NTT DATA

October 2026

AI economics: A leadership framework for sustainable net value

A practical guide to assessing the benefits, costs, risks and trade-offs that determine the enterprise value of AI



Author

Pawel Andre, AI Director, NTT DATA Belgium

Pawel Andre leads AI transformation and strategy, while advising client executives on data, AI and broader enterprise transformation. He helps organizations turn AI and data into measurable business value and lasting transformation. Previously, he served on the EMEA board at a leading AI and analytics company. He holds three degrees from ICHEC and ESSEC, including an executive master's degree in digital transformation, awarded summa cum laude.

Contributors

Midialis Gonzalez, Head of Strategic Partnerships, NTT DATA Belgium

Elke Willems, Open Universiteit

Mark Van Achter, Professor and expert in human capital, AI and economics, KU Leuven

Contents

04 **The role of AI economics in AI transformation**

05 **AI is only part of what it takes to create value**

06 **AI Net Value: Toward a broader measure of AI returns**

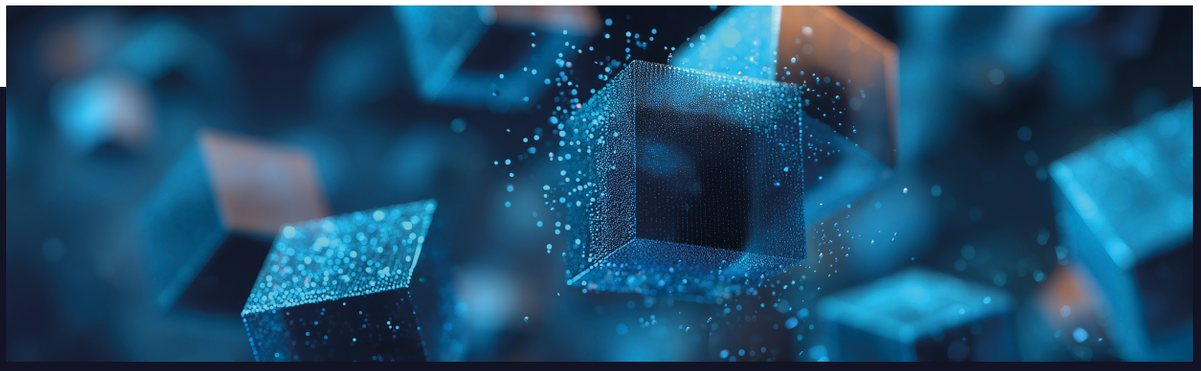
19 **AI economics is a management discipline**

The role of AI economics in AI transformation

AI has crossed an important threshold. Leaders know the technology can perform useful work, but they may be less certain about how to convert the rapidly expanding capability of AI into sustainable value.

NTT DATA's 2026 Global AI Report, which surveyed 2,567 senior executives across 35 countries and 15 industries, classified around 15% of organizations as AI leaders. These AI leaders tightly align AI with business strategy, focus on high-value domains that unlock disproportionate value, redesign workflows from end to end, and create a cycle where initial investments fuel early success that drives reinvestment for further growth.[1]

This is the starting point for AI economics — looking beyond productivity gains, model performance, adoption figures and visible technology costs to understand the broader economic system surrounding AI.



This guide proposes a framework of six dimensions that, when taken together, transform the role of the AI business case:

AI Net Value = Realized Value + Future Value + Protected Value – Lifecycle Cost – Residual Risk – Economic Trade-offs

Note: This equation is an executive synthesis rather than an accounting standard. It is not intended to create an illusion of mathematical precision but to help leaders identify the factors that determine whether an AI investment ultimately creates sustainable net value.

AI economics is not a one-time calculation made before funding is approved. As models, processes, people and providers change, value, cost and risk continue to evolve. Because of this, the business case needs persistent ownership. Business leaders should remain accountable for the outcome, while Finance helps establish credible baselines and verify that the benefits attributed to AI are supported by evidence.

This also explains why AI economics ultimately leads to AI transformation. Understanding what creates net value is only the start. Organizations also need to build the capabilities to repeatedly improve that result. Business direction, talent, operating models, technology, data, adoption and scaling all influence how efficiently technological possibility translates into organizational performance.

As the frontier-model race continues to expand the capabilities of AI, the more enduring difference between organizations will increasingly be determined by how well they convert those capabilities into value. The question for leaders now is not how much AI they can deploy but how much sustainable value they can repeatedly create from it.

[1] NTT DATA. 2026 Global AI Report: A Playbook for AI Leaders. December 2025.

 © 2026 NTT DATA, Inc.

AI is only part of what it takes to create value



Yesterday, we were asking: "Can AI make us more productive?" Today, the question is: "Can AI make us more productive *and* more profitable?"

We already know that AI can accelerate task completion, analyze information, generate content, facilitate decision-making and increasingly automate complex tasks. It's becoming easier to answer the technical question. The same cannot be said for the economic question.

For businesses, the ultimate determining factor is whether AI improves growth, margins, productivity or capital efficiency. For public agencies, it's the measurable value of their mission — better services, more informed decisions, reduced costs, fewer losses or increased capacity.

In both cases, the challenge for leaders is the same: To what extent does AI's rapidly expanding potential actually translate into value?

This question requires leadership attention as AI capabilities advance faster than organizations' ability to integrate them, and costs become an increasingly important issue. Executives are hearing more and more about model pricing and token consumption, particularly as generative AI (GenAI) and agentic AI move into production. But model and token costs are only part of what it takes to run AI at scale.

The 1–2–3–4 rule

Abhijit Dubey, CEO of NTT DATA, illustrates this principle with a simple 1–2–3–4 rule:

$1 to create the technology

$2 to drive adoption

$3 to build the ecosystem

$4 to prepare the data[2]

The message for leaders is this: Model and token costs are only part of what it takes to run AI at scale, and AI itself is only part of what it takes to create value.

[2] Semafor World Economy. Interview with Abhijit Dubey. April 2026.

Focusing solely on model prices or token costs can be misleading

The real economic question isn't simply how much AI costs but what it takes to turn AI into a reliable business outcome. The same applies to productivity. An employee who completes a task 30% faster does not automatically reduce the company's costs by 30%. That increase in capacity must be captured.

This is where many AI business case analyses fall short. They measure the immediate benefit and the visible technological cost but overlook the organizational work required to capture it. The economic challenge therefore extends beyond the technology itself.

Sometimes, the constraint is not the intelligence of the AI at all. Instead, it comes from fragmented data, contradictory content, inconsistent processes or poorly defined accountability within the organization. NTT DATA's work on enterprise knowledge underscores the same fundamental point: AI requires reliable data and content, context, business meaning and clearly defined accountability to function reliably at the enterprise level.[3]

Models open up new possibilities. The organization transforms these possibilities into value. This report calls that broader conversation "AI economics":

AI Net Value = Realized Value + Future Value + Protected Value – Lifecycle Cost – Residual Risk – Economic Trade-offs

The equation is not intended to create an illusion of mathematical precision, nor to suggest that these six dimensions are exhaustive or fixed. When applying the framework, it is important that the dimensions remain conceptually distinct so that the same economic effect is not attributed to more than one component.

