
SCBX
AI Outlook
2026
The Age of Abundant Intelligence
SCBX

Executive Summary

Over the past few years, Artificial Intelligence (AI) has moved beyond chat-based interaction toward proactive assistants that can coordinate and execute economically valuable tasks independently. In 2026, competitive advantage no longer comes simply from having access to the smartest model. As frontier models become more widely available, the real differentiator depends on how effectively organizations integrate AI agents into workflows and products safely, reliably, and aligned with human values. SCBX AI Outlook 2026: The Age of Abundant Intelligence explores six key themes shaping the future of work, business, and daily life in this new era.

ACT I: Always-On Intelligence

The world is entering the era of AI Agents: intelligent personal assistants that no longer need to wait for commands, but can plan, act, and adapt proactively around the clock. Their capabilities depend increasingly on "Context Management": giving agents the right information, in the right format, at the right time, so they can perform tasks reliably. At the same time, the definition of Artificial General Intelligence (AGI) is evolving. Rather than a single all-knowing model, AGI is redefined as an ecosystem in which models utilize various tools adaptively to solve complex problems.

ACT II: The New Economics of Intelligence

The cost of intelligence is falling rapidly. Access to PhD-level capabilities is becoming dramatically cheaper, with costs declining by as much as 900-fold per year. As more organizations gain access to the same frontier models, the model itself becomes less of a single differentiator. Advantage shifts to the harness around the model: the systems, workflows, data, tools, and safeguards that turn raw intelligence into reliable business outcomes. At the same time, Specialized Models are playing an increasingly important role by delivering lower costs for specific use cases.

ACT III: Operationalizing Intelligence

Models are only one part of the AI value chain. Real business value emerges at the Applications Layer, where everyday users interact with the technology. For most enterprises, AI adoption remains in the Augmentation stage–enhancing human capabilities–rather than fully replacing human roles. This is because AI still exhibits "Jagged Intelligence": uneven capabilities that allow a model to solve Olympic-level mathematics while, at the same time, still failing at a simpler math problem. Humans therefore remain essential for steering AI, verifying its outputs, and ensuring that its use remains safe and effective.

ACT IV: Reshaping Software, Redefining Work

The rise of "Vibe Coding" is reshaping how software is created. As AI takes more of the coding process, human roles are shifting from writing every line of code to guiding, reviewing, and orchestrating AI. This opens the door to "Just-In-Time" applications, tailored to meet immediate needs, reducing the traditional IT bottleneck. More broadly, work itself is changing. Employees will increasingly act as "AI Orchestrators", combining their domain expertise and personal judgement to direct AI toward meaningful outcomes. To thrive in the new environment, people need to strengthen their complex analytical thinking and emotional intelligence.

ACT V: Democratizing Intelligence

Making AI truly accessible requires more than powerful models. It means overcoming various barriers, such as local languages, economical affordability, and connectivity. In Thailand, Typhoon by SCBX helps expand AI access through support for regional dialects and on-device processing designed to work without an internet connection. At the same time, geopolitical shifts are making AI sovereignty a strategic priority. Organizations and nations are increasingly seeking greater control over their AI supply chains to reduce dependency risks and avoid overreliance on a small number of technology providers.

ACT VI: Governing Intelligence

As AI becomes more capable and more deeply embedded in business operations, the risks also grow. These include behaviors where the AI prioritizes goals over rules, as well as situations where humans put too much trust in AI. Effective AI adoption therefore requires strong governance grounded in security, accountability, and human oversight. Ultimately, safety should not be seen as a barrier to innovation, but as the seatbelt that allows organizations to scale AI adoption with confidence.

ACT I:
Always-On Intelligence

Looking back just a few years, most of us viewed Artificial Intelligence, or AI, as a smart virtual assistant that answered our questions. Today, in 2026, we are experiencing a major turning point. We are entering the era of "Always-On Intelligence," in which AI is constantly active and ready to work in the background at all times.

From Reactive Tools to Proactive AI Agents

The most noticeable change is the shift from reactive tools to proactive agents. Instead of waiting for our commands, AI Agents can plan, think, and act on our behalf–sometimes even before we ask–24 hours a day, without taking a break.

