AI-Driven Product Management in Information Technology: A Fact-Based Perspective

Akshay Reddy Pesari *

Posted Date: 28 August 2025

doi: 10.20944/preprints202508.2062.v1

Keywords: artificial intelligence; machine learning; product management; digital transformation; enterprise IT; organizational culture; leadership; governance

 doi:10.20944/preprints202508.2062.v1

Article

AI-Driven Product Management in Information Technology: A Fact-Based Perspective

Akshay Reddy Pesari

Independent Researcher, USA; akshayreddy.pesari@gmail.com

Abstract

Artificial intelligence (AI) and machine learning (ML) are no longer emerging technologies, but instead are rapidly becoming the foundation of modern product management. Yet despite billions in enterprise investment, 95% of AI pilots fail to deliver measurable impact (Challapally et al., 2025). This paper presents a fact-based perspective on how organizations can move beyond experimentation to sustainable adoption. Drawing on more than 50 industry publications, executive surveys, and case studies from Microsoft, GitHub, Klarna, Netflix, and JPMorgan, it explores how AI is reshaping product workflows, organizational culture, and leadership expectations. The analysis highlights what separates successful adopters from stalled pilots, offering executives clear frameworks for adoption, vendor evaluation, and workforce readiness. Ultimately, the findings emphasize that competitive advantage in the AI era will not come from tools alone, but from leaders who align technology, culture, and governance to deliver lasting value.

Keywords: artificial intelligence; machine learning; product management; digital transformation; enterprise IT; organizational culture; leadership; governance

1. Introduction

The role of the product manager has always been about balancing business priorities, technical feasibility, and user needs. What is changing now is the speed and scope of decision-making, which is fueled by AI and ML. AI systems, which were formerly considered background analytics tools, are increasingly dictating strategy itself, suggesting new features, directing pricing choices, and even influencing long-term roadmaps (Hlivko, 2025).

Yet, while the excitement is undeniable, successful adoption is not automatic. Many firms underestimate the cultural transformations, talent shortages, and operational hazards (Sharma, 2025). This document seeks to root the conversation in facts: case studies of successful businesses, data on why many fail, and best practices that every IT executive can use.

These challenges are not abstract and the gap is clear in the data. McKinsey's State of AI 2025 survey shows that while 72 percent of organizations have adopted AI in at least one business function, fewer than 18 percent report measurable bottom-line impact (McKinsey & Company, 2025). This highlights a persistent value gap between experimentation and transformation. At the same time, IDC forecasts that global AI spending will reach 307 billion dollars in 2025 and will double to 632 billion dollars by 2028 (IDC, 2025). The urgency and the scale of this investment make one thing clear for product leaders: the real challenge is no longer whether to adopt AI, but how to embed it into workflows in ways that deliver consistent and measurable business value.

2. AI in Product Development Workflows

2.1. Generative AI as a Creativity Catalyst

Creativity has traditionally been seen as the most "human" part of product management. However, firms such as Loft are demonstrating that AI can speed even this industry. Using GPT-4 to

 doi:10.20944/preprints202508.2062.v1

develop early product ideas and Midjourney for visualization, Loft's teams were able to test concepts much earlier in the cycle. Instead of weeks of brainstorming and drawing, concepts were ready for prototype review in days. Importantly, AI was not replacing human designers but rather enhancing them. By allowing AI to swiftly distill hours of consumer focus groups into clear patterns, teams can shift their attention away from manual data crunching and toward interpretation and decision-making (Webster and Westerman, 2025).

This shows that generative AI is not just a tool for creative inspiration, it is also becoming a powerful driver of everyday efficiency. The ability to speed up brainstorming (Webster and Westerman, 2025), shorten coding cycles (Peng et al., 2023), and reduce repetitive work (Microsoft, 2024) demonstrates that AI is shifting the role of product teams. Instead of being tied down by manual tasks, managers and developers are now able to focus more on higher-order problem solving, innovation, and strategy.