AI economics will continue to evolve as the technology, its uses and its risks evolve. Its purpose is simple: It gives leaders a more complete view of the components that determine whether AI creates net value and a practical framework for asking better questions.

AI economics should therefore be treated as a living view of value, refined as assumptions are replaced by evidence, and tested against the best realistic alternative.

The question is no longer whether AI works but whether it creates sustainable net value once benefits, costs, risks and trade-offs are considered. That is the shift from promised AI value to proven net value.



[3] NTT DATA. 8 ways to improve your enterprise knowledge strategy for AI. July 7, 2026.

AI Net Value: Toward a broader measure of AI returns

1. Realized Value

A pilot cuts processing time in half. A model produces highly accurate recommendations. Employees say the new assistant saves them hours every week. Has the organization created value?

Possibly. But not necessarily.

The first component of AI economics is Realized Value. This is the measurable business or mission value of AI that the organization is capturing.

It is important to note that the potential value of AI can far exceed the value eventually reflected in business performance.

Examples of Realized Value

Realized Value for businesses

- Higher revenue, stronger margins or lower operating costs
- Greater throughput, shorter cycle times or a more efficient use of capital and assets
- Better quality, stronger decision-making and improved customer outcomes
- A better employee experience, as reduced administrative work creates capacity for higher-value activities
- …

Realized Value for public-sector organizations

- Shorter backlogs
- Greater service capacity
- Better resource allocation
- Stronger policy outcomes
- Greater value for money
- …

The important word is "realized."

A complementary way to understand how AI creates economic value is to start at the level of tasks. The European Commission's Joint Research Centre distinguishes between tasks, the skills required to perform them and the broader competencies through which these skills are applied.[4] This perspective is particularly relevant to AI, since jobs are bundles of tasks that can be deconstructed and recombined as technology evolves.[5]

Task-based economics also differentiates between the various ways that technology can change work. AI can automate certain tasks, augment individuals performing other tasks, and contribute to the creation of new human tasks. These mechanisms can have different effects on productivity, capacity, quality, skills and, ultimately, economic value.[6]

[4] European Commission, Joint Research Centre. A Unified Conceptual Framework of Tasks, Skills and Competences: JRC Working Papers Series on Labour, Education and Technology 2021/02. 2021.

[5] Rodrigues, M, Fernandez, M.E. and Sostero, M. A unified conceptual framework of tasks, skills and competences. European Commission. JRC121897. 2021.

[6] Acemoglu, D and Restrepo, P. Automation and New Tasks: How Technology Displaces and Reinstates Labor. Journal of Economic Perspectives 33 (2): 3–30. 2019.

 © 2026 NTT DATA, Inc.

From improvement to economic value

Let's consider productivity. If AI reduces the time required to complete a task from ten hours to seven, three hours of capacity have been released. However, the organization has not automatically generated the financial equivalent of those three hours.

Economic value appears only when this capacity is used to process more work, improve services, take on higher-value activities, reduce external expenditure or otherwise affect a measurable outcome.

NTT DATA's work with Damm, a leading Spanish brewing group, illustrates this distinction. We designed a proof of concept that combines automated reasoning with tools such as geolocation applications and simulated database queries to detect duplicate records in hospitality establishment data. The solution achieved a reported decision accuracy rate of 97% and reduced database improvement effort by 60% to 70%, representing approximately 6,000 to 7,000 hours annually.[7] The accuracy rate tells us something about the quality of the AI. The improvement rate tells us something about the operational capacity released. The economic question is what the organization can do with that capacity.

These effects are not necessarily uniform, however. Productivity gains can vary significantly depending on the task and the level of experience.[8] In some cases, AI can even reduce performance when a task falls outside its current capabilities.[9] Organization-wide averages can therefore hide important differences in where value is actually being created and captured.

This is why **we distinguish between potential value and realized value.**

Realized value equals potential value multiplied by the Realization Factor. The Realization Factor describes how much of the opportunity identified in the business case actually translates into measurable business or mission performance.

Realization Factor = Adoption × Effective Utilization × Output Quality × Process Integration × Scale Coverage × Value Capture × Persistence

These conditions depend on one another and can evolve together over time. For example, Output Quality can influence trust, which can then influence Adoption. Effective Utilization can also strengthen or weaken the skills and organizational routines that affect future performance.

