

TECH FUNCTION

How AI Is Reshaping the Economics of Technology

By Michael Grebe, Joppe Bijlsma, Djon Kleine, Vladimir Lukic, Daniel Baker, Eike Kock, Juuso Soininen, Marc Roman Franke, Sabine Döschl, and Felix Sippel

ARTICLE OCTOBER 02, 2026 15 MIN READ

AI demand is driving up technology costs fast. Server prices are peaking at 80% above 2025 rates¹ and software inflation is running four to five times the market rate.² Meanwhile, average token costs for the top US large language models (LLMs) rose from below \$5 per million at the start of 2025 to \$20 per million in 2026³—a point-in-time snapshot of a market where model mix and pricing continue to change rapidly. Price, however, is only part of the story; as adoption accelerates, token volume becomes an equally important cost driver.

This is the era of token-based competition (see article [“The Era of Token-Based Competition Is Here”](#)), and it is only adding to the CIO’s responsibilities. The question is no longer only how much technology costs but also how much business value the intelligence being purchased generates (see article [“Return on AI: What CEOs Need to Know About the True Cost of Artificial Intelligence”](#)). On cost, the answer has been stable: from 2010 to 2025, IT budgets averaged 2.8% to 3.7% of revenue across industries (see article [“How CIOs Can Prove the Value of Technology in the Age of AI”](#)). AI now tests that stability by calling on CIOs—and on CDOs, CTOs, and other senior technology leaders—to govern enterprise AI demand and tie technology investment to business value.

This article examines how AI is reshaping enterprise technology economics: which cost elements grow, which decline, and how total spending is likely to evolve. Our scenario-based model combines two effects across low, medium, and high AI adoption scenarios—the productivity AI creates inside the [Tech/IT function](#) and the demand it creates in the business. Our modeling is not focused on AI-native companies but on organizations that are progressively embedding AI into existing operations.

The Budget Dynamics of AI

AI moves technology costs in three directions at once. It *raises IT cost* in absolute terms through new platforms, data, integration, governance, licenses, and tokens. It *lowers IT cost* by making the Tech/IT function more productive. And it *adds AI spend* outside Tech/IT in business and support functions and increasingly in product IT and operational technology where AI is built into what a company sells. Where the cost-to-income ratio lands depends on how much each cost lever moves, and on AI’s impact on revenue.

A company’s starting point matters as much as its AI ambition. Organizations with large enabling functions and low Tech/IT intensity, such as industrial goods and transport, must plan for heavy business-driven demand even as AI improves Tech/IT productivity. Those with high intensity and digitally enabled operations—such as media and digital services, financial services, and telecommunications—have a larger cost base to work on and more room to capture AI savings. Their challenge is reinvesting those savings into growth.

Companies fund AI in different ways. Some treat it as part of the IT budget, while others allocate it, one-time and recurring, to the business functions that capture the benefit—for example, when AI replaces personnel cost in sales or operations. Hybrid models are also common. This matters because the formal IT budget captures only part of an enterprise’s technology spending.

While IT spending has remained relatively stable as a share of revenue over the past 15 years, shadow IT (such as technology spending in regions, markets, and business functions outside the

central IT budget) has long added to the total. AI now appears to be expanding the technology footprint further, with much of the new investment occurring outside Tech/IT. Our Applied AI 2026 survey indicates the scale of the shift: respondents reported that AI's share of the IT budget has more than tripled, from 5% to 17%. The survey also suggests that AI spending outside the Tech/IT budget is rising significantly as companies move from pilots to enterprise-wide programs. Based on respondents' estimates, companies spend on average 3.3% of revenue on enterprise IT plus 2.8% of revenue on AI initiatives funded outside the Tech/IT budget. But that split varies widely by industry.

As AI workflows scale into core operations, their costs—data platforms, model operations, integration, governance—belong in the enterprise technology stack, where reliability, security, and cost control must be protected. AI agents are enterprise assets too, and need the same visibility, governance, and life cycle management as any other part of the estate.

Not all AI spend is the same, and it should not be judged the same way. One-time costs, recurring costs, and production tokens land in different places in the P&L: tokens that build reusable capability are a capital expense, tokens that run internal work are an operating expense, and tokens inside a product a customer pays for are part of the cost of goods sold (see article "[Return on AI: How CFOs and CIOs Can Manage the Token Meter](#)").

What unifies them is the return on AI (RoAI). For a CIO, the practical test is simple. Every scaled AI investment needs an explicit value thesis and evidence appropriate to its purpose: banked efficiency, measurable growth, or staged strategic value. If it repeatedly fails that test, the investment should be rescope or stopped.

