a structured approach – one that connects use cases to value, embeds governance from the outset, and ensures agents are integrated into the way the business operates.

And the gap isn't capability. It's process.

The following framework outlines how to move from early ideas to operational impact. Follow the seven steps laid out below, or click to jump to a relevant section for you.

Only

11%

of organisations are running AI agents in production, despite

79%

having initially adopted them.

saasultra

Define the outcome, not the agent

Identify the right use case

Prepare your data and context

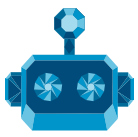

Build the agent

Deploy with governance on day one

Enable your people

Measure, refine, and scale

STEP ONE

Define the outcome, not the agent

Start with the problem, not the technology.

The most common reason AI agent initiatives fail is simple: they begin with the technology, rather than the outcome. Organisations are often drawn to the potential of AI agents and ask where they can be applied, rather than stepping back to consider what problems actually need solving.

This leads to solutions that are technically impressive but commercially insignificant.

To create meaningful impact, AI agents must be rooted in clear, measurable business objectives. That means identifying friction within existing processes—where time is lost, effort is duplicated, or outcomes are inconsistent—and focusing agent design around removing those constraints.

AI agents are not valuable because they exist. They are valuable because they solve problems.

In practice, the most effective starting points are processes that are:

- Repetitive and time-consuming
- Dependent on structured data
- Prone to delays or bottlenecks
- Already well understood within the business

This could include:

- Lead response times in sales
- Approval workflows in finance
- Query handling in customer service
- Internal reporting cycles

The key is clarity. If the outcome is not clearly defined from the outset, success cannot be measured. **Without measurable success, the initiative will not scale.**

STEP TWO

Identify the right use case

Start by thinking about high impact, low complexity.

Once a clear outcome has been defined, the next step is to identify the right use case to deliver it. This is where many organisations lose momentum—either by choosing use cases that are too complex, or by attempting to solve too many problems at once.

Successful organisations take a different approach. They prioritise use cases that allow them to demonstrate value quickly, while minimising risk and dependency.

This is not about limiting ambition—it is about sequencing it effectively.

Early AI agent deployments should focus on areas where:

- The process is already stable and understood
- The data required is accessible and reliable
- The busines impact can be clearly measured
- The risk of failure is low

By starting with focused, high-value use cases, organisations are able to:

- Build confidence internally
- Develop capabilites required to scale
- Create a clear business case for further adoption

Three in four businesses are seeing ROI in the first 12 months of deployment, and this initial success is critical. It shifts perception from experimentation to value—and creates the momentum needed to expand into more complex use cases over time.

The objective is not to build the most advanced agent... it is to build the first one that works, **and can be proven to work.**

74%

of organisations deploying AI agents report seeing ROI within the first year.

saasultra

STEP THREE

Prepare your data and context

Agents are only as good as the information they use.

Once a use case has been defined, attention must shift to the foundation that underpins every successful AI agent: data. While agent capabilities continue to advance, the quality of their outputs remains entirely dependent on the quality of the information they are given access to.

This is where many initiatives begin to falter.

Organisations often attempt to deploy AI agents without fully understanding the state of their data estate—resulting in outputs that are inconsistent, incomplete, or unreliable. In these scenarios, trust is quickly lost, and adoption stalls before value can be realised.

Effective implementation requires a deliberate approach to preparing both data and context.

This includes:

- Identifying the specific data required to support the use case
- Ensuring that data is accurate, consistent, and up to date
- Establishing clear access controls and permissions
- Defining how data will be used, updated, and governed

Importantly, this is not about achieving perfection. It is about ensuring that the agent is grounded in relevant, trusted information that enables it to perform its role effectively.

Without a strong data foundation, even the most advanced AI agent will **struggle to deliver meaningful outcomes.**

STEP FOUR

Build the agent

It's time to connect systems, define actions, and enable execution.

With a clear use case and a prepared data foundation in place, the next step is to design and build the agent itself. While the technology required to do this is becoming increasingly accessible, the emphasis should remain on simplicity and clarity – particularly in the early stages.