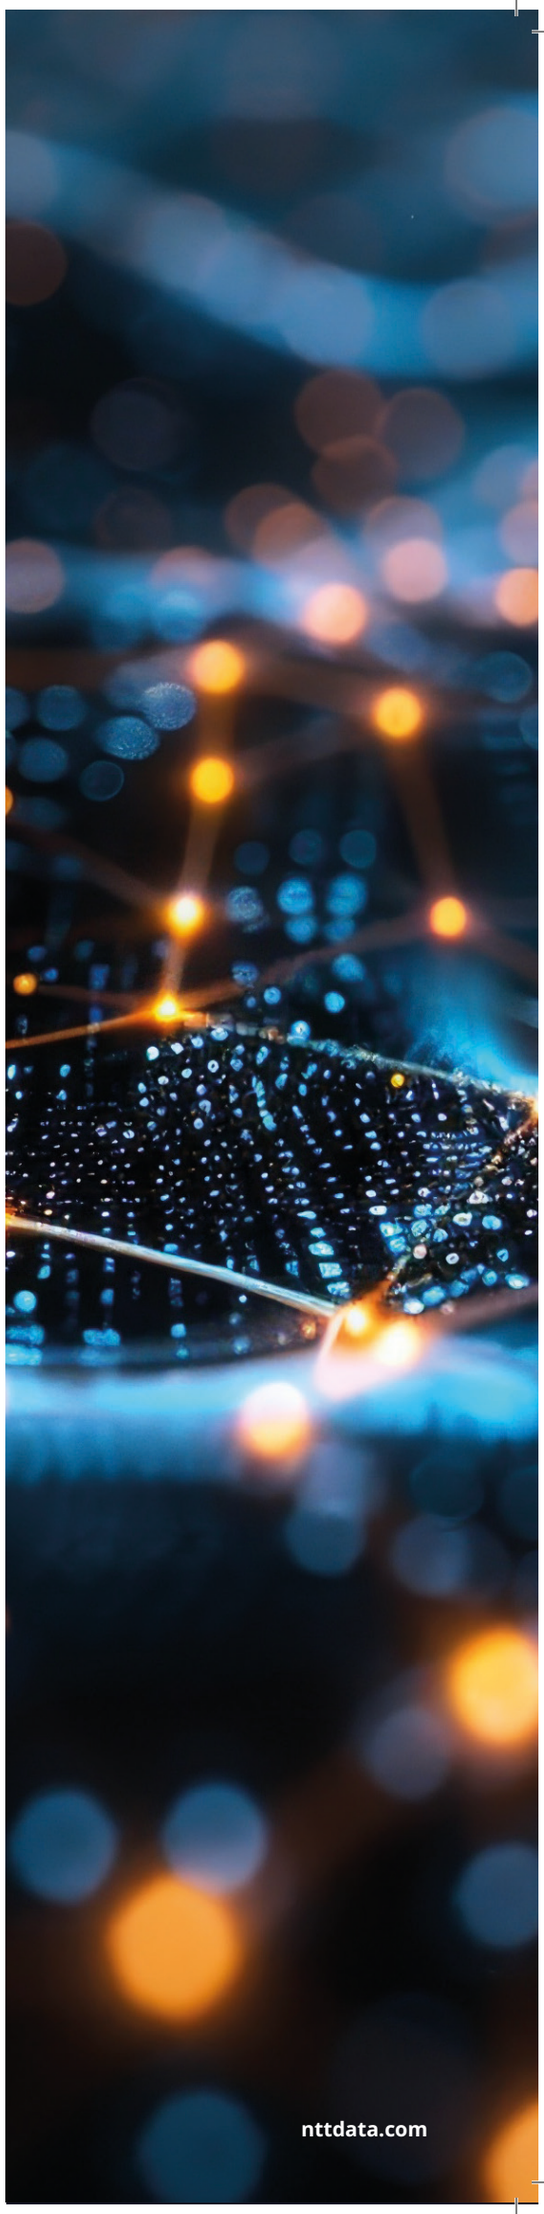
Even if an AI solution performs exceptionally well, it can still produce disappointing economics if people do not use it, if it is applied to low-value work, if the surrounding process does not change or if the improvement never reaches sufficient scale. The formula is not intended to suggest that every factor can be measured with mathematical precision. Rather, it is a practical way to identify where potential value is being lost.

[7] NTT DATA. Generative AI agents to optimize data at Damm.

[8] Brynjolfsson, E, Li, D, and Raymond, L. Generative AI at Work. The Quarterly Journal of Economics, Volume 140, Issue 2, May 2025.

[9] Dell'Acqua et al. Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of Artificial Intelligence on Knowledge Worker Productivity and Quality. Harvard Business School Technology & Operations Management Unit Working Paper No. 24-013. The Wharton School Research Paper. Revised March 17, 2026.

 © 2026 NTT DATA, Inc.



7 areas where potential value can get lost

Each component of the Realization Factor tests where an AI initiative could succeed or fail, for different reasons.

01

Adoption asks whether the intended people or systems will actually use the solution.
A predictive maintenance model may correctly identify equipment at risk of failure.
If technicians continue to follow the old maintenance schedule and ignore its recommendations, little value will be realized.

02

Effective Utilization asks whether AI is being used where it can materially improve an outcome.
People may use GenAI to answer basic questions that could be handled just as well by a web search, or generate additional content that nobody ends up reading.
While activity increases, economic value may remain close to zero. More AI-enabled work does not necessarily result in more valuable work.

03

Output Quality asks whether the result is good enough for the decision or process it supports.
A fraud detection model identifies more suspicious transactions.
If it generates too many false positives, investigators may spend more time clearing alerts than pursuing actual cases.

04

Process Integration asks whether the surrounding work has changed enough to utilize the improved output.
A demand forecasting model predicts demand more accurately.
If replenishment decisions are still made using the same weekly process and fixed ordering rules, inventory and availability may barely change.

05

Scale Coverage asks how much of the addressable opportunity has actually been reached.
A claims model used by one team handling 10% of total claims may prove that the solution works.
But most of the potential value remains unrealized at that scale.
The economic impact grows when the model reaches a meaningful share of the relevant claims volume, markets or operations.

06

Value Capture asks what the organization actually does with the improvement.
If AI reduces the time required to prepare a sales proposal by 30%, it frees up capacity.
But Value is captured only when this additional capacity results in more proposals, a higher conversion rate, lower external expenditure or another measurable outcome. Time saved does not equate to cash saved.

07

Persistence asks whether the result lasts.
While a predictive model may initially improve production performance, it can deteriorate as equipment, behavior or data changes. Similarly, an AI assistant can initially attract strong adoption that later falls away.
Realized Value should reflect the benefits that survive real operating conditions, not just the performance observed during the launch period.

 © 2026 NTT DATA, Inc.

These seven factors provide a practical way to identify where value is being lost. They also help to explain why organizations may report high adoption rates, strong model performance or promising pilot schemes, yet still struggle to demonstrate meaningful economic results.

For executives, the test can be very concrete:

- If processing speeds up, does throughput increase or does the backlog fall?
- If forecasting becomes more accurate, does this improve inventory or availability?
- If employees save time, is that capacity actually being redeployed?
- If decisions improve, do conversion rates, losses, service quality or other relevant outcomes change?



The economic value of AI does not lie in isolated improvements but in the improvement the organization succeeds in capturing.

When the organization becomes the constraint

This is also where another constraint becomes apparent. Sometimes, the organization may wish to capture the value, but the information, processes and rules on which the AI relies are not yet in place.

Structured data may be fragmented across systems. Unstructured content, such as policies, contracts, reports, emails, procedures and presentations, may exist in multiple versions and lack clear ownership or be unsuitable for reliable machine use. The real process may differ from the documented one. Important exceptions may remain implicit. Decision rules, permissions, responsibilities and controls may be unclear.

The broken organizational-intelligence layer

Work by NTT DATA on enterprise knowledge has highlighted the dependencies between structured data, unstructured content, context, business meaning and clear ownership. We refer to this as the "broken organizational-intelligence layer." Many organizations do not lack information. They lack usable organizational intelligence.

A controlled pilot can hide these weaknesses because the data is curated, the content is selected and experts are close at hand. But in production, there is nowhere to hide. What initially appears to be an AI problem can quickly turn out to be a problem of data, content, processes and accountability.

The answer is not to stop and repair the entire enterprise before implementing AI. Start with the economic leverage point or mission outcome that matters most. Transform the relevant business domain, redesign the end-to-end process and repair the structured data, content, decision logic and controls required to make that transformation work. This changes the role of the use case.

The use case is not the unit of transformation

A use case remains a useful building block. However, a better approach is to begin with the economic outcome, followed by the business domain, the processes that need to change and the organizational intelligence and decisions required to implement them. Only then should the organization decide which combination of AI, data, technology and other solutions is needed to deliver the result.