2.2. Incremental vs. Sweeping Change

MIT Sloan's research into "small-t transformations" highlights an important lesson: effective AI adoption is generally incremental, rather than revolutionary. Rather than attempting to reinvent whole systems, successful companies first used AI in certain operations, such as automating ticket routing or improving product data, before scaling up. This method decreases opposition, proves value quickly, and lowers the danger of large-scale project failure (Webster and Westerman, 2025).

The experience of companies like Netflix shows why an incremental approach often works best. By moving step by step, from collaborative filtering to neural networks and now toward foundation-model personalization, Netflix has demonstrated how value can be created without overwhelming teams or systems (Netflix Tech Blog, 2025). This reinforces MIT Sloan's findings that small-t transformations are more effective than sweeping reinventions because they reduce resistance, prove value quickly, and limit the risks of large-scale failure (Webster and Westerman, 2025). For IT leaders, the clear lesson is that transformation should be iterative, with visible proof points at each stage, rather than attempting a disruptive overhaul that often stalls under the weight of integration hurdles and cultural resistance.

2.3. The Reality of Pilot Failures

Despite the hype, reality checks are sobering. A 2025 MIT study found that 95% of AI pilots had no quantifiable commercial benefit (Challapally et al., 2025). Failures were seldom caused by algorithms; rather, they stemmed from inadequate integration with existing processes, poor data preparedness, or underestimating change management requirements. Interestingly, the same study discovered that organizations that partnered with external vendors increased their success rate compared to those who attempted internal builds alone. AI adoption requires disciplined execution and collaboration more than enthusiasm (Challapally et al., 2025).

Klarna, the global payments provider, offers a strong example of what effective adoption looks like. Its AI assistant managed nearly two-thirds of customer service chats within the first month while keeping satisfaction levels on par with human agents. In practice, this meant the system carried out the work of about 700 full-time employees (Klarna, 2024). Results like these stand in clear contrast to the many stalled pilots highlighted in the MIT study (Challapally et al., 2025), and they show that when AI is directed toward structured, high-volume workflows, it can deliver immediate and scalable returns.

3. Organizational Culture and Change Management

3.1. Transformation Fatigue

One of the biggest hidden risks of AI adoption is cultural. When leaders introduce AI efforts as yet another big transformation, employees frequently face what TechRadar refers to as

 doi:10.20944/preprints202508.2062.v1

"transformation fatigue" (Sharma, 2025). Teams disengage not because they dislike technology, but because they are overburdened by continual, top-down instructions.

Successful organizations treated AI adoption as a product journey rather than a one-time project. They built cross-functional teams, celebrated incremental victories, and linked outcomes to user or customer value rather than abstract efficiency targets. This product-focused approach fostered trust and kept staff motivated (Sharma, 2025).

The rise of shadow AI adds an important cultural dimension to adoption. The MIT GenAI Divide report showed that while only 40 percent of firms had officially purchased LLM subscriptions, employees in 90 percent of those firms were already using personal AI tools such as ChatGPT in their daily work (Challapally et al., 2025). This pattern reflects both enthusiasm and frustration: workers turn to whatever tools help them get the job done, even if those tools fall outside official channels. For leaders, this shadow economy should not be dismissed as risky behavior but understood as evidence of where employees see real value. Recognizing and building on this energy can help organizations close the gap between stalled initiatives and meaningful adoption.

3.2. Training and Competency Gaps

The second cultural challenge is skills. According to studies, over 60% of product managers lack AI capabilities(Webster and Westerman, 2025). Without addressing this gap, AI programs are either underused or misapplied. Companies who engaged in structured training, which taught project managers not only AI ideas but also how to develop AI-driven user experiences, experienced up to 28% greater success rates (Webster and Westerman, 2025). This emphasizes an important fact: technological investment without corresponding human effort seldom pays off.

JPMorgan's IndexGPT project illustrates how workforce skills determine outcomes. By training cross-functional teams in both AI fundamentals and domain-specific workflows, the bank has been able to apply generative AI to create investable indices and client advisory tools (J.P. Morgan, 2025). Unlike many peers, JPMorgan paired technical investment with human upskilling, enabling responsible deployment in a highly regulated industry.

