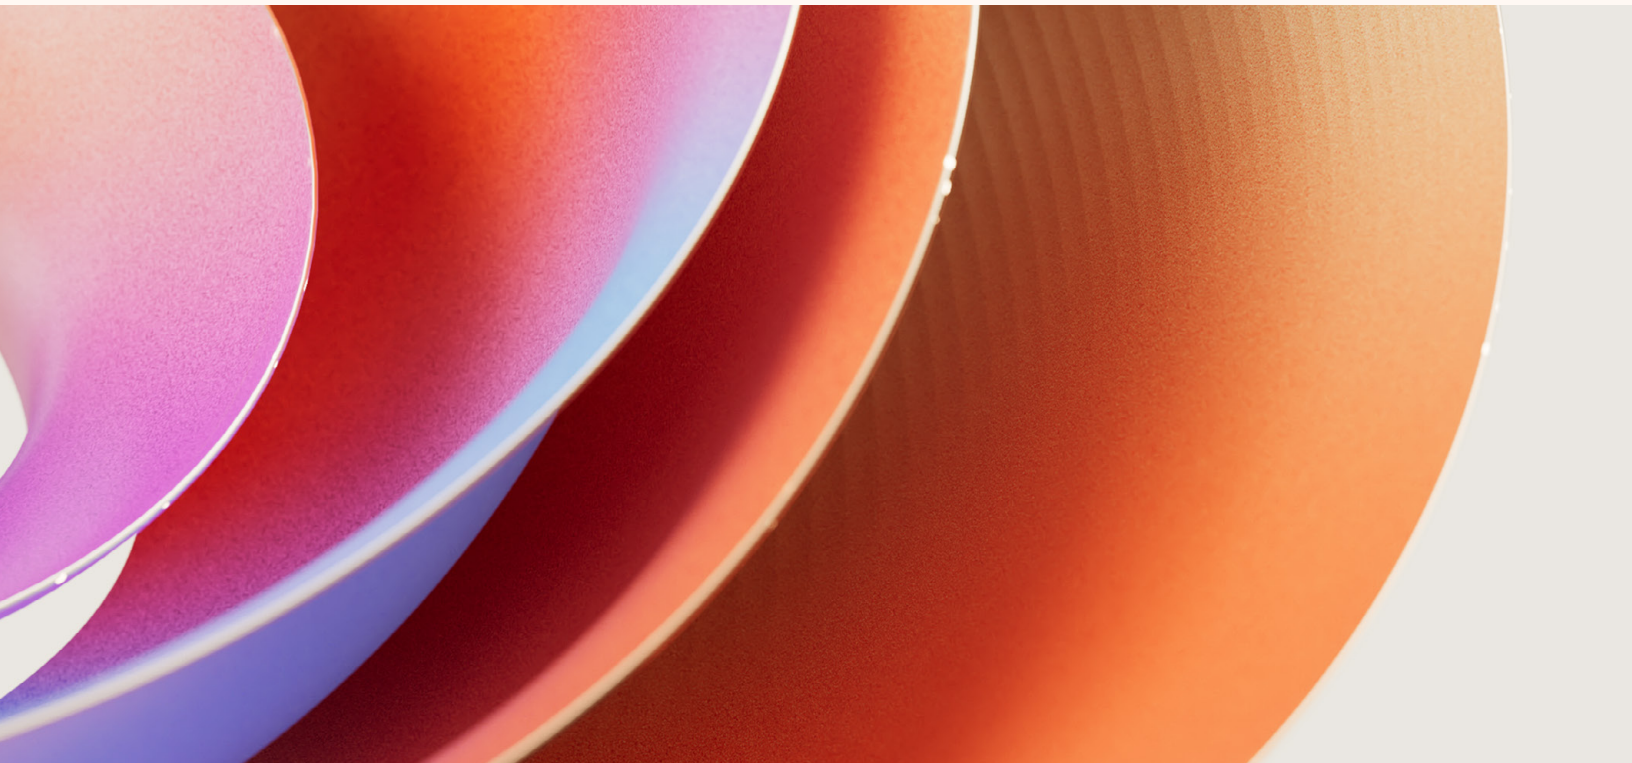


Microsoft

The AI Strategy Roadmap

How organisations are achieving Frontier Transformation



Executive summary

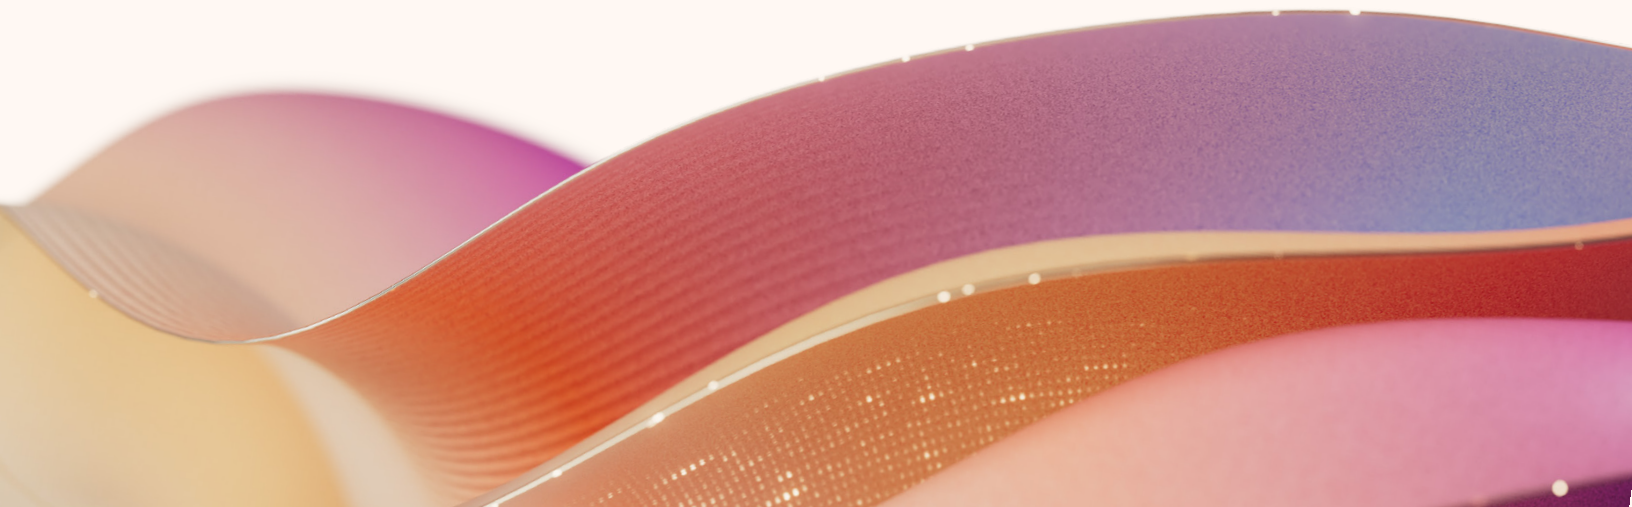
AI is reshaping how work gets done, helping organisations optimise operations, gain competitive advantages and deliver more value faster. Reaping these benefits requires a foundation for responsible, scalable and outcome-driven AI adoption. It also depends on workforce capability – equipping people with the AI and human skills needed to work effectively with AI systems.

Organisations are embracing the Frontier Transformation process to build that foundation and drive faster innovation and growth across their organisations. They're making AI systems a core part of their business processes – changing how their organisations do work and empowering their employees to focus on the tasks that require human judgement, knowledge and creativity.

What is Frontier Transformation?

Frontier Transformation is the next evolution of digital and AI transformation: a holistic reimagining of business that aligns AI with human ambition to help organisations achieve their highest aspirations and growth potential. It requires moving beyond isolated AI use cases focused on productivity and efficiency to integrate AI throughout key business processes.

This roadmap translates research-backed patterns and experiences from successful Frontier Firms into practical strategies and actions that can help you become an AI-powered business.



Lessons learned through experience

The recommendations and guidance provided here are informed by detailed interviews with 70 business and IT decision-makers from leading organisations, all actively engaged in their transformation journeys.[1] They're also informed by our internal AI transformation at Microsoft and by the work we've done to support the transformation efforts of many customers.

Across these experiences, a consistent set of challenges emerges.

Organisations struggling with AI transformation generally have good intent and ideas, but often struggle to control scope, establish trust, manage unanticipated complexity and build the workforce capability.

In contrast, successful Frontier Firms deal with challenges by acting early, planning for iteration and staying clear eyed and realistic about AI capabilities, requirements and scope. They transform by tackling real problems, building confidence through an iterative series of small wins.

Learning how successful Frontier Firms deal with these challenges – and more – can help you accelerate your organisation's Frontier Transformation journey.

Issue	Potential challenge	How Frontier Firms approach
Ideas abound, but there's no clear starting point.	Overwhelming volume of opportunities makes prioritisation a significant hurdle.	Anchor on one measurable use case before expanding.
Return on investment (ROI) is abstract, so leaders hesitate to fund scale.	It's difficult to convince leaders that concrete evidence of ROI isn't the goal.	Create transparency and show measurable progress against key performance indicators (KPIs), instead of engaging in isolated black-box experimentation.
Data is fragmented/not ready, outputs are unreliable.	Siloed systems with poor data stewardship result in inaccurate or hallucinated outputs.	Treat data readiness as the critical path. Start by building patterns, and then expand from there.
Compliance and security reviews delay expansion.	Regulatory and review processes are segmented across a wide variety of stakeholders, and testing needs aren't aligned.	Prove safety with humans in the loop and layered sign-off.
Pilot works, but scaling feels risky and undefined.	Responsibility is ambiguous for resolving exposed gaps in governance, operating models and cost modelling.	Scale step by step, and validate one use case before moving on to the next.
Users/customers don't trust the agent, so usage stalls.	Early encounters with unpredictable performance, bad data, governance issues or cultural rejection of AI reduce trust.	Offer choice (opt-in), test safely and prove results to end users.
Teams lack the skills and confidence to adopt AI broadly.	Limited AI fluency, uncertainty about how to use AI responsibly and over-reliance on a small group of experts can slow adoption and make scaling difficult.	Invest in role-based skilling, hands-on learning and change support that helps employees build confidence, adopt AI responsibly and apply it in real workflows.

[1] Emerald Research Group: Sponsored by Microsoft, AI Transformation Strategy Qualitative Research Report, May 2026.

Contents

02
Executive summary

05
How to approach Frontier Transformation

09
AI value driver #1: Business strategy

20
AI value driver #2: Technology and data strategy

30
AI value driver #3: AI strategy and experience

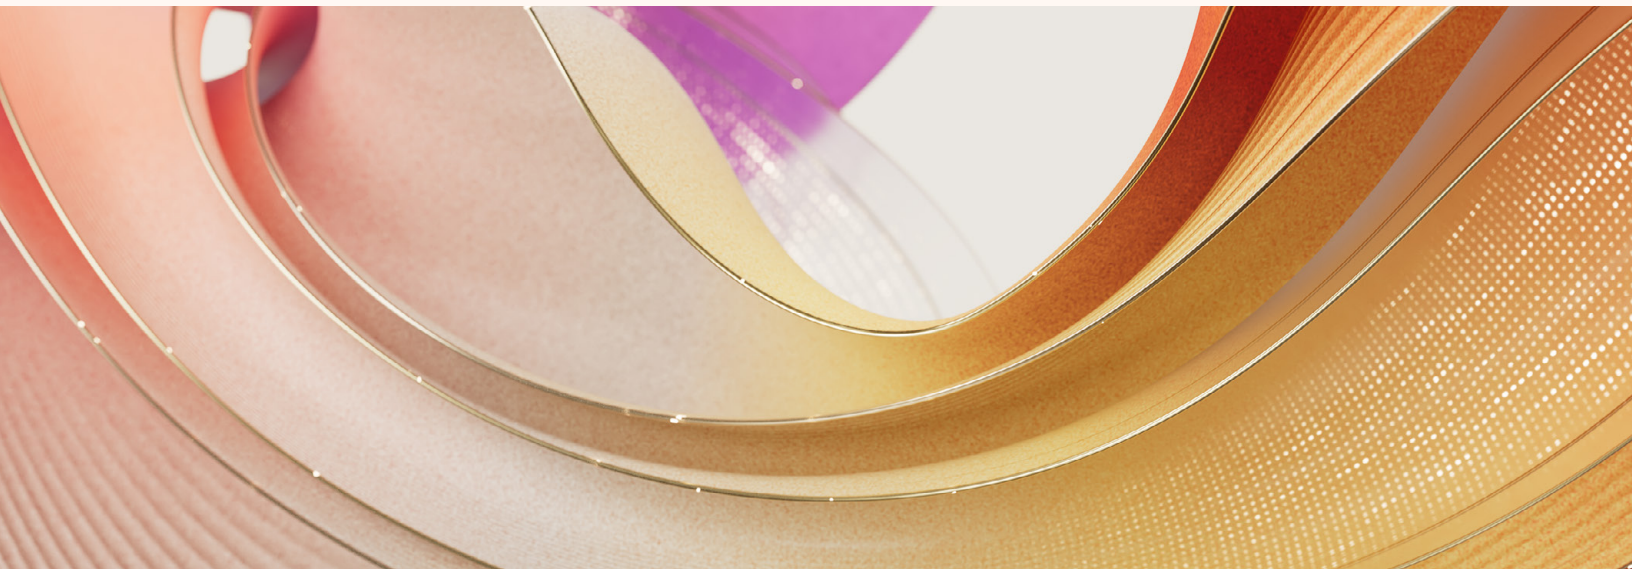
38
AI value driver #4: Organization and culture

47
AI value driver #5: AI governance and security

58
Microsoft as customer zero

63
How can Microsoft help?

65
Conclusion



How to approach Frontier Transformation

Most organisations see the advantages of Frontier Transformation and the benefits it offers: scaling AI across the organisation and harnessing it to improve business processes.

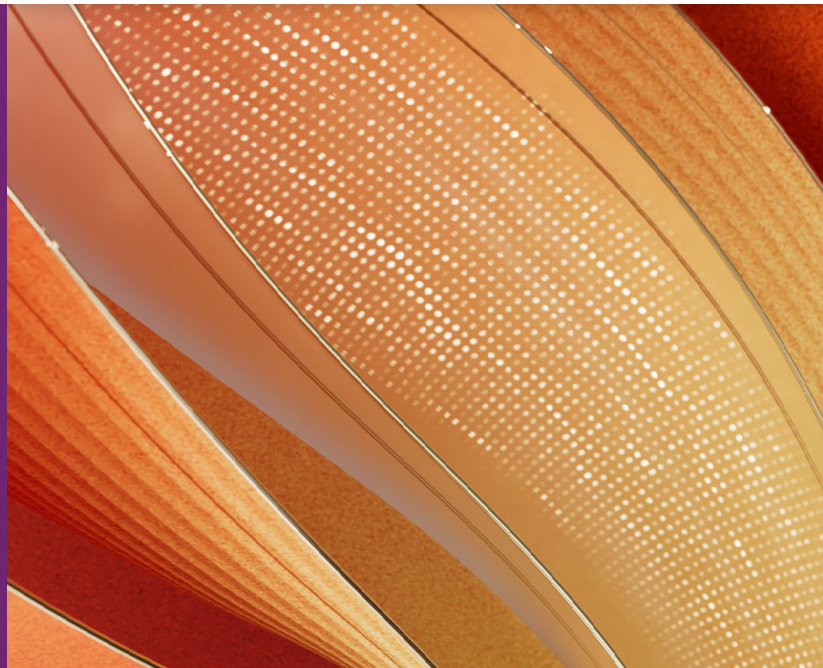
To get the most out of transformation, you need a thoughtful, iterative approach to AI that can address potential challenges ahead of time and align your AI ambition with technological and organisational reality.

What changes with Frontier Transformation

- **AI becomes foundational** and is embedded across all business functions.
- **Humans and AI work together**, with employees increasingly collaborating directly with agents.
- **Processes get rebuilt**, with AI integrated into workflows.
- **AI is integrated into continuous business intelligence**, so decision-making is faster and more widely distributed.
- **Experimentation is replaced by enterprise-scale execution**, with success less dependent on pilots and more on scaling well-governed AI systems across the entire business.
- **Skills evolve alongside AI**, as employees develop both AI literacy and the human skills (judgement, critical thinking, collaboration) needed to work effectively with AI systems.

Ultimately, transformation means AI does more than simply improve your staff productivity or automate workflows. It introduces new capabilities to your organisation, so you can achieve outcomes that weren't possible before.

> It really started with asking, 'Where are the biggest inefficiencies in our workflows?' We were looking at things that were taking a lot of manual effort and asking if there was a better way to approach them.
>
> – Vice President (VP) of Tech Solutions, Financial Services



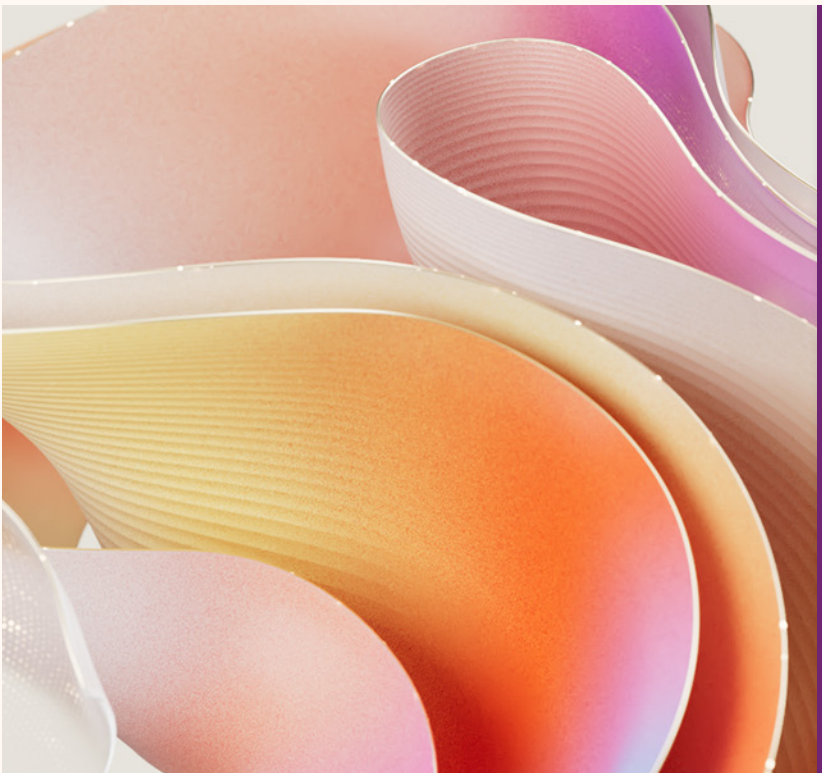
Think of Frontier Transformation as a journey, not a destination

Frontier Transformation is not just about technology. It's a fundamental change to how your business creates value and gets work done.