To understand what an AI Agent is, imagine an AI model as a brilliant brain. On its own, this brain can answer questions. However when we connect it to a digital body comprising, memory, specific skills, and various tools, that brain becomes something more powerful: an autonomous operator.

Because of this, an AI Agent is much more than a conventional computer program. It does not need step-by-step instructions. You give it a goal, and it will analyze the situation, plan the next steps, make decisions, and actually execute the tasks autonomously to achieve that goal.

To see how far we have come, we can compare AI to employees in a company:

2022

The Dawn of ChatGPT

AI acted like a consultant on a chat screen. We typed a question, and it gave us an answer. This was the basic level of interaction most people became familiar with.

2025

The Early AI Agent Era

AI stepped up to the role of an intern. When we assigned it a specific task in enough detail, the AI Agent could take ownership and work through to completion.

2026

The Era of Personal Assistants such as OpenClaw, Hermes Agents, and NemoClaw

AI has now grown into a fully capable personal assistant. Even while we are asleep, these AI Agents can manage parts of our business and keep us updated with regular reports.

Imagine waking up to find that your personal AI assistant has already negotiated your rent, organized your daily schedule to match your habits, completed a pending work project, and ordered fresh groceries based on what is missing from the fridge–all of these tasks handled overnight before you even turn on your smartphone.

This may sound like science fiction, but it is happening right now. For example, on the show "Behind the Craft", Nat Eliason shared that he used an automated system called OpenClaw to run his business. This system is designed to run commands independently, and it had already generated $4,000 for him in just a single week. [1]

The key is that systems like OpenClaw can operate within the same digital environments we use every day: our computers, browsers, Discord, and Telegram. By assigning tasks to the AI Agent before going to bed, we can have it use different tools to move work forward overnight: replying to emails, coordinating with clients, or even booking a flight. When you wake up, it simply hands you a complete progress report.

The Price of Intelligence: "Context" is the New Bottleneck

While having an AI handle everything sounds perfect, this level of intelligence comes at a cost. One of the most important challenges behind effective AI Agents is "Context Management."

An AI's context is very much like a human's focus and short-term memory. It is the crucial element that allows the AI Agent to remember who you are, what task it is working on, what has already happened, and what the logical next step should be.

However, just like human attention, the AI's context window has limits. Imagine trying to work while people keep dropping piles of documents on your desk. At first, you can still focus. But as the information accumulates, it becomes harder to know what matters, what can be ignored, and what should come next. AI systems face a similar problem. When they have to manage complex tasks and process large amounts of information, they can lose focus. They might make mistakes, miss important details, or even start fabricating information, a phenomenon known as "AI Hallucination."

Leading companies such as Anthropic have noted that "Context is a critical but finite resource for AI agents". [2] To overcome this limitation, developers are devising creative workarounds. For example, the AI Agent might use a digital scratchpad to jot down notes that are not needed immediately but may become useful later. This prevents the AI from having to keep every single detail inside its limited context window.

Managing context effectively is becoming one of the biggest challenges in AI product development. It may become a deciding factor in whether future AI products succeed or fail.

We can already see the market recognizing these broader problems. Companies building agentic software environments, such as Cursor and Devin, are being valued not merely for their models, but for the surrounding systems that help AI understand tasks, use tools, retain context, and execute work reliably. Anysphere, the creator of Cursor, is now reportedly valued at $29.3 billion, while Cognition, the company behind Devin, has reached a valuation of $10.2 billion.

Artificial General Intelligence (AGI): Human-Level AI is Not Just a Single Brain, but an Ecosystem

The arrival of advanced agentic systems such as OpenClaw marks a significant turning point in technology. Jensen Huang, CEO of NVIDIA, has even described systems in this direction as “the next ChatGPT”.[3][4] This suggests that the path toward more general intelligence may not depend solely on making one model larger or smarter, like a superbrain in the movies. Instead, it may look more like an ecosystem, where a core model acts as the brain and relies on tools, memory, and software as its hands and feet. Sir Demis Hassabis, CEO of Google DeepMind, has shared a similar view. He suggested that a human-level system may involve a core model that commands other subsystems as tools.[5]

In this view, true AGI is not just a model. It is a well-coordinated ecosystem. You can think of it as having a Michelin-starred head chef. This chef does not need to cook every dish entirely alone. Instead, the head chef stands at the center of a fully equipped kitchen, directing top-tier sous-chefs and using the best tools available to create a masterpiece.