Realized Value therefore asks a harder question than "Does AI work?" It asks how much measurable value survives the journey from technical possibility to organizational performance.

And even that is only the value that is visible today. Some AI investments create assets and capabilities whose most significant economic impact will be realized in the future.

2. Future Value

What if the most valuable outcome of an AI investment is not the solution itself but the new capabilities that the organization gains as a result?

Realized Value considers what the organization captures today. Future Value considers what is created today that can generate additional value tomorrow. This includes reusable assets, organizational capabilities, learning and strategic options.

Building what can be reused

Imagine an insurance business builds its first AI solution for claims processing. In order to make it work, the organization prepares the relevant data and content, develops integrations with core systems and creates evaluation methods, controls and monitoring procedures. A team is trained and new ways of working are established. Much of this groundwork can be used for the next AI solution. For example, data products can be reused, APIs and connectors can serve new workflows and evaluation methods, or guardrails can be applied again. A shared AI platform can support additional solutions without the need to rebuild the same foundations each time.

The organization has therefore done more than create an application; it has also increased its ability to build the next one. This is where the economics of reuse become important.

The economics of reuse

Some of these assets are easy to identify. These include platforms, models, agents, reusable components, prepared datasets and authoritative content. Common integration patterns, monitoring tools and security controls can also be reused.

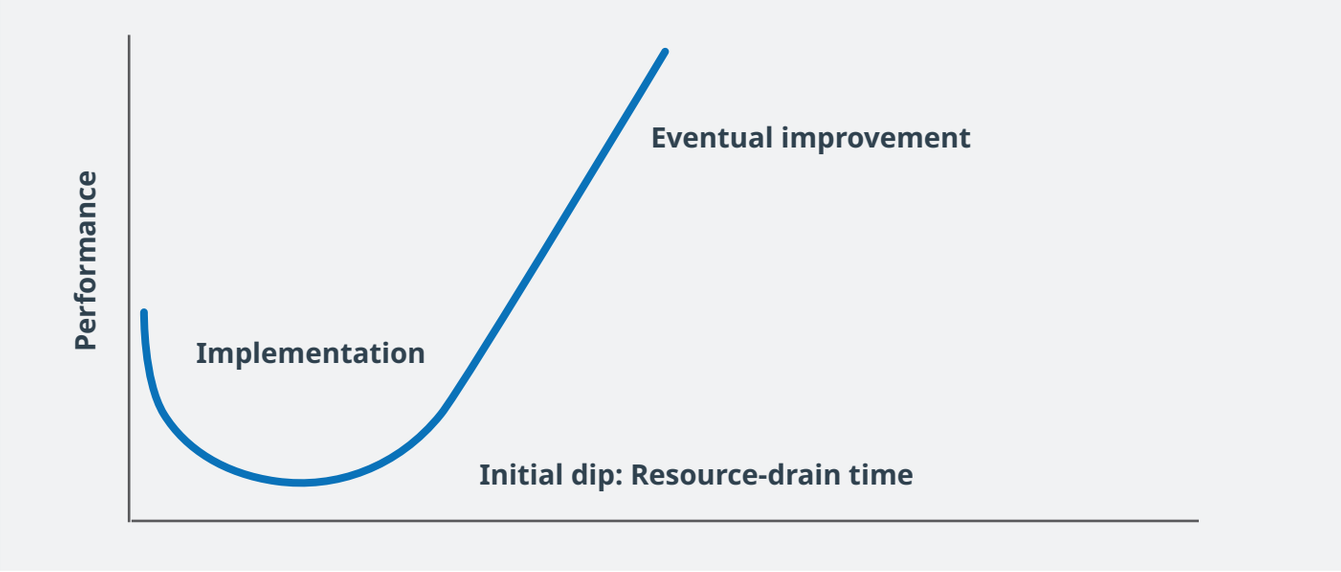
Other forms of Future Value are less visible but can be equally important. Teams learn how to take an idea into production. Product owners become better at managing adoption and value, while leaders improve their judgment about where AI is worth pursuing. Governance teams learn how to apply controls without recreating the process for every new solution.

The operating model can also improve

Decision-making becomes clearer, funding mechanisms become more effective and teams accumulate experience instead of starting again with every project. NTT DATA's 2026 research points to the same pattern: AI leaders combine disciplined scaling with stronger foundations and continued investment in the capabilities required for future waves of AI.[10]

Human and organizational capital play an economic role

The productivity J-curve described by Brynjolfsson, Rock and Syverson provides a useful lens. General-purpose technologies can require complementary intangible investment before their productivity effects become fully visible.[11] Training, job redesign, accumulated experience and trust should therefore not be viewed only as implementation costs. When they create durable capability, they are also investments in Future Value.



J-curve of Digital Transformation, adapted from Brynjolfsson, E. Rock, D. & Syverson, C. (2020). The productivity J-curve: How intangibles complement general purpose technologies. Cambridge: National Bureau of Economic Research.

There are two issues to consider here: first, whether people perform better when using AI and, second, what happens to human capability over time. Where independent judgment remains important, organizations should consider whether people continue to learn and transfer their knowledge to new situations, and if they remain capable when AI support is unavailable. The same question applies to how expertise is renewed. Agoria's Belgian AI Skill Shift study points to a related pipeline risk as routine tasks that traditionally helped junior profiles build experience become increasingly automated.[12] An immediate efficiency gain may therefore need to be considered alongside whether new learning pathways are replacing the experience that automation removes.

[10] NTT DATA. 2026 Global AI Report: A Playbook for AI Leaders. December 2025.

[11] Brynjolfsson, E, Li, D, and Raymond, L. Generative AI at Work. The Quarterly Journal of Economics, Volume 140, Issue 2, May 2025.

[12] Agoria, in collaboration with Indiville, The AI Skill Shift: AI and the Future of Talent in the Belgian ICT Sector.

From reuse to repeatability

Reuse is only part of the story. A reusable component can be applied again. A repeatable capability goes further, improving the organization's ability to reproduce a successful outcome.

Each deployment can refine the selection of opportunities, estimation of effort and design of solutions. Production experience reveals which assumptions were incorrect, while performance data creates better baselines for future decisions. Successful approaches can then be documented in playbooks and applied again.

Over time, one deployment improves the next. In this way, AI value creation becomes more repeatable rather than just reusable. The organization is both accumulating assets and developing a stronger system for converting those assets into outcomes.

The economic effects can be very practical. The next solution may require less engineering and be deployed more quickly. Data preparation and assurance may become easier while estimates become more accurate. As future deployment costs and efforts fall, opportunities that were previously too expensive may become economically viable.

Valuing what comes next

Future Value can also come from strategic options. Proprietary data, intellectual property or new organizational capabilities may enable the creation of products and services that do not yet exist. These assets could also facilitate the adoption of future AI improvements without the need for complete rebuilds.

