



Mayfield IT Leadership Network

Enterprise AI in 2025: The CIO's Roadmap

Insights from 200 IT Leaders Driving AI Transformation

January 2025

Inside...

The State of Enterprise AI: Looking Ahead to 2025	<u>4</u>
Early AI Success Stories: From Pilot to Production	<u>6</u>
The AI Investment Landscape: Build vs. Buy.....	<u>9</u>
Security Challenges: The Path to Enterprise AI	<u>11</u>
The AI Leadership Question: Evolving Ownership Models.....	<u>13</u>
IT Budget Outlook: Overall Budget Sentiment	<u>14</u>
From Experimentation to Execution: Real-World Insights from IT Leaders	<u>18</u>

Navigating AI's Enterprise Reality: A View From the Field

2024 represents a turning point for AI in the enterprise. No longer a future prospect, AI has become a present reality—with most Fortune 500 companies deploying it, half their workforce reporting productivity gains, and C-suites ranking it among their top three priorities. Yet this swift adoption presents IT leaders with a complex challenge: how to integrate AI initiatives—some mandated from above—with existing strategic priorities and infrastructure.

This report draws from our annual survey of 200 Fortune 2000 IT leaders and broader insights from the Mayfield IT Leadership Network. Through regular forums, meetings, and events, this community of over 3,500 CIOs, CISOs, and CTOs collaborates with Mayfield to explore emerging technologies and share practical insights. Their collective experience offers a unique view into how leading organizations are navigating this transformation.

We present these findings with gratitude to our survey participants and the broader IT leadership community whose candid insights make this report possible.



Gamiel Gran
Chief Commercial Officer
ggran@mayfield.com



Shelby Golan
Director of Business Development
sgolan@mayfield.com

The State of Enterprise AI: Key Findings

Executive Summary

The enterprise AI landscape has fundamentally shifted in 2024. While our [2024 CXO Report](#) focused on AI's potential and pilots, our survey reveals a clear transition to production deployment and revenue generation. This shift brings new challenges: massive infrastructure investments and innovative partnerships with hyperscalers (U.S. hyperscalers alone have collectively [invested \\$200 billion on infrastructure](#) and may invest \$1 trillion in 2026), unprecedented demands on computing resources, and complex questions about data integration and ownership.

The story that emerges from our survey isn't just about technology adoption—it's about organizational transformation. IT leaders are navigating a complex balance between leveraging existing vendor relationships and exploring new solutions, all while managing expanding budgets and evolving governance structures.

7 Key Insights for 2025

1 AI Has Moved Beyond Experimentation

68% of surveyed organizations have at least one AI use case in production.

Production AI deployments have reached critical mass. More telling: Over one third of these deployments directly impact revenue or customer experience, marking a clear shift from experimental to business-critical status.

2 Data Remains the Primary Barrier for AI Adoption

42% of IT leaders say governance and security are a top concern. Before CIOs can deploy, they must deal with two substantial burdens: data security and data readiness. Within this are data security, privacy, access, pipelines, labeling, and more.

3 The Decline of Custom Solutions

Organizations are abandoning custom applications and private data centers in favor of public clouds and pre-built solutions. This shift reflects both the growing maturity of available solutions and the increasing costs of maintaining custom infrastructure.

4 Legacy Vendors Face New Scrutiny

While established vendors (e.g., Salesforce, ServiceNow, Microsoft, Workday, and SAP) still dominate AI implementations with pre-existing services agreements, CIOs are underwhelmed by these types of offerings. **Declining user satisfaction, rising costs, and complex implementation requirements are creating openings for new market entrants.**

5 AI Governance Remains Decentralized

Only 10% of organizations have appointed Chief AI Officers, with CIOs leading AI initiatives in 43% of cases. This decentralized approach often results in business units driving use cases and bringing in IT leadership for execution.

6 IT Budgets Are Increasing for 2025

The budget outlook for 2025 is notably positive, with **77% of leaders expressing optimism about budget increases.** Gartner's projected 9.3% growth in IT spending appears well-aligned with our findings.

7 AI Investment Tracks Overall IT Growth

AI spending is increasing proportionally with overall IT budgets. Most organizations are funding AI initiatives through existing budget allocations to fund AI pilots and prototyping, with **76% planning increased investment in generative AI and LLMs** specifically.

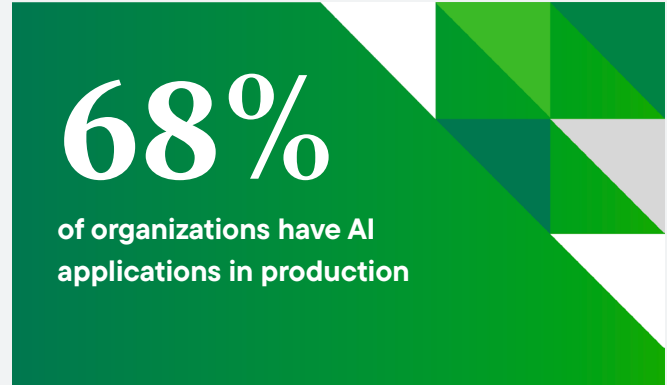
Looking Ahead

These findings suggest a maturing AI landscape where the key challenges are shifting from technical feasibility to organizational integration and scale. Success in 2025 will likely depend less on selecting the right AI technology and more on building the right organizational structures and data foundations to support it.