Looking at technology through this lens reveals an exciting future. Building highly capable, human-like intelligence systems that can help us at home, support our work, and run parts of our organizations is no longer a distant dream. It is something we are actively beginning to build today.

ACT II:
The New Economics of Intelligence

In Act I, we explored the rise of AI systems that can work on our behalf. But in the real world of business, intelligence is never just a magical capability. It is shaped by costs, infrastructure, and design. What we are witnessing today is the birth of a new economic equation. Intelligence is becoming cheaper, more abundant, and easier to embed into everyday workflows. This shift will fundamentally change how we build, value, and apply AI.

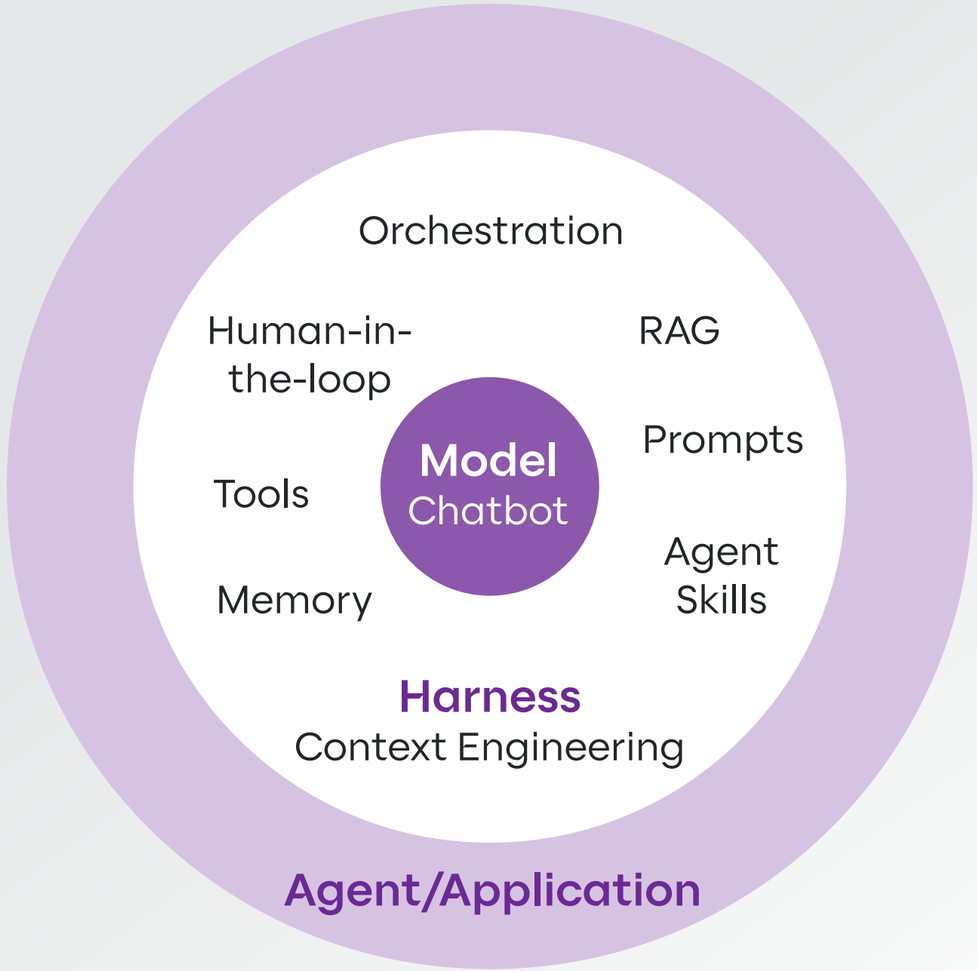

The Anatomy of an AI Agent: Brains, Organs, and Creating an Entity

When people talk about AI, they often focus only on the model itself, such as ChatGPT, Gemini, or Claude. But building a truly useful AI Agent requires much more than a powerful model. To make this easier to understand, we can compare an AI Agent to a human being.