In some sectors, these capabilities may be necessary simply to remain competitive. This is sometimes referred to as a "right to exist."[13] However, strategic necessity should not be used as an excuse to abandon economic discipline. Sound discipline recognizes that part of the return may lie in preserving future options rather than producing immediate financial results.

This makes Future Value harder to measure than Realized Value.

Measuring Future Value

Rather than assigning arbitrary financial values to concepts such as learning or capability, we should look for observable evidence. Did the second deployment cost less than the first? Was it delivered faster? How much technology, data or assurance work was reused? Did estimates become more accurate? Did an existing capability create a new opportunity?

Where the future benefit remains uncertain, it is better to use ranges and probability-weighted scenarios than to attempt to achieve false precision.

In short, Future Value asks what an AI investment leaves behind. It considers what an AI investment delivers today as well as what it will make possible tomorrow.

However, value does not come only from creating more. It can also be created by preventing something costly from happening, and this is what we term "Protected Value."

3. Protected Value

If AI prevents a loss of €5 million, has it created value even though revenue has not increased and no existing cost has disappeared? Yes.

Some of the most significant economic outcomes of AI arise from events that never occur. For example, a fraudulent payment is stopped. A machine is repaired before it breaks down. A quality issue is identified before products reach customers. A cyberattack is contained before it can spread. An incorrect decision is identified before it has financial or operational consequences.

These outcomes may never generate additional revenue, but they can still generate substantial economic value — and this is Protected Value.

Protected Value is the value created when AI helps to prevent or reduce losses that the organization would otherwise have faced.

Protected Value = Expected Loss without AI − Expected Loss with AI

[13] Davenport, TH and Srinivasan, L. 7 Factors That Drive Returns on AI Investments, According to a New Survey. Harvard Business Review. March 17, 2026.

The economics of prevention

Protected Value can appear almost anywhere in an organization. For example, banks may use AI to identify fraud or improve credit decisions. Insurers may use it to detect suspicious claims or reduce claims leakage. Manufacturers can use it to anticipate equipment failures or identify quality problems earlier. Cybersecurity teams can use AI to detect threats, prioritize incidents and respond more quickly.

The same principle applies to customer relationships. Earlier detection of service issues can prevent customer churn and compensation costs. Better compliance monitoring can reduce the likelihood of regulatory intervention. More reliable operational decisions can prevent disruption, rework and contractual failures.

For public organizations, prevention can be equally important. AI and analytics can help to detect fraud, waste and administrative errors. They can also improve eligibility decisions, strengthen oversight and safeguard scarce public resources.

In each case, value is created by reducing either the probability of an undesirable outcome or the damage caused by it. This is important because organizations often prioritize visible improvements to productivity, revenue or costs when developing their AI business cases. Prevention may be less visible, but it can be economically just as important.

Measuring what did not happen

Protected Value creates an obvious measurement challenge. If AI generates an additional $1 million in revenue, for example, the result can often be observed directly. However, if AI prevents a major operational failure, the organization sees precisely the opposite: Nothing happens.

Nevertheless, the absence of an event does not mean the absence of value. It simply means that the organization needs a credible view of what probably would have happened without the intervention.

Historical experience can provide a useful starting point:

- How frequently do fraud losses occur?
- What is the cost of an hour of production downtime?
- How much is usually spent correcting quality failures?
- What are the typical consequences of a service disruption or cyberincident?

Once this baseline is clear, the next consideration is whether AI materially changes this exposure.



Putting a value on prevention: A practical example from cybersecurity

IBM and the Ponemon Institute studied 600 organizations that had experienced data breaches. The study found that "Security teams using AI and automation extensively shortened their breach times by 80 days and lowered their average breach costs by USD 1.9 million compared to organizations that didn't use these solutions."[14]

Because this research is observational, it does not show that AI alone caused these differences, but it illustrates how prevention can be expressed in economic terms.

The same thinking applies to smaller events that happen much more frequently. Even a modest reduction in payment errors, defects, false claims or service failures can create substantial value when repeated across thousands or millions of transactions.

Where the event is uncertain, the estimate should reflect that uncertainty. A potential loss of $10 million is not automatically worth $10 million if there is only a small probability of it occurring. Ranges and probability-weighted scenarios provide a more accurate picture.

Protected Value therefore broadens the way leaders think about AI returns by focusing on the losses AI can help the organization reduce in addition to what it helps produce more of.

Preventing these losses requires investment. Detection systems require data and integration. Models need to be tested and monitored. People may still need to investigate alerts, review exceptions and make important decisions. These resources lead us to a crucial component of AI economics: Lifecycle Cost.

[14] IBM and Ponemon Institute. Cost of a Data Breach Report 2025: The AI Oversight Gap. 2025.

4. Lifecycle Cost

What is the true cost of AI once it moves from experimentation to production?

This is an increasingly important question as AI becomes more widely adopted, and the answer is much broader than the price of a model, license or token.

Generative AI has been adopted on a massive scale very quickly, both in people's private lives and at work. Much of that initial uptake was through free access, fixed subscriptions, enterprise licenses or AI capabilities bundled into existing software, which meant usage could increase significantly without the organization seeing a proportionate increase in its bill.

This is beginning to change. While fixed licenses remain important, they are increasingly coexisting with consumption-based pricing. For example, OpenAI now offers a token-based billing system for certain new enterprise agreements,[15] while Microsoft allows Copilot Studio customers to pay according to the credits their agents actually consume.[16]

As AI becomes more deeply embedded in business operations, two things happen: The relationship between usage and cost becomes much more visible, and the unit cost of intelligence continues to fall. Better infrastructure, more efficient inference and intense competition are making many AI capabilities cheaper to use.

Not all intelligence is priced the same

Industry analysis suggests that a frontier model designed for complex tasks can cost roughly 5 to 25 times more per token than a simpler model performing basic tasks.[17] Greater context, reasoning effort and retrieval, and repeated agent interactions can further increase the cost of completing the actual work.

Today's prices should also be treated with some caution. The market is still being shaped by extraordinary levels of investment in models, computing power and infrastructure, as well as intense competition and strategic pricing that have helped to accelerate adoption. This does not mean that every model is being sold below its underlying cost. However, it does mean that organizations should not assume that today's prices represent the long-term economics of intelligence.

This means AI can become cheaper per unit of intelligence even as total organizational spending on AI continues to rise. It also changes one of the familiar economics of software. Once traditional software has been developed, serving an additional user often involves relatively little additional cost. AI inference is different. Every additional interaction consumes computing power, and more demanding interactions consume more of it. Although scale can improve the unit economics, the marginal cost of usage does not disappear.