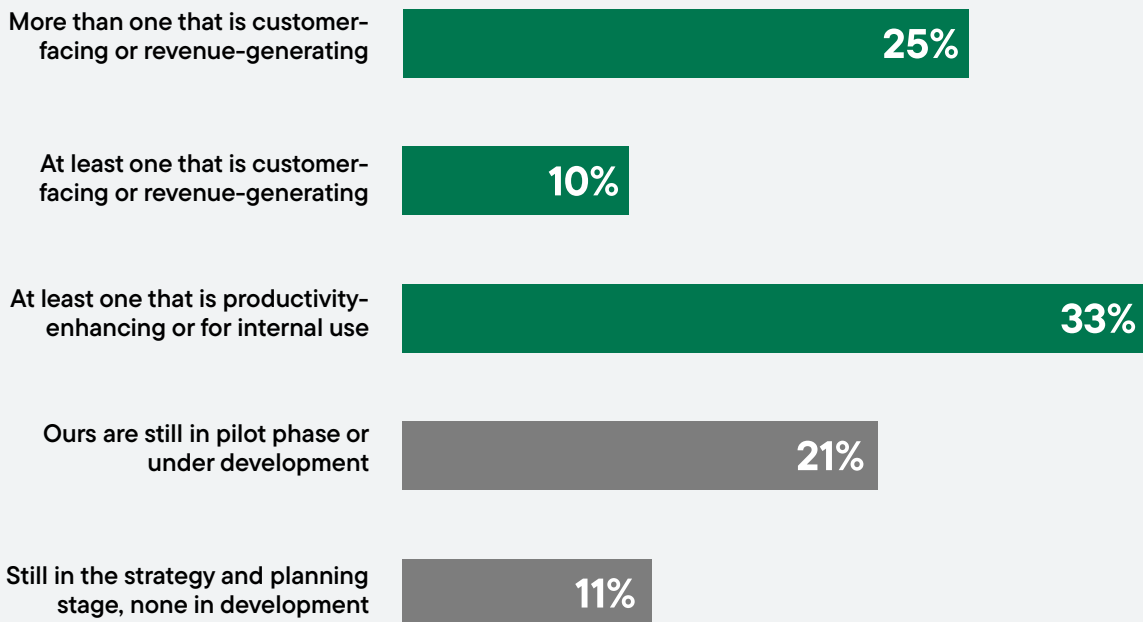
Early AI Success Stories: From Pilot to Production

The Shift to Production-Ready AI

The enterprise AI landscape marks a fundamental shift from traditional automation to intelligent adaptation. Our survey reveals a decisive shift: 68% of organizations now have AI applications in production, with 35% already deploying customer-facing or revenue-generating solutions. This rapid adoption reflects a key difference from previous automation waves—today’s AI solutions offer adaptable, context-aware capabilities that can process natural language inputs, making them accessible to non-technical users across the organization.



How many Gen AI applications do you have in production?



Four Core Implementation Areas

The range of early use cases is remarkably broad, though several core areas are emerging as the go-to applications. Below are four horizontal use case areas and emerging areas that are making significant progress.



Customer Support

Self-Service Chatbots

Provide automated support for common queries and simple issues.

Agent Assistance

Help service agents quickly find relevant information within knowledge bases.

Sentiment Analysis

Monitor and analyze customer interactions across social media, reviews, and communications.

Intelligent Routing

Direct inquiries to appropriate teams based on expertise, availability, and complexity.

Vertical Solutions

Deploy industry-specific support like claims processing and patient care assistance.



General Productivity

Task Automation

Automate routine work in scheduling, email management, report generation, data entry, and analyzing reports.

Employee Self-Service

Enable staff to independently access resources, policies, and handle basic IT requests.

Collaboration

Provide automated task tracking, deadline management, smart search, and content organization.

Content Generation

Create first drafts of sales materials, marketing content, and communications.

Design Assistance

Offer design suggestions, automated layouts, and intelligent component recommendations.



Developer Productivity

Code Auto-Complete

Suggest contextually relevant code snippets during development to improve accuracy and avoid rework.

Natural Language Programming

Transform plain English descriptions into functional code to ease the learning curve.

Automated Code Review

Identify bugs, security vulnerabilities, and style issues to streamline reviews before deployment.

Test Generation

Create comprehensive test cases based on code changes to improve test coverage and reduce manual testing.

Documentation Assistant

Generate and maintain technical documentation from code comments and sources.



Synthesizing Content

Document Analysis

Summarize complex reports, research papers, and legal documents so users can quickly interpret large volumes of information.

Transcription

Transcribe and extract key points, actions, and decisions from meetings, interviews, doctor visits, or other forms of audio content to summarize key points, actions, and next steps.

Report Generation

Compile and structure data from multiple sources into cohesive reports.

Language Processing

Provide real-time translation and interpretation of multilingual content (audio and video).

Media Analysis

Process and analyze images, video, and audio content for insights like vehicle inspection photos, drone video on damaged infrastructure, manufacturing, construction, and facial recognition.



Emerging Areas

Predictive Analytics

Forecast demand, assess risks, predict outcomes, inventory management, labor demand, HPC optimization, price adjustments, and customer behavior analysis (including churn).

Security Operations

Detect threats (phishing, malware, other anomalies), automate incident response, conduct digital forensics, fraud detection, and SOAR.

Sales Ops and Enablement

Identify opportunities, generate competitive insights, and automate outreach.

Healthcare Innovation

Support precision medicine, pattern recognition, drug discovery, clinical trial optimization, medical imaging, and robotic-assisted surgery.

IT Automation

Optimize cloud configuration and costs, CI/CD, automate incident response, and infrastructure-as-code.

AI in Action

Klarna

Their **AI customer service agent** handles two-thirds of customer chats (refunds, payments) saving \$40M yearly.



Reduced third-party coding needs by 70% through **AI-assisted development**. AI is now handling a significant portion of the company's coding. Human developers validate the final 30%.

Walmart 

Deployed Wallaby, **a retail-specific LLM** for customer-facing experiences.

JPMorganChase

Implemented an AI tool for wealth management called **LLM Suite** based on OpenAI technology.