Begin with clear business goals:

- **Build early momentum**, focusing on AI outcomes that are easy to track. That usually means beginning with low-risk, high-volume workflows or processes.
- **Make progress through iteration, not perfection.** With differing levels of AI adoption and knowledge, ongoing, hands-on experience builds trust in AI systems and strengthens the skills and confidence employees need to use AI tools effectively.
- **Choose use cases that show ROI early.** High-value use cases have big benefits, but they often take a lot of time and resources to tackle.
- **Treat each project as a learning vehicle**, documenting both the work and the lessons learned as you go. As systems scale and evolve over time, keep documentation up to date in a shared knowledge base.

As your transformation matures, pilot and validation cycles shorten, governance becomes more integrated, AI literacy increases and confidence grows. With tools that integrate clear guardrails, more of your employees become makers, building their own AI tools and agents safely. Over time, your AI efforts shift from a collection of individual projects to repeatable systems that deliver business value throughout your organisation.



> We had been trying to solve this problem in more traditional ways for a long time, and it just wasn't getting us where we needed to go. That's what pushed us to start looking at new approaches.
>
> – VP of Client Services, Professional Services

Evaluate AI readiness

Your organisation's AI readiness, or your ability to move beyond isolated AI experiments and reliably deploy, scale and govern AI as a foundational business capability, is key to successful Frontier Transformation.

It's useful to evaluate your AI readiness based on these five stages.

Exploring

In this early stage, you're exploring what AI can do, generating ideas and running low-stakes tests. Isolated pilots and proofs of concept can show promise, but without a centralised effort and without building early understanding and confidence among employees, they won't bring value to your wider organisation.

Planning

At this point, you're making a deliberate shift from experimentation to identifying ways that AI can bring value to your core business. This includes identifying and prioritising use cases, defining success criteria, aligning business and technical stakeholders, establishing initial governance guardrails and beginning to identify the skills and change support needed for broader adoption.

Implementing

Here, you start building and deploying AI solutions as systems in progress: instrumented for measurement, with defined ownership and clear success criteria for decision-makers. A small number of prioritised use cases move into production, allowing you to prove reliability, safety and usefulness and to build employee confidence in using AI in real workflows.

Scaling

After confidence and trust in AI is established, you expand to repeatable, governed, enterprise-wide delivery of AI solutions. This includes adopting shared platforms, reusable assets, consistent lifecycle management and broader skilling so new use cases can be delivered faster, adopted more widely and scaled more effectively.

Realising

AI is fully embedded as a core business capability, with value continually measured and improved. AI-enabled workflows are the standard, leaders trust results, governance is mature and auditable and the workforce has the skills, confidence and operating rhythms to continuously iterate and compound value.

Embrace the five drivers of AI readiness

We've identified five core drivers that consistently help organisations use AI to deliver value across their entire businesses:

Business strategy	AI efforts are anchored to a measurable business strategy and clear vision that helps leaders prioritise, fund and sequence work on AI initiatives.
Technology and data strategy	A strong technological and data foundation helps you transform experimental systems into scalable core capabilities.
AI strategy and experience	Intentional and iterative adoption builds user and leadership trust, fuels ongoing adoption and drives sustained business impact.
Organisation and culture	Intentional change management, leadership vision and the development of new operating models help drive AI adoption and build value.
AI governance and security	Strong governance and security guardrails create trust, encourage adoption and build confidence to scale and automate using AI systems.

Together, these drivers form a practical, interconnected framework for assessing where you stand today and what you need to do to move forward. When they're developed in parallel, your teams can move faster, manage and mitigate risk and realise the full value of AI.

"

We're in the implementation phase, but challenges around internal skills, data security and especially unclear ROI are slowing progress. We need clearer, business-aligned use cases to justify scaling across different units.

– Global VP of Sales, Manufacturing

AI value driver #1

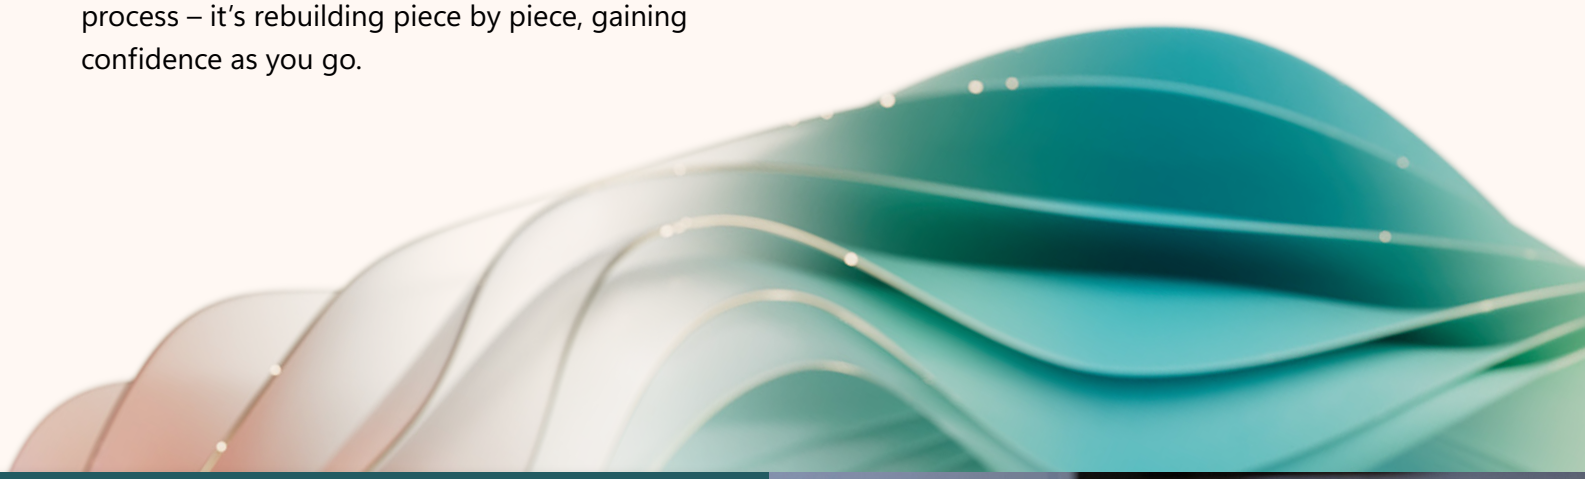
Business strategy

Frontier Firms align AI use cases with business strategy, IT and governance teams from the start. They begin with a clear business problem, design AI tools and agents to address it and measure ROI.

Early AI projects are often agents that automate low-risk, high-volume tasks, with easy-to-measure ROI. These projects evolve into integrated workflows through incremental wins, governance and team development. It's not a rip-and-replace process – it's rebuilding piece by piece, gaining confidence as you go.

In alignment with your AI investment strategy, use cases iterate with increasing success, ultimately developing into systems integral to your business strategy.

When AI investments are aligned with strategic business outcomes and measurable impact, value is the result.



"

We didn't start with a big, fully baked AI strategy. It was much more about identifying a single area where we could realistically make a difference and testing it there first. Once we had something working and could show tangible impact, that's what helped us expand and bring others along.

– Senior VP of Growth Marketing, Financial Technology



Evaluate your AI business strategy

Do you have an end-to-end AI operating model and clear prioritisation of high-value use cases? Understand your current level of readiness and look at the next stage of maturity to identify where you should be investing time and resources.

Unfocused experimentation	Foundation for enterprise AI strategy	AI in business planning	Leadership-driven AI strategy at scale	AI investment and portfolio management
• Leadership engagement limited • Experimental AI projects in business groups • Minimal investment • Unclear business value	• Clear AI goals • Priority use cases • Measurable outcomes	• AI projects driven by business plans • ROI and KPIs tracked at the organisation level	• AI initiatives scaled across organisation • Performance reviewed regularly by leadership	• AI fully integrated into business model and decision-making • Measurable, enterprise-wide ROI • AI as a business capability

Determine your next steps

Based on your current state, what are your next steps to arrive at an end-to-end AI operating model with prioritised, high-value use cases?

Create your AI vision and identify use cases

Without them, AI efforts tend to fragment. Among the organisations we interviewed that struggled with AI execution complexity, 15% cited lack of executive buy-in and strategic alignment as an AI barrier.

You don't need a polished AI manifesto to get started, but even in the early phases of your Frontier Transformation journey, you'll benefit from a vision that provides clarity, consistency and follow-through and AI use cases with measurable ROI.

> We know AI has value, but we're still early and figuring out how to measure it. Right now, efforts are spread across siloed teams without a central strategy, leading to duplication and lack of alignment.
>
> – VP of IT, Retail

Address key elements of an AI vision

Your strategic vision for AI should include these elements.

Executive sponsorship. In organisations making real progress, senior leaders do more than approve AI initiatives. They stay involved – actively championing initiatives, reinforcing priorities and holding teams accountable for outcomes. This sponsorship creates permission to innovate and offers clarity around what matters most.

Clear connections between AI and business priorities. Over 70% of organisations with advanced strategy and execution readiness report that leadership communicates an AI vision clearly aligned with business strategy.[2] They know why they're investing in AI and which business outcomes it supports, and they balance near-term ROI initiatives with longer-term, higher-risk bets that can transform how the organisation operates and competes.

Metrics that matter to the business. Frontier Firms clarify outcomes before they begin building, concentrating on clear and measurable metrics and keeping efforts focused. The desired outcome can be simple; for example, reduced call volume or faster internal processes. When value is visible and tracked in business terms, AI stays on the leadership agenda.

Workforce capability and skilling. Organisations define how they'll build AI literacy, role-based skills and ongoing learning programmes to support adoption at scale.

Deliberate funding decisions. Rather than spread AI investment thinly across disconnected initiatives, mature organisations secure budget for prioritised use cases that align with business-critical workflows and deliver clear impact. To show short-term success while sustaining AI strategies over time, they balance investments in proven initiatives that show rapid ROI with AI investments that have higher risks and take more time to mature, but can reshape how organisations operate and compete.

Proactive risk management. An effective AI vision explicitly addresses the potential risks introduced by AI technologies. These range from data privacy and security to model bias, compliance and operational reliability. Frontier Firms keep risk management accountability in-house, anticipating these challenges early and defining clear risk management frameworks, ownership and mitigation strategies. They change the conversation from *Should we do this?* to *How do we do this safely?*

Commitment to responsible AI (RAI). Leading organisations embed RAI principles directly into their strategy from the outset, building fairness, transparency, privacy, security and accountability into every AI project. By aligning AI initiatives with clear ethical standards and risk management practices, they create the trust needed for broad adoption, reduce regulatory and reputational risk and build the confidence to innovate at scale.

With these elements in place, your AI vision helps you prioritise AI initiatives for business impact.

[2] Microsoft Agent Readiness Survey, September 2025.

Start small, with a narrow and easy-to-measure use case

Although they start in different places, all Frontier Firms converge on a single, achievable and measurable use case to move from experimentation to execution. They don't build out a large-scale AI development plan immediately. Instead, they start with a narrow use case that's easy to measure, helps teams build practical skills and confidence and creates a foundation for broader adoption.

Evaluate internal processes to prioritise use cases

To choose this first use case and prioritise others, start with identifying and evaluating internal processes. Processes that create the most friction, cost or risk, such as long handling times or inefficient workflows, can be good candidates for AI use cases that can yield a significant ROI in a short time frame.

You should be confident that these use cases are feasible. Consider the complexity of data and system interactions, security and compliance requirements and how specific outcomes will connect to ROI, efficiency, quality or compliance tied to business priorities.

For your first use case, you're looking for a quick win with meaningful outcomes that you can build on.

As your AI initiatives expand, evaluating AI use cases using these criteria can help you modernise workflows and unlock value for agent deployment.



> "Let's pick a very specific use case where we know what good outcomes absolutely look like and build towards that.
>
> – VP of Tech Solutions, Financial Services

Embed AI in business plans and track results

Your first use case sets the stage with small wins and visible ROI. Even as it progresses, start planning to embed AI into enterprise-wide business plans and establish processes that will keep a steady stream of high-priority use cases in the queue and provide guardrails to help you scale safely.

Develop a use case prioritisation framework built on business criteria

Establish a use case prioritisation process that addresses feasibility, risk and business value, and apply it company-wide. Evaluating use cases consistently across the organisation supports broad participation and portfolio-level visibility, which will help you manage and scale your AI strategy.



Every selected use case should meet these criteria:

- It addresses one business challenge.
- Risk is identified and manageable.
- Use case aligns to strategic priorities.
- It's well suited for AI.
- You can validate success before expanding.
- You can create a repeatable expansion cycle without losing control.
- Funding and staffing is available.
- You can obtain executive sponsorship and visibility.

If you'd rather adopt an existing prioritisation framework, the Microsoft AI Decision Framework, sometimes called the Business, Experience, Technology (BXT) decision framework, can help you evaluate agent use cases based on three dimensions: business impact, experience and technology feasibility. By scoring across these factors, you can pick use cases that deliver value, are embraced by users and can be implemented securely.

Making a business case for AI

- For high-value use cases, look at repetitive, rules-based work that still needs judgement.
- Evaluate existing tools and processes, where AI can yield opportunities to integrate data and streamline workflows in ways not currently possible.
- Integrate use cases into existing workflows.
- Look for existing processes that can be reimagined, piece by piece.

As use cases are approved, secure executive support, building approved use cases into your business strategy and adding them to your organisation's AI investment portfolios.

Agentic AI use cases and metrics across business functions

AI value concentrates where work is repetitive, time intensive and rules based, but judgement aware, delivering the most value when it's applied to business-critical workflows and scaled deliberately. Agentic AI use cases can deliver compelling ROI across core business functions.

	Agent	ROI
Customer service	• Self-service AI assistant (support assignment) • Always-on AI assistant (problem resolution)	• Reduced call volumes • Reduced time required to address
Sales	• Sales AI assistant (lead generation) • Request for proposal (RFP) response automation (customer engagement) • Always-on sales AI assistant (customer engagement)	• Boosted cross-sell and upsell • Increased average sales ROI
HR	• Ask HR (employee engagement) • AI-enabled candidate search (recruiting) • Career development AI assistant (talent management)	• Reduced cost per hire • Reduced onboarding hours • Reduced HR hours per new hire
Legal	• Automated compliance management (compliance) • AI-optimised contract management (contracting) • Ask legal (advisory services)	• Reduced lawyer hours per contract • Reduced document review time
IT	• AI-enabled app modernisation (IT operations) • AI helpdesk agent (IT operations) • AI-driven assessment and mediation (security and compliance)	• Reduced effort in IT operations • Reduced unplanned downtime • Reduced time to issue resolution



Define measurable outcomes

You'll find that AI can take over a variety of tasks that deliver measurable benefits.

Increased productivity: Complete multistep tasks autonomously, reducing the need for human supervision and freeing team members for more strategic work.

Cost savings: Take over rigid workflows to reduce operational effort, manual exception handling and overhead, driving substantial and sustained cost efficiency.

Improved accuracy: Make decisions in variable or ambiguous scenarios based on interpretation, context evaluation and dynamic validation, minimising errors that traditional, rule-based automation can't detect or handle.

Enhanced compliance: Monitor actions continuously and apply policies contextually, escalating anomalies in real time. Interpret regulatory requirements to reduce compliance gaps that arise in manual or static automated processes.

Better collaboration: Act as intelligent teammates that interpret human intent, coordinate across systems and collaborate with other agents, improving communication flow and reducing friction between teams and functions.

Stronger security: Assess patterns proactively, detect deviations and enforce security rules autonomously, strengthening the security posture beyond static access controls or predefined rule sets.

Faster turnaround: Agents can self-orchestrate end-to-end workflows, resolving ambiguities and handling edge cases without waiting for human input to speed up completion times for complex processes.

Streamlined decision-making: Agents can automate research, analyse feedback and generate competitive insights, summarising the details your team needs to make informed decisions faster.

Scalability: Agents scale cognitive work – not just repeatable tasks. They coordinate tasks, distribute workloads and adapt to demand without proportional increases in staffing or automation engineering.





Expanding AI initiatives with engaged leadership

Successful organisations treat AI initiatives as a managed investment portfolio. They continually evaluate performance, funding and alignment. Through iterative small-scale wins, they build trust in AI and develop strong governance and operational practices.

Assess your AI portfolio

Between different teams and projects, organisations commonly span multiple stages. Businesses that scale AI successfully don't treat initiatives as isolated projects. They manage them as a portfolio of investments, with clear oversight, performance tracking and the ability to rebalance as conditions change.

Recognise that gaining confidence to scale is incremental. Data readiness is the baseline you need to invest in use cases. Successful use cases drive more volume. Change management builds adoption, leading to further expansion. Visibility into this process builds trust.



Be ready to evaluate the current state of your AI portfolio

- **Is there an oversight and rebalancing strategy?** To allocate resources strategically, you should regularly consider performance, alignment with organisational strategy and time frame to ensure that you're maximising ROI.
- **Do you have an overall AI investment plan?** To ensure that prioritised AI projects get the funding and support they require to succeed, you should have a concrete plan for allocating budget and resources across projects. An investment plan accelerates delivery, demonstrates executive commitment and minimises the risk of spreading resources too thinly.
- **Do you review the portfolio on a regular schedule?** Regular review builds agility into your portfolio management, providing an opportunity to respond to changing business needs, technology developments and incremental project milestones with refinement, redirection, renewed vigour or even a pause.
- **Are you capturing and sharing lessons learned?** This is particularly crucial early in your AI journey, although it's important at every stage to use successes and setbacks to improve your path forward. Responding to lessons learned can help your organisation reduce missteps, invest wisely and accelerate business transformation.

Integrate AI across the enterprise

As your AI use expands, you'll likely see common themes emerge across AI use cases and maturing projects. Successful scaling requires ongoing executive support, committed and well-informed sponsorship, clean and trustworthy data, reliable governance, consistent and strategic funding and trust built on visible results.

Strengthen trust in AI to expand adoption

As AI adoption increases, growing trust becomes the critical foundation for scaling impact across your organisation.

Internally, trust builds and adoption grows as AI systems demonstrate accuracy and take over more of the day-to-day workload. Employees have more time to spend on strategic and creative work.

As trust is proven, apps and agents are rolled out to customers, improving customer satisfaction, lightening the customer service load and opening up time to address more complex customer issues.

Throughout the enterprise, AI initiatives begin to scale across teams and functions.

