The AI Strategy Playbook 2026

Real strategies enterprise AI leaders are using to win

The deeper you get into AI transformation, the more in-depth — and evolving — you'll realize your AI strategy needs to be. From paving the way for agents to zeroing in on the right things to automate to building a tool stack that unlocks without adding costly bloat, your AI strategy needs to cover it all.

So we gathered 6 enterprise leaders to share how they are thinking about AI strategy at scale. In this working playbook, you will learn from them directly - and gain strategies to take back to your own company.

The top 6 strategies for a successful AI deployment

1. **Turn inference into business value.**
   The winners will not simply buy more AI. They will manage AI like a scarce resource and convert it into faster decisions, better work, and lower friction for employees and customers.
2. **Unfreeze the middle managers.**
   The biggest AI adoption blocker is often the managers whose expertise, workflows, and leadership identity were built in a pre-AI world.
3. **Clean the workflow before you automate it.**
   AI can accelerate bad processes just as easily as good ones. Map the work, identify owners, and remove broken steps before layering agents on top.
4. **Build a citizen developer layer.**
   Empower business users with approved tools, data access, and guardrails so they can build automations and agents closer to the work.
5. **Treat token costs like cloud sprawl.**
   AI spending is following the same path as cloud adoption: early excitement, rapid growth, and eventual cost scrutiny. Put governance around token usage before the bill arrives.
6. **Become a better wisher, not a better prompter.**
   As AI gets better at execution, the competitive advantage shifts from writing clever prompts to defining the right outcomes, context, and business problems worth solving.

How to Build an AI Supercompany

with Greg Shove & Scott Galloway,
Section & Pivot

AI transformation is not a tool rollout, it's a shift in how companies create business value. The next generation of winners will be Supercompanies: organizations that reduce friction for employees and customers while learning to manage inference as a core business input. Leaders need to experiment aggressively while capital markets are still subsidizing the cost of AI.

"Never a cheaper time than now to be experimenting with AI."

The mandate for leaders is to make AI a management discipline. That means setting expectations for usage, funding experimentation, managing token spend, and getting every employee closer to an AI-native way of working.

1. **Manage inference like a business resource.**
   The companies that win will not treat AI usage as free. They will give teams access, measure value, and learn where expensive inference actually creates return.
2. **Reduce friction on both sides of the business.**
   Supercompanies create less employee friction and less customer friction. AI should make work faster, products easier to consume, and service delivery more responsive.
3. **Give everyone an AI chief of staff.**
   Every knowledge worker should have agentic support - not only executives. The future work day starts with an AI system that helps prioritize, draft, search, summarize, and act.

The practical takeaway: Write the AI manifesto, run the hackathon, find the catalysts, put agents into production, and build enough trust that teams are willing to change how they work.

Choosing the Right Things to Automate

with Lashorelle Morgan,
NBCUniversal

Governance is the thing that allows innovation to scale safely. At NBCUniversal, the goal is not to build a permission bureaucracy. It's to help employees know which tools are safe, which use cases need review, and how to move from usage to building without creating unmanaged risk.

"AI is really, really good at blowing up a bad process."

Before adding AI, write down the workflow — use pen and paper if you have to. If the workflow is unclear, ownerless, or broken, AI will not fix it — it will amplify the problem faster.

1. **Start with the workflow, not the tool.**
   Ask workflow owners to write down exactly how the workflow operates today. You need to know this so you don't amplify the problems in a confused or broken workflow with AI.
2. **Use intake to say yes faster.**
   Track every tool and use case, then decide which categories are high risk. If 10 out of 100 requests need deep review, the governance committee can quickly approve the other 90 instead of slowing everything down.
3. **Keep humans accountable where the blast radius is high.**
   Calendar automation does not need the same controls as consumer-facing messages or legal filings. The higher the risk, the more explicit the ownership and review process must be.

The practical takeaway: Governance should be designed like guardrails on a mountain road. It should help the business climb faster without driving off the cliff.

Building the Enterprise Infrastructure for AI Agents

with Scott Likens,
PwC

Much of the AI agent conversation focuses on models, agents, and infrastructure, but the biggest barrier to adoption is often human, not technical. Executives typically understand the urgency of AI, and frontline employees are often eager to experiment – the bigger challenge sits in the middle manager layer, whose expertise, processes, and leadership identity were built in a pre-AI world. These leaders are responsible for driving adoption, but they are also the group most likely to feel threatened by changes to how work gets done.

"The frozen middle is the real adoption blocker."

The infrastructure question is not just which model to pick, it's how to create a trusted AI ecosystem that supports model access, orchestration, memory, data connections, security, and governance.

Key insights:

1. **Design for flexibility over longevity.**
   Models, standards, and orchestration patterns change constantly. Build an architecture that can hot swap models, onboard new capabilities quickly, and avoid disrupting the whole stack.
2. **Good data matters more than perfectly centralized data.**
   AI does not always need data organized the way humans need it organized. The bigger questions are where the data lives, who can access it, and whether it is legal and appropriate to use.
3. **Capture context, not just data.**
   The most valuable enterprise context often lives in people's heads: the process around the data, the judgment applied to it, and the tacit knowledge behind the work.
4. **Build bilingual teams.**
   Central teams can build reusable AI capabilities like voice interaction or document extraction. Business teams need to apply those capabilities in context so the same AI asset can serve different use cases.
5. **Optimize model cost at the architecture layer.**
   Not every task needs the most expensive model. Route simple work to efficient models, test quality, and preserve frontier models for the work that justifies the spend.