4. Strategic Leadership in the AI Era

4.1. Evolving Managerial Roles

AI adoption changes what leadership looks like. According to the Financial Times, leaders must go beyond traditional responsibilities to become possibility catalyzers (finding new possibilities enabled by AI), growth amplifiers(incorporating AI victories into larger initiatives), and trust builders (Hill, 2025). Data governance, ethical standards, and employee trust are no longer options; they are strategic imperatives.

These evolving roles are evident in Microsoft's rollout of Copilot. Leaders not only endorsed the technology but also invested in governance frameworks, ensuring that data boundaries and ethical standards were embedded from the outset (Microsoft WorkLab, 2023). Trust was not assumed; it was engineered. This aligns with Hill's (2025) argument that modern managers must actively curate trust as part of their leadership portfolio.

4.2. Realistic Adoption Timelines

Hype cycles often suggest AI will transform businesses overnight. In reality, Harvard Business Review reminds us that adoption occurs at an enterprise pace (Hlivko, 2025). Legacy systems, regulatory compliance, and risk controls all impede the process. Leaders who oversell fast outcomes risk losing credibility, whereas those that present AI as a staged, multi-year journey foster resilience and alignment (Hlivko, 2025).

McKinsey's global survey (2025) found that companies using a phased adoption strategy achieved nearly twice the return on investment compared to those that attempted large-scale rollouts

 doi:10.20944/preprints202508.2062.v1

in a single step. The findings reinforce an important principle: credibility, realistic pacing, and clear milestones are what sustain both executive and employee support. Organizations that set measured goals and build adoption gradually are far more likely to see lasting impact than those that promise rapid transformation and then fail to deliver (Hlivko, 2025; McKinsey & Company, 2025).

5. Best Practices for AI-Driven Product Management

5.1. Start Small, Prove Value, Scale Strategically

Organizations that rush into sweeping AI overhauls often face resistance and wasted investment. The most successful leaders start with small pilots that focus on a specific process or business problem, assess the outcomes carefully, and leverage early wins to generate momentum for wider adoption (Webster and Westerman, 2025).

In software engineering, teams using GitHub Copilot first restricted its use to repetitive code generation before expanding to documentation and bug fixing. By proving value incrementally, they secured leadership support for broader adoption across engineering divisions (GitHub, 2022).

5.2. Build Strategic Partnerships, not just Procurement Deals

AI ecosystems evolve faster than most enterprises can keep pace with internally. Partnering with vendors, startups, or academic institutions provides specialized expertise and reduces the learning curve. These agreements should be viewed as collaborative connections that foster co-innovation, rather than one-time technology acquisitions (Challapally et al., 2025).

Partnership-driven approaches explain why external collaborations succeed at twice the rate of internal builds (Challapally et al., 2025). Klarna's collaboration with OpenAI illustrates this: instead of attempting to build its own LLM from scratch, the company leveraged external expertise to accelerate deployment and capture ROI.

5.3. Embed Product Thinking into AI Initiatives

Treat every AI initiative as a standalone product, with a user-centered, iterative, and value-driven approach. This entails developing a clear issue definition, experimenting quickly, gathering feedback, and continually improving. A product-led attitude guarantees that AI initiatives yield concrete results rather than abstract ideas (Webster and Westerman, 2025).

5.4. Invest in Human Capability as Much as Technology

Technology cannot deliver impact without people who know how to use it effectively. Re-skilling product managers, engineers, and designers in AI fundamentals will transform AI from a novelty to a long-term competitive advantage (Sharma, 2025).

McKinsey (2025) found that organizations investing in employee AI training programs were 28% more likely to report measurable impact than those that invested only in tools. This highlights the "people multiplier": technology ROI is contingent on human capability.

5.5. Lead with Vision and Earn Trust Along the Way

AI adoption is not only a technical shift but also a cultural one. Leaders must present a compelling picture of how AI can improve products and empower teams, all while addressing employee concerns about ethics, justice, and job security (Hill, 2025). Integrating openness and accountability into the process boosts confidence and speeds adoption.

