

The AI ROI Playbook

SPRING 2026



How enterprise leaders are getting to value with AI initiatives

AI ROI is no longer hypothetical - it's just not always obvious how to unlock it. Getting value from AI requires a deep understanding of business levers, customer pain points, and areas of friction and opportunity within margin-driving processes.

We gathered 11 enterprise leaders to demonstrate how they found high-value ways to implement AI in their organizations - and how they're proving that value to key stakeholders. In this playbook, you'll get the high-calorie download from each one, and gain strategies you can fast follow on in your own company.

The top 11 strategies for unlocking AI ROI

1/ War-game how AI-native startups will attack your business

(SCOTT GALLOWAY & GREG SHOVE)

Don't wait for consultants to tell you where you're vulnerable - put your best internal people in a room and have them build the case for how an AI-native competitor would undercut you. Incumbents that don't do this risk losing 15-20% of revenue before they even see it coming.

2/ Structure ROI around workflows, not tasks

(IBM)

Individual productivity gains matter, but real ROI shows up at the workflow layer. IBM mapped 12 top-level workflows into hundreds of components. That's the level of specificity CFOs actually respond to.

3/ Track revenue per employee, not just headcount reduction

(EVERCORE)

The most telling AI metric isn't how many roles you eliminate - it's whether your teams are generating more revenue with the same or more people. Meta is the clearest example: AI helped drive 30% revenue growth, not just cost cuts, which is a very different and more valuable story.

4/ Reframe what ROI means in your context

(NORTHWELL HEALTH)

Northwell measures clinician time reclaimed and burnout reduction rather than headcount cuts, because in a tight labor market, retaining doctors is worth far more than eliminating roles.

5/ Treat governance as an accelerant, not a speed bump

(CISCO)

Security and compliance slow AI down only when they show up too late. Cisco's "discover, detect, protect" framework - built upfront - removed the fear that was actually blocking deployment.

6/ Embed AI into workflows, not sidebars

(UKG)

AI tools that live in a separate tab get abandoned. UKG found that adoption only stuck when AI was integrated directly into the processes people were already doing.

7/ Lead by example, not mandate

(NAPSTER)

Napster's CEO spent a week coding alongside his engineering team to show what AI could actually do. That hands-on demonstration converted skeptics faster than any top-down directive would have.

8/ Define your 'why' before building an agent

(SALESFORCE)

Before approving any agent, Salesforce requires the business owner to document who it serves, why it's needed now, and what success looks like. That baseline is what makes ROI measurement possible later.

9/ Run company-wide hackathons to build AI muscle

(ZAPIER)

Zapier paused operations for a week so every employee could experiment with AI tools hands-on. Weekly AI users jumped from 10% to over 50% - no passive training program comes close to that result.

10/ Focus agents on vertical, not horizontal use cases

(SIEMENS)

Horizontal agents tend to produce incremental gains that don't move financial needles. A Siemens customer built a narrowly focused virtual engineer agent that cut production downtime fix time by 85% and saved 6,000 hours annually.

11/ Move from one-size AI to a portfolio of value models

(OPENAI)

Different parts of your business need different AI strategies - broad adoption, process reengineering, new customer channels, deep domain expertise, and expanded coding capabilities. Treating AI as a single initiative is how companies underinvest in the ones that actually move revenue.



The biggest questions on AI today

WITH SCOTT GALLOWAY, PIVOT & GREG SHOVE, SECTION

Scott Galloway and Greg Shove opened the AI:ROI Conference by debunking alarmist predictions about AI-driven unemployment and analyzing what AI ROI actually means in practice.

"The only way they can justify these valuations is to offer up EBITDA through cost cutting. The only way they can justify that is to project a world of labor apocalypse."

- SCOTT GALLOWAY

While doomsday narratives about AI destroying jobs make for compelling headlines and help justify inflated startup valuations, the reality on the ground tells a different story. Executives are seeing real productivity gains in software engineering and knowledge work, but these gains haven't yet translated into widespread labor displacement or consistent revenue growth. The key distinction is between genuine productivity improvements and narrative-driven cost-cutting dressed up as AI transformation.

