

The AI Strategy Roadmap

The path to successful
AI implementation

Contents

AI readiness looks different for every organization. This roadmap highlights actionable strategies for successful AI transformation based on over 100 in-depth interviews with IT leaders and business decision makers who figured out how to deliver positive business outcomes using intelligent solutions.

03
Executive summary

05
Define your business strategy
to help ensure your AI projects serve strategic objectives.

12
Create a technology and data strategy
to establish the foundation needed to deploy AI solutions.

20
Build an AI strategy and gain experience
to empower your organization to support AI at scale.

26
Refine your organization’s structure and culture
to encourage the widespread adoption of AI-powered tools and workflows.

33
Govern AI across teams and departments
to easily maintain security, privacy, and regulatory compliance.

38
Conclusion

Executive summary

AI is reshaping how work gets done, enabling organizations to optimize operations, gain competitive advantages, and deliver more value faster. Yet, the ability to benefit from AI requires more than just access to the technology. It requires a foundation that supports responsible, scalable, and outcome-driven adoption.

The key to a successful AI transformation is AI readiness—or your organization's ability to effectively integrate AI within business operations and workflows. This readiness is shaped by five essential drivers: business strategy, technology and data strategy, AI strategy and experience, organization and culture, and AI governance.


Together, these drivers form a practical framework for assessing where an organization stands today and what it needs to do to move forward. When developed in parallel, they enable teams to move faster, manage and mitigate risk, and realize the full value of AI.

We talked to over

100+

real-world business and IT leaders who have successfully deployed intelligent solutions to bring you their top insights for realizing business value from AI at scale.

Figure 1: Five drivers of AI readiness



To better understand these drivers and their impact on AI readiness, we interviewed business and information technology decision makers from leading organizations who have successfully deployed prebuilt or custom-built AI solutions to generate significant business value at scale—and we’re sharing their insights with you.

The following findings are based on 104 in-depth interviews conducted by Emerald Research.¹ They identify emerging best practices for organizations to improve each aspect of their AI readiness and provide insights that help you turn AI initiatives into scalable solutions that deliver tangible business value.




Driver 1: Business strategy

Ensure your AI projects serve strategic business objectives

Successful AI initiatives aren't pursued in isolation as technical experiments. Instead, they should be tightly aligned with the overarching objectives of your organization. This helps ensure that your AI solutions are designed to solve clearly defined business problems—such as improving customer experiences or increasing operational efficiency—and can deliver measurable business value.

According to decision makers from leading organizations, this alignment requires more than just enthusiasm for emerging technologies. It demands leadership buy-in and a shared vision of success that helps everyone understand the goals of your project. As your organization prepares for AI transformation, you should begin to view intelligent technology as an essential way of achieving strategic objectives.

Key takeaways:



Get support from leaders and executives

Involve department heads and senior leadership in your AI projects to get necessary resources.




Clearly define the problem and the solution

Identify specific AI use cases that deliver clear business value by improving how your business operates.



Align everyone on one vision for success

Clearly articulate what your AI projects aim to achieve and make sure every person involved agrees on that direction.



Get involvement from business leaders and executives

One of the most critical enablers of AI readiness is leadership buy-in. This means more than just verbal support—it involves active sponsorship, visible endorsement, and a willingness to allocate resources to AI initiatives. When senior leaders champion AI, it signals to the rest of the organization that these efforts are not just experimental or optional, but central to the company's strategic direction.

Business leaders from organizations that have successfully adopted AI say that leadership support also plays a pivotal role in securing funding, aligning cross-functional teams, and setting the tone for risk tolerance and innovation. Without it, AI projects often struggle to gain traction, remain siloed within technical teams, or fail to scale. In contrast, when executives are engaged and invested, AI becomes a shared priority—one that's integrated into business planning, performance metrics, and long-term vision.

“I'll tell the leadership team 'It costs this much, it needs to do this, and I know we have the funds. What's the reason you would push back on purchasing this product?' They don't see the actual value. That's the more important thing.”