Modeling How AI Will Reshape Enterprise Technology Cost

To understand how these dynamics affect total IT cost, we built a scenario-based model that assesses nine high-value AI areas, seven of which sit inside Tech/IT and share one goal, which is to raise the productivity of internal technology work:

- **Software development life cycle (SDLC).** AI improves coding, testing, and documentation; speeds up delivery; and reduces rework.

- **Legacy migration.** Legacy analysis and code conversion are enhanced with AI to speed up modernization.
- **IT operations.** Run operations benefit from AI-enabled monitoring, triage, and remediation.
- **Application management.** Smarter diagnosis, support resolution, and proactive maintenance are enabled by AI.
- **Enterprise Resource Planning (ERP) transformation.** Analysis, design, migration, and testing are supported by AI to accelerate delivery.
- **Tech sourcing.** Spending analysis, supplier insight, and negotiation preparation are improved with AI.
- **IT service-desk automation.** Ticket intake, knowledge retrieval, and routine resolution are automated with AI.

Two high-value AI areas sit in the business: using AI in support functions such as finance, HR, and legal to reduce workforce cost and using AI to grow revenue through personalization, pricing, and new business models.

The assumptions behind the model are explicit. We model a five-year horizon, with savings and costs fully ramped by year three. One-time AI implementation cost is phased over the first two years, and each of the nine high-value areas carries its own cost and savings profile. Together they form the enterprise AI one-time and run costs. Returns differ by scenario: the higher the rate of adoption, the larger both the investment and the benefit. We assume baseline growth without AI of roughly 6% nominally each year—3% organic revenue growth compounded with around 3% price inflation. We also assume an evolutionary approach to AI adoption rather than a greenfield one. For more on the model, see our Methodology sidebar.

– Methodology

Our model calculates AI cost and savings for each high-value area separately and assigns both to the specific parts of the IT budget they affect: internal staff by discipline, contractors, software and token spend, infrastructure and hosting, hardware, and managed services. It therefore shows not only how much technology cost moves but also where.

Three inputs drive the ongoing cost: how many staff are given AI tools and actually use them, how much routine work is automated, and how many AI use cases run in each function. From these we build up six cost lines: AI licenses (licensed staff × annual price), token cost for staff use (staff with AI × token

volume × token price), token cost for agent use (roles automated × token volume × token price), hosting the platform, maintaining and refreshing the models, and the AI specialists each use case needs.

The one-time cost of building AI works in a similar way. Multiplying the number of use cases to build times the amount of effort per use case gives the size of the build team. That sets the cost of labor, platforms and data, and the tokens consumed while building—with training, change management, and program governance scaled from those. For modeling purposes, we assume all AI-related spending ends up in IT's cost-to-income ratio, regardless of whether it is booked there.

One caveat is central: the model assumes that AI productivity gains inside Tech/IT are realized, meaning freed capacity leaves the cost base rather than being absorbed elsewhere. Experience says that this is the hard part, since roughly 10% of value comes from technology, 20% from data and algorithms, and 70% from changes to people, process, and operating model. Tooling alone rarely bends the cost curve; without those changes, savings weaken while demand arrives on schedule.

We applied the model to two exemplar companies: an industrial goods company representing lower Tech/IT intensity and a media and digital services company representing higher intensity. Each company's size, revenue, and operating model differed, so outcomes were varied.

Highlights for the Industrial Goods Company

The industrial goods company has comparatively low Tech/IT intensity, a large operational and support-function footprint, a \$1.0 billion IT cost baseline, and an IT cost-to-income ratio of 2.6%. Model assumptions are described above, with 2% annual traditional IT cost streamlining (for example, by capturing AI benefits also in external contracts) and one-time AI investment phased over the first two years.

While absolute technology spend rises, the cost-to-income ratio does not necessarily follow because AI also lifts revenue. Where the Tech/IT cost-to-income ratio (ITCIR) lands depends on the productivity AI delivers inside Tech/IT, the pace of AI adoption in the business, and the revenue AI generates. ITCIR provides a measure of technology affordability relative to revenue; the ultimate test of AI economics, however, is the enterprise value created. Three scenarios show the pattern: (See Exhibit 1.)