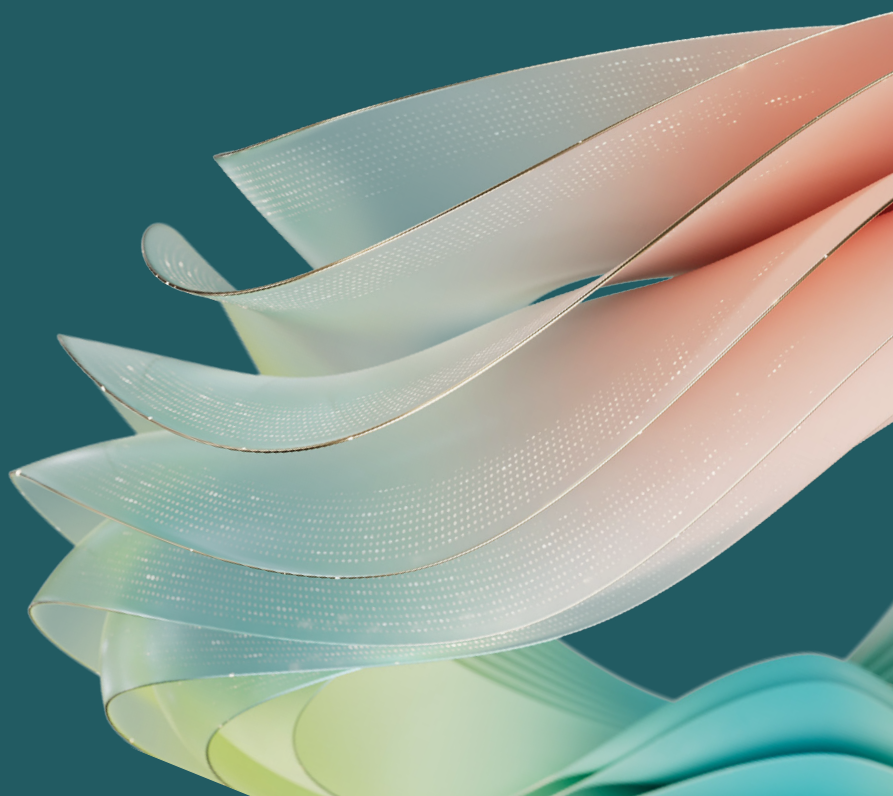
AI becomes fully embedded in organisational strategy and value creation. As a competitive differentiator, AI becomes core to the company vision. AI initiatives deliver measurable, repeatable value.

Communicate your success as AI use grows

This is a good time to define how you'll measure the value of AI initiatives across the business. Consider implementing a scorecard or dashboard to track KPIs for AI initiatives, including risk-reduction and governance, at the executive level.

As you continue to scale, the AI processes you've established become part of the heartbeat of the organisation.

- Executives provide ongoing support, reviewing AI portfolios and results on a regular cadence.
- AI strategies are funded, monitored and rebalanced as progress is made and business and competitive needs change.
- Use cases are prioritised regularly using a consistent framework.
- Lessons learned are integrated into the portfolio and process.



Scale safely and effectively

To maintain control as they scale, Frontier Firms use the following strategies:

Expand one business challenge at a time. Start with narrow, well-defined use cases. Avoid broad rollouts that increase risk and complexity.

Validate success before expanding. Measure performance and outcomes early, using clear metrics to build confidence before moving to the next use case.

Use transparency to maintain trust. Clear guardrails, human oversight and visible performance help align stakeholders and reduce resistance as agents scale. Involving legal and compliance in early conversations smooths out barriers.

Create a repeatable expansion cycle. Each successful agent establishes patterns for data, governance and workflows that make subsequent deployments faster and easier. The more you document, the easier it is to repurpose.

Anticipate roadblocks

Manage the flow of ideas. Too many ideas with no clear starting point can be overwhelming, resulting in stagnation.

Be transparent about measurement and progress. When ROI goes unnoticed or seems abstract, leaders hesitate to fund black-box experimentation at scale.

Align IT, governance and business teams to ensure approvals and deployments. The sooner you get all the key players aligned, the better.

Identify reluctant stakeholders and address concerns early. Risk-averse stakeholders can include legal, security, finance and other employees who might be worried about AI safety, job impacts and ROI.



"

You have to walk through architecture, security and risk up front and make sure everything is aligned. Once that foundation is in place, it's much easier to move solutions forward and get the business on board.

– VP of IT, Retail



Reaching the goal: AI embedded into strategy and value creation

You've reached the point where AI is an enterprise capability, firmly integrated into your vision, your business model and your decision-making.

- AI as a factor in ROI, measured across the enterprise, delivers repeatable value and might be a competitive differentiator.
- As part of your enterprise strategy, AI investment and objectives are embedded in planning processes across the organisation and AI initiatives are reported at the executive level.
- An AI investment plan, which is regularly reviewed, defines budget and resource allocation based on priority and expected returns for projects in your portfolio.
- AI initiatives are managed, reviewed and rebalanced using a portfolio approach. (Organisations often double down on AI initiatives that show results and might pause or refine those that underperform.)
- All AI implementation is governed, compliant and secured. The processes and guardrails are in place for you to continue expanding your AI capabilities to support new business endeavours, markets, teams, products and more.
- Your growth is no longer constrained by head count. Instead, you can increase capacity using software, handling higher volumes of work and distributing workloads as you adapt to demand.

Replace hesitation with confidence-led execution

- **Don't wait for perfection to get started.** Frontier Firms move forward with available skills and infrastructure, learning and evolving as they go.
- **Act despite uncertainty.** Frontier Firms proceed with use cases to surface and resolve risk, governance and integration issues rather than waiting to eliminate them up front.
- **Show, don't tell.** Focus early rollouts on users and teams most open to AI. Resistance decreases when employees see agents reducing manual work and supporting their roles.
- **Turn early AI efforts into repeatable systems.** Make the most of early wins, indirect ROI, executive sponsorship and agent champions to sustain momentum and scale. Then move forward with new, more complex use cases and scale them faster.

AI value driver #2

Technology and data strategy



When data is trusted and technology infrastructure is designed for growth, AI becomes a repeatable capability rather than a series of isolated experiments. With this foundation, you can deliver consistent results, embed AI across your organisation and generate measurable business impacts.

Keep these key points in mind when approaching your technology and data strategies:

Address data readiness as a prerequisite. Data quality, identity and access controls and governance need to be addressed before expanding AI capabilities.

Build and connect infrastructure incrementally. Early AI use cases may require minimal resources and integration between systems. Integrate AI tools and agents into your wider infrastructure only after their reliability is proven.

Stabilise foundations before scaling AI capabilities. After you've established data pipelines, governance patterns and monitoring systems, you can scale your agents and systems faster and more safely.



"

It's really important to get your data, infrastructure and your cybersecurity architecture right from the get-go. You have to consider what you want to plug into, what those data sources are going to be.

– VP of Technology and Cyber Security, Retail

Evaluate your current technology and data readiness

Are your data foundations, infrastructure and IT teams ready to scale AI responsibly? Understand your current level of readiness and look at the next stage of maturity to identify where you should be investing time and resources.

Emerging data foundation	Foundational build-out	Production-capable platform	Enterprise-wide infrastructure	Continuous optimisation
• Data siloed or limited in quality • No AI-specific infrastructure beyond individual experiments • Starting to assess data, platforms needed for agentic AI	• Efforts underway to implement data readiness, establish initial AI infrastructure • Starting to set up cloud, hybrid and on-premises resources, data lakes, sandbox environments for agentic AI use cases	• Robust, scalable data and AI platform in place to deploy AI solutions at scale • Data pipelines feed production AI models • Tools and environments support model development, deployment, monitoring reliably	• Data platform unified, serves multiple teams • Agentic AI solutions integrated with enterprise systems, embedded into business processes • Infrastructure handles large-scale, real-time AI workloads securely, efficiently	• Data architecture and AI infrastructure improved continually • New AI technologies readily adopted, supported by data platform and tools • Data treated as strategic asset, fuelling ongoing AI innovation

Determine your next steps

Based on your current state, what are your next steps to creating a scalable well-governed data and technology infrastructure that can support your growing AI adoption efforts?

Assess your data quality and AI readiness

During the exploration stage, it's common to find data that's fragmented or requires substantial manual preparation before AI systems can use it. If you can't systematically count on data quality, consistency or governance, AI outcomes can be unreliable, difficult to scale and hard to trust.

With a trusted, unified and well-governed data foundation, your AI tools and agents can move beyond isolated pilots and scale reliably into secure, repeatable, enterprise-wide capabilities.

To prepare your data and systems for AI:

Perform a formal data maturity assessment.

- Investigate what data you have, how it's accessed and how it's used.
- Develop a structured and thorough assessment that identifies gaps in data quality, availability and governance before they cause problems with your AI use cases.

Identify and address gaps in data needed to support your use cases.

- Document any data you'll need, including the systems, records and attributes required to produce reliable outputs.
- Assess that data and identify gaps, such as missing fields, siloed sources, limited access or inconsistent standards.
- Identify the data and technology touchpoint dependencies for inputs, outputs and supporting systems.
- Create a plan to close the gaps, including actions for data cleanup, integration, governance and ongoing maintenance.

> "On paper, we had all the data we needed, but it was spread across systems and not always aligned. We had to create some structure around what data to use and when, otherwise the outputs just weren't reliable.
>
> – Senior VP of Growth Marketing, Financial Technology

Rationalise and modernise legacy data.

- Identify legacy data sources and technical debt. This can include data in outdated platforms, fragmented architectures and inconsistent data structures that are difficult for AI systems to work with.
- Consider modernisation options, such as replatforming, refactoring, archiving or retiring legacy datasets and systems.
- Consolidate your data into scalable, unified modern platforms that support standardised architectures, strong governance and AI-friendly cloud-native capabilities.

Build a unified enterprise semantic model for your data.

- Define a consistent semantic model for core business entities; for example, ensuring common definitions for concepts like *customer, order* and *region* across your systems and teams.
- Identify and resolve inconsistencies in how your data is defined and used across sources, eliminating conflicts that can lead to unreliable AI output.
- Standardise and govern data definitions centrally so AI systems can correctly interpret data, scale across use cases and deliver consistently accurate output.



Start developing an AI-ready infrastructure

If you've reached the planning stage, your efforts to establish data readiness are underway. Your next priority should be laying the groundwork for an AI-ready infrastructure designed to support continuous AI development, deployment and operations at scale. This infrastructure should provide:

Scalability, so AI workloads can grow as demand increases.

Security and identity by default, so data and actions are continuously protected.

Reliable data access, with clear boundaries and governed integrations.

Operational visibility, making performance, usage and costs transparent.

Flexibility, supporting a mix of low-code, pro-code and packaged solutions.

Successful organisations build capabilities incrementally, proving reliability and governance at each stage before expanding. It's important to plan early in your transformation journey for your infrastructure to support widespread AI adoption.

As you start building out your AI infrastructure:

Address data governance early.

- For AI that scales securely, reliably and with trust, establish governance that spans the full data lifecycle, including ingestion, model training, inference and operations.
- Approach data governance as a way to enable safe innovation, rather than as something that constrains or blocks your staff's use of new technology.
- At the organisational level, establish clear ownership, standards and data controls that make creating and deploying AI tools and agents faster and easier, while also ensuring safety.

Adopt a managed platform.

- A shared managed AI platform provides a common infrastructure for building, deploying and governing AI models, applications and agents across your organisation.
- The platform reduces duplication, and you can apply central controls and governance of all agents, tools, and data connections.
- It also helps you scale: your models, data and insights are reusable across multiple use cases, while enforcing consistent, enterprise-wide governance and lifecycle management.



Develop AI optimised 'cloud landing zones'.

- Use standard patterns and configurations to create deployment-ready environments with built-in networking, identity, security and policy controls.
- Standards provide a structured starting point for scalable AI, so you can create safe, scalable infrastructure for your AI solutions.
- Accounting for AI requirements in your landing zone design helps you minimise the risk of latency issues, cost spikes, inconsistent model behaviour and security gaps.

Standardise on a consistent AI integration framework.

- Adopting a reusable integration framework (for example, model context protocol [MCP], OpenAPI or similar enterprise connectors) drastically reduces time and effort to integrate new AI tools into systems.

Update your cost management practices.

- AI usage can scale rapidly and pricing varies by model and workload type, so small design choices can have outsized financial impact. Design for cost efficiency early, evaluate usage patterns continually and refine design decisions as systems evolve. Optimise your operations continually and intentionally, rather than after budgets are exceeded.

Shared AI platform features

Common capabilities instead of custom features. Don't reinvent the wheel for every pilot. Instead, minimise duplication and operational overhead so teams can focus on data and features.

Standardised security and identity patterns build trust by providing consistent protection and access controls that can scale alongside usage.

Common integration frameworks make it easier for agents and applications to connect to data and other systems consistently and securely.

Reusable operational practices, including monitoring, cost management and lifecycle management, help you scale AI systems more safely and reliably.

With these shared foundations in place, teams can move faster over time, because each new AI initiative builds on what already exists rather than starting from scratch. As your AI usage grows, this infrastructure grows with it.



Acquire and scale AI capabilities

As you proceed to integrate AI into your organisation, decide how best to build and scale AI capability – both by developing the expertise your teams need to adopt and deliver AI solutions and by making deliberate decisions about when to buy, extend or build.

Plan for the AI expertise you need

You've got several options for acquiring and building the AI expertise you need for a successful transformation journey.

Partner strategically

External partners that bring proven patterns, deep technical expertise and lessons learned from other transformations can help you accelerate yours. Partners are especially useful for initial platform design, governance set-up and complex integrations and deployments that internal teams may not be ready to handle.

External suppliers can accelerate progress, but it's important that ownership of decisions, governance priorities, risk assessment and accountability remain internal. The goal is to jump-start momentum and build a foundation, while continuing to build internal skills, confidence and leadership over time.

Consider bringing in external partners when:

- You've identified promising AI use cases, but your internal teams aren't sure where to start or how to turn your ideas into practical business value.
- You're facing fixed timelines or your transformation is significantly behind the competition, and you lack the expertise or skills to get up to speed quickly.
- Your business processes allow suppliers to make progress quickly using an out-of-the-box approach.
- To deploy or scale your AI solutions, you need dedicated support and extensive resources that only experienced partners can provide.

Train from within

The most durable path to long-term AI success involves building deep technical expertise in designing, deploying and operating AI systems. Ideally, upskilling should span technical roles and business leaders. It should also cover not only how AI tools work, but also when and why to use them, how to interpret outputs and where human judgement remains essential.

Consider building AI expertise in your existing staff if:

- Your culture is focused aggressively on AI and you want existing staff to embrace experimentation widely within your organisation.
- You're in a competitive position that gives you adequate time to upskill staff to support broader adoption.
- Your business is complex or unique, so it's easier to train your staff on AI than to teach outsiders about your business.
- Budgetary limits restrict additional head count or heavy investment in external consultants.



Hire experts

Experts can help you navigate early design decisions, avoid common pitfalls and establish strong technical and governance foundations. New hires are most effective when integrated with existing internal teams, transferring and building knowledge rather than simply operating in isolation.

Consider hiring experts if:

- Your culture is open to integrating new expertise and you can share expert knowledge more widely.
- You need skills quickly to separate your organisation from competitors.
- You can get a newcomer up to speed on your business fast enough to start making real contributions in a relatively short time.
- You can afford the additional head count.

Ultimately, your Frontier Transformation succeeds when expertise is embedded throughout your organisation. With the right mix of people and external support, you create the conditions to fully benefit from AI.

These choices are never global, exclusive or permanent. Needs evolve. Each use case, agent or AI system might need a different mix of staff, training and partner assistance.



Know when to buy, extend or build custom AI solutions

Custom development can be powerful, but it's expensive and difficult to sustain at scale. Similarly, relying exclusively on off-the-shelf tools can limit customisation that would allow AI to materially improve how your business operates. Successful organisations buy what doesn't need to be unique, extend platforms where they already have trust and integration and build custom solutions only where the returns justify the investment.

AI services generally fall along a spectrum of effort, control and responsibility:

Bring your own model. Infrastructure-as-a-Service (IaaS) approaches are effectively fully custom builds. Organisations select models, manage data pipelines, design architectures, implement security and governance and operate everything from end to end. This approach offers maximum flexibility and customisation, but also carries the highest cost, complexity and operational overhead.

Managed AI development platforms. Platform-as-a-Service (PaaS) solutions provide managed models, data infrastructure, orchestration and operational tooling, while still allowing teams to build custom AI applications. This balances flexibility and speed by simplifying model, infrastructure and governance management and still supports customisation of application logic, integrations and user experience.

Packaged commercial AI solutions. Software-as-a-Service (SaaS) solutions offer ready-made AI capabilities where most features, infrastructure, security and updates are managed by the platform provider. These are typically fastest to deploy and easiest to scale, making them ideal for standardised workloads that don't require unique customisation.

Early in your transformation, buying or extending proven tools may be the best way forward. As AI becomes more deeply integrated with your business, some capabilities might justify deeper customisation or full custom development.

Key questions to guide buy, extend or build decisions

To inform decisions about whether to buy, extend or build AI solutions, develop a clear decision framework that includes the following questions:

- Where does custom development matter for competitive advantage or strategic differentiation?
- Do existing services and platforms offer integrations and customisation features you can extend rather than building a solution from scratch?
- Can solution accelerators, templates and reference architectures help you move faster and more safely?
- Do you have the skilled staff, operating model and governance needed to build and run custom solutions sustainably?
- Does it make sense to acquire low-code or no-code AI tools that allow subject matter experts to build simple AI solutions that streamline or improve their day-to-day work?

Build optimisation and resiliency in your AI infrastructure

As AI tools and agents move from development and testing to production use, performance can become unpredictable. Costs can drift if they're not actively managed. Availability becomes critical as AI systems move from experimental or developmental use to real business processes.

Design for performance at scale. Unlike traditional applications, AI workloads generate bursts of activity that can spike latency and increase resource demands. Address optimisation from the start, so use cases can expand with predictable, repeatable performance as usage increases. Adopt standardised deployment patterns, test performance under realistic workloads and monitor continuously to detect issues before they affect users or business operations.

Build resilience into your AI systems. Resilient AI infrastructure accounts for failure through fallback behaviours, isolation between critical and non-critical components and clearly defined recovery paths. The aim isn't to eliminate every possible issue ahead of time, but to maintain safe, predictable operations during disruptions.ADesign your AI tools and agents to deliver consistent performance at production scale, your architecture to handle agentic and multimodal workloads reliably and your infrastructure to maintain availability during outages or model failures.

Establish a foundation of security and governance. As AI systems scale, security and governance need to scale with them. Establish a baseline that offers consistent identity controls, access policies, monitoring and compliance guardrails across your AI environment. As AI usage grows, make sure this baseline is continually augmented to ensure that you're protecting data, managing risk and operating systems safely and predictably.

Ultimately, optimisation and resiliency are about trust. Your organisation can trust AI systems that perform consistently. Your staff can trust platforms that remain available and predictable. And business leaders can trust solutions with costs they can predict and benefits they can see.



Anticipate roadblocks

High-quality data and a scalable infrastructure are critical to ensuring that your AI systems are useful and gain acceptance. Widespread AI adoption across your organisation accelerates when you focus on these common issues early.

Access to quality data determines AI system reliability. Without data that's accurate and applicable to your use cases, your AI systems might produce inaccurate, biased or unreliable outputs. Ensuring that you have high-quality data that's fit for purpose should be an essential first step in building any AI system.

You can only trust your AI as much as you trust your data. When your data isn't clean, consistent or complete, your AI systems can't be relied on to produce accurate or relevant outputs. If users get inconsistent or low-quality results, their trust in new AI systems rapidly declines, leading to low adoption and stalled business impact.

Lack of a mapped data hierarchy can trip you up. It's common for data to exist in multiple places. When it's inconsistent, which version should your AI system use? Which sets do you trust more? A clear, complete hierarchy of which datasets to prioritise leads to more consistent responses from your AI systems.

Infrastructure needs to meet the changing demands of AI workloads. AI adoption can stall if new AI solutions aren't available, performant and secure. As you deploy more AI systems, infrastructure needs to scale to meet demand and handle errors or outages gracefully.

Carefully consider how you expand AI access and connect datasets. The more datasets an agent can access, the more it can do. However, you need to balance this against the risk that excessive access can inadvertently lead to exposing sensitive or protected information.