TESLA

Advanced autonomous vehicle development with **Cybercab**.

Investment Signals

The market's confidence in enterprise AI is reflected in 2024's investment patterns, with over ten billion-dollar funding rounds. This level of investment suggests we're entering a phase where clear market leaders will emerge, backed by both venture capital and corporate strategic investments.

The AI Investment Landscape: Build vs. Buy


Our survey responses suggest a clear preference in the enterprise AI market: Most organizations plan to buy rather than build their AI solutions, primarily leveraging capabilities embedded within existing enterprise platforms. However, respondents also indicate an emerging hybrid approach—while they rely heavily on established vendors, many are beginning to supplement these core platforms with specialized AI providers for specific use cases.




The Enterprise AI Platform Race: Microsoft Leads

Microsoft has established clear dominance in the enterprise AI space, with 72% of respondents utilizing their AI offerings. This commanding lead over Amazon (17%) and Google (10%) stems largely from their strategic OpenAI partnership and the successful deployment of Copilot, which has created a robust ecosystem of AI-powered applications across their product suite.


Frequently-Cited AI Vendors


 **Claude** Emerging as a leading alternative to OpenAI


 **databricks** Data analytics and AI infrastructure

watsonx Enterprise-grade AI solutions





















scale AI development and deployment platform

 **Moveworks** AI development and deployment platform

 **Glean.ai** Enterprise search and knowledge management

 **perplexity** Advanced search and analysis

Emerging Solutions by Category

 Developer Tools  Dosu Code maintenance and optimization  prisma.ai Visionary Intelligence Intelligent database interactions  Katalon Advanced test automation	 Sales and Marketing  Outreach AI-powered sales engagement  copy.ai Marketing content generation
 Search and Productivity  YOU AI-enhanced search and assistance  Orby Personal workflow optimization	 Data Operations  surge^{AI} Intelligent data labeling  Braintrust AI-powered talent acquisition
 Customer Service  CRESTA Service optimization  Embrace.ai Customer experience  Forethought Support automation  glia Digital customer service  kore.ai Conversational AI  yellow.ai Customer engagement	

Strategic Implications for CIOs

Platform Integration: The dominance of established vendors suggests CIOs should evaluate how to maximize AI capabilities within existing enterprise platforms before seeking specialized solutions.

Specialized Solutions: While core AI capabilities may come from primary vendors, specialized providers are proving valuable for specific use cases where deep expertise matters.

Vendor Selection Criteria:

- Integration capabilities with existing systems
- Enterprise-grade security and compliance
- Scalability and performance track record
- Total cost of ownership, including implementation and training
- Vendor stability and market position

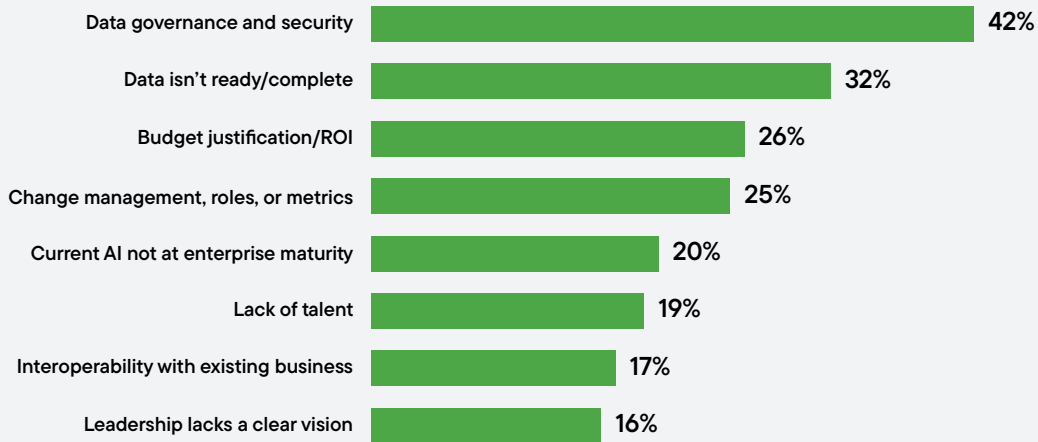
Risk Management: How do you balance the risk of vendor lock-in against the complexity of managing multiple AI platforms? Keep proof-of-concept projects ready with alternative providers for critical functions.

Security Challenges: The Path to Enterprise AI

Despite the AI adoption surge in 2024, enterprise leaders face a complex set of deployment challenges. Our survey paints a clear picture: 42% cite security and governance as their primary concern, while another third struggle with data readiness. Beyond these technical hurdles, leaders report significant gaps in internal expertise, strategic uncertainty, and infrastructure limitations—all while facing pressure to justify AI investments.



What's holding you back from getting AI use cases into production?



Security remains the foremost challenge in enterprise AI adoption, driven by three key factors:

Data Exposure Risks

- AI systems require extensive training data, often containing sensitive information
- Traditional security measures may not adequately protect AI-specific vulnerabilities
- Growing concerns about intellectual property protection in AI systems

Model Integrity

- Potential for training and input data manipulation
- Risk of adversarial attacks on production models
- Challenges in validating model accuracy and reliability

Operational Complexity

- Difficulty in monitoring AI system behaviors
- Challenges in identifying security vulnerabilities
- Need for specialized security expertise

Additional Implementation Challenges

Survey respondents highlighted several emerging challenges:

- Difficulties integrating AI into a broader data pipeline
- Lack of employee readiness, or even employee pushback
- Concerns around the regulatory environment
- A belief that the accuracy LLMs provide today is not good enough for external use cases
- Troubles evaluating these new use cases given that standard tests such as MMLU are not effective

Looking Ahead

Organizations need a comprehensive approach to AI security and implementation. Check out [Gartner's AI TRiSM](#) (Trust, Risk, and Security Management) framework for proactively identifying and mitigating the risks in AI models and applications.