Director of Sales and Business Development,
Technology

“The more you can make something visible and show that it's valued, the better. I wanted to have a holistic program that involved leadership This gives the perception that leadership values it. There's social pressure to adopt these tools. Alignment and messaging at multiple levels accelerates adoption because it's clear that everyone values it.”

Head of Mobile and Engagement Engineering,
Enterprise SaaS


**Navigate the AI landscape
with tips from Microsoft leaders**

[Read the e-book](#)

“I think [we] need to have that buy-in from the CEO and the investors to really accelerate what we’re doing in AI.”

Chief AI Officer, Insurance





Identify the specific problem you plan to solve with AI

A critical step in building AI readiness is understanding the problem you're solving. Although every organization's strategies are unique, the business and IT decision makers we interviewed took a similar approach to AI deployments—they begin by clearly defining use cases, the business problems that need solving, what success would look like, and whether AI is the right tool for the job. Regardless if they're designed to reduce churn, improve forecasting, or automate repetitive tasks, the most successful AI initiatives are practical and grounded in measurable value generation.

Equally important is identifying and aligning on a vision for success. This means articulating both qualitative and quantitative outcomes and ensuring that all stakeholders—including leadership and delivery teams—are on the same page. A shared vision prevents fragmentation, keeps teams focused, and provides a benchmark for tracking progress and ROI. Without it, even well-intentioned AI efforts can drift off course or fail to scale.

“You have to have something real. You have to have something that you can show ... and explain in very detailed ways. It can't be just a fluffy kind of idea—it needs to be a very certain thing [It needs to have] some hint of an ROI. Like, we're solving this problem faster. We're improving our users' lives by doing x, y, and z. We are accomplishing this task much faster. Something like that needs to be stated.”

Head of CTO, FinServ




“I think there's probably a million use cases that you can think of. The important thing is you need to identify the business problem.”

SVP of Portfolio Construction
and Client Analytics Technology, FinServ

“We need to make sure that we are not [just] getting a solution. We have to have a problem and an enhancement opportunity that we are getting a solution for.”

VP of Business and Technology Transformation, FinServ






Establish a shared vision of what success looks like

A clear, shared vision for success is an essential part of any AI initiative. This includes defining what success looks like—both qualitatively and quantitatively—and ensuring that everyone involved is aligned. Whether it's improving customer satisfaction scores, reducing operational costs, or launching a new AI-powered product, having a well-articulated vision helps teams stay focused and coordinated.

Without this alignment, AI efforts can become fragmented across departments, leading to duplicated work, conflicting priorities, or unclear outcomes. A shared vision not only keeps stakeholders moving in the same direction but also provides a benchmark for measuring progress and return on investment. Business and IT leaders from leading organizations stress the importance of metrics and key performance indicators (KPIs) in gaining the clarity needed to scale AI beyond isolated pilots and into enterprise-wide transformation.

Success looks different from one organization to the next—and what it looks like may change over time depending on your unique business challenges and objectives.



“I’m sure you’re familiar with business cases, but you need to include certain metrics. What does your total cost of ownership look like? What does technical debt look like? What is the opportunity cost?”

VP of Business and Technology Transformation, FinServ

“Is it a productivity play, an efficiency play, a quality play, or a combination of these? It’s really about understanding the ‘why’ and the use case.”

VP of Operations and Business Transformation, Health Insurance





Driver 2: Technology and data strategy

Establish the foundation needed to deploy AI at scale

Turning AI initiatives into practical business solutions that drive value across your organization is only possible with the right technology and high-quality data. While aligning projects with business goals and objectives helps set the right direction, technology and data strategies determine whether organizations can move from AI experimentation to scalable business impact. By turning data into a strategic asset and treating technology like a catalyst for growth, your organization becomes better positioned to innovate responsibly, adapt quickly, and realize measurable value from AI.

Key takeaways:



Effective AI solutions need high-quality data

Ensure your data is consistent and up-to-date to help improve the accuracy and reliability of your AI solutions.



Preparing data for AI is an ongoing process

Continuously refine your data—and understand that high-quality, AI-ready data takes time and effort.



