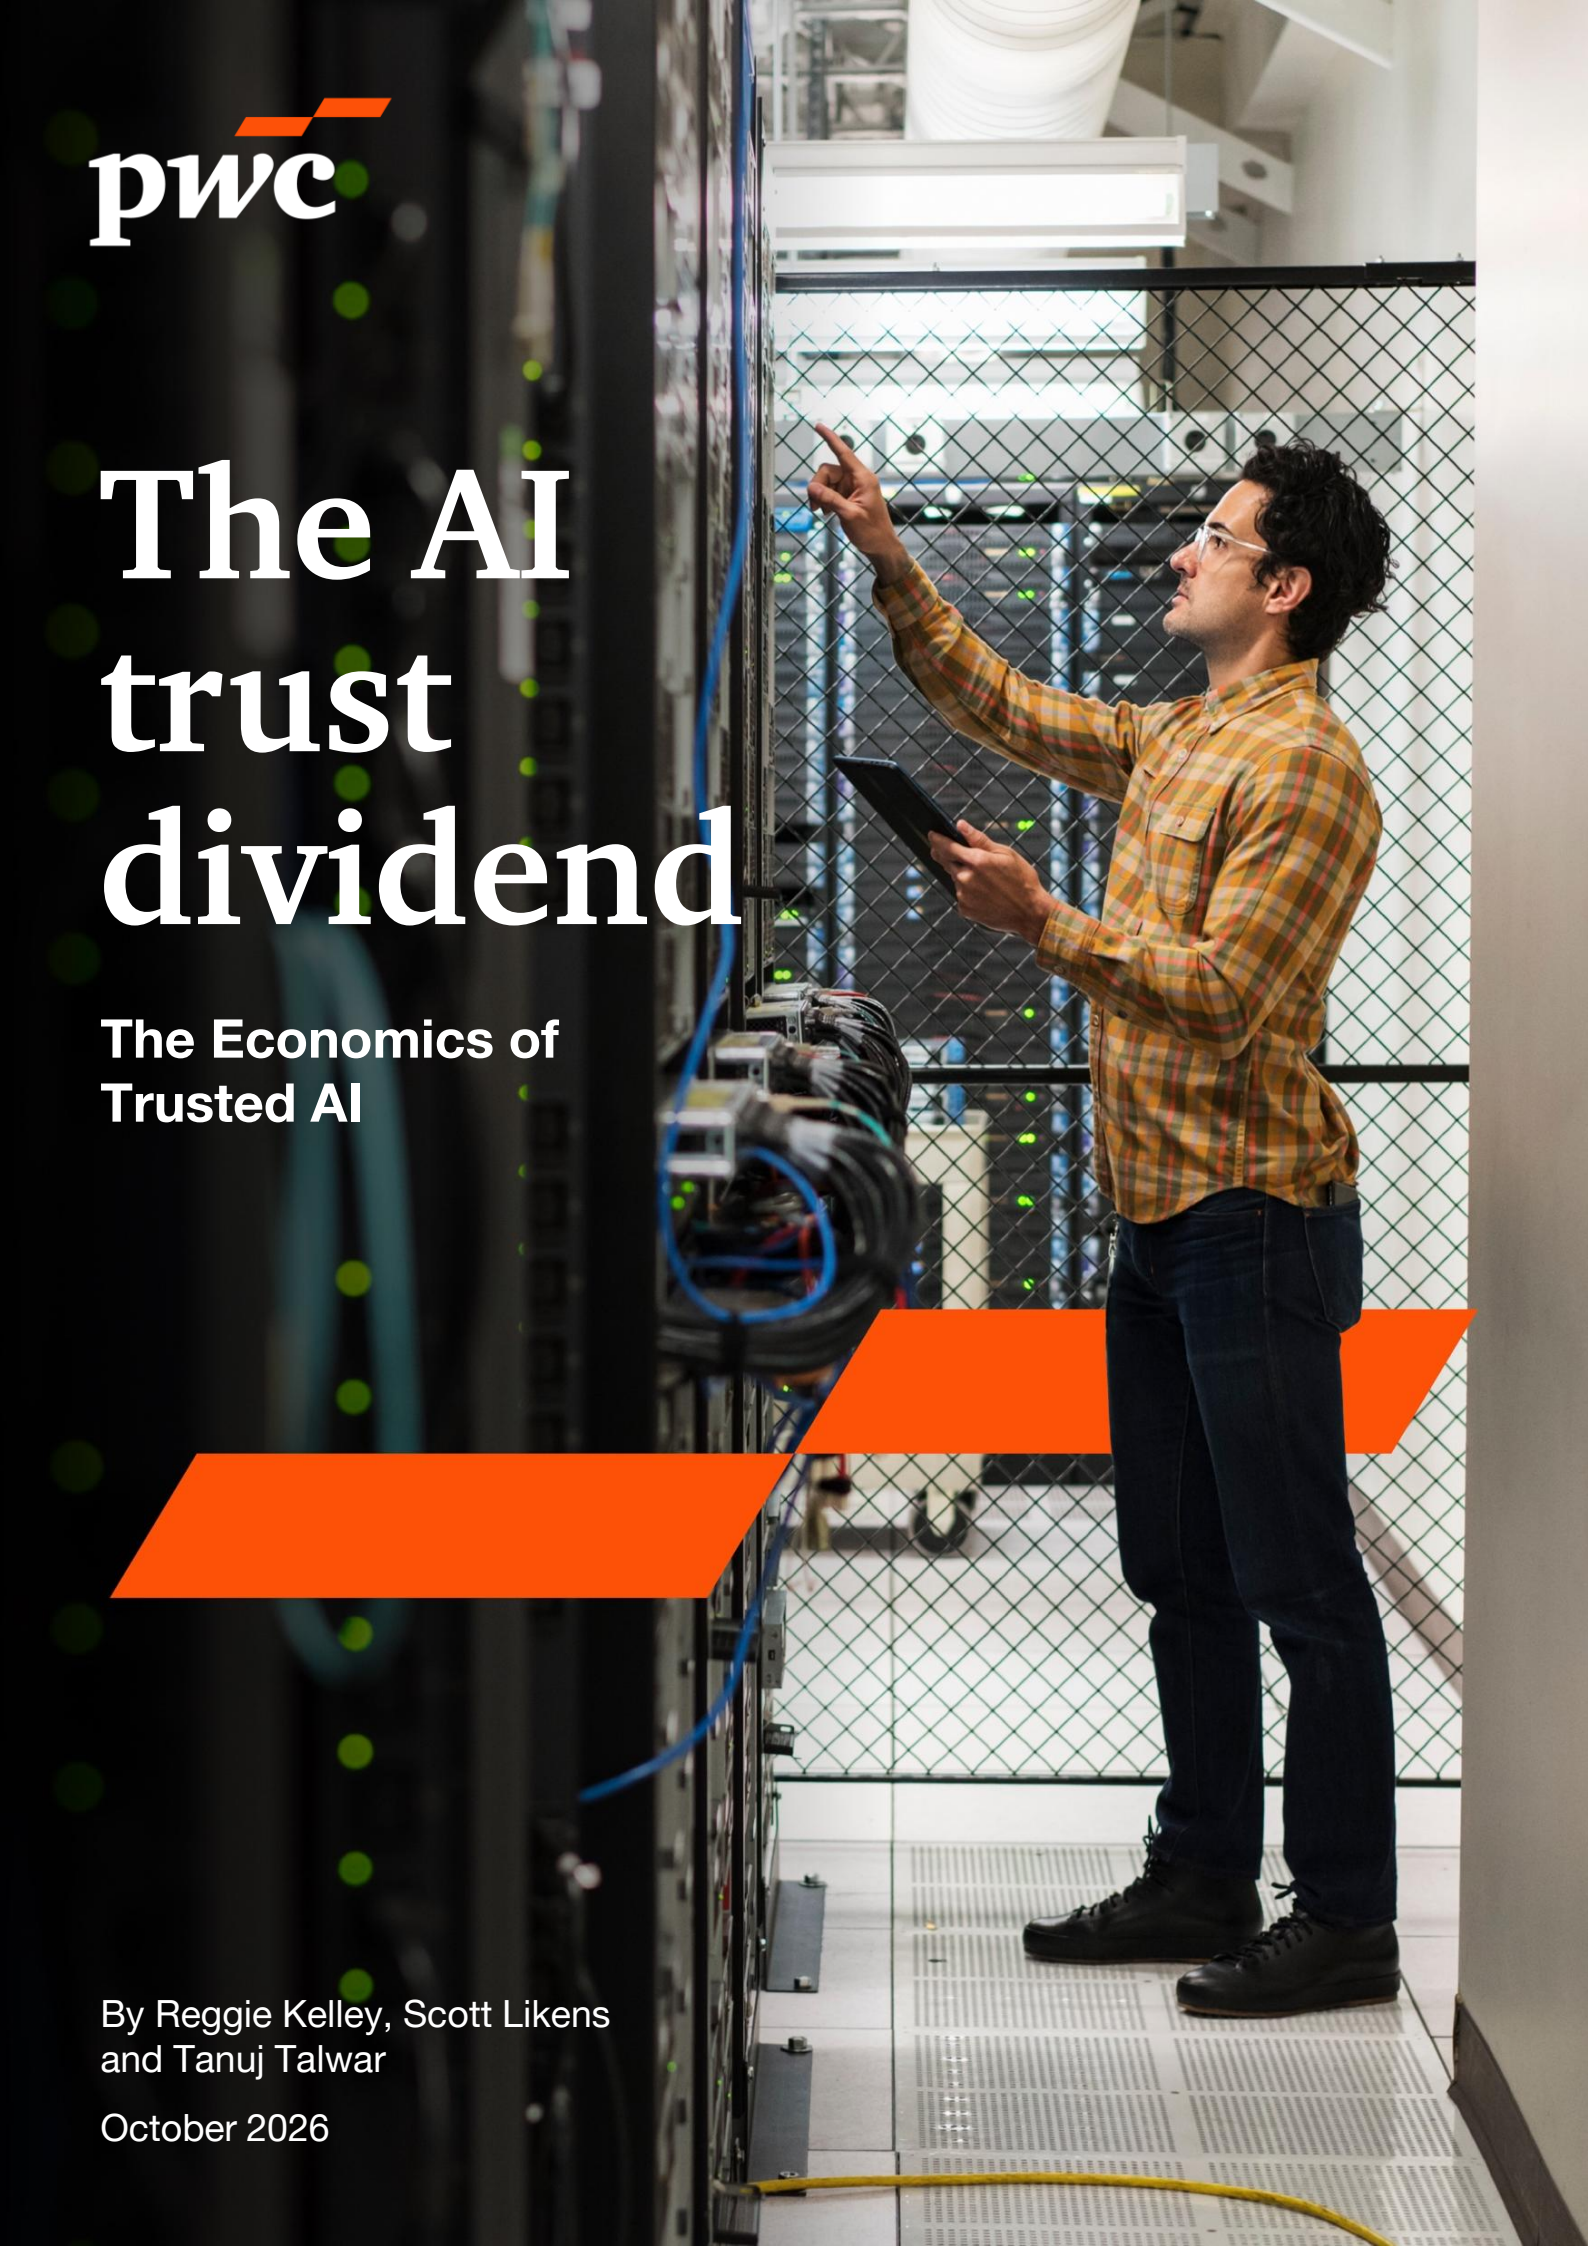
pwc

The AI trust dividend

The Economics of Trusted AI



By Reggie Kelley, Scott Likens and Tanuj Talwar

October 2026

Table of contents

Introduction 03

Three realities of trusted AI economics 04

Building AI economics by design 07

Three questions the Chief Executive Officer and board should ask 08

Authors and contributors 09



Introduction

As AI adoption grows, AI economics cannot be separated from technology and trust risk. Choices about models, infrastructure, performance, and controls shape what AI costs, how reliably it performs, and whether people can trust it. The organisations that lead will be those that govern these dimensions together, continuously optimising the relationship between economic performance, technological resilience, trust, and business value.