Key insights:

1 AI is corporate Ozempic, not a magic bullet.

Companies like Meta demonstrated you can grow revenue 23% with 20% fewer people, but this is the exception. Most executives haven't yet seen the promised productivity gains despite massive investments. The market is in "no hire, no fire" mode as companies wait to prove AI delivers.

2 Incumbent companies must move fast with internal talent.

Stop relying on external consultants and instead empower your best internal people to war-game how AI startups will attack your business. Without fast, committed internal transformation, incumbents risk losing 15-20% of revenue to AI-native competitors.

3 Job destruction fears are overstated; loneliness risk is understated.

There's little evidence AI is disrupting labor markets yet, following historical precedent where automation destroyed roles but created net new work. The real understated risk is cognitive outsourcing and social isolation as AI sequesters people from meaningful human relationships.

The path forward for AI ROI isn't catastrophizing or hype - it's pragmatic, internal execution. Leaders must stop waiting for certainty, move fast with their best people, and prepare for a 3-5 year transition where entrepreneurship will flourish.



Integrating AI into team workflows

WITH MATT LYTESON, IBM

Matt Lyteson, IBM's CIO for Technology Platform Transformation, detailed how IBM systematically drives AI ROI across individual, team, and enterprise-wide workflows using a tiered framework approach.

"From a CIO perspective, I need to be able to enable a series of capabilities across the enterprise."

IBM's framework recognizes that enterprise AI requires fundamentally different thinking than personal productivity tools. The company organized its approach into four layers: individual agents, team-level sharing, core workflow automation, and a centralized enterprise agent.

Key insights:

1 Structure ROI around workflows, not just tasks.

While individual productivity gains matter for employee happiness, real ROI emerges at the workflow layer. IBM identified 12 top-level workflows decomposed into hundreds of components. CFOs care about measurable impact at this scale, not just personal time savings.

2 Create a unified "front door" to manage AI proliferation.

Matt's team deployed "Ask IBM" as a central concierge agent to prevent confusion about their many AI tools and hundreds of agents. It functions as a single entry point for employees to discover the right AI tool, access data interactively, and receive role-specific analysis.

3 Build small, cross-functional teams around specific workflows.

IBM uses 7-9 person fusion teams - mixing technologists with business domain experts - aligned on a single outcome like contract review or inventory management. That blend of AI skill + business context is what turns pilots into production.

The takeaway - start with clear outcomes, small empowered teams, and 90-day sprints - then layer in enterprise-wide orchestration as you prove value.



The state of the AI market

WITH MARK MAHANEY, EVERCORE

Mark Mahaney, a veteran tech analyst at Evercore with 25 years covering Internet and tech companies, shared his perspective on AI market investments and ROI, contrasting the present moment with past tech bubbles.

"As an investor, you look for supply-constrained industries - where you should have invested 100 billion dollars two years ago."

Mahaney emphasized that the trillion-dollar capital expenditure by tech giants is fundamentally different from the dot-com bubble. Today's spending is funded by current cash flow from highly profitable companies with seasoned management teams.

Key insights:

1 Supply constraints signal a bullish market.

Major cloud providers are explicitly stating they would generate more revenue if they had more compute capacity. This supply constraint is a positive indicator for investors, not a concerning sign of overcapacity.

2 AI margins are still lower than SaaS.

Return on invested capital for AI services is lower than traditional cloud/SaaS businesses due to the cost of compute and inference. Companies must drive scale and cost efficiency to maintain competitive pricing while preserving margins.

3 Speed and latency are a competitive moat.

Reducing latency in AI responses will drive adoption, similar to how Google's search speed improvements increased usage. Even shaving seconds off response times creates meaningful user experience improvements and increases volume.

"Companies that do more with more - generate more revenue per employee with more employees - those are the most interesting companies."

4 Revenue per employee is a key metric.

Rather than pure cost-cutting, investors track whether companies can generate more revenue per employee through AI. This productivity metric is more valuable than simple headcount reduction, indicating genuine innovation driving economic value.



The biggest opportunities for AI ROI in healthcare

WITH KRISTIN MYERS, NORTHWELL HEALTH

Healthcare organizations face significant pressure to adopt AI while maintaining regulatory compliance and ensuring patient safety. New York's largest healthcare system, Northwell Health, has demonstrated how to scale AI implementation across thousands of clinicians in a highly regulated environment.