Know when to buy and when to build

Consider whether it makes more sense to purchase prebuilt AI tools or build custom AI models tailored to your specific needs.



Focus on the quality and consistency of your business data

Business and IT leaders from leading organizations say that prioritizing data and ensuring consistent data quality is foundational to AI readiness. This means treating data as a strategic asset from the outset—identifying the most valuable sources, breaking down data silos to make them more accessible, and maintaining their integrity over time. It also involves cleaning, annotating, and validating data to eliminate errors and building data dictionaries or semantic data models to ensure your data is well aligned to business objectives.

Poor-quality data leads to unreliable predictions, biased outcomes, and failed deployments. Clean, consistent data enables AI models to learn from trustworthy inputs, which improves their accuracy and reliability. Plus, it lays the groundwork for scaling AI across the organization—without this foundation, even the most promising AI initiatives risk being built on shaky ground.

“ We need to make sure that we have the technical capabilities. Do we have the data? Because in some cases, yes, we have the data, but the data is not right [In] terms of the data format [or] structure ... we need to have a full audit ... because garbage in, garbage out.”

Orchestration Leader of Digital, AI and Innovation,
Automotive



“ [My best] advice would be to prioritize your data. Make sure that you have consistent data quality [and] your data infrastructure represents the AI stack.”

CTO, FinServ



Key techniques for structuring and managing business data include:

Labeling and annotating data

Label and annotate data to mitigate bias and improve the reliability of your AI solution’s outputs.

Breaking down data silos

Connect different data sources together to streamline data usage and retrieval for your AI model—and decrease the likelihood of inaccurate outputs.

Incorporating real-time data

Incorporate real-time data integration to improve the relevance and effectiveness of your AI models.

Transforming data for AI

Convert raw data—such as text and images—into structured formats that can easily be processed by your AI solutions.

“The first step was massaging our data [and building] a solid semantic data model, a data dictionary, which I fed. Obviously, the better your underlying data is, the better the answers get [Building a data dictionary] is about taking a business view of our data and mapping it to our system implementation.”

SVP, Portfolio Construction &
Client Analytics Technology, FinServ, US



“We are cataloging and putting in definitions for EVERYTHING. That effort is not an overnight effort. It’s a data governance initiative that runs for years.”

Head of Data and Digital Transformation,
Retail eCommerce

“So many things can go wrong. So you have to come into it like, ‘Let’s start basic. Let’s not try to throw in a dozen pieces of data all at once. Let’s be iterative about this.’”

VP of Digital Analytics Services Engineering,
Banking, FinServ, and Capital Markets

Preparing business data for AI takes time

Getting data ready for AI isn’t a one-time task—it’s a continuous process. As organizations mature in their AI journey, the data that supports intelligent systems must evolve too. This means regularly refining, labeling, and enriching datasets to improve their quality and relevance. It also involves building feedback loops that use AI outputs to uncover gaps or inconsistencies in the data and feeding those insights back into the system.

This ongoing effort requires a data foundation that can adapt and scale with your organization’s needs. Business conditions shift, new data sources emerge, and expectations change—data that was once considered AI-ready may no longer be sufficient. Teams that treat data readiness as a continuous discipline—not a project with a finish line—develop more resilient systems and gain the business agility needed to support innovation over time.

“Data literacy is extremely important, and we lack that today. A data literacy mission is also something that we’ve started in parallel so that we can train more people to be more data-aware and help us improve the quality of data, whatever they’re managing with.”

CIO of Mobile, Manufacturing





Know when to buy AI solutions and when to build them

Deciding whether to buy or build your AI tools is a strategic decision that shapes the pace, cost, and impact of adoption. Buying prebuilt solutions offers speed, lower upfront investment, and minimal technical overhead. These tools are ideal for general use cases such as summarization, content generation, or productivity enhancements. Building your own AI models allows for deeper customization, total control over data, and increased alignment with your specific business needs—but building also requires higher budgets, more time, and greater technical expertise than buying does.