Reaching the goal: AI technology and data maturity

With a thoughtfully crafted AI technology and data strategy behind you:

- Your data is recognised and managed as a strategic asset and fuels ongoing innovation.
- Strong governance, quality and access controls are in place, leading to data that can be trusted and used effectively across AI use cases.
- Your scalable, flexible platforms adapt to new workloads, improve performance over time and remain reliable as AI becomes more deeply embedded in operations.
- Rather than being static, these foundations are optimised continually through feedback, monitoring and iterative enhancement.

In this mature state, you can adopt emerging AI technologies rapidly and with minimal friction as your data platforms, pipelines and tooling are ready to integrate new models and capabilities. You'll be able to support faster experimentation and deployment and remain competitive as new technologies and processes emerge.

AI value driver #3

AI strategy and experience

Over time, your approach evolves into an AI factory model, which supports the repeatable, scalable creation and deployment of reliable AI systems your organisation depends on. Trust expands from individual tools to the systems, governance and teams behind them, creating the conditions for sustained AI adoption and usage.

As you expand the use of AI in your organisation, a number of deliberate steps can help your staff and leadership build experience and trust in new systems and ways of working:

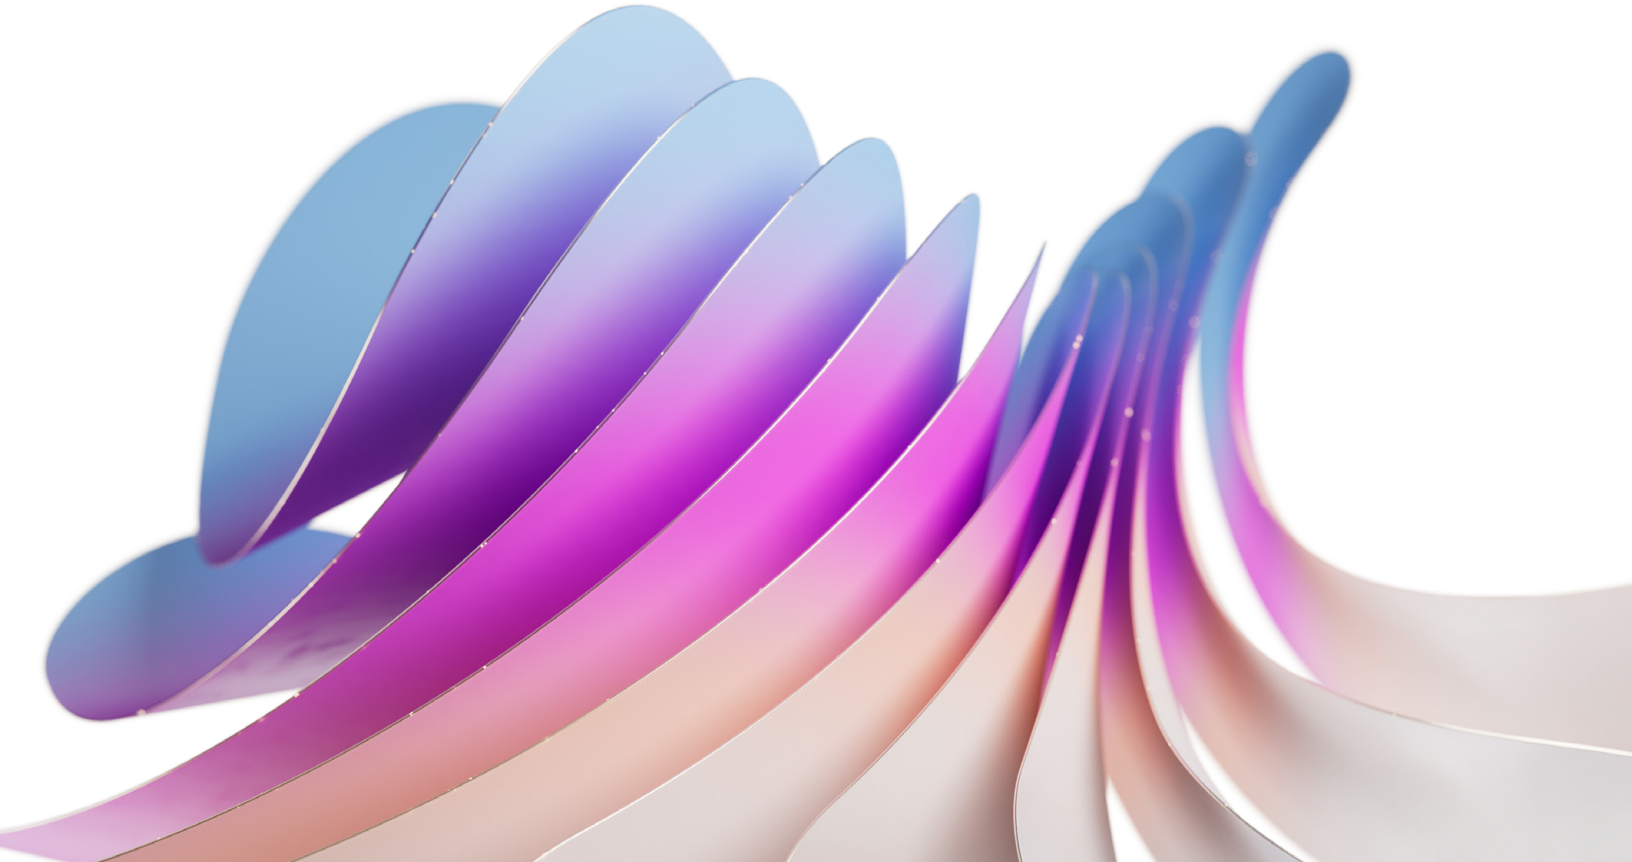
Secure leadership buy-in. By keeping leaders engaged from experimentation to wide-scale adoption, you build trust and justify investment.

Define guardrails and compliance before expanding. Involve the right teams from the start of every project. Ensure that early learnings and infrastructure can be reused at scale.

Build confidence through visible AI wins. Start by proving your AI systems can handle straightforward, low-risk tasks effectively.

Set the end-user change narrative through honest communication during broad rollouts, framing agents as tools that improve work and growth.

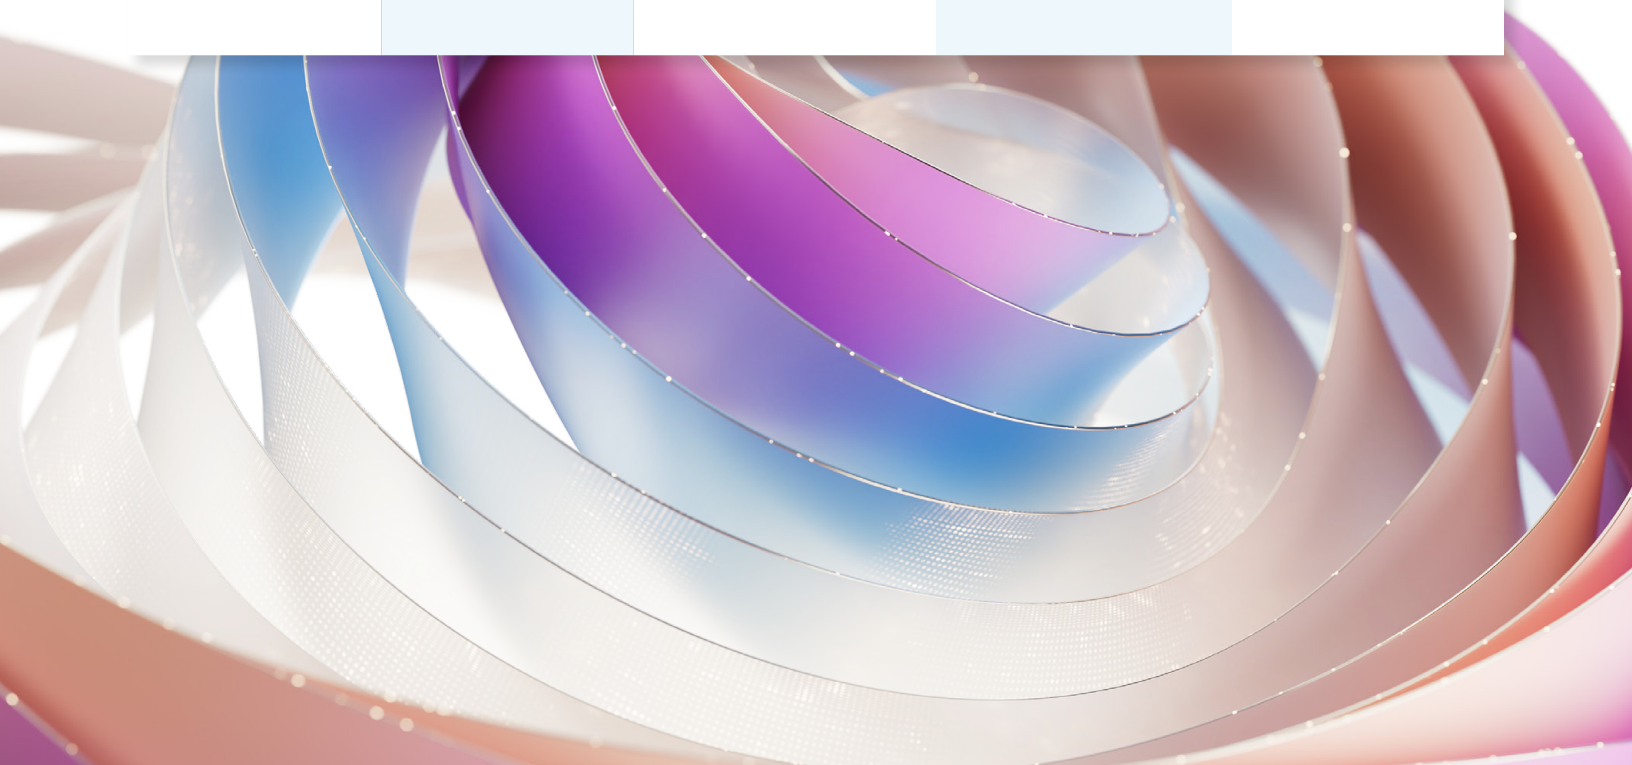
Build on what you have. After you've established trust and broadened adoption, apply agents to improve additional steps in business processes until they're completely transformed.



Evaluate your AI experience readiness

Is your organisation building experience deploying and using AI systems in ways that increase your leadership's trust and confidence in AI systems, while encouraging adoption organisation-wide? Find out which readiness stage you're at currently to help identify what your next goals should be.

Intentional experimentation	Structured piloting	Operational AI practice	Repeatable and reusable AI	Innovative and advanced
• A few individuals or teams running pilot experiments (proofs of concept) • AI expertise limited to specific teams or individuals • One-off or uncoordinated development • Anecdotal success	• Teams actively building multiple AI use cases • Initial AI roadmap or set of prioritised use cases • Repeatable processes emerge from initial projects	• Organisation moves from experimental pilots to operational AI • AI systems in production, teams follow defined workflows • AI operational practices introduce versioning, testing and monitoring of models • AI expertise includes dedicated data scientists/machine learning engineers	• AI development standardised and efficient • Best practices and templates widely used • IT role transformed from guardians to enablers • Many teams have AI experience, knowledge-sharing common • Organisation takes factory approach to agentic AI	• Deep AI expertise, innovation capabilities across organisation • Past projects used to tackle new, more complex AI challenges • AI delivered with agility and scaled reliably across the enterprise



Determine your next steps

Based on your current state, what should you do next to build trust in your AI experiences and encourage wider AI adoption across your organisation?

Create cross-functional AI teams

Effective AI solutions are not just dependent on the right technology choices. Successful AI projects also involve participation and support from various teams across your organisation. These teams bring business, technical and risk expertise together at the start of any AI project. They help create a single point of alignment where priorities, trade-offs and constraints can be addressed holistically rather than sequentially. This shared ownership and coordinated decision-making helps you convert technical experiments into valuable production systems.

To be effective, your AI teams should include a balanced mix of perspectives; for example:

A business or product owner responsible for value, outcomes and prioritisation.

Data and machine learning engineers to design, build and evaluate AI components.

Software or integration engineers to embed AI into real workflows and systems.

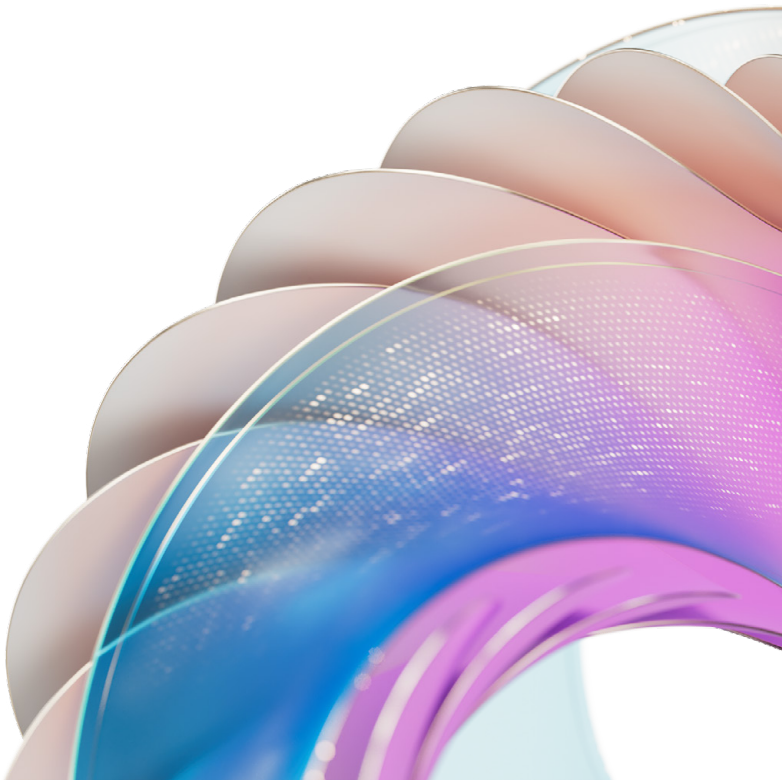
Security, compliance, privacy and RAI leads to help AI systems stay in alignment with your trust, safety and regulatory requirements.

Support intentional experimentation

Experience is the best teacher, and teams learn the most by experimenting with small, safe pilots and proofs of concept. These early experiments help teams build the skills, confidence and practical understanding to adopt AI more broadly.

How you experiment matters. If you don't have defined goals or know how to measure success or what lessons to learn from pilots and proofs of concept, you risk wasting effort and delaying progress.

Intentional experimentation is key. Define what success looks like up front, use experimental pilots to create learning opportunities and put guardrails in place to keep this experimentation safe. Human-verified results in controlled environments help reduce fear and build trust with stakeholders and users. The goal is to create fast feedback loops that turn ideas into organisational learning and tools that deliver real business value.



To get the most out of your experiments, keep these key points in mind:

Design experimental pilots to maximise learning, with clear hypotheses, success criteria and mechanisms for capturing and sharing lessons learned.

Balance speed with responsibility, ensuring that rapid iteration doesn't bypass safety, security or RAI practices.

Establish governance early, so experimenters know the boundaries and can innovate confidently within them.

Provide solution accelerators, such as templates and design patterns, to create AI solutions using trusted architectures.

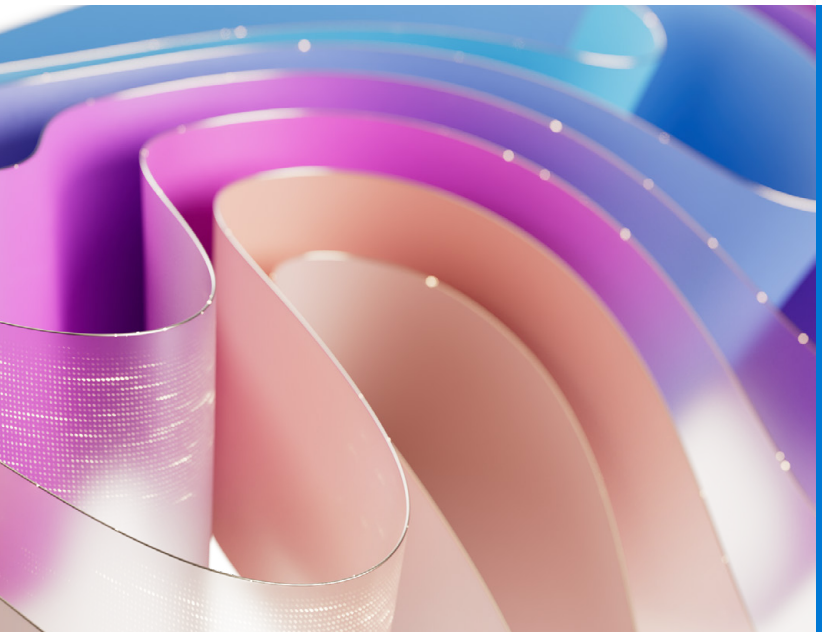
Evaluating your experiments in a safe environment is critical. To limit risk, validate AI tools and agents with trusted internal users before you deploy more widely. Smaller internal deployments give teams space to identify common issues, like hallucinations, prompt vulnerabilities and integration gaps, while containing risks to reputation, compliance and customer trust.

Support experimentation with controls and governance

When nontechnical staff can build internally useful AI tools and agents using well-governed low-code and no-code platforms, adoption and innovation increase dramatically. Facilitating widespread user experimentation works best when you maintain strong controls while providing users with the flexibility and capabilities to build solutions.

Problems emerge when IT imposes bottlenecks that make safe, approved paths slow or difficult, leaving your staff looking for faster alternatives. Frustration can lead to experimentation with unauthorised public tools or moving data outside governed environments in a misguided attempt to see progress, unintentionally increasing risks.

Avoid this pitfall by making the *safe path* the *easy path*. Shift the role of IT from owning every solution to governing the systems and data that users experiment with. If you integrate security controls and responsible AI policies into the tools that make AI experimentation easy, you end up strengthening governance rather than weakening it.



> "The risk to introducing friction internally is a lot lower than starting directly with customers where we have SLAs [service-level agreements] and could impact performance.
>
> – VP, Financial Services

Encourage reusability

Instead of starting each use case from scratch, build continually on what you've learned. As your AI capabilities compound over time, you can transform isolated successes into organisational wins.

To encourage reuse and shared learnings from your AI projects:

- **Create and maintain a library of reusable AI assets.** A curated asset library containing patterns, prompts, components, validated models and governance workflows helps spread AI knowledge across your organisation.
- **Document best practices in a structured, version-controlled format that's required reading before building or deploying AI solutions.** Quality documentation is critical for scalable and responsible AI development and operations.
- **Establish mandatory feedback loops that feed directly into your reusable asset repository.** Feedback mechanisms, such as postmortems, retrospectives and pattern write-ups, turn hands-on, project-specific lessons into institutional knowledge.
- **Embed standardised governance, safety, data handling and evaluation requirements in all reusable templates and pattern libraries.** Embedding guardrails and best practices directly inside reusable assets helps prevent accidental non-compliance.
- **Implement scalable generative AI operations and agentic development and operations (DevOps) practices.** See a detailed breakdown of implementing GenAIOps and agentic DevOps on page 35.