As AI adoption accelerates, organizations must elevate data security from a technical consideration to a strategic priority. This requires:

- Development of robust AI governance frameworks
- Investment in specialized security tools and expertise
- Regular assessment and updating of security protocols
- Balance between innovation and risk management

The path to secure AI implementation requires organizations to think beyond traditional security measures and develop comprehensive approaches that address both technical and organizational challenges.


The AI Leadership Question: Evolving Ownership Models

Current State of AI Leadership

While CIOs lead **43% of enterprise AI initiatives**, our survey reveals a fragmented leadership landscape. Only 10% of organizations have appointed dedicated Chief AI Officers, and 14% report no clear owner at all. The remaining organizations split ownership across various C-suite roles, from CTOs (14%) to Chief Digital Officers (7%).



Role and Ownership Percentage



* Chief AI and Technology Officer, Chief Data and AI Officer, CEO, CFO, Chief Analytics Officer, Chief Architect, Chief Innovation Officer, Chief Scientist, CISO, Chief Product Officer, VP Advanced Technology

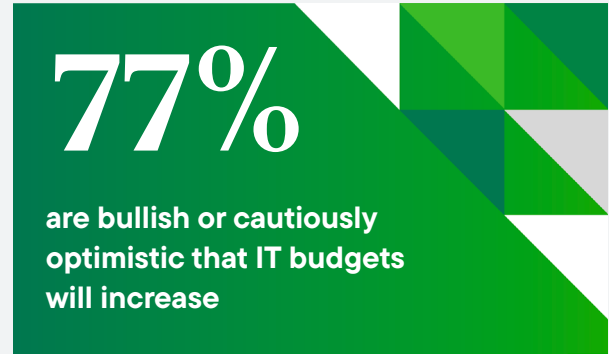
Teamwork Makes the AI Work

Effective AI deployment emerges from enterprise-wide collaboration. Data science teams craft the models, IT provides the foundation, product teams create user experiences, and legal/security teams ensure protection and compliance. Success depends not on any single group, but on their cross-functional coordination.


IT Budget Outlook

Overall Budget Sentiment

Seventy seven percent of IT leaders express optimism about 2025 budgets, aligning with **Gartner's projected 9.3% increase** in global IT spending to \$5.74 trillion.




How likely is your IT budget to increase in 2025?



Investment Priorities for 2025


High-Growth Areas for Increased Budget Spending

As in 2024, Generative AI and LLMs are the top categories earmarked for increased spend in 2025.



Areas Facing Cuts

Where will cuts come from? Outsourced services such as contractors and management consultants. Traditional service providers may not yet have the talent required to address complex AI-related needs, leaving startups that focus on services with an opportunity to get a foot in the door at the enterprise level.



Where and how will your spending change?

Decrease
 Flat
 Increase


Cybersecurity




Business Productivity Apps (With or Without Gen AI)



Data and AI



Services




AI Budget Projections

The share of IT budgets allocated to AI is set to grow. While Gen AI projects represented a modest 1–2% of IT spending in 2023, our survey reveals plans for significant expansion. Organizations expect to more than double their AI investment by 2025, reaching 4–5% of total IT budgets—a clear signal that AI is moving from experimental to essential.

The scale of AI investment reveals sharp contrasts across the enterprise landscape. [TheCUBE's research](#) exposes the divide: While 33% of organizations invest less than \$100,000 annually in AI, market leaders—representing 6% of companies—deploy over \$10 million. This gap appears most pronounced among large enterprises, where 26% of companies with \$5B+ revenue plan to commit more than 10% of their IT budgets to generative AI by 2025. These investments align with broader IT spending trends, as 43% of our survey participants expect AI budget increases of 10% or more.



How much do you expect your IT budget to increase?



ROI Expectations

With a real budget comes real expectations. One-third of survey respondents require hard ROI metrics (in terms of cost savings or top-line revenue growth) to justify AI investments, while others accept softer metrics or are still developing measurement frameworks.

What do you need to see to measure an AI project's success?



Looking Ahead

As AI moves from experimental to essential, organizations must:

- Develop clear ROI frameworks
- Balance innovation with operational efficiency
- Shift resources from legacy systems to AI-enabled solutions
- Build measurement capabilities for AI initiatives

The 2025 budget landscape suggests a clear pivot toward AI-driven transformation, with organizations reallocating resources from traditional services to new capabilities. Success will depend on balancing ambitious AI initiatives with pragmatic implementation and clear performance metrics.

From Experimentation to Execution: Real-World Insights from IT Leaders

As we step into 2025, AI has transitioned from a field of experimentation to a key catalyst for enterprise transformation. In this report, we showcase interviews with IT leaders from diverse industries who share how their organizations are successfully adopting AI while overcoming its challenges. Their insights provide a valuable, practical roadmap for navigating this transformation—covering everything from AI integration to measuring ROI.

Art Hu Global CIO and Chief Technology and Delivery Officer, Lenovo

Ben Davis EVP, Information Technology, Cambria

Paul Heard Former Chief Information Officer, Mimecast

Jonathan Hodges Chief Architect, AI & Machine Learning, Bolt

Shyam Bhojwani Senior Director, Business Technology & Cybersecurity, Workato

Jon Harding Senior Vice President, Global Chief Information Officer, Conair

Mark Dash VP and CTO, Information Management & Technology, The Penn Mutual Life Insurance Company

Sathish Muthukrishnan Chief Information, Data and Digital Officer, Ally Financial

Howard Miller Chief Information Officer, UCLA Anderson School of Management

Mike Brown Chief Data Officer and Co-Founder, Strider

Leo Griffin Former Vice President, Global Head of Consumer Technology, Hanesbrands

With AI, Less Can Be More

Lenovo

Art Hu, Global CIO and Chief Technology and Delivery Officer,
Solutions and Services Group

A multinational technology company that develops, manufactures, and sells personal electronics and enterprise solutions.