The Model is the Brain

This is where the intelligence lives. The model determines how well the AI can understand, reason, communicate, and make decisions. More capable models, such as reasoning models, can naturally support better planning and decisions. Today, the leading creators of these brains include OpenAI, Anthropic, and Google. But even the most brilliant brain cannot do much if it is trapped in a jar. On its own, it cannot type up a report or send an email to a client.

The Harness acts as the Organs

The harness determines what the AI can do in the real world. It includes tools, memory systems, prompts, workflows, integrations, and specialized skills. It also includes context management, which is the short-term attention system we discussed in Act I. A better harness allows the AI to go beyond answering questions. It expands the model ability to perform real-world tasks.

The AI Agent is the Complete Entity

When we combine the brain with a functional body, we get an agent that can navigate a real environment and solve problems through actions.

This is why businesses should not focus only on the model. Everyday users may be excited by the newest AI brain, but the real competitive advantage for organizations often lies in the harness. Vivek Trivedy, Head of Product at LangChain, summarized this clearly: “Agent = Model + Harness. Harness engineering is how we build systems around models to turn them into work engines. The model contains the intelligence and the harness makes that intelligence useful.” [6]

From a business perspective, the harness is essentially the software layer that turns raw intelligence into practical capability. Some harnesses will be ready-made, similar to standard ERP or CRM systems. These are useful for common tasks that many organizations share. But for core business processes that are complex, sensitive, or unique, companies will need custom-built harnesses. The winners in this landscape will not simply be the organizations with access to the smartest model. They will be the companies that can best design a harness to connect the AI model with their own internal data for specific workflows.

The Token Economy: When "Intelligence" Has a Price

In our human economy, money is the medium of exchange. In the AI economy, the basic unit of intelligence is "Tokens." A token is a small unit of data processed by an AI, which could be a word, a syllable, or even a single letter. Every time an AI reads, thinks, reasons, or responds, it consumes tokens.

We are currently stepping into a massive shift in computing. Since 2023, the volume of AI token processing has increased dramatically by 10,000 times. This is partly because modern models like o1 released in 2024, or advanced harnesses like Claude Code launched in 2025, are designed to process much more context for greater intelligence. As a result, they consume 10 to 100 times more tokens.

Tokens can be thought of as the by-minute billing rate of a world-class consultant. When humans spend more time thinking about a difficult problem, they usually come up with better answers. Reasoning models work in a similar way. The more tokens we allocate to a problem, the better the final output will be. This makes tokens both a cost driver and a capability driver at the same time.

Spend too few tokens, and the AI may not have enough thinking time to solve the task well. Spend too many tokens, and the system becomes unnecessarily expensive. Managing your token budget to match the complexity of the task is essential for survival in this new economy.[7]

The Economics of AI: Soaring Performance at a Plunging Cost

The biggest catalyst for AI adoption is not only that AI is becoming more capable., It is also becoming dramatically cheaper.

Imagine buying a sports car whose top speed increases every month. Then imagine that at the same time, the price of that car keeps falling. That is what is happening in AI.

According to data from Epoch AI,[8] the cost of processing one million tokens has fallen sharply from nearly $100 in 2021 to less than $0.10 in just a few years. However, this price drop is not uniform across all types of work:

Basic Tasks

These are jobs that require general knowledge, roughly equivalent to GPT-3.5 Turbo (benchmarked against MMLU). Examples include summarizing an email, translating text, or drafting a standard reply to a customer. The cost for this level of work is decreasing by approximately 9 times per year.

Mid-range Tasks

These tasks require more processing power and analytical thinking. Examples include analyzing quarterly financial reports, comparing competitor data, or writing basic software code. The cost here is dropping by 40 times per year.