As models change, data shifts, tooling updates and your business goals evolve, keep your documentation and assets updated. Ensuring that these resources stay in line with the latest learnings and experience will improve the baseline quality of new projects, so your AI capabilities improve consistently over time.



> "Each successful implementation gave us a blueprint we could reuse. Over time, that created a kind of playbook that made expanding into new use cases much more straightforward.
>
> – Chief Procurement Officer, Healthcare

Implementing GenAIOps and agentic DevOps

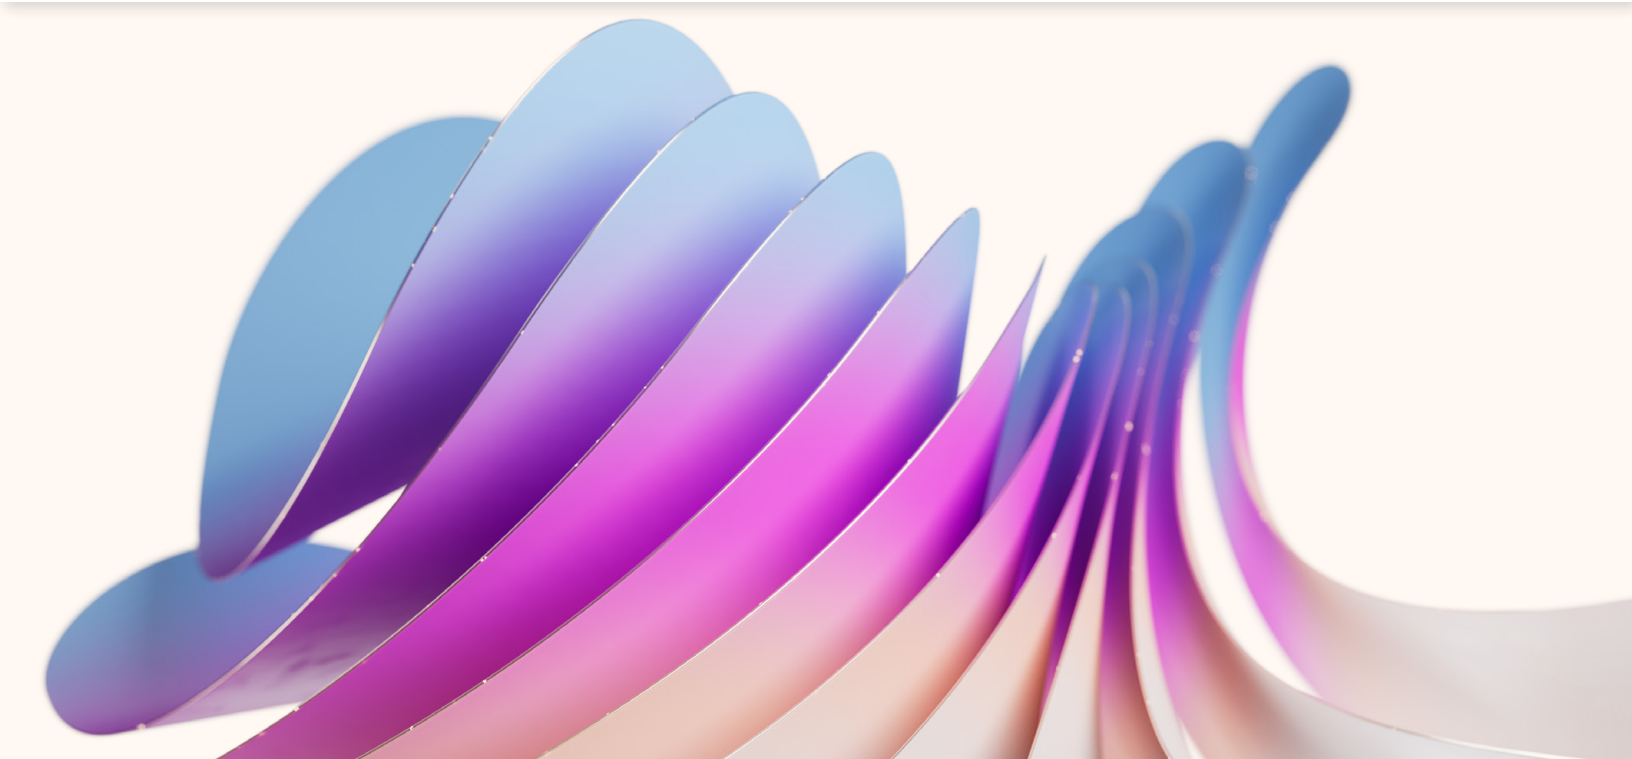
Generative AI operational practices (GenAIOps), tools and operating processes help you reliably deploy, manage, monitor and evolve generative AI systems. They extend traditional machine learning operations (MLOps) practices to account for the unique characteristics of generative and agentic AI, covering topics such as prompt management, model updates, data drift and evaluation of non-deterministic outputs.

Agentic DevOps builds on this by integrating AI as a core component in running and optimising technology operations. It uses agents to accelerate software delivery, improve operational efficiency, reduce downtime and automate complex IT workflows.

Operational practices that account for AI enforce versioned data, models, prompts and pipelines. They enable capabilities like continuous testing and model drift monitoring.

Agentic DevOps extends this to agent logic, workflows and integrations, bringing more predictability to change. Together they enable reusable templates, patterns and orchestrations that give teams a head start on development and accelerate delivery across use cases.





Measure what matters

In AI projects, measurement is focused on demonstrating business impact, building confidence in new systems and steering the direction of your future efforts. In practice, that means looking beyond ROI alone to include adoption, trust, system quality, risk exposure and operational efficiency.

Align metrics to real outcomes, not assumptions. Rather than tracking everything, define a small number of evaluation criteria that are specific enough to make progress visible over time. With narrow, well-chosen metrics and KPIs, stakeholders can see improvement clearly, support informed investment decisions and create shared confidence that AI initiatives are delivering value.



Evaluate AI pilots as systems in progress, not disposable experiments. Instrument them from the start with telemetry, user feedback mechanisms and performance monitoring, reflecting system behaviour in actual use.

Decision-oriented metrics and KPIs can help you evaluate system impact and potential. These often target business outcomes (such as time saved, quality improvements or adoption rates), system performance (like latency, reliability and accuracy), risk and safety signals and cost or efficiency measures. The right mix depends on your use case, but the intent is always to demonstrate value, identify issues and inform future efforts.

Insights benefit your next iteration, guiding you as you refine models, adjust workflows, improve prompts and modify guardrails. This feedback loop allows you to build ongoing support from leadership while growing small initial pilots into substantial systems that support core business functions.

> "What really moved things forward was being able to show results early and keep leadership engaged in that progress. At the same time, we had to make sure the data foundation was solid so the outputs were reliable. Once people saw consistent performance and real impact on day-to-day work, that's when trust built and we could start scaling more broadly.
>
> – VP of Customer Engagement, Healthcare

Anticipate roadblocks

AI systems that lack clear goals, don't address user needs or haven't been tested properly can increase scepticism and slow adoption.

Build trust through small wins. Each incremental success builds experience as you increase AI capabilities gradually over time. Intentionally narrowing scope and delivering clear results shows stakeholders and users the value of AI and encourages further adoption.

Set achievable goals to ensure success builds trust. Start projects to solve a single, well-defined use case that you're confident will bring benefits in production.

Only add new functionality after you've proven that a system is stable and valuable. A narrow scope and disciplined approach make benefits clear, so leadership can appreciate the value of new systems.

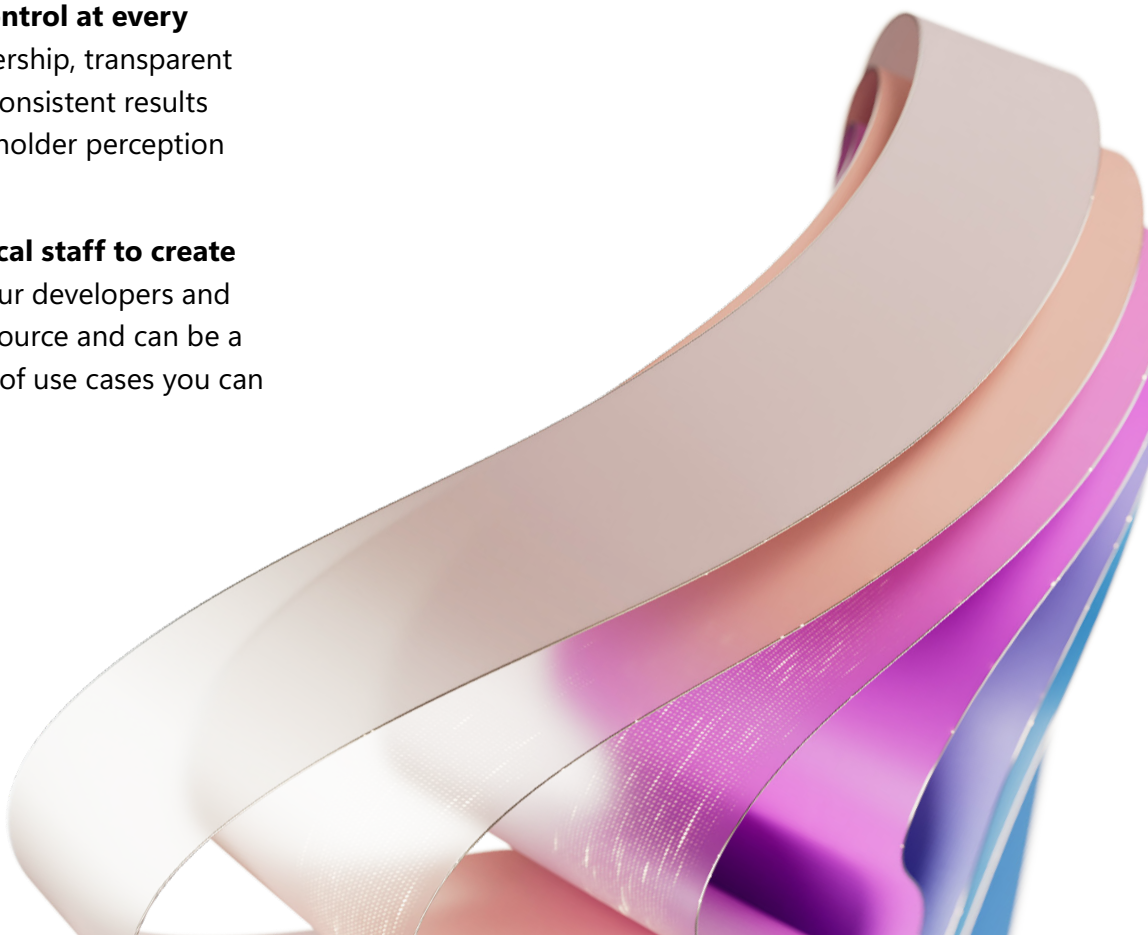
Maintain visibility and control at every step. Oversight, clear ownership, transparent performance metrics and consistent results reduce fear and shift stakeholder perception from uncertainty to trust.

Find ways for non-technical staff to create AI-powered solutions. Your developers and AI experts are a limited resource and can be a bottleneck on the number of use cases you can tackle at any one time.

Reaching the goal: An advanced and innovative AI ecosystem

As experience compounds, your organisation will develop deep expertise in building AI solutions and establish proven deployment and operations patterns. Lessons learned from past projects, along with the trust built gradually with business owners and stakeholders, allow you to take on more complex challenges over time with greater confidence and lower friction.

As a result of this process, you can deliver new AI solutions with greater speed, discipline and consistency across your organisation. Teams can reuse what works, adapt it to new contexts and scale it through trusted processes, so you can innovate continually while maintaining reliability, agility and control.



AI value driver #4

Organisation and culture

According to the Microsoft 2026 Work Trend Index Annual Report, organisational factors account for twice the impact of individual behaviour, driving 67% of the realised value from AI. For many organisations, organisational and cultural factors represent a much bigger impediment to scaling AI adoption than technical capacity.

Successful transformation depends on rebuilding your operating model to focus your teams less on step-by-step task execution and more on setting direction, defining standards and evaluating outcomes using human judgement. Without this shift, even the most skilled employees can fail to fully harness AI or turn new technology into institutional advantages.

Introducing new ways of working into your organisation is a key challenge for leadership, as is finding ways to make changes stick. Even the strongest executive vision will stall if teams lack the skills, confidence, support and clarity needed to work effectively alongside AI systems. Success depends on ensuring that the metrics, incentives and expectations are in place to reward your workforce for adopting new ways of doing their job.

This journey isn't linear, and readiness within your organisation will be inconsistent. Teams, functions and workloads might be at very different stages of AI adoption and acceptance, with some groups building advanced agents and automations while others are just beginning to understand how AI fits into their daily work. Recognise and accommodate these differences, and introduce processes that encourage ongoing change within your organisation.

> It all comes down to culture. Without organisation-wide adoption, it's difficult to push AI initiatives forward. Scaling AI depends on broad alignment and buy-in.
>
> – HR Head, Technology



Evaluate your organisational and cultural readiness

Are you building the skills, confidence and support structures your teams need to work confidently and effectively with AI? Identify where you are today to understand what organisational and cultural changes you need to make going forward.

Uncoordinated efforts	Forming AI governance	Organised and supported	AI-empowered culture	AI-first culture
• AI happens in silos driven by enthusiasts • No formal AI roles or governance structure • Culturally, awareness just beginning • Few employees with AI skills • Leadership interest cautious, with no unified vision	• Leadership articulates AI vision and sponsors initial projects • AI working group or Centre of Excellence (CoE) established to coordinate AI efforts, share learnings • Culture shifting from scepticism to curiosity as people see practical results from initial use cases	• Formal Frontier CoE, clear operating model in place across departments • Leadership communication, change management drive user adoption • AI champions support adoption • Create training programmes that build AI skills across your organisation	• Organisation embraces AI across all functions, executive roles or councils steer AI strategy • Most employees have some AI literacy • Many roles redefined to work alongside AI, and new roles (AI leads, data scientists, machine learning engineers) integrated into your teams • Employees actively seek AI-driven improvements in their work	• AI ingrained in how your company operates • Employees use AI to support decision-making, underpinned by strong change management, leadership vision • Robust AI CoE oversees portfolio of AI solutions, with feedback loops • Every function considers AI in its strategy • Workforce is agile and continually reskilled

Determine your next steps

Based on your organisation's current state, what are your next steps to enabling an AI-ready culture that embraces change, continually builds skills and has enthusiasm for harnessing the power that AI systems offer?



Establish centralised guidance, standards and support

When leaders clearly articulate how AI connects to business outcomes, invest in building workforce capability and actively encourage experimentation, they create momentum for change. However, these efforts can just as easily lose momentum without a unifying structure to guide them. Guidance, standards and training support give teams direction, tools and confidence to adopt AI responsibly and at scale.

A Frontier Centre of Excellence (CoE) team is a powerful tool to meet these needs. The Frontier CoE establishes central guidance and documentation for designing, deploying and scaling AI. It isn't tasked with owning solutions; instead, it sets the conditions for success and supports responsible AI usage across your organisation.

A central Frontier CoE team can help your organisation by:

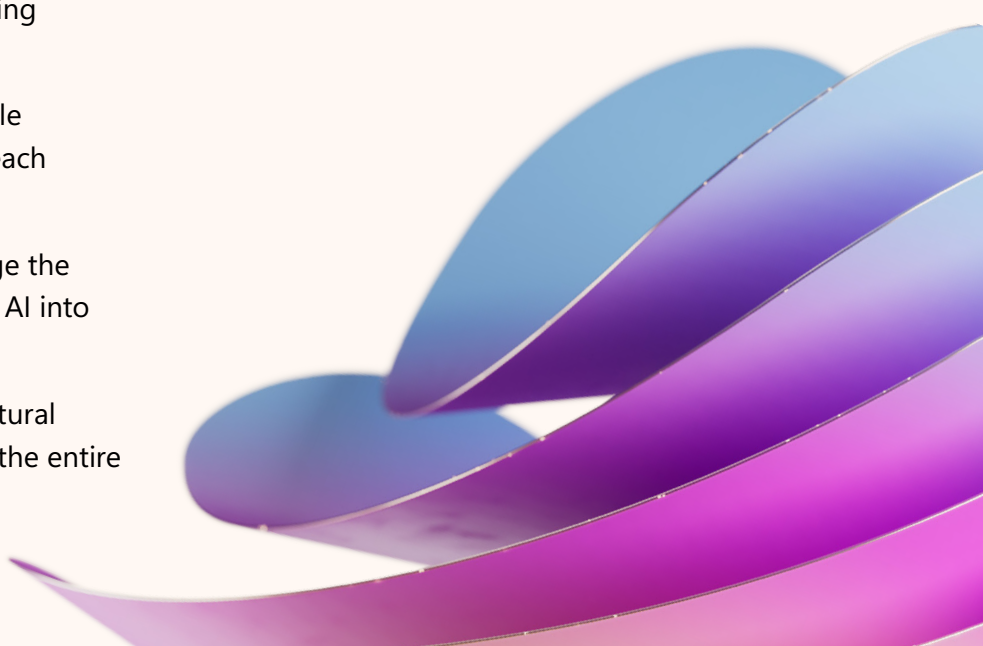
- Defining standards, governance and best practices.
- Establishing shared patterns, guardrails and platforms.
- Creating and encouraging skill-building efforts.
- Identifying and documenting reusable patterns so you don't need to start each project from scratch.
- Establishing processes that encourage the integration of responsible, governed AI into all AI solutions from the start.
- Encouraging faster adoption and cultural change by involving stakeholders in the entire process.
- Keeping individual projects aligned with your broader AI strategy.

The Frontier CoE won't have all the answers from day one – it will expand iteratively as AI adoption grows. To identify challenges and address them early, integrate feedback from employees and end users continually, keeping guidance, standards and support resources up to date as AI use increases.

With strong executive sponsorship, a dedicated cross-functional team and adequate resourcing, your Frontier CoE can balance oversight with empowerment. With Frontier CoE support, your teams can create solutions more easily while maintaining trust, security and alignment with organisational goals.

Provide organised support for change

As you move beyond initial experimentation, your success depends on putting the right foundations in place. You need to communicate with employees about your AI vision, equip them with AI skills and encourage progress. By managing how AI is introduced, understood and adopted across teams, you can reduce resistance, align your workforce with your standards and encourage adoption.



Establish a change management team

Leadership defines and communicates your organisation's AI vision, which should help employees understand how they'll adapt to new technology and new ways of working. Change management teams work with leadership to communicate that vision and establish the critical narrative that manages fears and resistance.

This process helps your staff understand the purpose behind AI initiatives and can alleviate their AI concerns or fears. Previously reluctant staff start to see how AI improves their work and delivers business value.

Change management can help you encourage AI adoption and build acceptance of ongoing change in several important ways:

Make sure people understand that agents are enablers, not replacements. Clear messaging reduces fear by positioning agents as productivity tools, not job threats.

Use experience to change minds. After employees see their productivity and capabilities grow through hands-on usage and experiences, scepticism and resistance fade.

Focus on people, not technology. Acceptance and adoption depend more on communication, reskilling and encouragement than on AI technology.