[15] OpenAI. ChatGPT Rate Card (Enterprise token-based pricing). Accessed August 2026.
[16] Microsoft. Billing Rates and Management, Microsoft Copilot Studio. Accessed August 2026.
[17] Bijlsma, J, Kleine, D and Scognamiglio, F. Return on AI: What CEOs Need to Know About the True Cost of Artificial Intelligence. Boston Consulting Group. July 1, 2026.

This is particularly important when AI is embedded directly into products and customer interactions. In these cases, AI usage is not just an IT expense. It can become part of the cost of serving every additional customer, thereby affecting the product's own economics and margin.

When AI becomes industrialized

Industrialization changes both the volume and intensity of AI consumption. An individual interaction with an assistant may be relatively simple, but a production process can be very different: AI may retrieve information, reason over it, interact with business systems, call external tools, validate its answer, and repeat parts of this process before reaching an outcome.

Agentic systems amplify this effect, because one request from a user can trigger many actions behind the scenes. Multiply that over employees, customers, transactions and continuous business processes, and it's clear that the economic challenge moves beyond access to AI. Organizations increasingly need to understand how much intelligence is required for each useful outcome.

The choice of intelligence is important

Not every task requires the strongest model available. Some tasks do not require AI at all. A deterministic calculation, look-up or predictable business rule, for example, can be handled more efficiently by conventional software. Routine AI tasks may work well with smaller models, while more complex analysis or reasoning may require more capable ones.

The objective is not to use the cheapest or most powerful model. Rather, it is to find the right balance between cost, capability, quality and the reliability of the outcome produced. A lower-priced model can still be more expensive if it requires more reasoning, retries or human correction to achieve the desired result.

This is where routing becomes economically relevant. Rather than automatically using frontier AI for everything, different tasks can be directed to different levels of intelligence.

The economic question is not how cheaply the organization can buy intelligence but how efficiently that intelligence can contribute to a useful outcome, and this is where Lifestyle Cost becomes much broader.



The cost beyond the model

Lifecycle Cost represents the complete cost of creating, capturing and sustaining AI value.

The model, token and license costs make up only a portion of the Lifecycle Cost. Determining the full cost starts with identifying the opportunity and the business processes that need to change. Data may need to be acquired, cleaned, engineered and governed. Unstructured content may need to be curated, structured and connected to authoritative sources.

This is where the broken organizational-intelligence layer also becomes economically visible. Fragmented data, inconsistent content, unclear processes and poorly defined rules reduce AI performance while creating additional work to prepare the intelligence on which the solution depends. The AI has to become part of the enterprise. Applications and agents need to be engineered. APIs and connectors must link them to existing systems. Retrieval and orchestration must work effectively. The economics of the solution are affected by platforms, cloud or private infrastructure, storage, network capacity and data movement.

A controlled experiment can conceal some of this effort. Experts can manually prepare data, select the right content, and compensate for weaknesses in the process or stay close to the system in case of issues. Industrialization reveals what the organization actually needs to make the solution operate reliably without exceptional support.

However, getting the technology into production is still not enough. AI requires evaluation, testing and monitoring. Security, privacy, legal and governance requirements also need to be addressed. Higher-consequence applications may require stronger controls and human oversight. People may be involved in reviewing outputs, approving actions, handling exceptions and responding to incidents.

Then there is the issue of adoption. People need to understand how the technology will change their work and how to use it effectively. Processes and roles may also need to evolve. Training, communication, user support, leadership attention and change management all consume resources. These costs are important because value is realized only when the technology becomes integral to how the organization operates.

The cost continues after deployment. Data and content must be kept up to date. Performance can deteriorate. Providers update their services and prices. Integrations and applications require maintenance. The system also needs to be improved as the organization learns from real usage. For AI, improvement is continuous. Ensuring the system remains effective as conditions and requirements change can involve changing the model, routing logic, prompts, retrieval tools, agent behavior, workflows, evaluation methods, data or content.

Scaling up also introduces further costs. Solutions may need to be replicated, localized or adapted for additional business units, countries or customer groups. Infrastructure may also need to be expanded. Regulations change. Ultimately, organizations may also encounter the expense of switching providers, migrating systems or decommissioning them entirely. The lifecycle view should also include externalities such as material, energy, water and hardware resources, where relevant.

From consumption to useful outcomes

Shifting focus from consumption to useful outcomes creates a broader perspective that changes how the cost of AI should be measured.

Token economics are important, but they are not AI economics

Knowing what a model consumed, and what a model call costs, is useful. However, this does not tell the organization how much was spent to turn intelligence into something valuable. This is why measuring cost per useful outcome is a more meaningful approach.

The relevant question here is what the complete system costs to produce an outcome at the required level of quality, reliability and control.

This includes the intelligence consumed as well as the infrastructure, data, integration, testing and controls, human effort and organizational change required to produce the outcome.

The important word is “useful.”

An interaction that is cheaper but produces more retries, escalation and rework may lead to a more expensive outcome. Similarly, a sophisticated agent that performs hundreds of actions without materially improving the result may simply consume more resources. Conversely, greater expenditure on intelligence can be economically rational when it improves the outcome or reduces costs elsewhere.

 © 2026 NTT DATA, Inc.

The cost picture becomes clearer over time

During experimentation, many assumptions remain uncertain. Production reveals actual consumption, integration effort, human intervention and support requirements. Continuous operation reveals maintenance and improvement costs. Scale shows what happens when these factors are multiplied across the organization.

The business case should evolve as this evidence replaces the original assumptions.

Lifecycle Cost asks a much broader question than how much the model, license or token costs. It asks what the organization must spend over time to create, capture and sustain the outcome that justified the AI investment in the first place.

Understanding this cost provides leaders with a much more comprehensive view of the investment. However, cost is still not the final deduction from AI value. Even after funding has been allocated for the necessary controls, assurance and human oversight, some economic exposure remains. This is the Residual Risk.

5. Residual Risk

An AI system is 99% accurate. Is it safe enough to deploy?

The answer depends less on the 99% accuracy than on what happens in the remaining 1%. A mistake in an internal summary may be easy to correct, whereas an error in a payment decision, insurance claim or critical production process could result in significant losses.

Protected Value considers what AI helps the organization avoid losing. Residual Risk considers the other side of that question: Once the AI has been deployed and the necessary controls are in place, what exposure remains?

Residual Risk is the economic exposure that remains after controls have been put in place, including risks introduced by the AI system itself.



The economics of being wrong

Risk cannot be understood through accuracy alone
Two systems with the same overall performance can have completely different economic implications because the cost of an error depends on how and where the system is used. For example, a false positive can block a legitimate transaction or prompt an investigation into the wrong case. A false negative can result in fraud, defects or security threats remaining unnoticed. A generative system can produce an answer that seems correct but is not, and a predictive model can gradually deteriorate as data, behavior or operating conditions change.

Overall accuracy can also mask significant weaknesses. A system may perform well on average while being poorly calibrated or consistently less reliable for particular types of cases or groups.