Gain a deeper understanding of key AI solutions and strategies

[Read the e-book](#)

According to businesses that have successfully deployed AI, the right choice depends on your specific goals, resources, and readiness. Buying can accelerate early wins and support rapid scaling across teams, while building may be necessary for differentiation or specialized use cases. Treating this decision as a deliberate tradeoff helps organizations allocate time, talent, and budget more effectively. It also ensures that AI investments are aligned with business priorities, not just technical ambition. Whether you choose to buy or build, leading organizations agree that AI technology is evolving rapidly—so your willingness to fully explore your options is important.

“You can buy now and stick with it, or you can take some time to build it once you understand the nuances.”

Sr. Director of Product, eCommerce

“So many different vendors are offering solutions in this space, but [my] company is doing a pretty good job in evaluating and trialing several of these tools, utilizing several ... vendors to figure out what works best.”

Direct of Product Strategy, Telecommunications

“It’s still so early. A lot of GenAI is still experimental, right? It’s evolving so rapidly. It’s hard to know that the tool I bought today is actually going to be there tomorrow or actually meet my needs as the market and everything keeps changing.”

CISO, Infrastructure and Defense





Decision makers from leading organizations ask these questions when choosing whether to buy or build AI models:

“Do we have the talent to build this right?”

If you’d have to hire outside talent to make it happen, consider buying as third parties can likely build and implement AI tools faster. If you already have an internal team with the necessary knowledge and experience, consider building your own tools.

“Are we under any time constraints?”

If you need a solution or results quickly, consider buying to set up intelligent solutions faster. If you have time to develop AI solutions that meet your unique needs, consider building them yourself.

“Is this solution part of an existing app?”

If your solution is a standalone application, consider buying prebuilt solutions to save time. If your solution is designed to be a part of your internal programs, consider building since your team already knows the technology.

“Can our technology stack easily incorporate AI tools?”

If your data systems are compatible with third-party provider solutions, buying can enable greater impact faster. If your systems are completely unique or proprietary, building custom AI models can be a better option than buying and implementing prebuilt solutions.

Figure 2: Determining factors in the choice to buy prebuilt AI models or build your own

	Buy (Prebuilt)	Build (Custom)
Speed	Faster to deploy	Longer development cycle
Cost	Lower upfront cost	Higher investment in time, talent, and infrastructure
Customization	Limited flexibility	Tailored to unique business needs
Expertise Required	Basic understanding of AI capabilities	Requires skilled AI and machine learning teams
Compliance Needs	Limited to the vendor-defined policies and configurations	Can be designed to meet specific needs



Driver 3: AI strategy and experience

Gain the expertise needed to support AI at scale

Without the tools, processes, and knowledge needed to turn plans and strategies into tangible solutions, companies can fail to realize any significant business value from their AI initiatives. While defining a clear vision for success is important, business and IT leaders stress how organizations must also develop technical proficiencies and repeatable processes—making it possible for teams to develop and deploy AI solutions consistently and at scale.

Key takeaways:



Start small and learn as you go

Launch your AI initiatives with focused, low-risk projects that are easy to implement and adjust.



Prioritize AI use cases by their business value

Focus your implementation efforts on projects that offer the greatest potential for significant returns on investment.



Use metrics to guide iteration and scaling

Make data-driven improvements to your AI deployments by tracking KPIs.



Use early pilots as learning opportunities

Leading organizations say starting small is one of the most effective ways to build momentum in AI adoption. By focusing on low-risk, easy-to-implement use cases, businesses can quickly gain hands-on experience without being overwhelmed by complexity. These early pilots serve as test beds for refining data pipelines, evaluating model performance, and shaping governance practices. More important, they create space for iteration, enabling teams to learn from real-world feedback, adjust their approaches, and gradually expand AI capabilities as their confidence and expertise grow.

This approach establishes a strong foundation for sustainable AI growth. Early, well-scoped successes not only demonstrate tangible value but also build trust across the organization—earning executive sponsorship and energizing teams. These initial efforts offer a practical environment to test and refine responsible AI practices, where fairness, safety, and compliance can be closely monitored. Most important, they shift AI adoption from theory to practice. By learning through real-world application, organizations gain insights that sharpen strategy, inform governance, and facilitate smarter investments.