Start building skills

Lack of AI skills and knowledge can create uncertainty and scepticism about new systems and processes. Introducing training and skill-building programs can help your teams better understand what AI is, how it applies to them and how it will affect their day-to-day work.

People equipped with the skills they need to harness AI are far more likely to engage with new tools and processes, encouraging adoption and use of new AI systems.

Beyond basic AI literacy awareness and specific technical skills, training programmes also align staff with the standards and best practices defined by your CoE team. They help you establish a shared baseline of AI knowledge and practical capabilities, improve adoption and drive better, more informed decisions across your team about using AI responsibly and effectively.

To help your workforce build these capabilities, explore learning platforms, like **AI Skills Navigator.** These resources allow you to discover role-based learning paths, track progress and scale AI skills across your organisation.

> "There was definitely some initial concern from teams about what this meant for their roles. We had to be very clear that this wasn't about replacing people, it was about helping them offload repetitive work so they could focus on more meaningful tasks. That shift in messaging made a big difference in adoption.
>
> – VP of Strategy, Financial Services

Use AI champions to support change

A community of AI champions can address leadership and employee concerns and build enthusiasm for your AI vision. Choose champions who have enough hands-on experience with your AI projects to understand technical details and potential benefits. They should also understand leadership's business goals, so they can support your AI efforts in several ways:

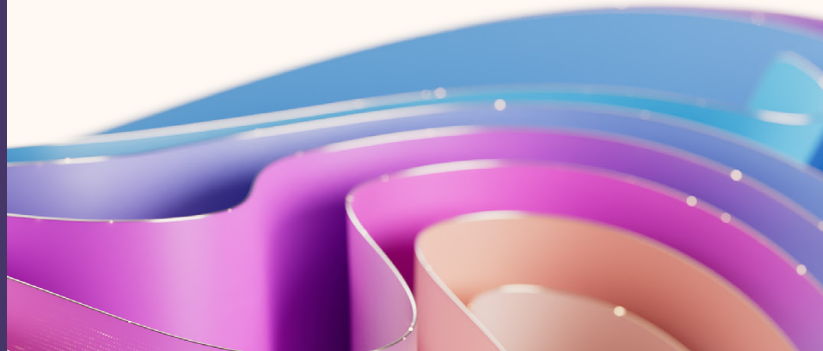
- **Translate progress into results leaders care about.** They connect day-to-day development with outcome-focused executive expectations.
- **Maintain momentum between projects and leadership.** They track wins, surface blockers and push decisions forward.
- **Prevent pilot drift.** By owning visibility and narrative, champions help maintain the focus of AI projects on their real business goals.

> "This isn't something where you just roll it out and you're done. You have to keep investing in training and helping people build confidence using it over time. That's what actually leads to sustained adoption.
>
> – VP of Client Services, Professional Services

Spread AI learning across your organisation

Critical thinking, communication, creativity and adaptability become even more important as AI takes on more routine tasks.

- **Expand training and upskilling efforts as AI use grows.** Focus on complementing employees' strengths as you grow proficiency in new AI skills.
- **Invest continually in employee skills and AI fluency.** People trained on new AI systems are more likely to trust them, to see how they reduce routine workload and to identify new opportunities to improve or redesign workflows.
- **Develop an AI learning strategy.** Don't rely on a one-size-fits-all approach. Assess each team for AI readiness, including awareness, mindset, skills and confidence. Adopt role-based training that aligns learning to different roles and groups. Incorporate hands-on projects to turn learning into practical capability that encourages adoption and delivers business outcomes.
- **Build a learning community.** Create internal communities where people can exchange best practices, document what works and learn together through activities such as hackathons, ideation sessions and shared resources. Involve leadership, business owners, developers and trusted AI champions to sustain momentum, reinforce credibility and spread best practices across your organisation.



Foster an AI-driven culture

In an AI-first culture, leaders visibly encourage and participate in AI-enabled workflows, instead of treating them as isolated experiments. AI-first cultures recognise and encourage teams that deliver real, measurable outcomes through AI and encourage widespread experimentation.

Innovation-driven organisations don't assume that AI is flawless. Instead, they understand that iterative process builds learning and experience leads to continuous improvement. The most important mindset shift is moving from *prove it works* to *help me get better at it*. When teams feel safe to experiment, refine and improve AI-enabled workflows over time, AI adoption and innovation accelerates rapidly.

Culture shifts fastest when people see AI succeeding in real work, supported from the top and reinforced through incentives, recognition and shared success. Frontier Firms share the following traits:

- **They set the pace based on readiness, not technology.** This includes expanding agent and tool use in manageable waves, while matching rollouts to your workforce's readiness and skills.
- **They maintain sponsorship from executives and other leaders.** When executives stay involved beyond kick-off, they signal that agent deployment is a priority. Appointed agent champions can help you further bridge the relationship between staff and leadership.

> "Once they saw it could eliminate repetitive work, they were much more open to it. It became clear this was about productivity, not replacement.
>
> – VP of Tech Solutions, Financial Services



Anticipate roadblocks

Transformation is often slowed by a combination of barriers related to leadership, culture and skills:

Without a shared vision or consistent executive sponsorship, organisations can struggle to align their efforts with overall AI strategy.

Risk-averse cultures, resistance to change and general fear surrounding AI technologies can discourage experimentation and the adoption of new tools.

Insufficient education or skill-building can further slow progress because employees lack the confidence or ability to use AI tools effectively.

Identifying, understanding and overcoming these barriers allows your organisation to adapt, learn and innovate using AI. When leadership is aligned, risks understood and managed and staff trained in the right skills, experimentation can evolve into usable production tools that bring real business value.

As you approach your own organisation's transformation, consider these barriers you might need to address:

Is AI adoption perceived internally as worth the cost and risk involved?

Stakeholders and leadership need to see clear value in AI to move from curiosity to conviction. Executives must understand the benefits of AI and justify the investment.

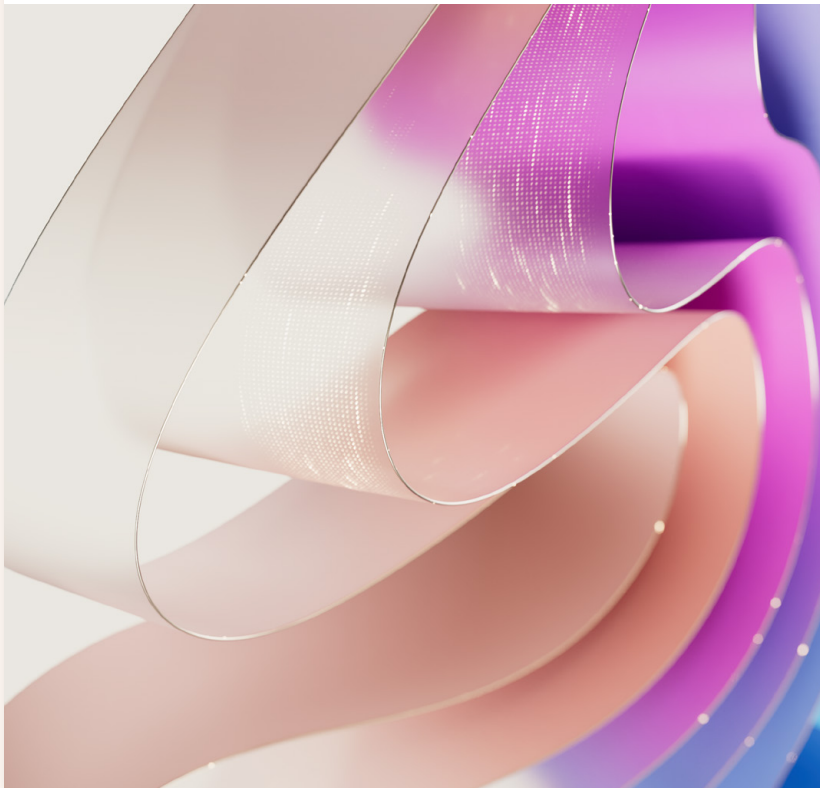
Risk prevention isn't an impossible challenge. Educate stakeholders on how AI risk management is a manageable process, not a massive barrier.

Build a culture that supports transformation

Employees are understandably cautious about changing core business systems or processes. However, as AI becomes foundational to productivity and day-to-day work, uneven adoption or resistance to change can create friction, slow progress and dilute the value of AI investments.

Getting people and teams to accept the new ways of working isn't just a matter of skilling and tools. It requires a culture that supports transformation. Clear leadership messaging, strong change management, visible success stories and accessible training programs can help your employees understand how AI supports their work. Creating opportunities for hands-on learning, peer support and safe experimentation allows sceptical teams to engage at their own pace without feeling left behind.

After your AI systems are proven reliable and essential to your business, these new ways of working aren't optional. Adoption becomes far smoother when your organisation actively supports staff through this transition.



Are there hidden or unspoken cultural barriers holding you back?

- Without ongoing executive support, momentum can stall. Engage with leaders to actively support your efforts, beyond just granting approvals.
- Competitors and customers can be forces of change. Leadership needs to be aware of how agentic systems can help strengthen your competitive edge and improve customer experiences.
- Enlist external partners and consultants to help you identify bottlenecks, align initiatives with business goals and reassure stakeholders.
- See your company's culture as an accelerator, not a brake. As leadership sees benefits, have them reward AI innovation, even if change is difficult.

Does your organisation believe that it has the right skills and capabilities?

- Your leaders and staff need to have confidence that they can succeed with AI technology. However, it's also important they understand that perfect readiness isn't a prerequisite for progress. Encourage forward movement even in the face of perceived skill or capacity gaps, knowing that these will be addressed as your teams gain experience.
- Risk management is critical, but insisting that 100% of risks be resolved up front will halt momentum. Leadership needs to give teams permission to act while your policy, governance and risk mitigation strategies are reviewed and improved.
- Ensure that you've put internal review, testing and support mechanisms in place that allow experimental pilots to safely grow into core business capabilities.

Reaching the goal: Frontier Transformation

When your organisation develops a mature Frontier culture, AI is fundamentally embedded in how you get work done. Employees rely on AI to inform decisions, augment judgement and accelerate outcomes. Change management supports a leadership vision, helping to ensure that AI adoption is purposeful, consistent and directly tied to business goals.

Both structure and culture sustain success at this stage. A robust AI CoE provides oversight, guidance and continuous feedback across a portfolio of AI solutions, ensuring that innovation remains aligned with governance, business priorities and user needs. AI becomes a shared capability that connects strategy, execution and outcomes across the enterprise.

Continuous upskilling allows your organisation to evolve alongside technology and opportunities. Employees have the skills and support to experiment, learn and improve how work is done, while leadership reinforces a growth mindset and confidence in AI.

AI value driver #5

AI governance and security

Governance isn't just about faster and safer innovation, and it's not optional or 'nice to have'. Regulatory expectations require controls and compliance mechanisms that are built directly into AI systems. Design proper governance into AI tools and agents from the start.

AI governance focuses on creating the rules, standards, systems and controls that drive trust. It's how you establish clear ownership, consistent guardrails and visibility across your business, technology and risk teams. With this foundation in place, you create trust with stakeholders and end users and establish what's allowed, what's expected and where the boundaries are.



"

We're actively developing governance guardrails, policies and controls for the rules and acceptable use cases for use of AI, but it's still a work in progress. These frameworks and standards are being developed alongside practical use cases and continue to evolve.

– IT Director, Retail

How is agentic AI governance different?

AI systems, and the content they generate, must be governed, reviewed and audited like any enterprise system. Agentic AI introduces some additional unique governance challenges:

- **Agent sprawl and resource access.** Maintaining clear ownership, consistent identity controls and well-defined access boundaries is important as agents multiply. Unmanaged sprawl can lead to over-permissioned agents, a larger attack surface and increased compliance risk.
- **Data oversharing.** Agents interact with many systems and users, raising the risks of exposing sensitive data if systems aren't properly managed and controlled.
- **Emerging AI-specific threats and vulnerabilities.** Attacks like indirect prompt injection and related vulnerabilities may not show up in traditional security models. Proactively protecting against malicious actors requires specialised attention.
- **Increasing regulatory pressure.** Governments and organisations are still trying to figure out their response to an increasingly agent-driven world, and laws change frequently. Staying on top of compliance is more important than ever.

Evaluate your AI governance and security readiness

Understanding your current governance state is the first step to identifying what's needed to make trusted AI possible. Are you implementing RAI guardrails, AI security and standardised dev and testing practices? Knowing where you are on the governance maturity journey lets you know where you should be heading.

No formal governance	Basic guidelines	Emerging governance	Integrated governance	Trusted and auditable AI
• AI experiments proceed without official guidelines • Little management of ethical, security or regulatory exposure • Employees embracing unsanctioned AI tools (shadow AI) • Lack of operational visibility (limited telemetry, monitoring, logging)	• Initial AI principles or guidelines drafted to steer projects ethically • Key risks (for example, privacy and compliance issues) identified for each pilot • Security reviews start for data and models in AI projects	• Governance processes in place for all new agentic AI deployments • Governance team reviews all major AI projects • Security and privacy controls integrated into agentic AI solutions • RAI introduced	• Governance built into agentic AI development lifecycle • Incident response plans specific to AI systems identified • Automated monitoring for security posture, not just model performance • Dedicated bodies (AI ethics committee, risk officers) oversee enterprise AI use	• Organisation has industry-leading AI governance • RAI part of the DNA • AI systems transparent, auditable and meet high ethical standards • Governance provides measurable business impact on trust, risk, business outcomes

Determine your next steps

Based on your current state, what should you do next to implement and improve governance and security practices that help you deploy safe, trusted AI at scale?

Define a RAI framework

A strong **RAI** framework establishes the 'rules of the road' for AI, defining clear expectations for fairness, transparency, privacy, security, accountability and human oversight in your systems. It lays out a structured set of principles, practices and governance processes to help you create AI systems that are ethical, trustworthy and aligned with your organisation's values, legal obligations and societal expectations.

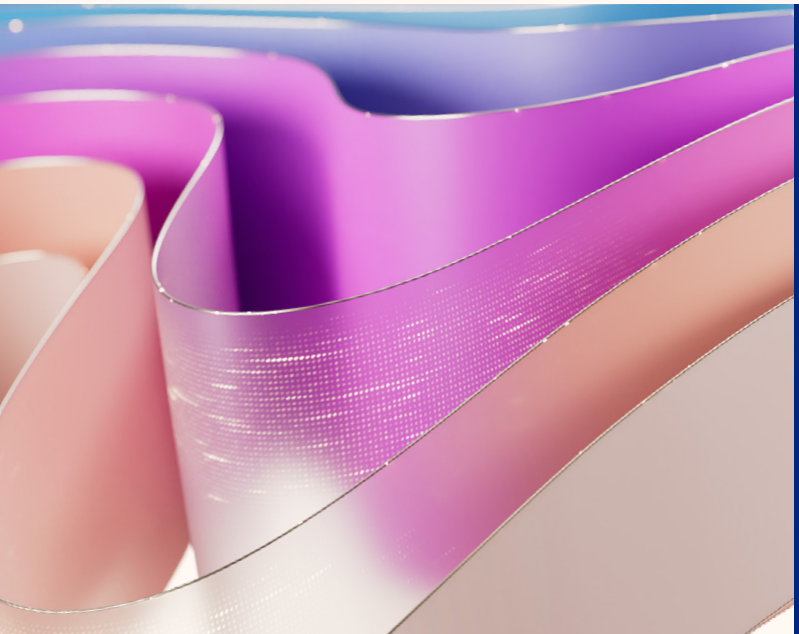
Defining your own RAI framework means moving beyond simple technical questions, like *Can you build an AI system?* Consider how your organisation's values affect AI development and how your AI use will affect your legal and ethical responsibilities.

Your AI CoE team and your RAI framework are complementary. The RAI framework defines what *responsible AI* means for the organisation, and the Frontier CoE embeds this understanding into policies, technical controls, review processes, tooling and day-to-day execution. Over time, it helps evolve the framework based on new regulations, lessons learned from deployments and changes in organisational risk tolerance.

For the framework to be effective, your leadership and AI champions must visibly support it and make clear to everyone that RAI is a shared responsibility rather than a compliance afterthought.

Establish AI governance

AI governance can exist on a spectrum from informal, poorly structured IT and security reviews to a well-established infrastructure that guides the entire AI lifecycle. As you move from experimental AI to broad, organisation-wide adoption, upgrade your governance infrastructure to trusted, observable and auditable AI systems.



> If you're going to use these tools, you still need the same level of responsibility as if you did it yourself. You can't rely on the system or an external solution to own that.
>
> – VP of IT, Financial Services

Understand the key components of AI governance

AI governance is based on a coordinated set of practices. As AI adoption grows, you'll address ethical expectations, security and risk management, data protection and consistency issues across teams and environments.

An RAI framework and principles define your ethical, legal and operational expectations for AI. Through a shared set of values, you define the guardrails you'll build and shape the associated policies, roles, review processes and technical controls. This helps to ensure that your AI systems behave safely, align with organisational values and regulations and build trust with employees, customers and regulators.

AI security and risk management are critical components to governance, providing the controls, visibility and accountability needed to support safe, compliant operations aligned with policy. Establish a central control plane for risk, along with continuous monitoring and threat detection systems that support consistent identity, access and policy enforcement at scale.

Data security and compliance mechanisms help to ensure that your AI systems handle information responsibly, protect sensitive data and meet regulatory and organisational requirements. As AI tools and agents move into production, data flows become more difficult to manage, so consistent data privacy controls are critical to avoid unintended exposure of regulated data. Data assessments and monitoring need to be ongoing, with continuous legal and compliance involvement as use cases evolve and regulations change.

Consistent, organisation-wide governance enforcement applies controls across all systems and teams. This helps to prevent fragmented or uneven controls from creating blind spots, inconsistent behaviour and actions that can't be traced or audited. Automation supports scalable governance, with continuous scanning and monitoring to surface issues early.

For all these components, the goal isn't perfection from day one. Instead, gradually develop a governance model that keeps up as your AI usage expands and that helps you avoid hidden risks along the way.

> At the end of the day, people will only trust the system as much as they trust the data behind it. If there's any doubt about the accuracy or completeness, adoption drops off quickly.
>
> – Global Head of Business and Marketing, Financial Services



Integrate good governance throughout your organisation

AI governance requires consistent and scalable controls across people, platforms and agent lifecycles. With your RAI framework, you set guidelines that determine how your organisation should develop, deploy and use AI solutions. The next step is to determine how you'll apply and manage these decisions in a practical way.

Establish a cross-functional AI governance committee. In alignment with your RAI framework and your Frontier CoE team's guidance, this group evaluates, monitors and approves AI initiatives. Their leadership and technical expertise inform decisions about projects based on ethics, risk, privacy, security and technical feasibility.

Enforce consistent governance, security controls and observability across your organisation. AI systems are multilayered, with models, prompts, data, APIs, agents and integrations all representing potential risk or compliance challenges. Make sure that you can evaluate risk and impose controls uniformly across systems and environments. Build safeguards and guardrails directly into your AI systems.

Implement automated security, compliance and safety scanning across all AI systems, prompts and agents. As AI scales and security and compliance staff are increasingly stretched, automated scanning solutions can identify issues like unsafe outputs, bias, data leakage, insecure configurations and non-compliant components.