Therefore, the relevant question is not simply how often AI is wrong, but what happens when it is wrong.

Decisions with higher consequences generally justify stronger validation, tighter controls and more human oversight. While this additional protection increases Lifecycle Cost, it can reduce exposure. Even then, risk rarely falls to zero.

This relationship is particularly important as AI moves from providing information to taking action. An incorrect recommendation that a person can review creates one level of exposure. An agent that can access systems, call tools, approve steps or initiate transactions creates another. A mistake can spread through a workflow undetected. Integration failures, incorrect actions, repeated agent loops or failed escalation can magnify a small error, turning it into a much larger operational problem.

This is why production AI increasingly requires continuous testing. It also requires observability, security, human oversight and clear limits on autonomy, based on the consequences of the decisions and actions involved. These are not simply technical safeguards. They also help to determine how much risk an organization is willing to accept in exchange for the value that AI creates.



The risk extends beyond the model

Some of the most significant exposure does not originate in the model itself. AI depends on the information, systems, people and organizations around it.

Incorrect or outdated data can influence decisions

Sensitive or confidential information may be exposed. Poor governance of content can cause an assistant to rely on the wrong policy or procedure. Intellectual property and licensing rights may be violated. Malicious inputs, attacks or compromised credentials can manipulate a system or grant an agent access to actions it should never have been able to perform.

Therefore, the broken organizational-intelligence layer affects both risk and value. Fragmented data, conflicting content, unclear decision logic and weak ownership make AI harder to control, as the system is operating on an unstable organizational foundation. The objective is not to repair everything before deploying AI. Instead, the aim is to ensure that the intelligence required by the relevant business domain is sufficiently reliable, governed and accessible, given the level of risk involved.

Exposure can also come from dependencies outside the organization

Providers may experience outages, change models or services or become unavailable. An example of this is Claude Fable 5. In June 2026, US export controls temporarily forced Anthropic to suspend access to the model; access was later restored.[18] This demonstrated a broader vulnerability: As frontier AI becomes more important strategically, access to intelligence may increasingly be shaped by national interests, regulation and geopolitics.

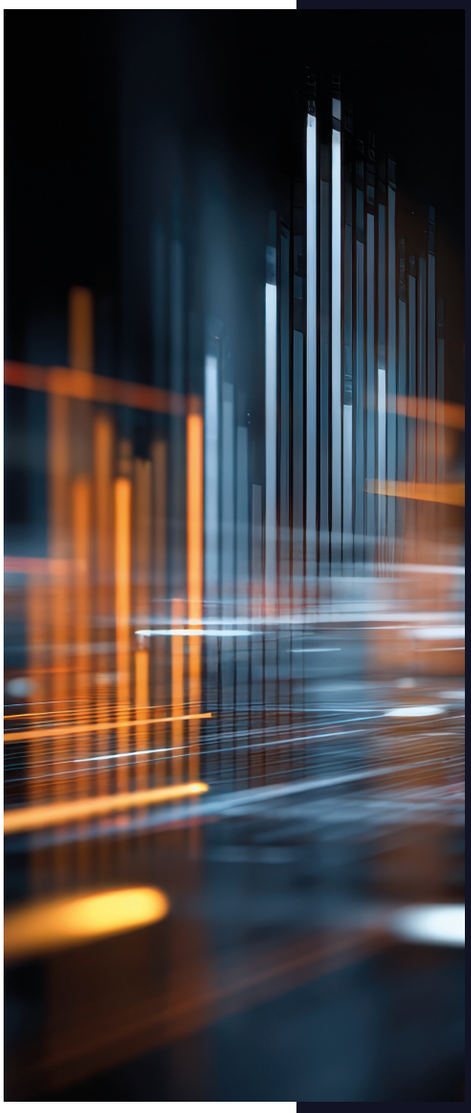
This raises the prospect of an era of AI inequality, where access to the most capable intelligence is unequal across countries, organizations and markets. In Europe, this reinforces the political and technological importance of sovereignty. For organizations, sovereignty is also an operational necessity. If critical processes become dependent on intelligence controlled elsewhere, continued access to that intelligence becomes a matter of resilience and business continuity.

[18] Anthropic. Statement on the US Government directive to suspend access to Fable 5 and Mythos 5. June 12, 2026. Anthropic. Redeploying Fable 5. June 30, 2026.

Human behavior introduces further complexity

Overreliance, weak supervision, shadow AI and the gradual loss of the skills required to challenge AI outputs when necessary can make AI adoption harder to govern. Concentrating many processes on the same models or providers can further amplify exposure by creating common points of failure.

A technical problem can then become a business problem. An incorrect action can result in financial loss, operational disruption, legal liability, regulatory intervention or a loss of trust among customers, employees, partners and the public. Current risk frameworks are increasingly recognizing this broader perspective. For example, the National Institute of Standards and Technology (NIST) treats generative AI risk as something that can emerge across the lifecycle and at model, system and ecosystem levels, rather than as a one-time model assessment.[19]



Understanding what remains

Residual Risk is ultimately about the exposure remaining after protections have done their work.
It is influenced by the likelihood of failure, the potential severity of the consequences of failure, the extent of organizational exposure and the amount of risk remaining after existing controls have been implemented.

Not every risk can or should be converted into a precise financial figure. Frequent events may provide enough historical evidence for a credible estimate, whereas less predictable risks may be better represented by ranges and scenarios. Rare events also matter when their consequences are significant — even a very low probability can be economically significant when AI is operating on a large scale or when a single failure could cause severe harm.

The objective is not to eliminate uncertainty but to understand the remaining exposure well enough to decide whether the value still justifies it.

Reducing exposure further is rarely free. More control can increase costs or slow processes down. Greater autonomy can improve speed but also increase the consequences of an error. Greater independence from one provider can improve resilience but requires additional investment and complexity.

Residual Risk therefore reveals more than what might go wrong; it also makes the underlying choices explicit.

How much additional cost and complexity is the organization willing to accept in exchange for greater control, resilience and independence?

Those choices are part of Economic Trade-offs.

[19] National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. July 26, 2024.

6. Economic Trade-offs

What if the most technically advanced AI option is not the best one economically? Every meaningful AI decision involves a trade-off. A frontier model may improve capability, but it will also increase cost and dependency. A private environment may provide greater control but require more investment and internal expertise. More automation may accelerate a process while making mistakes more consequential. A common platform may improve reuse and consistency while limiting local flexibility.

These are not just secondary implementation questions. They shape the economics of the decision itself.

Where Lifecycle Cost tells us what the chosen solution consumes and Residual Risk tells us what exposure remains, Economic Trade-offs asks a different question: What are we giving up by choosing this path over another?

Economic Trade-offs captures what the organization sacrifices, constrains or becomes dependent on when choosing one path over another.