“Irrespective of your investment, start small and then scale gradually. Do a certain level of [proof of concept], understanding what the feasibility could be, and the approach needed Once you have a rock-solid approach, scale and operate to accommodate the growth, which is important as the product evolves.”

CTO, FinServ

“To be honest, I started by doing the initial work ... myself. And how did I do it? It was a combination of reading online [and] trying different things. I think it was really a lot of [trial and error] learning.”

SVP of Portfolio Construction and Client Analytics Technology, FinServ

Focus on the most impactful AI use cases

Effective AI adoption starts with selecting the right problems to solve. Prioritizing use cases based on business impact means identifying how AI can help teams overcome real challenges, whether by improving efficiency, increasing revenue, or enhancing customer experiences. It also requires evaluating potential projects by their potential returns on investment, alignment with strategic goals, and implementation complexity. This disciplined approach helps teams focus their initial development efforts on high-value pilots that are both achievable and meaningful.

See today’s top AI use cases for creating business value

Read the e-book

“Start with small use cases and build on those over time. I would say ‘Don’t try to be perfect right away’, which allows that innovation to happen without huge expectations and then make it widely available. They don’t need to know why AI works or how AI works. They just need to know that it works and let them play with it and discover with it.”

COO, Commercial Real Estate Services



More than half of the decision makers surveyed saw a return on investment from custom-built AI applications in less than one year.¹

“Trying to build a generic solution would have been super hard. For the proof of concept ... try and first identify actual questions that business users have asked for Solve for them and then expand from there.”

SVP of Portfolio Construction
and Client Analytics Technology, FinServ





Across industries, the most common metrics that decision makers used include:

★ Quality of output

📷 Return on investment

👍 User satisfaction

👉 User engagement

↔️ Width of impact

🕒 Time to completion

👥 Number of users



“ROI? That’s really tough to measure. I have tried. Really, I think the best proxy that we could use is adoption and engagement—just usage of the tool.”

VP of Product Management,
Marketing Agency

Make data-driven improvements to your AI projects

Establishing a measurable indicator of success early in an AI initiative is essential, even if that metric is as simple as usage. Waiting for the perfect metrics can delay progress and obscure early signals of value. Instead, leading organizations suggest beginning by tracking what you can—such as how often a tool is used, who is using it, and for what types of tasks. These basic indicators provide a starting point for understanding adoption, identifying friction points, and guiding iterative improvements.

Measuring KPIs from the outset creates accountability and focus. It gives business leaders and individual teams a shared understanding of what AI success looks like and helps validate whether a solution is addressing the intended problem. For example, usage data can reveal patterns that inform future development and scaling decisions. In early-stage pilots, where learning is the goal, even simple metrics can provide the feedback loop needed to strengthen your strategies and build the case for broader investment.

“You measure return on investment based on what you are spending before and with the AI in place. The expectation is that [you] will be spending less, but it’s yet to be determined because we have just deployed these systems. They just went into production. This is all very new.”

Developer Decision Maker, Technology






Driver 4: Organization and culture

Encourage teams to drive value creation with AI

The success of your AI projects relies on more than just technology. Without the right organizational structure and culture in place, even the most promising initiatives can stall or fail entirely. Prepare your business for AI transformation by aligning people at all levels of your organization—including employees, teams, and business leaders—with the various demands of achieving AI-powered value creation at scale.

Key takeaways:



There's no one-size-fits-all approach to AI

Tailor your adoption strategy to your organization's culture and AI readiness.



Approach AI projects collaboratively

Build cross-functional teams to ensure your AI solutions are relevant, practical, and scalable.



Foster a culture of continuous learning

Invest in training opportunities that help teams maximize the business value of your AI projects.



Progressive rollouts are a key part of adoption

Gradually scale your AI projects to refine your strategies, train teams, and minimize disruptions.