Establish a unified lifecycle governance process for AI. Governing AI systems requires knowledge of the AI tools and agents that exist and what they do. A registration and lifecycle management system can track all your AI systems, from approval to deployment to retirement. Integrate this system into your agent creation processes and tooling to make it easier for teams to build compliant systems.

> "One of their biggest concerns was, does it violate any contractual agreements that we already have with the suppliers? Is the agent going to do something that breaks that? Well, no, I'll just program it not to.
>
> – Senior VP of Insights, Technology



Identify and mitigate AI security risks

AI and agentic systems face a unique threat landscape. Traditionally, your focus has been on protecting your systems by controlling access and validating inputs. Although these concerns still apply, AI systems must also protect against the manipulation of a model's reasoning and content generation capabilities, the potential for learned behaviour and the integration of unapproved AI systems (shadow AI).

Agents also introduce risks related to their autonomy and ability to act independently. They connect to various tools, APIs and data sources which, if abused, can lead to unauthorised actions or data leakage. Unlike traditional AI chat systems, agents often add persistent memory or context that malicious actors can potentially exploit to compromise systems and access sensitive information.

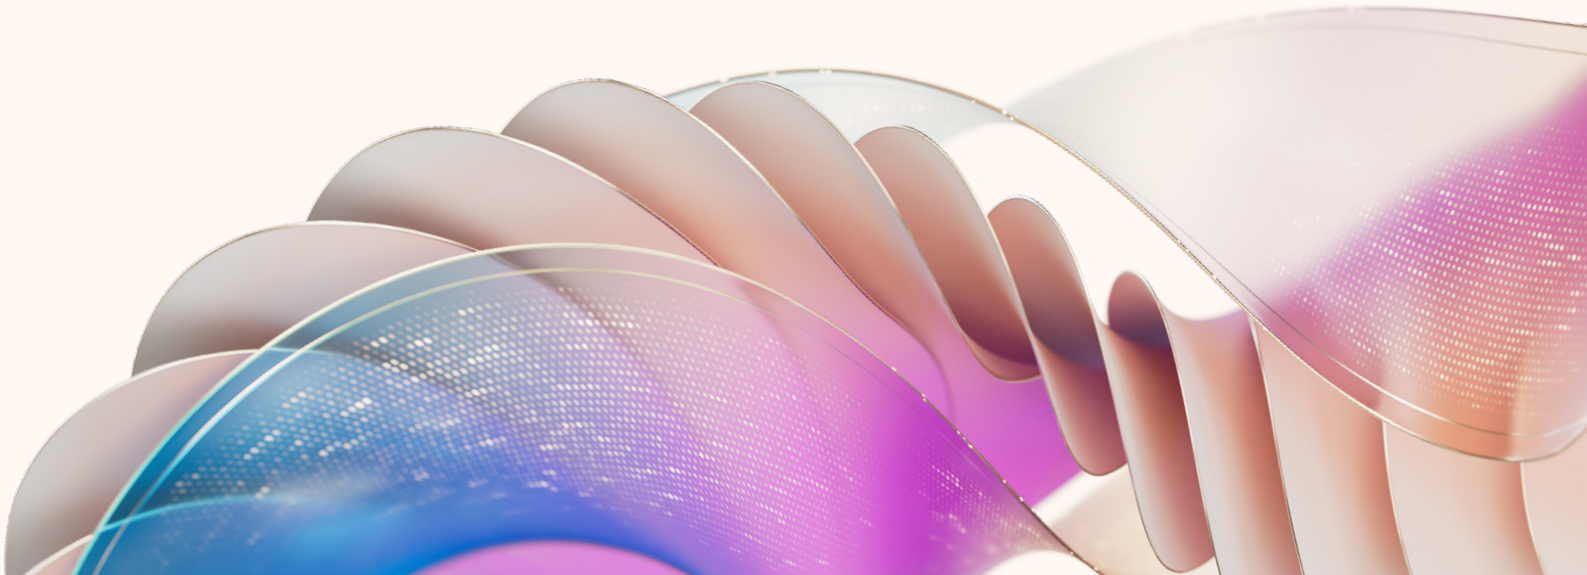
These risks are manageable, but security teams can't solve them on their own. As your AI tools and agents scale, governance, observability and shared controls matter even more and security needs to be a holistic, shared responsibility. Developers and IT teams play a critical role, alongside security, in setting boundaries, enforcing access and monitoring behaviour.

Consider these steps to address the unique risks introduced by agentic AI systems:

Apply identity management and access controls to agents. Agents share many of the same identity management and access control issues that people do. As agents interact with more of your systems, they need identities, defined permissions, access to resources, lifecycle controls and monitoring for suspicious activities. Integrating agents into your existing identity management systems is an easy way to better protect them.

Create data security controls. Implement consistent data classification and access mechanisms, automated monitoring and auditing of AI and agent behaviour and clear lifecycle processes for approving and retiring agents. Set up cross-functional oversight involving legal, security and compliance teams.

Implement a unified security control plane for all AI apps and agents. AI introduces new identities, tools and cross-system integrations. A centralised control plane prevents shadow AI, supports least-privilege access and gives security and compliance teams the visibility they need to monitor and govern agents at scale.



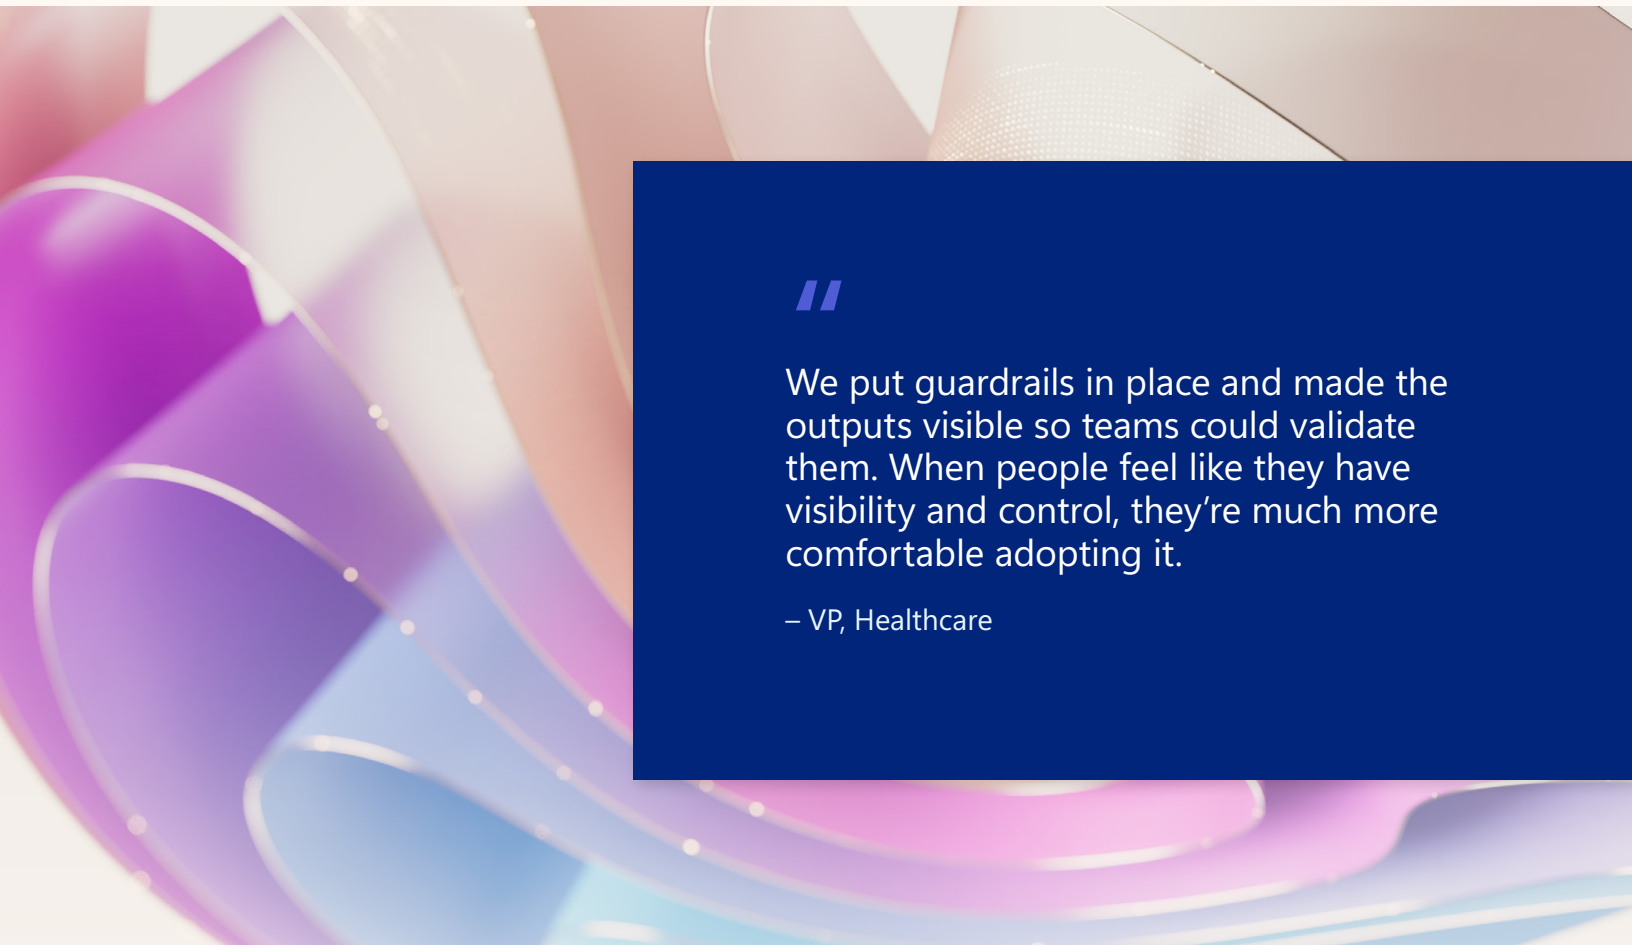
Ensure that you have continuous monitoring and real-time observability for agents and their actions. The sophisticated and complex capabilities of agentic AI can lead to unexpected actions, unsafe tool calls or behaviours that bypass standard security controls. Automated detection systems are essential for detecting prompt injections, chain-of-thought manipulation or tool exploitation. Implement continuous monitoring to catch suspicious behaviour early.

Conduct pre-deployment safety evaluations for every agent. Pre-deployment analysis and testing by security teams can prevent high-risk behaviour from reaching production. A consistent and thorough review process reduces security incidents and builds trust with internal audit teams, business stakeholders and regulators.

Address privacy and regulatory compliance

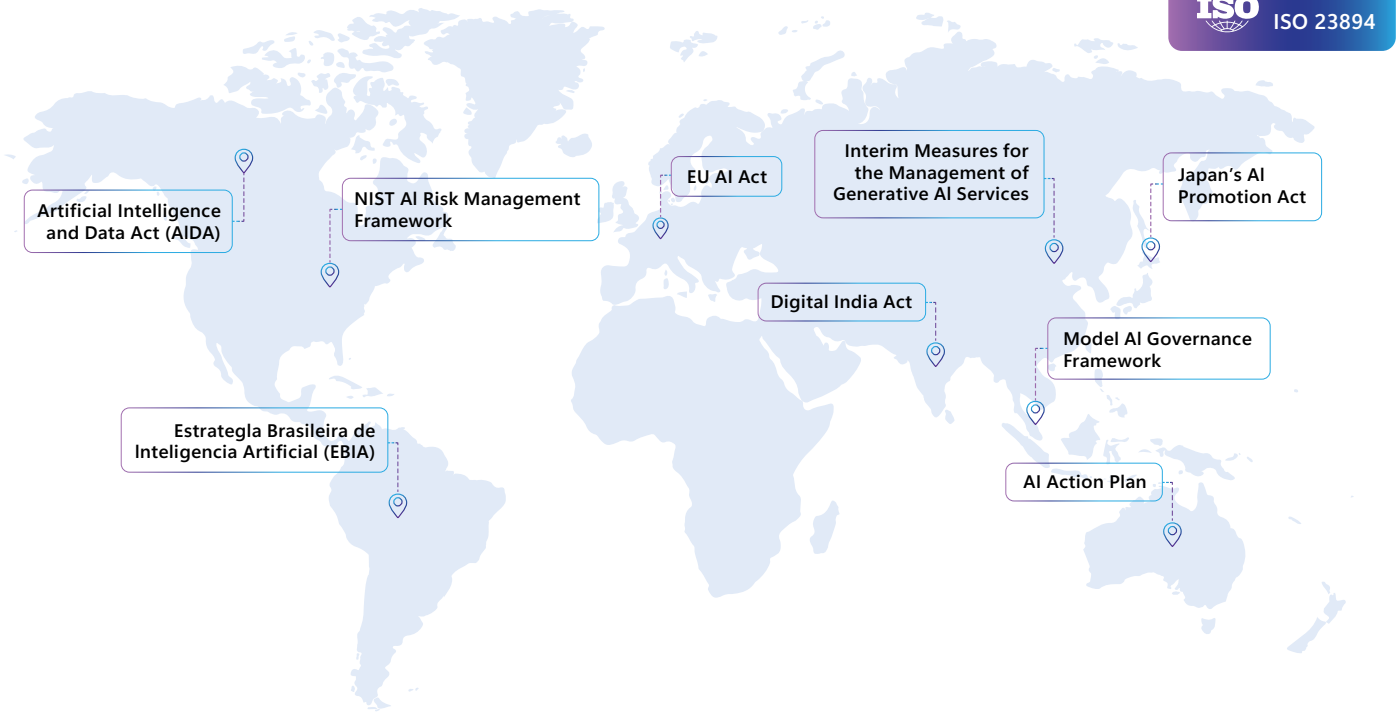
AI systems rely on large amounts of data, which often includes sensitive, personal or proprietary information. Agents also commonly interact with multiple systems and data sources. These complex behaviours introduce new challenges that your data privacy and security policies must manage.

At the same time, the legal landscape is rapidly evolving. Governments and regulators are expanding data protection with AI-specific regulations that emphasise transparency, accountability and responsible use. You're increasingly expected to not only protect sensitive data, but also demonstrate how your AI systems comply with regulations.



> "We put guardrails in place and made the outputs visible so teams could validate them. When people feel like they have visibility and control, they're much more comfortable adopting it.
>
> – VP, Healthcare

AI regulations are emerging across all regions



Multiple AI regulatory frameworks are emerging across the world.

To address these concerns, embed data privacy, security and regulatory compliance into your overall AI governance strategy.

- **Enforce rigorous data privacy controls for all AI and agentic workloads.** Implement controls that prevent unintended access and protect sensitive or regulated data in line with emerging AI regulations. Remain in compliance with existing privacy frameworks, like General Data Protection Regulation (GDPR) and Health Insurance Portability and Accountability Act (HIPAA).
- **Conduct regular security assessments of AI systems to confirm compliance with regulations.** Formal assessments, performed by legal, privacy and security teams, help ensure that your AI systems meet required standards before they reach production. Ongoing assessments help to keep them in compliance, reducing exposure to legal, reputational and operational risks.
- **Integrate legal and compliance teams throughout the AI lifecycle.** Include experts in regulation on your cross-functional governance teams to help you keep AI systems compliant over time and meet your legal obligations.
- **Deploy automated model monitoring tools.** If not detected, model behavioural drift can lead to privacy breaches, inaccurate outputs, biased outcomes or misuse of sensitive data. Implement continuous monitoring to help you catch issues early and respond quickly. This type of monitoring is a regulatory requirement in a growing number of jurisdictions.

Anticipate roadblocks

As you scale AI into more areas of your business, monitor system behaviour and create guardrails to manage risk – without slowing useful innovation to a halt. The challenge isn't just adding more oversight; it's designing AI governance that's transparent, measurable and accountable while still providing the flexibility to experiment, learn and improve.

Make your AI systems observable

Observability is the foundation of AI governance. You can't manage, review or protect systems you're unaware of, nor can you control systems you can't monitor or understand.

As AI use spreads, visibility needs to be a priority for every role:

- IT teams need to understand what AI systems exist and their operational state.
- Developers need visibility into how their AI tools and agents behave to operate and manage them responsibly.
- Security teams need to discover, test and protect these systems, no matter where they're running or who created them.

A shared set of security controls, like a centralised AI inventory, consistent access policies, global data protection and unified threat detection, can help bridge these requirements. As AI becomes part of your core business, you need these controls to provide the observability baseline that your governance strategy is built on.



Balance governance and innovation

Although security and governance are non-negotiable, overly restrictive policies can slow down experimentation and AI adoption. Striking a balance between innovation and control is key to maintaining momentum while improving trust and safety in your AI systems.

As AI use scales and your governance model matures, stakeholders have different priorities. Teams building AI solutions want freedom and flexibility to move quickly. IT wants scalable, cost-effective systems that follow policy. Security teams want agents to be visible, data protected and access controlled, with systems defended against continuously evolving threats.

Governance shouldn't slow down one of these groups to satisfy another. Instead, create shared guardrails so innovation can move forward without creating blind spots or unmanaged risk.

Your goal should be to turn your Frontier Centre of Excellence into a Centre of Acceleration, streamlining and encouraging innovation while analysing risk realistically. Balance your approach to prioritise the initial deployment of lower-risk use cases as you invest in the guardrails needed to safely support critical, higher-risk use cases.

Reaching the goal: Trusted and auditable AI

Strong AI governance and security practices create the conditions to scale AI with confidence. As you scale, RAI is embedded into how you design, deploy and manage AI systems. Governance is built in across the AI lifecycle, so AI tools and agents are auditable, transparent and aligned with ethical, regulatory and business expectations.

At higher levels of maturity, governance provides measurable business impact. It increases trust among leaders, employees, customers and regulators. It streamlines deployment, helping your organisation move faster without compromising control. The result is trusted and auditable AI that supports stronger business outcomes.

> "To scale AI with confidence, the main caps are usually less about the technology itself and more about operating model, governance clarity and enablement across the business.
>
> – IT Director, Retail

Microsoft as customer zero

Our own Frontier Transformation journey provides many practical, real-world examples of how the principles in this roadmap come together at scale. As both a technology provider and an enterprise operating across diverse business functions, Microsoft applied these strategies internally to advance our transformation.

Establish a Frontier CoE early

Microsoft Digital, our internal IT organisation, launched its first AI CoE in 2023 as an advisory group supporting internal AI efforts across the company. Teams across Microsoft Digital began building tools, workflows and AI-enabled solutions with enthusiasm.

Momentum accelerated quickly, but we soon saw different solutions addressing the same problems, uneven governance and inconsistent use of standards and reporting. To drive consistent and measurable AI outcomes at scale, we evolved the role of the Frontier CoE from guidance to coordination, helping to set priorities, define guardrails and connect AI work directly to business outcomes.