The value of the alternative

As we saw with Lifecycle Cost, not every task requires the strongest model, and some tasks do not require AI at all. Model choice and routing therefore affect the cost of achieving a useful outcome.

Economic Trade-offs takes this idea further. The question is no longer simply which technology can do the job. Rather, it is about the economic differences that arise when one option is chosen over another.

A smaller model may reduce costs but offer less capability on some tasks. A frontier model may improve performance, but this comes at the cost of increased consumption and dependence on a particular provider. Conventional software may be more reliable for predictable tasks. Greater autonomy may increase speed, but it requires stronger controls. Human judgment may add time and cost to a process, but it remains essential when the consequences of a wrong decision are significant. There is rarely one universally superior answer.

This is why AI should ultimately be compared with the best realistic alternative — another form of AI, for example, or conventional automation, process redesign, hiring more people, outsourcing, modernizing an existing system or simply waiting until the economics or technology improve.

It's important to make these comparisons because AI can appear attractive when doing nothing is the only alternative presented. However, organizations rarely face such a simple choice. Capital, people and management attention can usually be deployed elsewhere.

This means the economically relevant question is: Does an AI investment create more net value than the most credible alternative available to the organization?

The price of control

The same logic applies to how AI is sourced, hosted and operated.

Packaged AI can provide speed and simplicity, and external models can make sophisticated capabilities available without the need to build the underlying infrastructure. On the other hand, private environments or systems operated on-premises can provide greater control over data, infrastructure and operations — bearing in mind that greater control usually necessitates increased investment, enhanced internal capabilities and greater responsibility for operating the environment.

Provider strategy creates another choice. Focusing on one provider can simplify architecture, integration, governance and commercial management. However, it can also create dependency, increase switching costs and reduce negotiating flexibility. Using several providers can preserve options and improve resilience, but this increases integration effort and operational complexity. NTT DATA research on private and sovereign AI considers these decisions in terms of cost, control, resilience, flexibility, jurisdiction and provider dependency, rather than viewing them as purely technical choices.[20]

As Residual Risk showed, dependency becomes more important as AI becomes more embedded in critical processes. The issue here is not whether dependency creates risk — we already know that it can. It's how much additional cost and complexity the organization is willing to accept to reduce that dependency.

[20] NTT DATA. 2026 Global AI Report: A Playbook for AI Leaders. December 2025.

 © 2026 NTT DATA, Inc.

The same tension arises when solutions are scaled up. Standardizing models, platforms and workflows can improve reuse, reduce duplicated effort and simplify governance. However, countries, business units and customer groups may have different regulations, languages, data and operating processes. Too much localization undermines economies of scale. Too much standardization can result in a solution that does not fit the purpose.

A similar balance exists between achieving immediate results and developing future capability. Building shared platforms, reusable data assets, integration patterns and common controls can reduce the cost and time required for future deployments (this is part of the Future Value described earlier). However, laying too much groundwork before sufficient value has been demonstrated can consume capital and delay results. Similarly, optimizing every initiative for the fastest possible return may result in the organization repeatedly rebuilding the same capabilities.

What the organization gives up

Some of the most important trade-offs have little to do with the technology itself.

AI competes for scarce resources

Engineers assigned to an AI platform cannot also work on modernizing a critical legacy system. Business experts redesigning one domain have less capacity to transform another. Executives can effectively sponsor only a limited number of major initiatives. Employees can absorb only so much transformation at a time.

This means that the opportunity cost of AI includes capital as well as talent, leadership attention and organizational capacity for change. An organization pursuing dozens of AI initiatives may find that its real constraint is its ability to absorb change rather than access to technology or funding.

What you choose to focus on has economic implications. Opting to pursue many opportunities keeps more options open, but spreads people, investment and management attention across a wider portfolio. Concentrating on a smaller number of important business domains can increase the depth of transformation and accelerate value capture but will inevitably mean postponing other opportunities.

AI can also alter the underlying business's economics

An AI-enabled feature may improve the customer experience while adding a variable cost to every interaction. Automation may reduce the cost of an existing service while eroding another revenue stream. A new AI channel may improve accessibility while altering the economics of an established distribution channel.

These effects are important because the value of an AI investment cannot always be understood by considering the solution in isolation.

Economic trade-offs provide another perspective on AI economics by making the alternatives, dependencies and opportunity costs of a decision visible.

The question is not whether the organization can build the most advanced AI solution. Nor is it whether one particular option looks attractive on its own. Rather, it's whether the chosen path creates the strongest sustainable net value relative to what the organization could realistically have done instead.

Sometimes, this will mean investing more in AI. At other times, it will mean choosing a simpler model, accepting more human involvement, building greater resilience or delaying an investment. **And sometimes, it will mean using no AI at all.**

Economic trade-offs remind us that the quality of an AI decision is measured by the option selected and the alternatives deliberately left behind.

 © 2026 NTT DATA, Inc.

AI economics is a management discipline

AI is becoming more capable at a remarkable rate, and the race to develop frontier models will continue to determine the capabilities of AI. However, greater capability does not automatically make value easier to capture. Organizations now face a greater number of possibilities and investment choices, as well as more opportunities to confuse technological progress with business performance.

One distinction has remained central throughout this guide: The capabilities of AI and the value that organizations ultimately gain from it are not the same thing. Value emerges from the system around the technology, not from technical capability alone, and the economics of that system change over time.

Production reveals the real effort required to integrate AI into workflows, maintain data and content, support users, operate controls and manage exceptions. Scaling up reveals how consumption, infrastructure, organizational capacity and dependencies behave when multiplied across the business.

As evidence replaces assumptions, an initially compelling business case may become stronger, weaker or simply different.

AI economics should therefore be treated as an ongoing management discipline rather than a one-off calculation performed before funding is approved. The objective is not to eliminate uncertainty but to understand the economics well enough to make better decisions as the situation evolves.

If access to powerful AI becomes increasingly widespread and access to models is no longer the most significant differentiator, how will some organizations consistently create more value from it than others?

This is where AI economics and AI transformation converge: AI economics helps leaders understand whether AI is creating sustainable net value. AI transformation determines whether an organization can repeatedly and at scale create that value.

Success will be measured by an organization's ability to turn technological possibility into sustainable value, which means organizations will need to identify where AI can significantly improve outcomes, then redesign relevant business domains, establish the necessary data and technology infrastructure, implement appropriate controls and enable employees to work differently.

The purpose of AI economics is not simply to measure the value of AI. It is to help organizations become better at creating value.

Visit nttdata.com to learn more.

NTT DATA is a $30+ billion business and technology services leader in AI and digital infrastructure. We accelerate client success and positively impact society through responsible innovation. As a Global Top Employer, we have experts in more than 70 countries. NTT DATA is part of NTT Group.

NTT DATA


NTT DATA