The key challenge isn't finding AI tools - it's moving from isolated pilots to organizational value through structured governance, careful vendor evaluation, and thoughtful change management.

Key insights:

1 Rigorous vendor evaluation drives adoption.

Northwell conducted a nine-month pilot with 30 physicians testing multiple ambient documentation vendors to ensure the selected tool integrated with workflows and earned clinical trust. This extended timeline upfront prevented costly implementation failures.

2 ROI in healthcare requires reframing success metrics.

Rather than measuring traditional financial ROI, Northwell focuses on reclaiming clinician time, reducing burnout, and improving retention - metrics that directly impact competitiveness in a tight labor market. Time freed up enables growth investments rather than simple cost reduction.

3 Governance is an enabler, not a barrier.

Northwell embeds finance and compliance directly into AI decision-making through an AI executive committee and structured intake process. This disciplined approach actually accelerates value delivery by ensuring investments align with organizational strategy and clinical priorities.

Northwell Health demonstrates that healthcare organizations can move fast with AI while respecting regulatory requirements by combining physician-led pilots with centralized governance. The 1,100 physicians using Ambient in just five months proved structured implementation works.

"You can't just look at the hard ROI."



Governance as an accelerant for ROI

WITH DJ SAMPATH, CISCO

DJ Sampath, SVP of AI at Cisco, discussed how governance frameworks actually accelerate AI adoption rather than slow it down, framing security and governance as essential enablers of enterprise AI value.

"The enterprise needs to know that you can use these agents in a safe and secure manner. It's the number one blocker."

Cisco pioneered a pragmatic approach to AI governance that treats it as an accelerant rather than a barrier. By establishing clear frameworks for discovery, detection, and protection, the company has enabled widespread AI adoption while maintaining enterprise security.

Key insights:

1 Governance is foundation, not friction.
Effective governance starts with visibility into where AI workloads run, detection of vulnerabilities in models and data, and runtime protections that prevent unwanted behaviors. This three-part framework - discover, detect, protect - builds trust that enables faster, not slower, AI deployment.

2 Shared language drives cross-functional alignment.
Cisco released an open-source AI safety and security taxonomy that gives legal teams, engineers, security, and business stakeholders a common vocabulary to discuss AI risks. This shared framework transforms governance from a blocker into a conversation accelerator.

3 Cultural ownership beats mandates.
Leadership must cultivate a mindset of "we manage AI, not the other way around." Cisco's responsible AI team spans legal, product, engineering, and compliance; they meet regularly and adapt assumptions weekly as the technology evolves, allowing employees to innovate responsibly at scale.

Governance is not a speed bump - it's the foundation that makes speed sustainable. Organizations that front-load governance conversations, invest in shared frameworks, and empower cross-functional teams see faster time-to-value and lower risk.



The pilot-to-production process at UKG

WITH PRAKASH KOTA, UKG

Prakash Kota, CIO at UKG, discussed how to identify which AI pilots deserve to scale versus which should be abandoned, balancing rapid experimentation with structured governance at a 14,000-person SaaS workforce platform company.

"It's not just about launching. It's also about landing."

UKG learned that successful AI adoption requires integrating capabilities directly into existing workflows rather than creating isolated tools. The company developed a tiered portfolio approach that mirrors venture capital methodology, with clear outcome-based prioritization.

Key insights:

1 Eliminate silos with centralized discovery.
Create an AI hub where all ideas are populated and tracked to prevent duplicate experimentation across teams. Assign regional champions who are subject matter experts to prioritize use cases based on business outcomes rather than technology coolness factor.

2 Embed AI into workflows, not sidebars.
Adoption fails when AI tools exist as separate chat windows or detours from core work. Force integration into existing workflows or redesign business processes to make AI the natural path of work, not an optional playground.

3 Measure outcomes over activities.
Stop tracking prompt volume or tool usage rates. Instead, measure how many AI capabilities are integrated into workflows, how much code reaches production, and whether business processes are demonstrably more efficient and effective.