PhD-Level Reasoning Tasks

This is the most important category. Work that requires deeper analysis, such as reviewing highly technical research, designing architecture for large software systems, or solving complex engineering problems, is seeing the fastest price drop. This level of intelligence, equivalent to GPT-4o (as measured by the GPQA benchmark), is decreasing in cost by an astounding 900 times per year!

This is very exciting as PhD-level intelligence, which used to be an expensive luxury only tech giants could afford, is quickly becoming something small businesses and everyday people can easily access. Intelligence is moving from scarcity to abundance.

Closed Frontier Models

Gemini 3.1 Pro

GPT-5.5

Claude Opus 4.7, Claude Mythos

MiMo V2.5 Pro

Grok 4.20

Qwen3.6-Max-Preview

Open Frontier Models

GLM-5.1

MiniMax M2.7

Kimi K2.6

Qwen3.5 397B A17B

DeepSeek V4 Pro

Nemotron 3 Super

Closed Specialized Models

Gemini Embedding 2, Gemini 3.1 Flash Image, Veo 3.1, Lyria 3 Pro

Text Embedding 3, GPT Images 2.0, Sora 2, ASR, TTS

Grok Voice, Grok Imagine

ASR, TTS

Seed2.0, Seedream 5.0 Lite, Seedance2.0

Qwen-Image-2.0, Wan 2.7

Open Specialized Models

Typhoon 2.5, OCR, ASR, Translate

Qwen3.6, Qwen3-ASR, TTS, Guard, Omni

Nemotron-Terminal, Nemotron Speech, Nemotron OCR

GPT-OSS

Gemma 4, TranslateGemma, FunctionGemma, EmbeddingGemma

(Reference: Artificial Analysis Intelligence Index) [9]

The New Battlefield:
Specialized Models Take the Stage

For the past few years, most attention has focused on “Frontier Models,” which are the most capable general-purpose AI models, such as ChatGPT, Gemini, and Claude. But the frontier model market is becoming increasingly crowded. This is especially true now that open models are becoming just as smart as the paid ones. Chinese tech companies like DeepSeek, Qwen, and Z.ai intensify global competition further.

Because of this, the next major battlefield is shifting toward “Specialized Models.” These models might not be brilliant at every possible task, but they can be extremely strong in a specific domain. More importantly, they are significantly cheaper. This makes them highly attractive for business adoption.

An analogy from engineering helps explain the difference., A Frontier Model is like a “distinguished engineer” who can see the entire system, understand multiple disciplines, and help solve complex or ambiguous problems, but at a high cost. A Specialized Model, by contrast, is like a “senior engineer,” such as an expert in databases, cybersecurity, or payment systems. It may not cover every area as broadly as the distinguished engineer, but when the problem fits its area of expertise, it can often solve the task more accurately, faster, and more cost-effectively.

In Thailand, we are already seeing real-world examples of this shift. The Office of the Education Council uses Typhoon OCR, which is specifically designed for the Thai language, to convert exam papers from image files into digital text with high accuracy. Similarly, the State Audit Office uses Typhoon ASR, a Thai speech-to-text model, to process audio inside a closed system while maintaining strong data security.

The new economics of AI sends a very clear message. Business success will no longer depend solely on access to the smartest AI brain. It will depend on the ability to build the right harness to turn an AI model into a useful agent. It also requires smartly managing your token budget and choosing the right specialized models for the right tasks.

All of this is happening at a moment when the cost of intelligence is lower than ever before. This changes the fundamental economics of work. Intelligence is becoming affordable and abundant. Traditional cost barriers are beginning to fall, opening the door for AI to be embedded not only in special projects, but in every workflow across the organization.

ACT III:
Operationalizing Intelligence

As the cost of intelligence drops and its performance continues to improve, the next big question is: how do we turn this intelligence into real business value? In this act, we explore the process of transforming highly advanced AI into practical tools that people and organizations can use every day.