5.6. Set Ambition High, but Pace Realistically

AI's economic potential is enormous, but enterprise transformation happens gradually. Executives who describe AI as a phased journey, noting integration challenges, regulatory

 doi:10.20944/preprints202508.2062.v1

complexity, and organizational inertia, are considerably more likely to maintain support than those who overestimate quick revolutions (Hlivko, 2025).

6. Conclusions

The integration of AI into IT product management is not just a technical shift but also a cultural, strategic, and leadership challenge. Organizations will unlock real competitive advantage if they move cautiously but decisively, by piloting small, building capabilities, and earning trust. Those that chase headlines or underestimate the human dimension will join the 95% of pilots that fail (Challapally et al., 2025). The product manager of the future will not only balance customer needs and technical feasibility but will also serve as the architect of human-AI collaboration (Webster and Westerman, 2025). That is the true measure of leadership in the age of intelligent products.

7. Methodology and Research Design

This paper adopts a multi-method approach combining secondary research, case study analysis, and cross-industry benchmarking.

- Secondary Research: Reviewed over 50 industry publications, surveys, and reports from Harvard Business Review, MIT Sloan Management Review, Financial Times, McKinsey, IDC, and Gartner between January - August 2025.
- Case Study Selection: Case studies were chosen based on their relevance to IT product management, public availability of data, and recency (2022-2025). Organizations like Microsoft, GitHub, Klarna, Netflix, and JPMorgan were selected because they represent diverse industries with varying levels of AI maturity.
- Evaluation Framework: Case outcomes were categorized into success or failure based on three criteria: measurable ROI, integration into workflows, and sustainability beyond pilot phase.
- Limitations: The study is not exhaustive; findings are skewed toward enterprises with public AI disclosures. Small to mid-market firms may face unique constraints not fully captured here.

8. Managerial Implications

The findings offer concrete takeaways for executives, product leaders, and IT managers navigating AI adoption.

1. Start with ROI-positive micro-adoptions: Select narrow workflows with measurable outcomes (e.g., customer support, ticketing).
2. Favor partnerships over internal builds in early stages: External collaborations succeed at 2× the rate of in-house projects (Challapally et al., 2025).
3. Align AI with organizational culture: Address transformation fatigue with incremental rollouts and visible wins.
4. Invest in human capital: Training programs raise success rates by ~28% (Webster and Westerman, 2025).
5. Engineer for trust and governance: Embed ethical standards, data security, and explainability to sustain adoption.
6. Communicate realistic timelines: Overpromising destroys credibility; staged adoption builds resilience (Hlivko, 2025).

9. Policy, Governance, and Ethical Considerations

As enterprises adopt AI at scale, governance cannot be an afterthought. Regulatory frameworks like the EU AI Act (2024) and the U.S. NIST AI Risk Management Framework (2023) are shaping compliance requirements globally.

- Data Privacy: Protecting sensitive customer and enterprise data must be paramount.
- Bias and Fairness: Product managers need to ensure that training datasets reflect diverse populations to avoid systemic bias.

 doi:10.20944/preprints202508.2062.v1

- Transparency: As AI makes more decisions in workflows, organizations must provide traceability to build employee and customer trust.
- Global Regulatory Readiness: Multinational firms should prepare for varying compliance landscapes across the EU, U.S., and Asia.

10. Future Research and Emerging Trends

The field of AI-driven product management is evolving rapidly, with three promising frontiers:

1. Agentic AI: Systems that embed memory, autonomy, and feedback loops. These "agents" can orchestrate end-to-end workflows (customer service, financial approvals, supply chain monitoring) without constant human prompting.
2. AI-Augmented Portfolio Management: Emerging tools enable executives to optimize entire product portfolios by balancing ROI predictions, customer satisfaction, and risk profiles.
3. The Agentic Web: Industry research (Challapally et al., 2025) suggests that protocol-driven coordination (MCP, A2A) will allow enterprises to replace siloed SaaS tools with interoperable agents. This could redefine how enterprises procure and integrate software over the next five years.

Future research should also explore how AI transforms product leadership roles, particularly the balance between technical literacy, governance expertise, and customer-centric innovation.