The practical takeaway: Build for people before you build for agents. Unfreeze the middle managers, create flexible AI infrastructure, and capture organizational knowledge so employees can adopt AI faster than the technology itself changes.

The Operating Model for AI at Enterprise Scale

with Amol Phadke,
Tech Mahindra

AI strategy cannot be generic. It must be mapped to the service lines, industries, and work patterns where value is created.

Treat AI learning like a graduate program: everyone gets the foundation, then specialists move into deeper role-specific applications. Foundational AI literacy gets employees onto the same page, advanced AI capability changes how their actual work gets done.

Key insights:

1. **Create a common AI foundation.**

   Every employee needs shared fluency in AI foundations, capabilities, architecture, and responsible use before the organization can scale more advanced use cases.

2. **Make the advanced layer role-specific.**

   For software teams, AI changes the software development lifecycle. For contact center teams, it means humans working alongside agents. The roadmap must reflect the work, not just the technology.

3. **Federate execution into the service lines.**

   Once the shared architecture and guardrails are in place, implementation should move into the business units that understand the work and can apply AI in context.

3. **Move beyond pilots into scale.**

   The market has largely moved past AI skepticism. The harder question now is how to modernize systems, workflows, and operating models so AI benefits can compound into real ROI.

3. **Centralize AI economics, decentralize AI execution.**

   Business units should be empowered to build and deploy AI solutions, but token budgets, model standards, and vendor management should be managed centrally to avoid duplicated spend.

> **"Token costs are the new cloud sprawl. We all thought cloud would save money, then we got the bill."**

The practical takeaway: Standardize the foundation, decentralize the execution, and govern the economics. The organizations that scale AI successfully will share a common architecture and AI language while giving teams the freedom to solve their own business problems – without letting token costs spiral out of control.

Scaling a Safe Agentic AI Strategy

with Patrick Murta,
Centene

Before scaling agents, leaders need people, process, data, change management, and an "AI nervous system." With the right centralization in place, individuals across the business can build and automate within approved guardrails, accelerating innovation without sacrificing governance.

> **"Build your citizen developer layer between personal productivity and engineering-led solutions."**

Rather than forcing every automation request through IT, organizations should enable employees closest to the work to build low-code automations, workflows, and agents. This approach accelerates innovation while maintaining the security, governance, and compliance standards required in highly regulated environments.

Key insights:

1. **Build data products for AI consumption.**

   AI-ready data is actionable, governed, accessible, and supported by stewards. A standard semantic layer helps teams access data from different domains without losing control.

2. **Scale through a nervous system, not isolated projects.**

   People, process, data, and change management need to connect. Otherwise, agentic work remains trapped in pilots and disconnected experiments.

3. **Use centralized telemetry and auditability.**

   Highly regulated organizations need to see what prompts are being used, where policies execute, and how controls prevent risky actions.

4. **Create an enterprise gateway for AI.**

   Route prompts through an AI gateway for greater transparency, security, and evidence when stakeholders ask what happened and why.

5. **Let departments build inside the guardrails.**

   The long-term goal is for business units to run their own roadmaps with low-code capabilities, while core business processes are prioritized by engineering.

The practical takeaway: The fastest path to scale isn't a larger central AI team. It's creating the shared data, governance, and platform capabilities that allow every business unit to become an AI builder.

Setting the Right AI Tooling Strategy

with Ryan Asdourian,
Lumen

AI tooling is both an infrastructure and a culture question. Enterprises need to move massive amounts of data quickly, securely, and across a multi-cloud world - but that tooling only works when people adopt it and use it in the right ways.

> **"Become a better wisher, not a better prompter."**

Balance broad enablement with orchestration – encourage teams to use AI, share what they are building, and compete around great use cases. Just be careful not to let tool sprawl create shadow IT, uncontrolled data access, or unclear ROI.

Key insights:

1. **Don't let hundreds of tools bloom unmanaged.**

   Experimentation is good, unorchestrated sprawl is not. The governance committee needs visibility into usage, data access, licenses, and measurable business value.

2. **Make AI adoption visible.**

   Organizations should encourage employees to test new workflows, share results, and learn from one another. Competitions and recognition programs can reinforce the behaviors that lead to long-term transformation.

3. **Build better AI operators, not just AI users.**

   The organizations that gain the most value from AI teach employees how to provide context, define outcomes, and orchestrate workflows – not just write prompts.

4. **Make ROI unmistakable.**

   Specialized platforms need clear value so the business knows why a license, model, or workflow deserves continued investment.

5. **Re-architect the company for AI.**

   The bottleneck will keep moving – from employees to IT to systems to compute to storage to power. Winners will keep redesigning around the next constraint.

The practical takeaway: The right tooling strategy is not the biggest AI stack. It is the clearest system for giving employees access, teaching them how to use it, and keeping the enterprise in control.

Will your company scale AI or stall again?

Book a Meeting