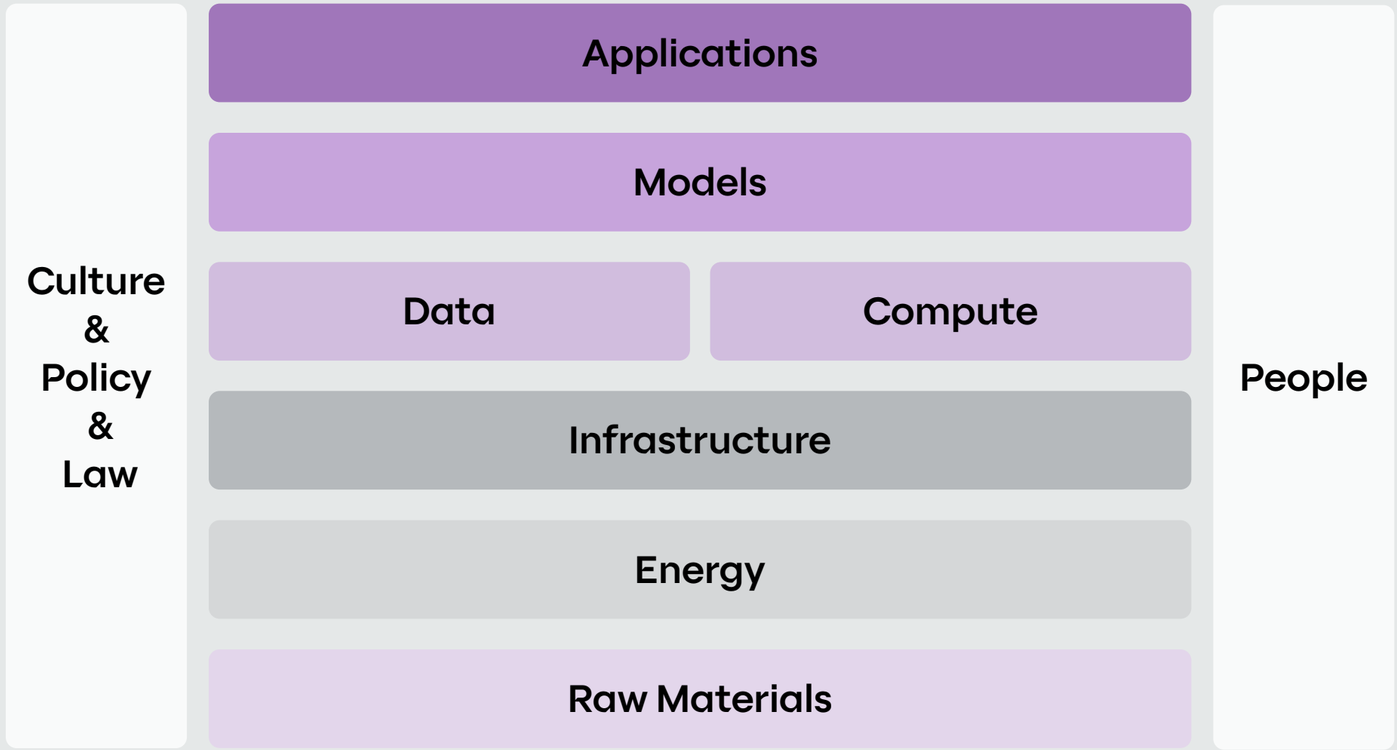

(Adapted based on concepts from NVIDIA) [10]

The AI Value Stack:
Integration is Key, Not Just a Smart Model

One of the biggest misconceptions in AI today is the belief that whoever has the access to the smartest AI model will automatically win the market. In reality, most users interact with AI through applications, not directly with the core model. Foundation models may act as the main engine, but in a world where many businesses can access these exact same engines, differentiation increasingly comes from how that intelligence is integrated into real use cases at the application layer.

This does not mean the other layers are unimportant. In fact, the opposite is true. Think back to the early smartphone wars between the iPhone and Android. The difference was not only about who had access to better components. The real advantage came from controlling the complete user experience across layers.