Shape your adoption strategy around your workplace culture

As you strategize your company’s approach to AI transformation, consider how teams and individual workers view AI technology—and whether there’s any resistance to incorporating intelligent solutions across business operations. When people distrust AI or feel worried it could take their jobs, they create change-resistant cultures where workers are less likely to try using new technology. However, when people see AI as an opportunity to grow, they create innovation-driven cultures that speed up deployments and enable faster returns on investment.

While most companies have a mix of characteristics from both change-resistant and innovation-driven cultures, understanding the specific nuances of your organization’s culture is an essential part of crafting effective adoption strategies.

Uncover 5 key considerations to make when training teams to use AI

Read the e-book

“The bottleneck always is on the adoption side—on the end user side.”

Director of Software Engineering, Healthcare, US

Figure 3: Key characteristics of change-resistant and innovation-driven cultures

	Change-resistant culture	Innovation-driven culture
Employees’ views on AI	Don’t replace me. AI sparks concerns over job security, slowing down adoption.	Don’t hold me back. AI garners excitement and is eagerly adopted.
Employees’ views on work	Don’t make me change. Long-established workflows and habits are difficult to transform.	Help me grow. AI is seen as an opportunity to improve performance—even if additional training is necessary.
Employee tendencies	If it doesn’t work, give up. AI solutions are quickly dismissed if they don’t solve problems immediately.	If it doesn’t work, try again. Users refine prompts and experiment with AI to achieve desired results when outputs fall short.
Past tech deployments	Yet another box to check. New technology is seen as burdensome.	Ready to embrace new opportunities. Technological innovations are embraced optimistically.



Encourage collaboration throughout the development process

Leading organizations say successful AI initiatives require contributions from various departments across your organization—not just your IT team. Rather than having developers and data scientists create AI solutions in isolation, consider building cross-functional teams—comprising talent across your legal, operations, HR, and IT teams—to ensure that your AI projects are developed with business goals, ethical standards, and user needs in mind.

Creating diverse teams brings together different perspectives which helps validate the effectiveness of your AI solutions and identify potential risks and biases early. While software testing is usually performed by IT teams exclusively, the business and IT decision makers we interviewed emphasized how including users in the AI testing process is essential because the same prompts or inputs can often generate different results.

“For every business unit, we’ve selected key people with a good knowledge of the business, good connections in their area of the business, good communicators ... to feed opportunities up to my AI team, but also for my AI team to feed down to them what some of the capabilities are and get people thinking about where they can use them.”

Chief AI Officer, Insurance



“Have a business sponsor ... that you’re working with. Don’t be doing development in a vacuum. This is not Field of Dreams [where] if you build it, [users are] going to come.”

Business Decision Maker, Retail

“Have humans go and check it out. There are companies coming to the market right now to do automated testing, but so far I haven’t seen one single one that does good testing in the output that’s not predictable. Until then humans must come in and do a little bit more rigorous testing.”

Managing Director,
Chief Data Architect, Professional Services





“Many employees are excited about AI [and machine learning]. When we offer training courses, we have oversubscription because people understand it’s good for their [curriculum vitae], makes them more competitive, and more resilient in terms of headcount.”

Senior Director of Strategy and Transformation,
Manufacturing

“Invest in the people. Invest in the talent. Invest in the training and hire ... skilled professionals in an incremental way. Not just a big bang hiring without understanding the use cases—without understanding what level of investment you can do.”

CTO, FinServ

Continuously provide workers with training opportunities

Preparing people for AI transformation is just as important as readying business systems. To deliver business value from AI at scale, teams must first acquire the technical skills and confidence needed to properly use intelligent solutions. By offering learning paths, certifications, or hands-on practice, workers gain a deeper understanding of how AI works and a greater appreciation for the ways it benefits them—which ultimately makes your organization more likely to embrace AI tools and drive adoption.

Discover 10 best practices that empower teams with AI skills

[Read the e-book](#)



Business and IT leaders suggest:



Starting with a group of power users

Find a group of employees willing to try AI, experiment with it, and share their feedback. A random group of users may not always engage with new technology.