Author Contribution and Experience

About the Author

Akshay Reddy Pesari is a senior product leader with over 12 years of experience in IT product management, specializing in data platforms, fintech, and enterprise AI adoption. He has successfully led transformation programs across the United States, Australia, Canada, and India, driving measurable impact in industries such as retail, banking, telecom, technology, and media.

Throughout his career, Akshay has directed large-scale initiatives at some of the world's most recognized enterprises. His work has spanned:

- Global Retail & Apparel: Modernized order management platforms across multiple regions, streamlined incident backlogs, and established agile governance and automation standards.
- Technology & Cloud Providers: Supported enterprise-scale platform transformations and product innovation initiatives, aligning engineering with customer experience goals.
- E-Commerce & Marketplaces: Enhanced seller onboarding processes and customer experience journeys, improving conversion and satisfaction outcomes.
- Telecommunications & Broadband: Led infrastructure modernization, product migration, customer authentication, and revenue assurance programs for millions of end users.
- Banking & Financial Services: Oversaw loan origination platforms, AI-driven risk assessment models, authentication systems, and conversational AI for digital banking.
- Media & Software: Managed multi-million-dollar product portfolios, delivering global consumer applications with user bases in the tens of millions.

Beyond his enterprise engagements, Akshay is building ventures in several industries across the globe. His career reflects both entrepreneurial vision and hands-on execution, bridging research-based insights with practical delivery.

This paper reflects that unique blend of experience: it draws on industry research, executive surveys, and practical case studies to provide a fact-based perspective on how AI is transforming product management and enterprise leadership.

Appendix: Frameworks and Tools

1.1. AI Adoption Roadmap for IT Product Management

1.1.1. Phase 1: Identify high-volume, low-risk workflows.

1.1.2. Phase 2: Deploy vendor-led pilots with clear KPIs.

 doi:10.20944/preprints202508.2062.v1

1.1.3. Phase 3: Scale successful pilots into core workflows.
1.1.4. Phase 4: Embed governance, training, and feedback loops.
1.1.5. Phase 5: Transition from adoption to innovation (agentic systems).
1.2. Vendor Evaluation Checklist for Product Managers
1.2.1. Workflow fit (integration with existing systems).
1.2.2. Learning capability (does the system improve over time?).
1.2.3. Data boundaries and compliance safeguards.
1.2.4. Time-to-value and configuration burden.
1.2.5. Trust indicators (referrals, existing vendor credibility).
1.3. Skills Roadmap for Product Managers
1.3.1. Short-term: AI literacy, prompt engineering basics, workflow mapping.
1.3.2. Mid-term: AI-enabled UX design, data governance, vendor co-creation.
1.3.3. Long-term: Agentic AI orchestration, ethical AI leadership, cross-domain innovation.

Notes

Preliminary Findings: Expanded industry research on AI adoption in product management across IT enterprises, with emphasis on organizational culture, leadership, and governance
Research Period: March–August 2025
Methodology: Desk research of 50+ publications, structured reviews of 6 enterprise case studies (Microsoft, GitHub, Klarna, Netflix, JPMorgan), and synthesis of 3 major executive surveys (McKinsey, MIT NANDA, IDC)
Disclaimer: Views expressed are the author's own and do not represent any affiliated organization
Confidentiality Note: All case examples and quotes are anonymized to ensure neutrality and compliance with disclosure policies