What does AI actually cost your organisation and how much value does it create? As AI spending accelerates, it's a question that can be difficult for leaders to answer. Some have been shocked to attention by rising token costs as model providers increasingly shift to usage-based billing structures. Yet token costs are only a fraction of AI cost, and cost is only one component in assessing the full economics of trusted AI. The greater priority is to understand the total cost and value of the business outcomes produced by the AI-powered agents doing the work.

As we explored in our recent article Creating strong foundations, trusted AI, and real returns, governing the complete AI lifecycle is essential to build trust and scale AI. This comes with an economic cost. Organisations that successfully connect cost, trust, and value will be better positioned to scale AI responsibly and realise returns over the long term.



Three realities of trusted AI economics

As business leaders address the costs of scaling AI, there are three realities to consider.

Getting to the heart of Total Cost of Ownership (TCO) and cost governance is challenging

Visible costs (token usage, API calls, licenses, and subscriptions) are the tip of the AI cost iceberg. While important, they comprise a small part of AI's total economic reality. The hidden costs beneath the surface are often more significant and considerably harder to measure.

PwC worked with one global company to identify 22 cost drivers shaping what AI actually costs at scale and built the framework to track them. Several had little to do with the model price itself. These drivers must be factored into any effort to establish a true TCO and establish the right cost governance for AI solutions.