EXHIBIT 1

Tech and AI Spending Scenarios for Industrial Goods Company



- **Applying AI in business high and Tech/IT side high keeps ITCIR at 2.6%.** With high adoption in both the business and Tech/IT, AI productivity of \$143 million offsets most of the \$194 million that volume and inflation add, holding the Tech/IT budget 3% below baseline at \$973 million. Business AI then adds \$277 million of Tech/IT run cost, taking year-three spend to \$1.25 billion. Revenue reaches \$48.4 billion including AI uplift, so the cost-to-income ratio holds at its 2.6% starting level.
- **Applying AI in business low and in Tech/IT high reduces ITCIR to 2.2%.** With high adoption in Tech/IT and low adoption in the business, the ratio improves to 2.2%. The Tech/IT budget falls 1% to \$985 million and total spend reaches \$1.06 billion, against \$1.0 billion today.
- **Applying AI in business high and in Tech/IT low increases ITCIR to 2.8%.** With low adoption in Tech/IT and high adoption in business, the ratio rises to 2.8%, around 8% above the 2.6% baseline. Total spend reaches \$1.34 billion: the business adds \$277 million of AI run cost while the Tech/IT function delivers limited productivity of its own.

Inside the Tech/IT function, the largest productivity gains come from the product and software development life cycle (PDLC/SDLC), ERP transformation, and application management. These are the biggest value pools, and they scale with adoption.

The pattern is consistent across scenarios. When Tech/IT transforms at least as fast as the business adopts AI, productivity gains within Tech/IT can partially offset the additional technology cost created by enterprise AI demand. But they are only one part of the equation: the overall economics also depend on whether productivity gains in the business are banked and whether AI generates incremental revenue. Without those benefits, higher AI adoption can leave the enterprise with a structurally higher technology cost base.

The lesson from this industrial goods example is clear: CIOs should not view AI as a cost-reduction program. AI fundamentally changes the economics and composition of enterprise technology spending. Productivity gains within Tech/IT can help fund enterprise AI adoption, but sustainable economics ultimately depend on converting AI investment across the enterprise into banked productivity, revenue growth, or both.

Highlights for the Media & Digital Services Company

The media and digital services company has a high Tech/IT intensity, the operating model is digital-first, the IT baseline is \$150 million, and the cost-to-income ratio 7.7%. Technology here is not a support function but a core element of the business, which is why the ratio starts higher. Other model assumptions are unchanged. (See Exhibit 2.)

EXHIBIT 2

Tech and AI Spending Scenarios for Media Company



Source: BCG AI Tech Costing Model, BCG project experience, Company data, BCG analysis.

As technology represents a much larger share of the cost base, Tech/IT productivity creates a particularly powerful lever for offsetting the additional cost of enterprise AI adoption. Three scenarios illustrate the pattern:

- **High AI adoption in both business and Tech/IT reduces ITCIR to 6.7%.** With high adoption on both sides, AI productivity of \$33 million more than offsets the \$29 million added by volume and inflation. Together with \$10 million in traditional IT cost savings, this reduces the year-three Tech/IT function budget by 9%, to \$136 million. Business AI adds \$25 million in run cost, taking total spending to \$161 million. With revenue increasing to \$2.4 billion, ITCIR falls from 7.7% to 6.7%.
- **High AI adoption in Tech/IT and low adoption in the business reduces ITCIR further, to 6.1%.** When Tech/IT transformation runs ahead of business adoption, the Tech/IT function captures the same strong productivity gains, reducing its budget by 9%, to \$137 million. With business AI adding only \$7 million in run cost, total spending reaches \$144 million. The combination of lower technology spending and higher revenue brings ITCIR down to 6.1%, the lowest ratio across the three scenarios.
- **High AI adoption in the business and low adoption in Tech/IT keeps ITCIR close to initial level, at 7.5%.** When business adoption runs ahead of Tech/IT transformation, lower internal AI productivity of \$13 million is insufficient to offset volume and inflation, leaving the Tech/IT function budget 3% above baseline, at \$155 million. Business AI adds another \$25 million in run cost, taking total spending to \$180 million. However, the associated revenue uplift to \$2.4 billion largely absorbs the additional spending, keeping ITCIR at 7.5%, close to its 7.7% starting point.

For CIOs, these two examples reveal a critical insight: the same AI productivity logic can produce opposite results. In the media company, it cuts the Tech/IT budget well below baseline, while in industrial goods, it only offsets growth and inflation because less of the cost base is addressable relative to the growth it must carry. Where technology is the business, revenue uplift can be sufficient to keep the ITCIR broadly intact even if AI adoption is unbalanced. A lower ITCIR is therefore not necessarily the better outcome; the ultimate test is whether the additional AI investment creates enterprise value.