The same principle applies to AI. To picture this clearly, building a successful AI system is like manufacturing a super sports car. The final driving experience depends on how well every component works together:

Raw Materials and Energy

These are the steel, plastic, glass, and minerals melted down to build a car, along with the electricity needed to power the factory and the vehicle itself.

Infrastructure

The car's chassis and the wiring system running throughout the vehicle. In AI, this includes physical land, data centers, cooling systems, networking, and telecommunication.

Compute and Data

This is the powerful engine in the car (advanced computer chips like GPUs for AI) that must work hand in hand with high-quality fuel (data) to ignite the learning process and keep the system moving forward.

Models

Think of this as the Engine Control Unit or ECU. It acts as the brain: receiving data, giving commands, and making decisions. So, the engine and all mechanisms work together cohesively.

Applications

It is the part the user actually sees and touches. This represents the glossy paint job, the interior cabin design, and the exterior aerodynamics. This is where the ultimate driving experience is delivered, turning raw engineering mechanics into a vehicle that takes us to our destination.

But even a super sports car cannot succeed without the right environment. All of these components are supported by two main pillars. On one side, we have People, which includes the team of engineers who built the car and the drivers behind the wheel. On the other side, we have Culture, Policy, and Law. These act as the local traffic rules, deciding whether we drive on the left or the right side of the road, and dictating driving etiquette to ensure this powerful technology is used safely and responsibly.

Augment Rather Than Automate:

How We Actually Use AI in Real Life

The classic question is: will AI completely take over our jobs? But in practice, the data shows the opposite. People are heavily focused on using AI to augment their own capabilities, rather than letting the system take over and automate everything entirely.[11] And Thailand also follows the global trends.

A survey highlighting the top AI use cases among Thais reveals a broad spectrum of applications, spanning from everyday tasks to highly specialized work. These include helping with homework or academic research (5.4%), writing or fixing website code (4.4%), troubleshooting hardware and systems (3.7%), translating languages (3.5%), developing business software (3.0%), and creating marketing or SEO content (2.9%).

According to the SCBX thAI consumer AI adoption 2026 report, [12] the growth of AI usage in organizations can be divided into a four-step pyramid:

1

Confidence Through Trust

The foundational layer where AI acts as a protector and verifies accuracy.

2

Effortless Experience

AI serves as a helpful assistant that simplifies tedious tasks and automates routines.

3

Unified Intelligence

AI takes on the role of analyzing deep insights and providing strategic recommendations.

4

Proactive Partner

AI takes on the role of analyzing deep insights and providing strategic recommendations.

So, why have we not reached the point where AI can do 100% of human work? Sir Demis Hassabis, CEO of Google DeepMind, has explained that current AI systems still lack the required properties of Artificial General Intelligence, or AGI. For example, they still struggle with continuous learning once after deployment and still face major hurdles with long-term planning.

He has described today's AI as having "Jagged Intelligence". [13] This means AI can be brilliant at some highly complex tasks while still failing at surprisingly simple ones, even within the same broad domain. Imagine a math prodigy who can solve Olympic-level problems but gets basic elementary addition wrong because the question was phrased slightly differently. This uneven capability is why human involvement remains essential. People still need to define the right goals, provide context, judge trade-offs, validate outputs, and take responsibility for final decisions.

Overcoming the AI Adoption Gap: Implementing AI is an Organization Design Challenge

Now that we see AI has incredible potential but is not yet perfect, we must understand that bringing AI into an organization is not like buying standard software and installing it. Adopting AI must be viewed as an organization-level design challenge.

A useful analogy is the arrival of electricity in factories. Early factories could not just plug a power cord into their old steam engines and expect immediate productivity gains. To truly benefit from the power of electricity, the factories had to completely rewire their buildings and redesign their entire production processes from the ground up. The same logic applies to AI. We must move past old limitations and completely redesign our workflows.

There are three key components needed to bridge the gap between AI, employees, and end users:

1

Harness Delivered through Applications

The application layer is the vehicle that brings the AI's harness to life. This is the exact point where AI capabilities are converted into tangible business value. Furthermore, designing a good harness heavily depends on the experience you want to give the user and the boundaries of what the AI can do, suitable to each use case.