The Often-Hidden Components of AI Cost

- Infrastructure and compute resources
- Agent orchestration
- Integration layers
- Ongoing maintenance and model optimisation
- Monitoring and observability
- Energy and other resources
- Governance and compliance activities
- Human oversight and trust controls

Chief Financial Officers (CFOs) and Chief Technology Officers (CTOs) have seen this TCO challenge before as they grappled to understand cloud economics. Despite years of investment in cloud Financial Operations (FinOps) disciplines, organisations still struggle with cloud cost transparency and optimisation. AI economics are substantially more complex than cloud economics because they are highly dynamic, usage-based, increasingly autonomous, and deeply influenced by architectural decisions. Model pricing structures and performance profiles continue to change rapidly and open-weight and frontier models offer different cost, operational, and trust trade-offs. The governance model therefore needs to span finance, technology, business, and risk teams rather than being treated as a technology-cost discipline alone.

There is also the question of the unit of measure. Business leaders rarely care about the cost of an individual database query or cloud transaction. They care about what those transactions produce. The same principle increasingly applies to AI.

Leading organisations are already focusing more on agent economics. With agentic AI, one human request can trigger multiple model calls, searches, tools, and downstream actions, all of which consume tokens. Organisations therefore need to govern not only human-triggered consumption but also potentially autonomous economic consumption. That creates a new requirement for limits, observability, and accountability.

Knowing the full TCO is important. The next maturity step is to have clear ownership, thresholds, decision rights, and controls around that cost. This moves the discussion from measurement into enterprise governance.



Three realities of trusted AI economics (Continued)

Strong governance turns architecture choices into a source of cost advantage

The same task can have dramatically different economics depending on engineering. Poorly designed agentic solutions can multiply costs through inefficient routing, agent loops, and inappropriate model selection.

Design engineers need to understand the cost decisions and implications of their work, which assumes they know the most efficient way to architect an agent and use the right model for the task. Given how rapidly AI is evolving, that training is rarely in place. From our experience, there is a more cost-efficient way to deliver the anticipated outcome in the vast majority of an organisation's existing AI use cases. Consider, for example, a routine finance task, such as pulling structured data out of an invoice. The same prompt, document, and expected output run 10,000 times through a flagship model might cost US$200 per batch. However, run the same workload through a smaller model chosen for the task and the bill could drop to about US$40 - an 80% savings on identical work. In one pilot PwC conducted, more than 90% of user queries could have been answered by the smallest model in the portfolio, yet 100% of users selected the largest.

To be clear, the goal is not always to choose the least expensive model, rather the one that delivers the best economics for the required outcome. A more expensive model may still be optimal if it materially improves completion, accuracy, customer outcomes, or avoids downstream human intervention. Ultimately the executive metric should increasingly be value per unit of AI consumption, not simply the lowest model cost.

Organisations have taken different steps to control the costs associated with model selection and token consumption with mixed success. A uniform token cap can result in unused capacity for some and a halt in productivity for others. Restricting access to certain models for certain users can also cap innovation.

This brings the need for economic governance into sharp focus. The same way we increasingly build trust and risk controls into the AI lifecycle, we should design economic controls into AI architecture: model routing, consumption thresholds, action limits, escalation triggers, and stopping conditions. Organisations must define clear boundaries around agent behaviour, including which tasks agents are authorised to perform, what data they can access, what resources they can consume, how long they can pursue a task, what economic thresholds apply, and when human intervention is required.



Three realities of trusted AI economics (Continued)

Value matters more than cost

Cost optimisation alone does not create value. An efficient AI system that does not create measurable business value is still a failed investment. While many AI costs are hidden, the metrics can be relatively straightforward once those costs are identified. Value metrics are more challenging. Organisations need broader value measures that include efficiency, accuracy, insight, business outcomes, and the sustainability value AI enables. Metrics such as cost per resolved case, completed transaction, customer interaction, or capacity created help translate the technology economics into business economics. Together these factors help organisations build a more complete picture of AI's overall net benefit, and link AI investment directly to measurable business impact.

Value may need to be measured differently depending on its incremental versus transformational impact. Many organisations are merely plugging AI into existing processes rather than redesigning them around AI-native ways of working. In this case value may be measured against old ways of working, including savings in people time, delay, rework, and opportunity cost.

Incremental automation produces incremental value, but transformational redesign produces transformational value. Model and infrastructure optimisation are important, but the larger prize comes from redesigning processes, roles, and organisational capacity around AI. Such reinvention requires research and development (R&D) spend, experimentation and some inefficiency before a use case scales. In these cases, early return on investment (ROI) may not be visible, but value and overall economics continue to stack up as adoption scales.