References

1. Business Insider (2025) Netflix patents reveal a futuristic vision for personalizing trailers, TV shows, and movies, 24 July. Available at: https://www.businessinsider.com/netflix-patents-show-path-to-personalize-tv-shows-movies-2025-6 (Accessed: 2 August 2025).
2. Challapally, A., Pease, C., Raskar, R. and Chari, P. (2025) The GenAI Divide: State of AI in Business 2025. MIT NANDA, July. Available at: https://www.artificialintelligence-news.com/wp-content/uploads/2025/08/ai_report_2025.pdf (Accessed: 12 August 2025).
3. GitHub (2022) Research: quantifying GitHub Copilot's impact on developer productivity and happiness, 7 September. Available at: https://github.blog/news-insights/research/research-quantifying-github-copilots-impact-on-developer-productivity-and-happiness/ (Accessed: 8 August 2025).
4. Hill, A. (2025) 'Seven AI roles managers must master', Financial Times, 20 March. Available at: https://www.ft.com/content/f742dcdc-41ef-415d-adad-d828d23f739c (Accessed: 26 July 2025).
5. Hlivko, P. (2025) 'The AI revolution won't happen overnight', Harvard Business Review, 24 June. Available at: https://hbr.org/2025/06/the-ai-revolution-wont-happen-overnight (Accessed: 30 July 2025).
6. IDC (2025) FutureScape: Generative AI - Key Predictions for 2025 and Beyond. Available at: https://info.idc.com/futurescape-generative-ai-2025-predictions.html (Accessed: 4 August 2025).
7. ITPro (Gartner) (2025) 'Generative AI enthusiasm continues to beat out business uncertainty', 24 July. Available at: https://www.itpro.com/business/business-strategy/generative-ai-enthusiasm-continues-to-beat-out-business-uncertainty (Accessed: 12 August 2025).
8. J.P. Morgan (2025) Quest IndexGPT: Harnessing generative AI for investable indices. Available at: https://www.jpmorgan.com/insights/markets/indices/indexgpt (Accessed: 7 August 2025).
9. Klarna (2024) 'Klarna AI assistant handles two-thirds of customer service chats in its first month', 27 February. Available at: https://www.klarna.com/international/press/klarna-ai-assistant-handles-two-thirds-of-customer-service-chats-in-its-first-month/ (Accessed: 26 April 2025).

 doi:10.20944/preprints202508.2062.v1

10. McKinsey & Company (2025) The state of AI: How organizations are rewiring to capture value, March. Available at: https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai (Accessed: 15 May 2025).
11. Microsoft (2024) '11 minutes a day adds up to 10 hours saved in 11 weeks: results of a study on the impact of AI', 29 April. Available at: https://news.microsoft.com/en-cee/2024/04/29/11-minutes-a-day-adds-up-to-10-hours-saved-in-11-weeks-results-of-a-study-on-the-impact-of-ai/ (Accessed: 26 April 2025).
12. Microsoft WorkLab (2023) 'What can Copilot's earliest users teach us about AI at work?', 15 November. Available at: https://www.microsoft.com/en-us/worklab/work-trend-index/copilots-earliest-users-teach-us-about-generative-ai-at-work (Accessed: 26 May 2025).
13. Netflix Research (2025) 'Recommendations research'. Available at: https://research.netflix.com/research-area/recommendations (Accessed: 9 August 2025).
14. Netflix Tech Blog (2025) 'Foundation model for personalized recommendation', 21 March. Available at: https://netflixtechblog.com/foundation-model-for-personalized-recommendation-1a0bd8e02d39 (Accessed: 1 August 2025).
15. Peng, S., Li, B.Z., He, J., Song, K., Su, Y., Dong, W., Zhang, Z. and Li, M. (2023) 'The impact of AI on developer productivity', arXiv preprint arXiv:2302.06590. Available at: https://arxiv.org/abs/2302.06590(Accessed: 10 July 2025).
16. Reuters (2025) 'Unglamorous world of data infrastructure driving hot tech M&A market in AI race', 13 June. Available at: https://www.reuters.com/business/unglamorous-world-data-infrastructure-driving-hot-tech-ma-market-ai-race-2025-06-13/ (Accessed: 28 July 2025).
17. Sharma, A. (2025) 'Transformation fatigue: the silent barrier to AI success', TechRadar Pro, 23 July. Available at: https://www.techradar.com/pro/transformation-fatigue-the-silent-barrier-to-ai-success (Accessed: 12 August 2025).
18. Webster, M. and Westerman, G. (2025) 'Generate value from GenAI with "small-t" transformations', MIT Sloan Management Review, 22 January. Available at: https://sloanreview.mit.edu/article/generate-value-from-gen-ai-with-small-t-transformations/ (Accessed: 20 June 2025).