The rapid influx of AI investment demands a focused, disciplined approach to implementation. Rather than pursuing hundreds of potential use cases, we're deliberately limiting ourselves to 5-10 high-impact initiatives that can scale across our entire organization. This focused strategy ensures we can properly manage change, allocate resources effectively, and achieve meaningful double-digit improvements in our chosen areas.

“The first 90% is easy, the last 10% is hard—there are no shortcuts when implementing AI at scale.”

Our approach to customer contact centers exemplifies this strategy. Instead of simply adding AI capabilities to existing systems, we're reimagining how AI can run entire contact center operations. This requires looking beyond superficial chatbot implementations to consider deeper, more meaningful integration. The same focused approach applies to our software engineering and marketing initiatives, where we're seeing particularly strong results in marketing applications that don't necessarily require binary right/wrong answers.

This selective approach yields practical benefits, particularly in resource allocation and financial planning. When you have a narrow subset of high-conviction use cases, developing clear resource requirements and how they're capturing value becomes much easier. This clarity helps secure buy-in from finance teams because we can commit to specific, measurable outcomes rather than making vague promises about future potential.

While maintaining focus on core initiatives, we still encourage employee experimentation through small-scale projects. This balanced approach ensures we don't stifle innovation while keeping our primary resources aligned with strategic priorities. It's crucial to recognize that there's often a significant

gap between rapidly evolving AI capabilities and an enterprise's ability to implement them effectively.

Data quality and management have emerged as

critical success factors in our AI initiatives. Estate management—the systematic organization and governance of data assets—often represents more than half the battle in successful AI implementations. Our experience shows that grounding yourself in data fundamentals is essential for AI to deliver value at scale over time. Poor-quality data inevitably leads to poor inference and results, regardless of how sophisticated your AI models might be.

Key Takeaways

- Focus on 5-10 high-conviction use cases rather than pursuing hundreds
- Commit to measurable outcomes when seeking resources
- Balance-focused execution with room for experimentation
- Recognize the gap between AI capabilities and enterprise readiness

AI as Copilot: Empowering Employees Through Intelligent Automation



Ben Davis, EVP, Information Technology

Provides custom software development and IT services.

The key to successful AI implementation lies in engaging business leaders and positioning the technology as an enabler rather than a disruptor. At Cambria, this approach begins with identifying champions across business units who deeply understand processes and opportunities to optimize and scale them.

The foundation of our strategy rests on three pillars. First, we focus on business engagement, working closely with unit leaders who understand their processes intimately. Second, we evaluate data readiness—examining dataset size, accuracy, and growth potential—applying the same rigorous prerequisites we use for any transformation initiative. Third, and perhaps most crucially, we emphasize employee adoption through careful branding and change management.

“Change management is so much more than just the technology, and having business leaders play a significant role is the key to widespread adoption.”

Our decision to brand AI projects as “copilots” reflects a deliberate strategy to position AI as a complement to human talent rather than a replacement. This messaging and active business leader involvement

have been crucial for widespread adoption. We measure success through quantifiable metrics, tracking time savings, and applying rate calculations to demonstrate concrete value.

Looking ahead, we’re exploring how AI can strengthen our partner network by enhancing data sharing and application capabilities. While we remain optimistic about AI’s potential, we recognize the importance of proper measurement frameworks as AI becomes integrated into daily business functions.

Key Takeaways

- Business unit engagement is essential for successful AI implementation
- Data readiness assessment should precede any AI initiative
- Strategic branding of AI as a “copilot” helps drive employee adoption
- Measuring concrete value through time savings helps justify AI investments

The AI Toolbox: Buying, Building, and Blending Solutions for Success

Paul Heard, Former Chief Information Officer, Mimecast

A cloud-native cybersecurity provider specializing in email and collaboration security solutions.

Unlike previous technology waves, where implementation paths were clear, AI requires a more exploratory approach. I've found that overall, traditional evaluation metrics don't perfectly fit AI initiatives. Instead of starting with rigid business cases, it's easier to invest in understanding and building organizational capability alongside achieving specific outcomes.

That being said, there's still a need to evaluate business cases for AI, and define the expected outcome and benefits, as well as the likely costs. There are many good off-the-shelf solutions to get started with. I recommend a mixed buy-and-build approach. If your current vendor has built-in AI, investigate that against specific AI add-in solutions. For some needs it is quicker and easier to build your own and partnering data engineers with prompt engineers gives you the development skills you need.

user groups to begin with. It's not the right time to tackle replacing a CRM or ERP, but it may be the right time to target smaller groups and enhance your existing platforms or add specific AI capabilities to your existing workflows. Having a small center of

excellence with the muscle to build Gen AI solutions is a good incubator to start things out, but equipping everyone in IT and as much of the business as possible is the ultimate goal. I recommend enabling your existing data warehouse team with the skills

to source data and ensure data quality and governance, as the data source for AI.

However, we've learned important lessons about operational costs. The speed of AI development can mask significant backend expenses, particularly in areas like Snowflake usage. Successful deployment requires careful engineering effort around optimization and tuning to manage these ancillary costs effectively.

“The iteration and speed of development is very high with AI—get somewhere in the ballpark to start, and business interest will drive iterative improvement.”