Transformational redesign does not have a benchmark to measure value against. There is a large amount of analysis, review, research, and problem solving that AI can do that businesses simply do not do today because the human cost makes it economically infeasible. This might not show up as a net ROI in the context of today's operations but the strategic value for the future business is invaluable.



Building AI economics by design

Organisations that want to scale trusted AI should focus on four priorities and manage AI as a business investment governed through a framework that connects cost, trust, and value.

01

Establish AI economic governance

Create transparency into the full cost of AI operations. Move beyond token tracking toward true TCO measurement. Build AI cost taxonomies and reporting dashboards and establish cost ownership and chargeback mechanisms. Create cost visibility, while also establishing ownership, budgets, guardrails, thresholds, exception processes, and clear accountability for cost versus value.

02

Design for trust through continuous evidence collection

Trust should be based on measurable evidence, not assumptions. Implement telemetry (continuous collection, processing, and analysis of data generated by AI systems), monitoring, observability, and control frameworks that make AI performance measurable and auditable.

AI telemetry can, in many cases, provide greater transparency than traditional human processes. Humans often cannot explain every step of their reasoning. AI systems can generate logs, decision traces, and performance metrics that provide visibility into how outcomes were produced and provide robust data to proactively monitor agent performance.

03

Optimise through architecture

Develop architectures that support intelligent model routing, workload optimisation, and evolving technology choices without lock-in. Workloads should use the most appropriate model for the task, recognising that the best model today may not be the best choice in 60-to-90 days as new models and upgrades are released. They should also be engineered for cost controls, such as consolidating multiple tasks into a single call and implementing escalation triggers when costs hit certain thresholds.

The ability to route workflows across models and providers and avoid model calls when possible is not just good technical architecture. It creates optionality as pricing and performance change and allows organisations to continually optimise cost, capability, and trust as the market evolves.

04

Measure value relentlessly

Create balanced scorecards that link AI investments directly to operational and business outcomes. If teams cannot articulate value, they cannot justify scale.

Value is not just one number when it comes to AI. Organisations should measure a balanced range of outcomes that may include cost savings, accuracy improvements, productivity gains, energy or resource efficiencies, new business insights, innovation outcomes, or customer experience improvements. Value, in other words, is multidimensional.

Three questions the Chief Executive Officer and board should ask

Governing and continuously optimising the relationship between AI cost, trust, and value is a shared responsibility of the CFO, CTO, and Chief Operating Officer (COO).

The CFO increasingly becomes a steward of the AI investment portfolio: understanding where spend occurs, whether it translates into business outcomes, and where capital should be reallocated as economics change.

The COO must know which processes use AI, what they cost, the value they bring, and how they are governed.

The CTO must build flexible systems that can adapt as models, pricing, and capabilities evolve.

Here are three questions to consider as the organisation assesses its trusted AI economics:

- Can we explain the full cost of AI across models, infrastructure, governance, and operations?
- Can we demonstrate business value from AI with the same rigour we measure cost?
- Have we built the architecture, controls, and governance needed to scale AI with trust?

Trusted AI advantage lies in managing the entire economic system

The challenge in scaling AI is not limited to funding AI cost. Organisations should manage the entire economic system: model selection, architecture, governance, controls, value realisation, and business outcomes. This requires building the governance, architecture, and operating model that allow AI investment and value to scale faster than its cost, risk, and complexity.

Token costs will continue to fall. Models will continue to improve. New architectures will emerge. What will differentiate organisations is not their access to AI. It will be their ability to govern it, measure it, trust it, and connect it relentlessly to business value.

This is the fourth article in ‘The AI trust dividend’ series, where we continue to go deeper on the challenges business leaders need to address at different points in their AI maturity to unlock its full potential. Other articles in the series include:

Creating strong foundations, trusted AI, and real returns

How trust by design accelerates AI value

Three governance shifts to build trust in agentic AI autonomy



Authors and contributors



Reggie Kelley
Partner, PwC UK
reggie.kelley@pwc.com



Scott Likens
Global Chief AI Engineer, Principal, PwC United States.
scott.likens@pwc.com



Tanuj Talwar
Partner, PwC India
tanuj.talwar@pwc.com

Thank you

Access the online version

© 2026 PwC. All rights reserved. PwC refers to the PwC network and/or one or more of its member firms, each of which is a separate legal entity. Please see www.pwc.com/structure for further details. This content is for general information purposes only and should not be used as a substitute for consultation with professional advisors