



# TRUST, TRANSPARENCY, AND THE COMMERCIAL AI PC

WHITEPAPER | 2026



Ever since AI burst on to the scene in late 2022, tech companies and software vendors have raced to explore how AI-enabled hardware and software could change and challenge the conventional boundaries of computing. Initial AI discussions and articles focused on what it was, how it worked, and its potential for various commercial, consumer, and scientific applications.

As more people and organizations have familiarized themselves with AI, conversations have evolved from more abstract explanations of the technology and its long-term potential to more concrete conversations about scaling, trust, governance, long-term business impacts, and ROI. This, in turn, has put pressure on both customers and vendors to move from pathfinding projects to real-world results.

## THE ADOPTION GAP: WHY EARLY AI ROI IS UNEVEN

Recent executive surveys on AI implementation outcomes, revenue impact, and costs have been mixed. While the exact figures differ depending on the particulars of a given report's methodology, relatively few companies report positive, enterprise-wide revenue growth and/or declining costs. Pilot projects remain difficult to scale into full production.

One observation that connects many of these analyses is that the problems that derail an AI project or limit its real-world value are often unrelated to the accuracy or quality of the underlying AI model. Non-AI technology issues explain some failures to launch, but many artificial intelligence pilots fall short for organizational, structural, or governance reasons.

The companies reporting the greatest success with AI have been generally those that deployed it with clearly defined use cases and well-articulated goals. Business leaders hoping to unlock new efficiencies with AI need to understand where production bottlenecks are and a concrete plan for how AI will address them. Skipping this analysis may result in minimal improvement if downstream choke points aren't alleviated simultaneously. The state of a company's current and archival data – how it's labeled, formatted, and accessed within the company – can also limit the success of an AI initiative, as can the age of its infrastructure or PC fleet.

Successful AI rollouts, however, aren't just a matter of technology, data hygiene, or a good vendor relationship. Trust is a common thread that connects these discussions together, both overtly and implicitly.

## THE IMPORTANCE OF TRUST

Trust as it relates to AI arises in many different contexts, including user trust in AI tools, between vendors and suppliers, employees and employers, and in company leadership. As work evolves, employees need to understand how they are being judged and under what criteria. Clear rules, transparent processes, consistent metrics and open feedback channels all improve outcomes, as can aligning employee recognition, rewards, and incentives with well-articulated AI adoption goals.

The importance of trust as a foundational element of successful AI adoption hasn't just surfaced in retrospective AI rollout surveys. The Gartner report "Predicts 2026: AI's Impact on the Future of Workforce" contains a series of key findings and strategic planning assumptions intended to help CIO's "create a symbiotic relationship between humans and machines at the workplace."<sup>1</sup>

The report characterizes the narrative of an "AI-driven job apocalypse" as "misplaced," noting that "while AI is profoundly reshaping work, it is not eliminating it." The report's key findings predict part of what that reshaping might look like, including the idea that AI will shorten the half-life of work skills from 2-5 years down from 8-12 years, and employers expect 39% of worker's core skills to change by 2030 due to "the automation of routine work, the erosion of legacy competencies, and the emergence of AI-orchestrated workflows, where machines increasingly execute tasks under human oversight." This will result in adaptability and learning velocity "being the primary metric for hiring."

The pressure on CIOs to appropriately balance the need for AI ROI with appropriate headcount adjustments will be high, autonomous AI agents will "redefine how users and systems interact" with 60% of digital products architected primarily for digital consumption. This means organizations will be navigating questions around the use of internally created digital avatars and how AI agents are represented on org charts. Trust is a factor in many of these discussions, sometimes unequivocally so.

Gartner identifies trust, security, and compliance as "decisive differentiators" for agent-first products and platforms, for example, and predicts "Buyers will prioritize vendors that provide zero-trust validation, auditability, and intent-verification guardrails for agent-driven actions. Markets will reward transparency and governance maturity over visual polish or legacy UX strengths." The predicted reduction in work skill half-life is explained partly by the increased use of machines to perform tasks under human oversight, and while Gartner does not explicitly point to this scenario as an example of trust, a certain amount between user and agent is arguably implicitly required.