Key Takeaways

- Balance quick wins with strategic capability building
- Structure teams to support both off-the-shelf and custom solutions
- Plan for optimization to manage operational costs
- Start with targeted user groups before broader rollout

Beyond Chatbots: Building AI into Your Core Products

BOLT

Jonathan Hodges, Chief Architect, AI & Machine Learning

A fintech platform that provides one-click checkout and shopper network solutions for e-commerce merchants.

Organizations making real progress with AI have moved beyond basic experimentation to integrate the technology directly into their core products. At Bolt, our approach focuses on identifying current AI capabilities and finding immediate applications while building toward deeper product integration. While there's plenty of low-hanging fruit available—chatbots, simple automations, and similar solutions—the real value comes from figuring out how AI can enhance your fundamental offerings.

We're seeing this in action by implementing multi-modal search capabilities, similar to Google Lens, which allows shoppers to search using photos they take. This technology also enables us to enrich our product catalogs, addressing a common challenge where merchant descriptions vary widely in quality. By using AI to standardize and enhance these descriptions, we've created a foundation for improved personalization and recommendations using traditional ML models.

However, this progress requires careful consideration of security and infrastructure. Significant investment in guardrails and prompt injection prevention is essential for models going into production. Even major tech companies face challenges in these areas, so robust security measures are critical. This necessitates a cross-functional approach, bringing together engineering expertise, strategic oversight, and operational knowledge to evaluate capabilities, risks, and alignment with existing roadmaps.

The path forward is becoming clearer as the technology stabilizes. RPA isn't becoming obsolete; it's evolving as organizations move from brittle interfaces to more sophisticated automation. We anticipate that security measures traditionally reserved for employees will be adapted for AI agents while organizations develop more comprehensive automation strategies. The role of the Chief Architect is also evolving, shifting toward process modeling and system integration as agentic systems mature and take over lower-level components.

Your competitive advantage will come from strategic differentiation. Simply investing in OpenAI and cloud services won't be enough—organizations need to identify where to build unique capabilities in-house. This differentiation will vary by organization, but it's becoming increasingly critical for long-term success.

“The real challenge isn't deploying AI—it's building it into your core products in ways that create lasting value.”

Key Takeaways

- Success requires moving beyond basic AI applications to core product integration
- Security and guardrails are essential from the beginning
- Cross-functional collaboration is critical for effective AI implementation
- Competitive advantage will come from differentiated in-house capabilities, not just vendor solutions

Orchestrating AI for Business Transformation



Shyam Bhojwani, Senior Director, Business Technology & Cybersecurity

An enterprise automation and integration platform that helps organizations automate business workflows across cloud and on-premise applications.

AI is at the top of my mind for Workato as a B2B SaaS provider and “customer zero” for our own product. Our experience implementing AI across product development and internal operations has revealed the power of what we call “AI orchestration”—the strategic coordination of multiple AI systems to enhance overall capabilities.

Our standout implementation demonstrates this orchestration principle in security operations. We’ve significantly reduced our security team’s investigation time by developing AI agents that analyze logs and alerts across platforms to provide remediation recommendations. This success has led to an even more ambitious approach: creating specialized AI agents that communicate with each other. For example, our security agent now collaborates with an infrastructure agent to gather comprehensive insights, showcasing the potential of orchestrated AI systems.

“The future lies in AI orchestration—creating specialized agents that communicate with each other to enhance overall capabilities.”

We’ve moved beyond traditional SLAs to implement XLAs (experience level agreements) to measure success, focusing on end-user-centric metrics. These measurements help us quantify AI’s impact on user experience and service outcomes, providing more meaningful insights than operational metrics alone. We’ve found this approach particularly valuable for internal benchmarking of AI initiatives.

Our “Workato on Workato” strategy—implementing our product innovations across our operations before sharing them with customers—creates valuable feedback loops. This approach,

enhanced by cross-functional engagement through initiatives like our internal hackathon, has generated ideas from IT, Development, Sales, Finance, and HR teams. As these projects scale, we’re focusing on performance, doing our best to reduce some of our SaaS applications, and putting those savings toward AI.

Key Takeaways

- Balance quick wins with strategic capability building
- Structure teams to support both off-the-shelf and custom solutions
- Plan for optimization to manage operational costs
- Start with targeted user groups before broader rollout

Making AI Pay Off: Quantifying Value and Streamlining Operations



Jon Harding, Senior Vice President, Global Chief Information Officer

Designs, manufactures and sells small appliances, personal care, and beauty products.

Despite the pervasive AI hype, our approach focuses on identifying projects with clear ROI potential while establishing robust governance frameworks. This strategy has proven particularly effective in customer service, where we've seen measurable results from our AI chatbot implementation through Yellow AI.

Our initial focus on customer service automation demonstrates how to balance innovation with practical value. Every interaction our AI chatbot handles represents direct cost savings by reducing call center volume. However, the implementation journey revealed important lessons about AI deployment. Training the system required careful attention to nuance, particularly in handling sensitive situations like product liability claims. The success of this initiative—now live for one brand and expanding to Cuisinart—proves that patience in training and testing pays off in reduced risk and improved customer experience.

“AI is becoming like cloud—people initially thought of it as a strategic end in itself, but now it’s viewed as a tactic. We’re already there with AI: focusing on where we can use it effectively.”

We've also found success with internal applications, particularly in digital content creation. Using Adobe Firefly for a quick first draft of digital advertising copy has significantly reduced our content backlog, addressing a critical business challenge of maintaining updated product information across multiple retail

platforms like Walmart, Best Buy, etc. While ROI is more challenging to quantify in this case, the impact on time-to-market provides clear business value.

The implementation process has revealed that AI requires even greater business unit involvement than traditional IT projects. Success depends on subject matter experts

spending time testing and refining the system, often in short bursts between other responsibilities. This “background task” approach to AI training has proven more sustainable than intensive dedicated sessions.