At this point, the goal is not to create the most advanced or autonomous system possible. It is to build an agent that performs a clearly defined function, reliably and consistently.

Within a Microsoft-led environment, this typically involves combining several components into a cohesive solution:

- **Copilot Studio** for designing agent behaviour and interaction
- **Power Platform** for connecting workflows and automating actions
- **Azure AI Foundry** for enabling reasoning and AI capabilities
- **Microsoft Fabric** for ensuring outputs are grounded in governed data

The key is orchestration – bringing these components together in a way that allows the agent to operate within defined boundaries, while still delivering measurable value.

At this stage, scope control is critical.

Agents should:

- **Focus** on a specific process or task
- **Operate** with clearly defined actions
- **Integrate** cleanly with existing systems
- **Avoid** unnecessary complexity

Successful agents are not those that attempt to do everything, but **those that do one thing well – and can be relied upon to do it consistently.**

STEP FIVE

Deploy with governance on day one

Control enables scale.

As AI agents move from insight to action, the stakes increase. At this stage, the focus shifts from capability to control—ensuring that agents operate securely, responsibly, and within clearly defined parameters.

This is where governance becomes essential.

Too often, governance is treated as a secondary consideration—something to be addressed once an agent has been proven to work. In reality, governance must be embedded from the outset. Without it, organisations are forced to limit access, restrict functionality, or delay deployment altogether.

Just one in five organisations has a mature governance model for its AI agents, and this gap represents one of the biggest barriers to scaling AI.

Effective governance frameworks should include:

- Identity and access management to control who and what agents can interact with
- Approval workflows for actions that carry risk or impact
- Monitoring and logging to track activity and performance
- Clear accountability for agent behaviour and outcomes

The objective is not to slow progress, but to enable it.

Governance, when implemented correctly, is what **allows organisations to move from controlled experimentation to confident, large-scale deployment.**

Just

23%

of organisations have the appropraite governance frameworks in place for their use of AI agents.

Deloitte

STEP SIX

Enable your people to work with agents

Adoption is a human challenge, not a technical one.

Even the most well-designed AI agent will fail if it is not used—or if it is not trusted by the people it is intended to support. While much of the focus in AI initiatives is placed on technology and data, long-term success is ultimately determined by people.

This is where many organisations underestimate the challenge.

AI agents introduce a new way of working. They shift responsibilities, automate familiar tasks, and require users to engage with systems differently. Without clear communication and support, this can lead to uncertainty, resistance, or disengagement.

Successful organisations address this directly, by:

- Clearly articulating what the agent does, and why it matters
- Providing training on how and when to use it
- Setting expectations around its capabilities and limitations
- Encouraging feedback and continuous improvement

Adoption is already happening – three quarters of knowledge workers are already using AI on a daily basis. The organisations that benefit most are those that actively shape how it evolves within their workforce.

AI agents should feel like part of the team – not an external system that people are expected to work around.

75%

of organisations have the appropraite governance frameworks in place for their use of AI agents.

Deloitte

STEP SEVEN

Measure, refine, and scale

Value is created through iteration.

The deployment of an AI agent is not the end of the process—it is the beginning of it. Real value is created over time, as organisations learn from how agents are used in practice and continuously refine their performance.

This is where the distinction between experimentation and operational success becomes most visible.

Organisations that struggle with AI tend to treat deployment as a finite point – launching agents without establishing clear mechanisms for measurement or improvement. In contrast, those that succeed treat each deployment as part of an ongoing cycle of optimisation.

This involves:

- **Tracking** usage and adoption across the business
- **Measuring** performance against defined outcomes
- **Collecting** user feedback and identifying areas for improvement
- **Refining** agent behaviour and expanding capabilities over time

And it's paying dividends. Organisations adopting AI agents are reporting increases of more than 170 per cent when deployments are scaled effectively.