UKG's three-tiered pilot portfolio allocates resources to scaling high-ROI use cases, growth opportunities with new customer value, and exploration for learning, with a maximum 90-day iteration cycle and transparent feedback loops.

Transparent communication about how AI transforms roles rather than eliminates them - helping employees see skill redeployment opportunities - dramatically increases adoption and removes organizational resistance to pilots and production deployments.



Bringing a legacy business into the AI era

WITH EDO SEGAL, NAPSTER

Edo Segal, CTO of the reimagined Napster, shared how he's transformed a legacy music brand into an AI-native company by building democratized intelligence for both consumers and enterprises using cutting-edge agent technology.

"Every day was like a month of productivity. Meaning, every day, I saw what we usually took a month to build."

Since Napster acquired his company, Segal has been rapidly implementing AI agents throughout the organization while managing two business lines: consumer-facing creative agents and enterprise solutions.

Key insights:

1 Lead by example, not mandate.
Rather than forcing adoption, Segal demonstrated AI's value by sitting with his engineering team for a week-long immersion. He showed that each engineer could achieve what 20 people previously did, creating empirical proof that shifted team mindset from resistance to enthusiasm.

2 Choose revenue growth over cost cutting.
While many executives focus on AI-driven headcount reduction, Segal argues this is a death spiral. The winning strategy is amplifying human capability to expand revenue and impact - hiring open-minded talent and enabling them with AI tools rather than replacing them.

The companies that win with AI won't be the ones cutting headcount - they'll be the ones using it to expand what their people can do. Segal's approach proves the point: show results, don't preach.



Managing your AI agent ecosystem

WITH DIANE IGOCHE, SALESFORCE

Diane Igoche, Director, Agentforce Governance & Strategy at Salesforce, discussed how to build governance frameworks and measure ROI when managing thousands of AI agents across an organization.

"I can't govern what I don't know."

Key insights:

1 Define your why before building.
Before green-lighting an agent, establish clear baseline metrics (leading indicators) and outcome metrics (what matters to leadership) to measure agent performance. Different agents require different ROI measurements - internal employee agents measure time savings, while external customer-facing agents measure cost savings and satisfaction.

2 Distinguish three types of metrics.
Use efficiency metrics (core metrics), activity metrics (leading indicators), and outcome metrics (what matters to leadership) to measure agent performance. Different agents require different ROI measurements - internal employee agents measure time savings, while external customer-facing agents measure cost savings and satisfaction.

3 Implement staged governance councils.
A multi-disciplinary council reviews use cases weekly, assessing technical feasibility, architectural alignment, and business strategy before approval. This prevents sprawl and ensures agents are built strategically rather than ad hoc across the organization.

Governance isn't just compliance - it's how you scale agents without losing control. The organizations that build structured intake, define clear metrics, and establish cross-functional councils from the start will move faster, not slower.



AI as a growth lever

WITH WADE FOSTER, ZAPIER

Wade Foster, CEO of Zapier, discussed how to leverage AI as a growth lever for your business, sharing strategies for company-wide adoption and practical implementation approaches.

"If this is how fast things are going to get better, then this has massive implications on our business. It can be a huge opportunity if we strike while the iron's hot."

Foster's journey at Zapier demonstrates the importance of recognizing inflection moments in AI capability and acting decisively. The company's Code Red response in March 2023 transformed AI from a novelty into a core operational strategy.

Key insights:

1 Run company-wide hackathons to build AI muscle.
Zapier paused operations for a week and conducted an all-hands hackathon where every employee experimented with AI tools. This single initiative increased weekly AI users from 10% to over 50%, proving that hands-on experience beats passive training.

2 Create a chief of staff agent as the universal starting point.
Foster recommends having all employees build a personal AI chief of staff using a simple prompt: describe your role and ask the AI to make you better at it. This creates daily briefs, task prioritization, and feedback loops that help workers discover additional automation opportunities naturally.

3 Give your agents a knowledge base.
Build repositories at the company, team, and individual level that agents can pull from - rather than relying on the AI's memory alone. This means your agents stay useful no matter which model you're running.

"You want people who are super duper curious and ready to go with the shovel. If you get that in your organization, you hold on to it. It is so precious."