How CIOs Must Manage the New Economics of Technology

As the economics of enterprise technology enter a new era, the CIO's mandate changes. The role is moving from running the technology function to orchestrating enterprise-wide AI capability (as we also discuss in the article “Why We Still Need a CIO in the AI-First Era”). Alongside that, CIOs become stewards of technology economics, making sure AI investments create sustainable business value rather than simply increase spending.

We believe that this requires CIOs to focus on three priorities.

Create full transparency over enterprise technology and AI spend. CIOs should establish a single view of technology and AI spending across the enterprise—not just what sits within the formal Tech/IT budget. That means bringing together core IT spending; shadow IT in business units, regions, and markets; and the rapidly growing AI investments made outside Tech/IT. This transparency should extend from one-time investments and recurring costs to token consumption and realized benefits, with clear ownership for every material workflow. As AI spending becomes increasingly distributed across the organization, this enterprise-wide view is essential not only to control cost but also to manage the technology complexity that comes with it.

Apply AI within Tech/IT to help fund enterprise AI. CIOs should lead by example, applying AI to the largest value pools within their own function to deliver measurable productivity gains and offset part of the growing cost of enterprise AI demand. They should optimize for return, not token consumption. RoAI—economic return divided by the combined cost of human intelligence and tokens—shifts the focus from AI spend to cost per successful outcome. Measured this way, the cheapest model is often not the cheapest answer. A more capable model that completes a task on the first attempt can cost less overall than a smaller one that retries, escalates, or produces work a human has to redo. The same logic applies to consumption. Total spend depends on the scale and intensity of AI adoption, the context agents carry through each loop, and the models teams use. That makes AI economics an architecture and accountability question, not simply a

procurement one. CIOs should use this lens to continuously optimize how people and AI work together.

Own the enterprise AI economics. As AI embeds across the enterprise, CIOs must keep technology and AI investment aligned with business value, measured by cost per outcome rather than activity. That means directing AI and token spending toward the opportunities with the strongest value potential rather than simply constraining consumption. They also become stewards of the emerging “synthetic” P&L—the economics of a workforce made up of people and AI agents. That means designing the operating model for human and digital labor, managing the life cycle of AI agents, and optimizing how work is allocated between the two.

Ultimately, AI will reshape technology economics, not simply increase technology costs. CIOs who steer their companies toward the highest-return AI opportunities while continuously modernizing the technology foundation can turn rising AI demand into an advantage—accelerating innovation, enabling new business models and avenues for growth, and strengthening margins. The opportunity is not simply to manage the cost of AI but to shape where and how it creates value.

Authors



Michael Grebe

Managing Director & Senior Partner
Munich



Joppe Bijlsma

Managing Director & Partner
New York



Djon Kleine

Managing Director & Partner
San Francisco - Bay Area



Vladimir Lukic

Managing Director & Senior Partner; Global Leader, Tech and Digital Advantage
Boston



Daniel Baker

Managing Director & Partner
Sydney



Eike Kock

Managing Director, BCG Platinion
Platinion – Cologne



Juuso Soininen

Partner & Director, Digital Transformation
Helsinki



Marc Roman Franke

Partner & Associate Director, AI and digital transformation
Berlin



Sabine Döschl

Senior Director - Tech Function & Cost Optimization
Munich

Felix Sippel

Manager, BCG Vantage
Amsterdam





ABOUT BOSTON CONSULTING GROUP

Boston Consulting Group bridges the gap between ambition and outcomes for the world's leading companies and organizations. We are built for this era of unprecedented change — bringing strategic clarity rooted in over 60 years of deep domain knowledge, combined with applied AI shaped by our practitioners. BCG works shoulder-to-shoulder with CEOs across industries and geographies to deliver transformative impact at scale: stronger returns, transferred capabilities, and change that sticks. For more information, visit bcg.com.

This publication is made available solely for your personal, informational, and non-commercial use.

You may download and retain a copy for such purposes. Except with the prior written permission of Boston Consulting Group (BCG), you may not use, reproduce, republish, distribute, sell, license, modify, create derivative works from, or otherwise exploit this publication or any portion of its content for any commercial purpose.

Without limiting the foregoing, the publication and its content may not be used for the commercial development, training, testing, evaluation, or improvement of artificial intelligence, machine learning, or other automated systems, or incorporated into any commercial database, dataset, product, or service, without BCG's prior written permission.

All rights not expressly granted are reserved by BCG.

© Boston Consulting Group. All rights reserved.

¹ Gartner (March 2026) Server Price Surge Survival Guide

² Vertice Inflation Index 2023–2026

³ BCG article [“The Great Divide: How the US and China Are Splitting the AI World”](#)