## TRUST, AI, AND THE MODERN ENDPOINT

This latest IT procurement and PC refresh cycle kicked off as AI PCs – PC’s with a dedicated neural processing unit for handling AI workloads locally – arrived in-market. Whether AI runs in the cloud, on-device, or both, the user’s experience of using it is always mediated through an endpoint device. Factors like responsiveness, performance, ease-of-use, flexibility, and software support define the boundaries of which AI-enabled products and services an end-user can access. This is true for all software – the experience of using an application cannot be entirely separated from the device it runs on – but the point has particular salience for AI given its rapid evolution.

The specialized NPU inside AI PCs can run AI workloads more efficiently than the CPU or GPU, saving battery power compared to using either of these processors and reserving them for other tasks. While many AI workloads run in the cloud today, hybrid AI, which uses both local and cloud rendering resources, is expected to become more common as AI PC adoption and ISV software support both rise. That’s the technical aspect of adoption and why the value proposition of AI PCs increases as software support improves.

Perceptually, adopting AI PCs as part of a fleet refresh cycle signals the importance of AI as a strategic goal while affirming the importance of the employee experience. AI PCs, especially Copilot+ PCs, offer the highest level of flexibility and performance for AI workloads by providing access to both local AI tools and cloud-based services.

But AI performance and flexibility are themselves only one aspect of the story. Standardized hardware platforms with robust software support and strong security are essential components of any AI ecosystem. That’s precisely what AMD delivers.

AMD led the x86 market with the introduction of neural processing units in 2023 and has the largest line-up of SKUs supporting Microsoft Copilot+ PCs available today. AMD Ryzen™ AI PRO 300, 400, and Ryzen AI Max PRO Series processors span a wide range of CPU and GPU configurations while standardizing on the AMD XDNA™ 2 NPU architecture. This gives customers the option of choosing processor SKUs based on their conventional CPU and GPU needs, while simultaneously providing software developers with a single NPU to target.

AMD PRO processors build trust by providing IT procurement teams with flexible AI performance in a variety of OEM partner-built form factors and include a number of features intended for the commercial market, including Supply Chain Security, industry-standard DASH profiles for straightforward manageability, 18 months of software stability and hardware availability, an additional five years of software support after the final ship date, and continuous firmware and driver validation for long-term reliability, security, and predictable support. AMD partners with major OEMs deliver a wide range of systems across a wide range of configurations and price points to support the business needs of every employee.

The AI era has jolted the business world and introduced a degree of uncertainty not seen in decades, but the success of many deployments depends significantly on the evolution of existing business practices, norms, and structures, not just the invention of better, faster, or more accurate AI. Human factors, including how effectively companies cultivate trust with their customers, partners, employees, and the general public will play an essential role in the future of AI and how quickly or slowly companies adopt it.

That, quite frankly, ought to be enough uncertainty for anyone. Introducing uncertainty in perceived IT strategy by ignoring the opportunities of AI PCs or only selectively deploying them may simplify AI deployments today but risks confusion and uncertainty in the longer term. Emphasizing broad AI availability across many roles is a simpler way to signal confidence, and AMD provides a flexible, scalable set of solutions that can meet commercial security and manageability requirements while addressing each organization’s unique interests and needs.

### ENDNOTES

1. Gartner, Predicts 2026: AI’s Impact on the Future of Workforce, By Arun Chandrasekaran, Helen Poitevin, Tori Paulman, Brent Stewart, Shawn Murphy, Afraz Jaffri, 4 November 2025

### COPYRIGHT NOTICE

© 2026 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective owners.

*GARTNER is a trademark of Gartner, Inc. and/or its affiliates.*