Over time, successful use cases can be:

- **Replicated** across other teams or functions
- **Extended** into more complex processes
- **Combined** into broader, end-to-end workflows

The organisations that succeed are not those that build the most agents, but those that **continuously improve and integrate them into the way the business operates.**

Average ROI increases of

170%

for organisations that scale agent deployments effectively.

saasultra

Working Alongside AI Agents

Managing a new kind of workforce

AI agents are often discussed in terms of capability—what they can do, how they are built, and where they can be applied. But their long-term impact is not defined by technology alone.

It is defined by how organisations choose to work with them.

As agents become embedded into workflows, they begin to take on responsibilities that were previously owned by people—handling tasks, making decisions, and coordinating actions across systems. This does not replace the workforce. It reshapes it.

The result is the emergence of a **blended operating model**, where people and AI agents work together to deliver outcomes. Current rates of adoption alongside human roles are set to triple in the next couple of years – this isn't a future state, it's happening now.

From tools to collaborators

The way organisations interact with AI is changing. Tools are used, systems are managed, but **agents are directed and overseen.**

That distinction matters, because AI agents are not passive. They act on instructions, operate across workflows, and influence real outcomes – and as a result, the role of the user shifts too. Intead of executing every task, people begin to:

- **Define** objectives
- **Delegate** responsibility
- **Review** outputs
- **Intervene** when neeeded

This is a fundamental change in how work gets done – and it requires a different mindset.

Agent adoption is expected to grow by more than

300%

in the next two years.

index.

Tools are used.

Systems are managed.

Agents are directed and overseen.

Redefining roles and responsibilities

As AI agents take on more operational tasks, roles within organisations begin to evolve. In practice, this means that routine and repeatable work decreases, oversight and decision-making responsibilities increase, and cross-functional collaboration becomes more important.

For many teams, it represents a move away from **execution**, and towards **co-ordination and optimisation**.

This shift is not always immediate or obvious. But over time, it changes:

- How teams are structured
- How success is measured
- Where value is created

The organisations that recognise this early are **better positioned to adapt.**

The importance of human oversight

Agent autonomy does not eliminate the need for control.

In fact, as AI agents take on more responsibility, the importance of human oversight increases. Organisation must decide:

- What agents are allowed to do
- Where human approval is required
- How exceptions are handled

Trust plays a central role – and right now, only a quarter of organisations express trust in fully autonomous Ai agents. This is why most successful implementations follow a human-in-the-loop model, where agents operate within defined boundaries and people retain accountability for their outcomes.

Over time, and as confidence grows, these boundaries can evolve – but **they should never disappear entirely**.

What this means for leadership

For senior leaders, the challenge is no longer whether to adopt AI. It's how to integrate it effectively, balance automation with control, and ensure that people and technology work together productively.

This requires a shift in perspective. AI should no longer be viewed purely as a productivity tool, but as **an operational capability that needs to be managed, monitored, and used to its full potential.**

The opportunity ahead is huge. From eliminating friction to creating capacity, the collective gains that can be harvested are what will truly set organisations apart.

The time to become an agent-ready enterprise is **now.**

Related content

Click the headlines or icons to read our Learning Hub content.

Agent AI: What, why, how and where?

The first of a three-part series covering everything you need to know about AI agents.

From AI anxiety to action

Your route to agent readiness, wherever your organisation is starting from.

The A-Z Guide of AI Agents

26 agent use cases to get you thinking about their application.

Enabling DIY Agents

Everything you need to know about governance, tools, and education.

Agent 365: Everything you need to know

Bring yourself up to speed on Microsoft's latest agent platform, released in 2026.

CloudDirect®

Building agent-ready enterprises

Cloud Direct is a leading Microsoft Partner that works to position organisations as **Agent-Ready Enterprises** by combining human expertise with responsible, scalable AI solutions. For more than 20 years Cloud Direct has supported its customers in maximising their technologies to realise meaningful business impact, with its focus now spanning cloud, data, AI and security to deliver human-led, AI-powered outcomes that matter.

hello@clouddirect.net **01225 300330** **www.clouddirect.net**