Collecting user feedback regularly

Set up review sessions or request written feedback to hold initial users accountable. An open-ended “just go try it” approach won’t get the feedback you need.



Collaborating and sharing knowledge

Empower people to collaborate using AI tools and share their experiences—and make it easier for users to build on each other’s learnings and strengthen AI use cases.



Defining a timeframe for deployment

Creating urgency pushes users to provide more direct input or ideas. By saying you have a week, month, or quarter to explore AI, you can get feedback sooner and iterate faster.

A progressive rollout is the key to a successful deployment

Breaking down AI adoption into manageable phases gives teams time to understand, test, and adjust to new tools and workflows. This strategy helps reduce disruption and build confidence while creating a space for feedback and iteration. Gradually rolling out your AI projects also highlights potential implementation challenges or cultural friction before they can become a barrier to broader adoption.

“We started off small, probably a couple hundred [users]. That was on purpose. Once they started [trying AI], the ideas started generating. There were things that people came up with on their own that we didn’t have in the plan.”

Chief Procurement Officer,
Telecom

“Pre-deployment testing and use case-specific deployments help prove out the value on a very small scale as a pilot to show leadership the value.”

Sr. Director of Technology –
Care Services and Delivery, Healthcare



Driver 5: AI governance

Deploy AI while maintaining security, privacy, and regulatory compliance.

AI governance is the framework that ensures AI solutions are developed and deployed responsibly and ethically—and used in a way that aligns with the values of your organization. As intelligent solutions become more embedded in business operations, governance provides the necessary guardrails to manage risks, ensure compliance, and build trust. However, effective AI governance is about more than just control—it's about coordination. It brings together legal, compliance, IT, and business leaders to define how AI should be used and monitored across your organization.

Key takeaways:



Gaining trust in AI technology takes time

Understand that early AI models may produce inconsistent results until they're exposed to more data, feedback, and real-world usage.



Collaborate with your security and compliance teams

Make sure the people responsible for protecting your business data and systems are involved in the development process.



Create guardrails and adapt them over time

Establish rules and systems that help ensure teams use AI safely, ethically, and effectively.



Gaining trust that your AI solutions will work effectively is an ongoing process. Most companies continually validate the accuracy of their AI outputs.

Building trust in your AI solutions takes time

AI systems are built to learn and improve through use. While early outputs may be inconsistent or incomplete, these imperfections are part of the process—not a reason to pause progress. Accuracy improves over time as models are exposed to more data, feedback, and real-world scenarios. Organizations that understand this dynamic are better equipped to move forward with confidence, even when results aren't perfect on day one.

Trusting technology doesn't mean ignoring its limitations—it means putting the right structures in place to monitor, evaluate, and guide its evolution. Leading organizations say early adoption should be paired with clear expectations, feedback loops, and accountability. Teams that accept and manage the learning curve are more likely to build internal expertise, avoid overreaction to early flaws, and realize value sooner. Progress comes not from waiting for perfection, but from engaging with the system and shaping it over time.

Explore ways to build trust in your AI solutions

[Read the e-book](#)



Partner with security and compliance teams during development

AI governance is most effective when it builds on the structures already in place. Security, compliance, and risk management teams bring deep knowledge of organizational policies, regulatory obligations, and operational controls. Their early involvement ensures that AI systems are integrated into existing frameworks for data protection, privacy, access management, and audit readiness. This collaboration helps address the unique risks AI introduces—such as model drift, data leakage, and bias—using tools and processes that are already trusted and tested.

Working with these teams also accelerates responsible AI adoption. It avoids duplicating effort by leveraging existing governance tools and ensures consistency across systems. As AI scales across departments, centralized coordination helps policies evolve in step with enterprise needs. By aligning AI efforts with established practices, teams can move faster, reduce risk, and maintain trust across the organization.



“We have an AI team, a data team, and also the ethics governance team to manage the AI developments.”