Key Takeaways

- Focus on projects with clear ROI potential while building governance frameworks
- Success requires significant business unit involvement in training and testing
- Start with use cases that have measurable cost savings
- Consider time-to-market improvements as alternative success metrics

Beyond the Hype: Differentiating Gen AI in a Competitive Insurance Landscape

**Mark Dash, VP and CTO, Information Management & Technology,
The Penn Mutual Life Insurance Company**

Provides life insurance, annuities, and investment products.

The distinction between traditional machine learning and generative AI marks a crucial evolution in insurance technology. While predictive machine learning has been foundational in insurance for years, our focus has shifted to identifying unique applications of generative AI that can create competitive advantages in an increasingly tech-driven industry.

Through our dedicated predictive analytics team, “Da Vinci,” we’ve taken a measured approach to AI implementation, initially focusing on R&D without immediate ROI expectations. This approach allows us to identify practical use cases for GenAI, which remains one of the most challenging aspects of implementation: determining how and where to apply it effectively.

Traditional machine learning has also seen extensive gains, particularly in underwriting which is traditionally a human-intensive process requiring linear scaling with application volume. By integrating AI, we’ve achieved remarkable results: Offer acceptance rates have risen from 73% to approximately 90%, while individual underwriter productivity has increased by 48%.

The impact extends beyond pure efficiency gains. Our semi-autonomous tools now pre-rate applications, fast-tracking qualified candidates while reducing the time required for human assessment.

Predicting the necessity for additional medical testing can save hundreds of dollars per application by minimizing unnecessary nurse visits or diagnostic exams, also improving customer experience.

For complex cases, AI assists by summarizing extensive medical documents and linking key information to original sources, enabling underwriters to focus on high-value decision-making rather than administrative tasks.

As we expand our AI capabilities, we carefully balance innovation and responsibility through our digital ethics committee. This is particularly crucial in life insurance, where our health-related underwriting requires greater caution than property and casualty insurance’s more straightforward applications.

“Building internal expertise has been critical—it’s less about developing models in-house and more about applying them strategically to solve business challenges.”

Key Takeaways

- Focus on differentiated AI applications rather than commoditized solutions
- Measure success through concrete metrics (90% acceptance rate, 48% productivity gain)
- Balance innovation with regulatory compliance and ethical considerations
- Build internal expertise rather than relying solely on external talent

AI at Work: Real-World Use Cases



Sathish Muthukrishnan, Chief Information, Data and Digital Officer

A digital financial services company offering banking, auto financing, and corporate finance solutions.

When evaluating new technology, the fundamental question remains constant: are you solving customer problems or creating organizational value? At Ally, we approach AI implementation with this pragmatic lens, focusing on concrete problems where AI can deliver measurable improvements.

Our remarketing use case demonstrates this practical approach. Traditionally, preparing a repossessed vehicle for transport to auction involves human review of dozens of vehicle photos to confirm that the key is present – a requirement for the

vehicle to be accepted at auction. With our Ally built Artificial Intelligence and Machine Learning (AI/ML) capabilities, we can now determine if the vehicle has a key with 95% accuracy and initial results demonstrate a significant reduction in vehicle storage time enabling faster auction cycles and inventory turnover, reduced depreciation and storage expense, and improvements in recoveries. In the small number of cases where the model can't make a determination, we default to human inspection - the same process we used before. This approach not only speeds up the transport of the vehicle but also optimizes our resource allocation.

“The collective notion of struggling to find articulated value in AI is real – the key is focusing on specific, measurable problems rather than general application.”

We're also seeing significant gains in software development and internal processes. We see that writing user stories is a repetitive task for our product owners. To streamline this, we use AI to capture conversations during brainstorming meetings and convert them into user stories. Incorporating generative AI allows us to take these user stories and

automatically generate test cases. This approach significantly enhances our efficiency and reduces the manual effort involved in writing user stories and creating test cases. And even better, it's an internally-facing, human-in-the-loop use-case.

Internal audit provides another compelling example of AI's practical value. Instead of internal auditor partners spending time understanding technology applications and filling out detailed forms from scratch, GenAI provides initial analysis, suggests risk parameters, and guides the planning process. The audit planner inputs a statement, goal, or scope, and GenAI helps to fill in the next steps, allowing auditors to focus on their core expertise – actual auditing and providing feedback – rather than preliminary investigation work.

Key Takeaways

- Focus on solving specific, measurable problems
- Start with internal use cases that avoid privacy concerns
- Look for opportunities where AI can augment existing processes
- Measure success through concrete operational improvements

Building the AI-Powered University

Howard Miller, Chief Information Officer

A graduate business school at UCLA, offering MBA, PhD, and executive education programs.

The transformation to an AI-powered educational institution requires a distinct strategy from corporate implementations. Through our AI task force—similar to a corporate center of excellence—we’re approaching AI adoption through two primary lenses: enhancing student success and demonstrating thought leadership as a top 20 business school.

Our initial focus on fundraising demonstrates the practical potential of AI in higher education. By implementing AI-powered email campaigns, we’ve seen improved open rates, click-through rates, and, most importantly, increased donations from a broader donor base. This eight-week pilot proved the value of targeted AI implementation. Our approach to pilots is deliberately focused and time-bound. We’ve found that eight weeks provides enough time to demonstrate value while maintaining momentum. Student experience represents our most ambitious initiative. We’re developing an AI system to guide

MBA students through their academic journey, from course selection to career preparation. For example, a student aspiring to become a product manager at LinkedIn receives AI-generated recommendations for core classes, electives, and club participation based on historical success patterns.

Teaching assistance represents another significant opportunity. While traditional ROI metrics might be challenging in an educational setting, we can still measure concrete outcomes.