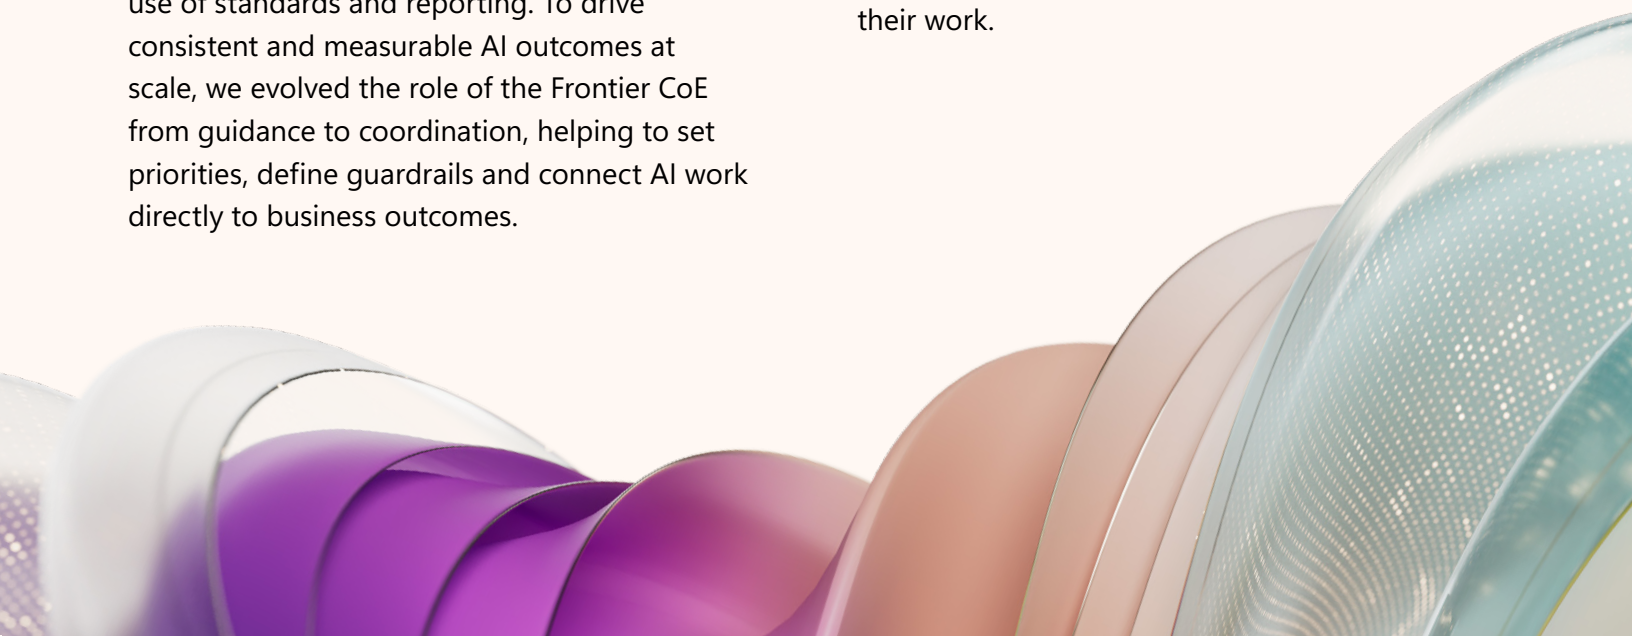
Our Frontier CoE operating model is intentionally simple. Employees submit ideas to the CoE using a common form. The Frontier CoE teams review initiatives against four shared pillars, which help teams to think beyond individual projects and to ensure that work aligns to business priorities, can scale safely, has a clear delivery path and supports responsible adoption:

The strategy pillar focuses on analysis of business value and implementation effort for AI projects.

The architecture pillar defines how we help teams scale AI solutions without creating security gaps, compliance issues or technical debt they'll have to unwind later.

The roadmap pillar focuses on how Microsoft employees interact with AI to ensure that user experiences are cohesive and align with our broader employee experience goals.

The culture pillar ensures that AI adoption across Microsoft is intentional, responsible and sustainable. It defines how Microsoft teams are expected to adopt AI and integrate it into their work.



Change the role of IT

Our own wide-scale adoption of agents forced a shift in IT mindset. We found that using IT staff as gatekeepers isn't productive. Instead, they need to act as trusted partners with developers and business owners.

To accelerate AI adoption safely and at scale, IT needs to be more proactive, strategic and platform-oriented, taking responsibility for the foundational capabilities agents need to operate reliably across the enterprise.

Our IT teams focused on several key areas:

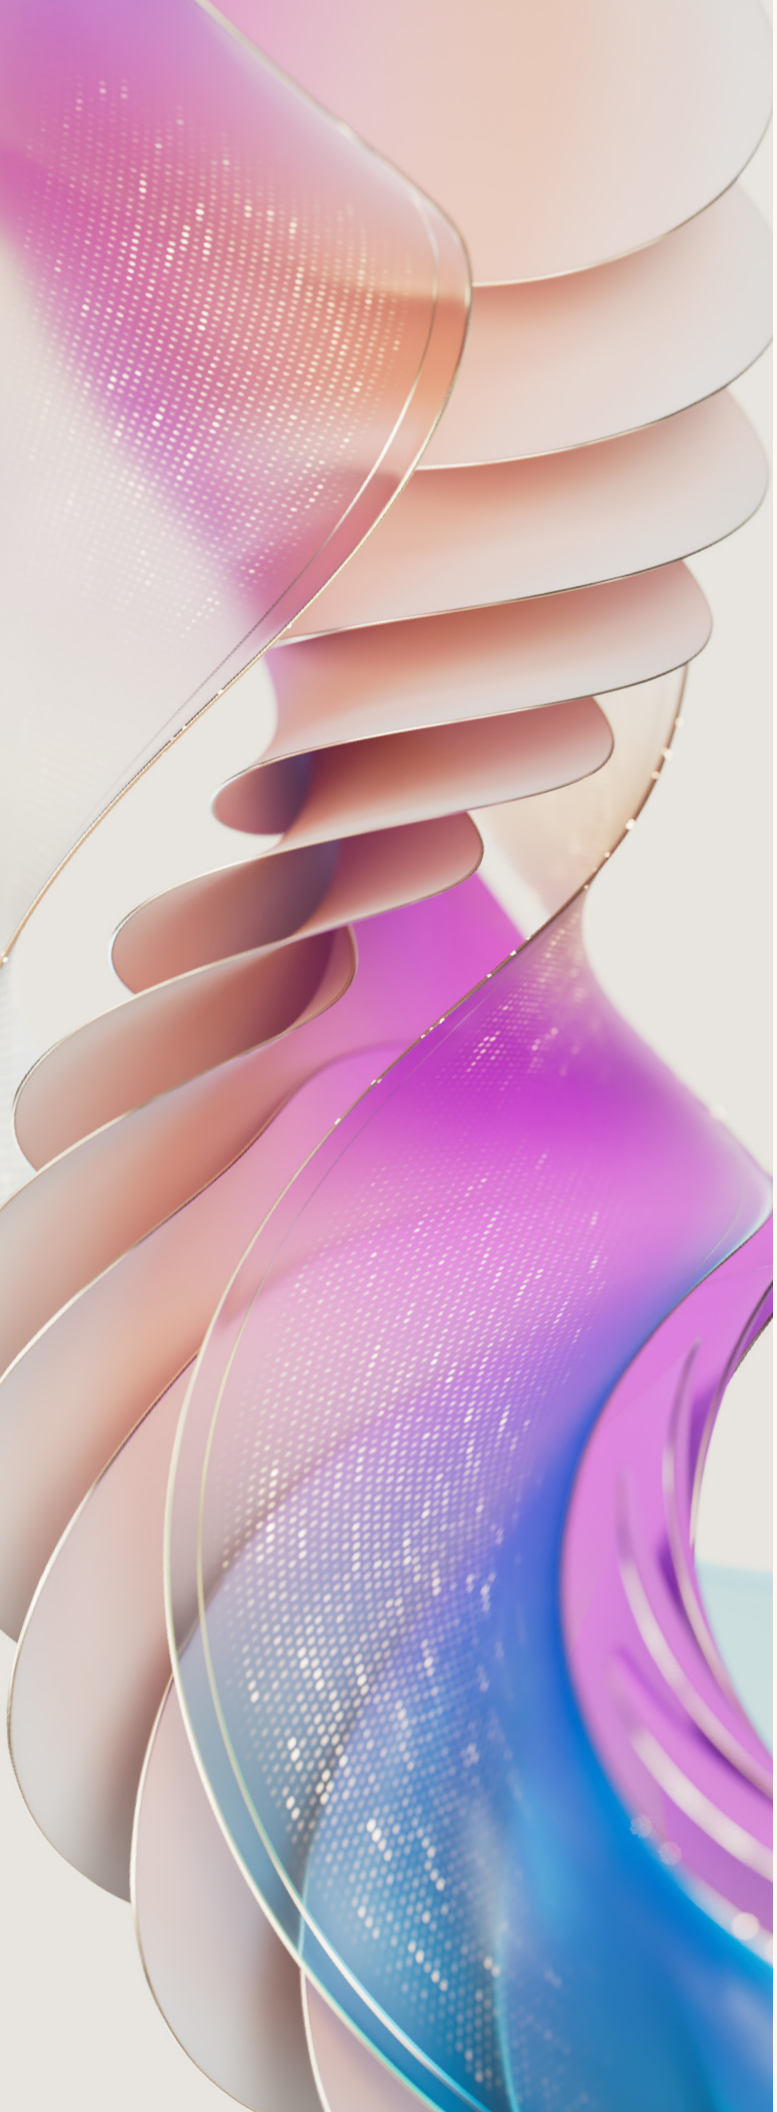
Developing and maintaining agent memory systems. Agents need persistent, contextual memory to function effectively. IT must ensure that the systems that store and manage this memory across workloads are secure, governed and consistent.

Orchestrating integration between agents and core systems. The real value of agents comes from connecting them to enterprise applications and business processes. IT builds and manages the infrastructure of APIs, MCPs, connectors and event-driven workflows, so agents can act safely and reliably.

Curating and exposing certified data sources. For agents, data quality becomes action quality. IT shifts from a gatekeeper role to a curator role, exposing trusted, well-governed data sources that agents can use with confidence. This reduces risk and increases reliability.

Enabling agents to 'graduate' from personal to enterprise-wide use. Many agents begin as individual or team-level experiments. IT manages the AI lifecycle to ensure that useful agents receive validation, security review, integration hardening and operational readiness checks before scaling them across the organisation.





Build skills, trust and adoption through experience

Microsoft recognised early that trust in AI is a prerequisite for transformation. We focused on creating clear guardrails, transparent communication and consistent user experiences. Responsible AI principles and compliance mechanisms helped us create AI systems that are safe, secure and aligned with business intent.

At the same time, we wanted to avoid overly restrictive controls that hinder AI use. We developed a 'fast train' approach for accelerating AI adoption in a controlled, governed way. It focuses on the following:

- A trusted subset of teams and employees are given **early, broad access to AI tools and capabilities**.
- This access is paired with **clear guardrails, policies and oversight**.
- Teams are encouraged to **experiment, learn quickly and iterate**.
- Risk is managed through **central governance, monitoring and shared standards**.

As adoption grew, so did trust in our systems. This allowed us to shift from creating individual productivity tools to building agents that worked in partnership with humans to produce new capabilities. Employees were empowered to create, deploy and refine agents within a governed ecosystem, combining human judgement with machine intelligence to transform workflows at scale.

Integrate AI into the culture

Our approach to Frontier Transformation required us to make organisational and cultural changes focused on five key areas:

Executive sponsorship. AI transformation is led from the top. Clearly defined business and skilling goals align to measurable outcomes. Leadership decisions are informed by telemetry that shows adoption rates, proves ROI and guides future resourcing.

Embedded governance and responsible AI. Our established Frontier CoE drives robust governance, RAI practices and compliance across initiatives. Staff can build agents using Microsoft Copilot Studio and Microsoft Foundry, which integrate our governance policy and guardrails. Microsoft Agent 365 provides management and observability across the organisation.

Enterprise-scale adoption and measurement. Microsoft 365 Copilot and agentic AI are deployed company-wide, with adoption tracked through real-time telemetry and usage targets.

Skilling and a gamified learning culture. To support ongoing learning, fun events and internal competitions encourage teams to compete for awards and recognition (such as pins and flair), reinforcing new behaviours and helping to turn learning into practical skills.

Champion networks and community scaling. Cross-functional champions, communities of practice and test-kitchen sessions help pressure test and scale successful practices.



Embed AI governance across the organisation

To embed governance throughout our organisation – as an accelerator rather than a constraint – our Frontier CoE drove collaboration between engineering, security, compliance, data and business leaders. Our approach combined human judgement with AI-driven risk analysis and monitoring, strict data governance and responsible AI and compliance principles built into design and deployment processes.

We learned some key lessons along the way:

A cross-discipline governance team is essential. No single team can govern AI effectively. Success came when leaders with diverse expertise collaborated on the AI governance team.

Human and AI systems work best when they work together. Humans still make the decisions, but AI systems augment governance. AI helps us analyse risk, map dependencies, flag potential violations and scale reviews across thousands of workloads.

Only the right data should be available to agents. We learned that agent productivity is directly tied to data hygiene. If data is poorly classified or if it's duplicated or overshared, agents amplify those problems.

Embed compliance and responsible AI principles from the start. Throughout our AI transformation, we aligned internal practices to global privacy frameworks and to our RAI principles. These principles are built into all our AI platforms as guardrails.



How can Microsoft help?

Microsoft's own Frontier Transformation journey, along with years of helping customers harness the power of AI, has generated a tremendous amount of practical knowledge. This exists as a wide variety of training and reference materials, examples, tools and systems that can help you move from early experimentation to enterprise-scale AI adoption more quickly and confidently.

If you're asking yourself *What do we do next?*, Microsoft can help with proven approaches, guidance, platforms and accelerators. Use these resources to help you build a layered intelligence foundation, establish repeatable governance and operating rhythms and skill up your teams to help them deliver solutions faster.



Enterprise AI guidance and frameworks

Microsoft offers comprehensive guidance, frameworks and resources to support every stage of your Frontier Transformation journey. Drawing from our own enterprise experience and work with customers, we share structured approaches, proven best practices and practical tools to help you move from experimentation to governed, scalable AI adoption.

- **Microsoft AI adoption guidance** offers a structured process that helps you break down complex transformation efforts into a set of more manageable phases.
- **Responsible AI tools and practices** resources provide a comprehensive set of frameworks, governance models, security controls, assessment tools and operational best practices designed to help you build, deploy, secure, monitor and govern safe and responsible AI systems at enterprise scale.
- The **AI Solution Accelerators** site offers templates, examples and guidance that can help you quickly build AI solutions.
- **Microsoft Inside Track blog** explores how Microsoft applies AI to solve problems, improve products and help customers.
- **AI Skills Navigator** helps you guide your workforce through structured, role-based learning journeys at scale, so you can track progress, build capability and sustain momentum over time.

Guidance for establishing your Frontier CoE

Microsoft **CoE guidance** can help you build a Frontier CoE to act as the central organising team for your transformation efforts.

This guidance outlines roles, responsibilities and engagement patterns for your Frontier CoE, which you can use within your business and technical teams. It includes best practices for change management and communications, along with advice for embedding governance, responsible AI and lifecycle management under the purview of your Frontier CoE.

The **AI Centre of Excellence** learning path is a good resource for teaching your staff about the CoE concept and its benefits.



Conclusion

This roadmap outlines the common drivers behind AI readiness and shows you many of the strategies and practices you can adopt to maximise your chances of significant, positive business impact from AI. We hope that these insights and guidance help you approach AI with greater clarity and confidence so you can innovate continually and see compounding value over time.

What's next?

Take the next steps on your Frontier Transformation journey.

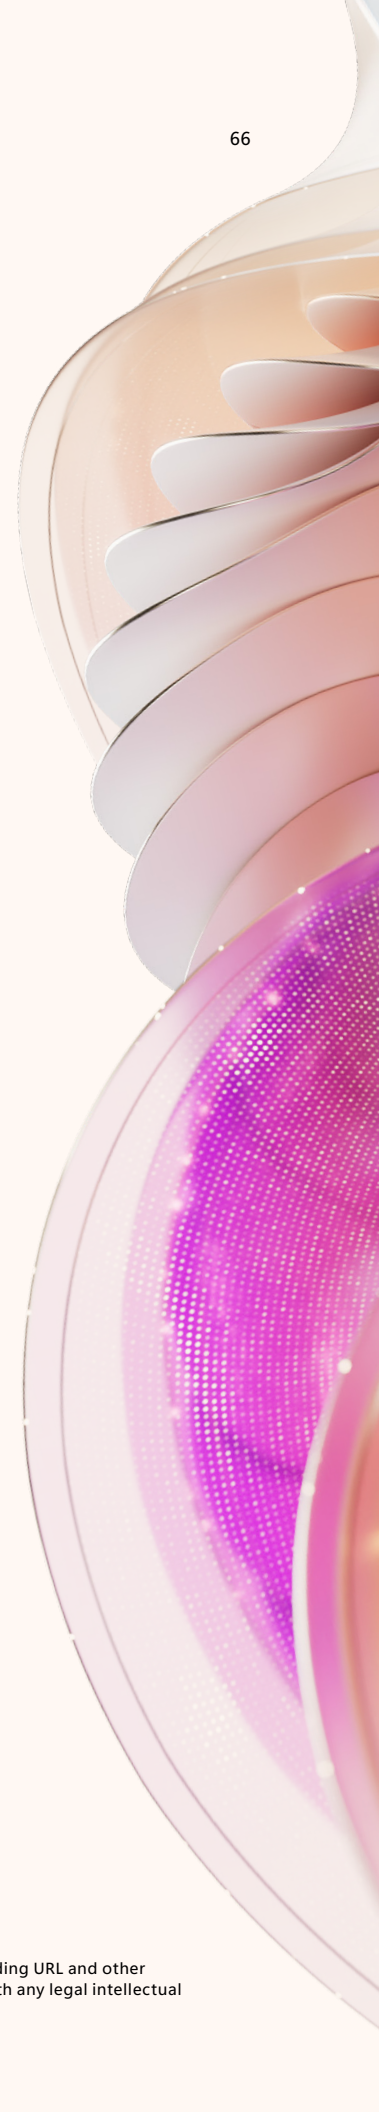
Explore **Microsoft AI**, and move from experimentation to enterprise-wide impact.

Build the skills to plan, strategise and scale AI projects on **AI Skills Navigator**.

Research and methodology

The research behind this eBook includes a study conducted by Emerald Research Group on behalf of Microsoft. The study consisted of 70 in-depth interviews with IT leaders and business decision-makers who are involved in bringing AI transformation to their organisations.

To ensure a variety of experiences in these studies, Emerald Research Group screened candidates from organisations that had already seen success in their efforts and those from companies that are deliberate or lagging in AI adoption and use. Emerald Research Group handpicked people from different disciplines and across various industries to increase the scope and value of the results. Researchers conducted a mix of 30–45-minute interviews between February 23, 2026, and March 13, 2026, to dive deep into their organisation's AI scaling journey.



© 2026 Microsoft Corporation. All rights reserved. This document is provided 'as-is'. Information and views expressed in this document, including URL and other Internet website references, may change without notice. You assume the risk of using said document. This document does not provide you with any legal intellectual property rights with respect to any Microsoft product. You may copy and use this document for your internal reference purposes.