Orchestration Leader, Digital, AI and Innovation,
Automotive

“The compliance team, with the data governance policies, want to make sure that we’re handling customer data appropriately. Just because we have customer data and it’s being used internally, doesn’t mean that everyone is entitled to see it or use it... That’s why all of these teams from compliance to legal to privacy to cybersecurity all have a role.”

Data Science Director,
Retail Technology





Establish a clear set of guidelines for using AI

AI governance is most effective when it reflects the specific needs, risks, and culture of your organization. Rather than relying on rigid, one-size-fits-all rules, leading organizations say teams should start by establishing practical guardrails—such as policies for responsible use, access controls, data protection standards, and model transparency. These initial safeguards create a baseline for safe experimentation and help ensure that AI systems are deployed with accountability from the outset.

As AI adoption grows, governance must evolve alongside it. Embedding guardrails into day-to-day workflows—through templates, automation, and feedback loops—makes compliance easier to maintain and scale. Over time, usage insights and lessons from pilots should inform updates to policies, tools, and training. This continuous improvement keeps governance relevant and responsive to new risks, technologies, and use cases. When done well, these guidelines protect against harm without slowing innovation, enabling teams to move quickly while staying aligned with ethical and regulatory expectations.

Common risks associated with AI usage are centered around:





“We know our process better than anybody else, and we needed to come up with our own restrictions ... to govern us by. [Legal] is very open and understanding that they’re going to be left behind if they don’t start to learn and work through solutioning with us.”

Senior Director Associate Productivity & Collaboration,
Auto Manufacturing

“We [created] an itemized list of do’s and don’ts. Do understand the tool. Do accuracy checks. Do collaboration enhancement training. Don’t solely rely on this. Don’t use this with bing.com or in Microsoft Edge. Don’t use for non–business-related activities.”

SVP of IT Delivery, Banking Fin Serv
and Capital Markets

Conclusion

This research is intended to share what we've learned from business and IT decision makers about the best practices for creating value at scale with AI, and how they differ depending on your organization's unique needs and level of AI readiness. We hope these insights help you chart your path with greater clarity and confidence.

Take the next steps on your AI transformation journey

01

[Explore](#) Microsoft AI solutions and see how Microsoft is empowering the world to achieve more with AI.

02

[Learn](#) how to plan, strategize, and scale AI projects on Microsoft Learn.

Research and methodology

The research behind this e-book includes two studies conducted by Emerald Research on behalf of Microsoft. The first research study consisted of 29 in-depth interviews with IT leaders and business decision makers who have built and deployed custom AI solutions. The second study comprised 75 in-depth interviews with IT leaders and business decision makers that have bought and implemented prebuilt AI solutions.²

To source participants for these studies, Emerald Research screened candidates for people who were already using AI tools, were involved in the adoption process from beginning to end, and viewed the experience favorably after achieving success. Next, they requested video responses to ensure each participant was articulate and a subject matter expert. Finally, Emerald Research handpicked people from different disciplines and from across different industries to increase the scope and value of the results—and ran 60-minute-long one-on-one interviews to dive deep into each participant’s AI journey.

Endnotes:
1. Please see the “Research and methodology” section for more detail on the research and analytical approaches that support this study.
2. Emerald Research Group: sponsored by Microsoft, GenAI Development Journey, November 2024. This research study was conducted by Emerald Research Group (a subsidiary of the Harris Poll) from October 24–November 22, 2024 and March 20–April 22, among sixty-seven (67) IT decision makers, fourteen (14) developer decision makers and forty-three (43) business decisions makers in enterprise organizations. Qualified participants have deployed one or more GenAI custom built solutions into a production environment or purchased a GenAI solution from a 3rd party provider and the participants must be actively involved throughout their entire GenAI development/purchase process. Those purchasing 3rd party solutions were required to work for companies that were already “scaling” or “realizing value” when it comes to GenAI, represented as the two most mature levels of a 5-point GenAI maturity scale.

©2025 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 bear the risk of using it. This document does not provide you with any legal rights to any intellectual property in any Microsoft product. You may copy and use this document for your internal, reference purposes.