AI can grade assignments according to rubrics more quickly than teaching assistants while providing more detailed feedback. If you give Gen AI the rubric, it can grade a lot faster than a human and provide way more qualitative feedback than a human who has just had to do this over 60 times repetitively. Today, I think of AI as a form of automation, and the technology shows real promise.

“The holy grail is personalization—using AI and machine learning to predict what students need before they know they need it.”

Key Takeaways

- Focus AI implementation on core educational missions: student success and thought leadership
- Start with measurable pilots before expanding to broader applications
- Use cross-functional collaboration to drive innovation
- Build on successful pilots to expand across the broader institution

Revolutionizing Language Analysis with Generative AI



Mike Brown, Chief Data Officer and Co-Founder, Strider

Provides intelligence software to help organizations protect intellectual property and talent from nation-state threats.

In the complex world of multilingual security analysis, AI is transforming how we process and understand vast amounts of data. At Strider, implementing LLMs has dramatically improved our ability to analyze billions of documents across multiple languages like Russian and Chinese.

We've nearly doubled individual analyst output by leveraging LLMs for translation, issue identification, and data review. This efficiency gain isn't just about speed—it's about enabling our analysts to focus on higher-level tasks requiring human expertise and judgment. They can concentrate on complex analysis and decision-making rather than spending hours on basic translation and initial review.

Our success stems from a pragmatic integration approach. Instead of completely rebuilding our

systems, we've added AI as a complementary layer to enhance existing capabilities. This strategy has proven particularly valuable in transforming unstructured data into structured formats, making it easier to visualize and analyze complex information. For instance, our systems

can now efficiently process and analyze content in right-to-left languages like Farsi, demonstrating AI's adaptability to complex linguistic challenges.

With this transformation, the role of language specialists has evolved. Rather than focusing solely on translation,

they now collaborate closely with engineers to optimize AI systems and ensure accurate, culturally nuanced analysis. This shift represents a broader trend in how AI is changing specialized roles, enhancing rather than replacing human expertise.

“Success comes from adding AI as a complementary layer, enhancing existing capabilities rather than replacing them.”

Key Takeaways

- Focus on augmenting existing systems rather than complete rebuilds
- Measure success through concrete productivity metrics (2x analyst output)
- Evolve specialist roles to combine domain expertise with AI capabilities
- Build systems that can handle complex linguistic challenges

From Learning to Earning: Scaling Gen AI for Marketing ROI

HANES
Brands Inc

Leo Griffin, Former Vice President, Global Head of Consumer Technology

Marketing stands at the forefront of AI transformation, balancing learning and execution. Drawing from venture capital experience, we see a clear parallel: just as startups must focus on learning before earning, enterprises must understand AI’s capabilities before focusing on ROI. However, we’re approaching a turning point where organizations must move from experimentation to scaling Gen AI with precise, measurable results.

“**Gen AI will drive the ability to create incredible, micro-segmented content at extremely low prices—transforming how brands engage with consumers.**”

The marketing function is particularly ripe for AI transformation, especially in content creation and imagery. For consumer brands and creative agencies, AI promises to transform how they develop and deliver content—imagine creating thousands of perfectly tailored product descriptions or visuals for different customer segments at a fraction of current costs. This shift has begun: While traditional television celebrities once dominated brand messaging, social media influencers now command audiences in the millions, often exceeding traditional media reach. The next frontier is even more transformative: AI-powered virtual influencers who can engage with audiences 24/7, speak multiple languages fluently, and maintain perfect brand consistency. This isn’t just automation—it’s a fundamental reimagining of how brands connect with consumers. While human voices won’t stop mattering, the availability and personalization of generic voices will make them highly compelling. As AI-generated content becomes more sophisticated,

it will create new opportunities for brands to scale their reach while maintaining authenticity.

The transformation extends beyond content creation. We’re witnessing the potential emergence of AI-mediated consumer interactions, where consumers

might deploy their own AI agents to interact with company AI agents. This shift could revolutionize how brands engage with customers, particularly as email becomes increasingly saturated and traditional digital marketing channels evolve.

However, technical debt and organizational resistance present significant hurdles. Many organizations still struggle with basic image management and formatting across platforms, while creative teams may resist AI adoption out of job security concerns. Success requires addressing these foundational issues while building the technical and organizational capabilities for AI implementation.

Key Takeaways

- **Balance learning and earning phases in AI implementation**
- **Focus on foundational issues before scaling AI initiatives**
- **Prepare for AI-mediated consumer interactions**
- **Address organizational resistance through clear value demonstration**



Building the Future of AI? Let's Go.

Mayfield is an early-stage venture firm obsessed with backing ambitious founders at the ideation stage. Since 1969, we've partnered with over 500 companies, including 120 IPOs and over 225 acquisitions. Today, with \$3 billion under management, we're deeply focused on founders building the next wave of Cognition-as-a-Service companies that are fundamentally changing how enterprises harness AI. Learn more at mayfield.ai and follow us on [LinkedIn](#).



Mayfield IT Leadership Network Benefits

Join our community of technology leaders and gain access to:

PEER LEARNING OPPORTUNITIES

Quarterly CIO Insight Calls: Deep dives into pressing topics with industry experts

Regional CIO Roundtables: Semi-annual regional gatherings for candid discussions with CIOs and other IT executives

Emerging Market Trends Briefings: Personalized 1:1 sessions on key areas of focus, with the goal of sharing new ideas we've seen that align to your specific use cases

2024 THEMES

Practical AI engagement and deployment insights from **BCG**

LLM application observability and evaluation with **LangChain**

UPCOMING IN 2025

Erik Brynjolfsson & Andy McAfee - Workhelix and Rationalizing the AI Agenda. [RSVP here](#)

Mayfield
People First.