4 Hire for curiosity and slope, not pedigree.
In the AI era, maximal curiosity outweighs experience or credentials. However, experienced people combined with high curiosity become super-performers. The key is identifying and protecting individuals who are genuinely interested in experimenting with AI, regardless of age or background.

5 Measure ROI through departmental KPIs and iteration cycles.
Move beyond efficiency metrics to track meaningful business outcomes like cost reduction or revenue impact. Zapier's support team now handles 50% of tickets via AI, and their SDR automation saves 10-20 hours weekly - proof that vertical AI deployments drive measurable results when tied to specific departmental goals.



Siemens' matrix for high-value AI initiatives

WITH HANS DE VISSER, SIEMENS

Hans de Visser, Chief Product Officer at Mendix within Siemens, described how to build high-impact AI agents by focusing on business outcomes, designing for architectural flexibility, and emphasizing human oversight throughout the agentic system.

"The AI is not just gonna do it right. There's a misperception about broad autonomy where the AI knows everything."

De Visser emphasized that successful AI agent deployment requires moving beyond the fantasy of fully autonomous systems. Organizations must think systematically about when agents deliver value, how to integrate human oversight, and what architectural patterns ensure long-term sustainability.

Key insights:

1 Vertical agents, not horizontal ones, drive business value.
Research shows horizontal agents improve productivity incrementally but don't move financial needles, while vertical agents specialized in specific domains substantially reduce costs and optimize operations. Focus resources on domain-specific agents that solve measurable business problems.

2 Design for architectural obsolescence from day one.
Build agents with pluggable components so you can swap LLMs and techniques as the field evolves rapidly. De Visser's team switched between models within half an hour because they decoupled architectural decisions from customer experience, ensuring flexibility as new breakthroughs emerge.

3 Human-in-the-loop is essential, not optional.
Agents have token costs and probabilistic outcomes - the fantasy of fully autonomous systems is a misconception. Design agents to ask permission before acting, enable human course-correction, and create feedback loops so agents improve over time through human guidance.

"If you can develop 10 times faster, you can also create a mess 10 times faster."

4 Start small with departmental use cases, then scale.
A Brazilian glass manufacturer customer built a virtual engineer agent that reduced production downtime fix time by 85% and saved 6,000 hours annually. Begin with narrow, high-impact problems like email categorization or equipment monitoring - not wholesale department automation - then expand as you gain experience.



The 5 AI value models

WITH BRICE CHALLAMEL, OPENAI

Brice Challamel, head of AI strategy and adoption at OpenAI, articulated a fundamental shift in how organizations should think about AI value creation - moving from a single utility model to a portfolio approach with five distinct value creation mechanisms.

"We have left the era of AI that speaks with us - to enter the era of AI that works for us."

Understanding where your organization sits in its AI adoption journey - from access and adoption through how to measure, ways of working, and reorganization - determines which value model applies and how proficiency success.

The five AI value models:

1 Broad adoption: the foundation layer.
Wall-to-wall workforce empowerment remains the essential first step, enabling people across all functions to use AI as a thought partner and assistant. This creates immediate productivity gains but requires intentional change management to sustain momentum.

2 Process reengineering: the revenue multiplier.
Business model transformation starts happening when you stop copying old workflows and instead reinvent how work gets done. This moves you from incremental cost-cutting to expanding revenue by doing multiple cycles of work or creating entirely new products.

3 Engagement channeling: new customer vectors.
Organizations must develop new ways to meet customers through AI conversations, integrated applications in ChatGPT and other platforms, and emerging advertising channels. This requires your team to live and breathe these capabilities to authentically engage with your audience.

"We need to think about AI in terms of portfolio management."

4 Deep expertise: vertical AI capabilities.
AI is moving beyond general-purpose use into domain-specific depth - running clinical trial research in pharma, modelling financial health in financial services, and coordinating complex patient-doctor-caretaker workflows in healthcare. These verticalized capabilities are where growth happens - rather than cutting costs incrementally, they let organizations do fundamentally different work than what was possible before.

5 Expanded coding capabilities: dependency management.
Coding agents now understand dependencies across complex systems - not just software, but manufacturing procedures, legal contracts, and business processes. This systems-level intelligence lets organizations automate work that previously required deep institutional knowledge.