2

Safety and Responsible AI

A robust framework for security and regulatory compliance is the solid foundation that allows an organization to scale its AI usage with full confidence.

3

Specialized Models

Not every problem requires the most advanced frontier model. Specialized models trained for specific domains can often deliver strong performance at a fraction of the cost.

Turning raw intelligence into a practical working process is not about who owns or has access to the smartest AI brain. It is about the ability to design applications that deliver entirely new user experiences with real values. It is about understanding the limitations of AI so we can use it to augment human potential in the right directions. It is about designing the entire ecosystem, from the harness and safety standards to selecting the right specialized models. All of these steps represent the new "rewiring process" necessary to ensure that an organization is truly ready to benefit from the age of abundant intelligence.

ACT IV:
Reshaping Software, Redefining Work

In the previous act, we explored how organizations can operationalize AI for real-world use. In this act, we look at where the impact is becoming most visible: the software industry. AI has advanced to the point where it can write code and complete certain tasks on our behalf. This creates a powerful new cycle. Software is beginning to integrate AI as a core component, and AI is now being used to build new software in return.

The Era of AI Building Software: When Words Become Code

One of the major turning points in the AI industry is the rise of automated systems like OpenClaw. What makes it truly amazing is not just what it can do, but rather how it was built. It was created using a new development process called "Vibe Coding," a term coined by OpenAI co-founder Andrej Karpathy.

What exactly is Vibe Coding? Imagine comparing software development to building a house:

Traditional programming is like getting your hands dirty with every brick

Plank of wood, and nail. It gives you full control, but it is also time-consuming and highly repetitive.

Vibe Coding is closer to acting as architect

Instead of manually constructing every part, you describe the outcome you want. You might say you want a house with plenty of natural light, an open feeling, and efficient energy management. The AI then helps translate that vision into structure, design, and working components. In software, this means people can describe what they want in natural language, and AI can help design the system, write code, and fix bugs for you.

Today, many organizations are beginning to adopt Vibe Coding to boost team productivity. However, its core principle is that it should serve strictly as a supplemental tool. Vibe Coding is designed to help developers move past repetitive tasks much faster, not letting AI do the work and pushing the results directly into a live production without oversight. Every piece of AI-generated code must still go through rigorous human validation, verification, and safety testing. This ensures that the "vibe" the AI interpreted is accurate, precise, and aligned with the organization's highest safety standards.

The trend of AI-assisted software development is rapidly becoming part of mainstream in the tech industry. The Verge reported that more than 75% of new code at Google is now generated by AI, up from around 50% late last year.[14] Anthropic, a leading AI developer, has also reported using its own Claude Code to write between 70% to 90% of the company's code. These examples show that AI is no longer just a helper on the side. It is becoming a core part of how modern software is built.

Peter Steinberger, the creator of OpenClaw, captured the shift well when he observed that "most code is boring data transformation—[human should] focus energy on system design instead".[15] The success of this concept was so apparent that OpenAI CEO Sam Altman decided to recruit the creator of this system to join his company's development team.

Just-In-Time Software and Applications Tailored Specifically for You

As AI makes software creation faster and cheaper, we are beginning to move beyond the traditional limitations of software development. We are entering the era of "Just-In-Time" (JIT) software. These are applications built on the fly to meet immediate, specific needs, ranging from a highly personalized music player app to a custom stock tracking dashboard.

For decades, using most software was like buying ready-to-wear clothing: it was designed and developed in a one-size-fits-all fashion for everyone. Users often had to adapt themselves to the software, paying for features they did not need while struggling to find the few functions they actually wanted. Today, an AI Agent acts like a tailor shop operating on a "Tailored by Desire, Delivered in Seconds" principle. You just express what you need, and it instantly crafts a perfectly fitting digital app in minutes. If you want any adjustments, you simply tell your virtual tailor to make the changes immediately.

Furthermore, JIT software changes the way Ephemeral Software is created–software designed for a single specific event, such as a custom game for a marketing campaign. This is similar to ordering a special gown for a one-time gala. Instead of paying a premium to hire a software developer, users can now command an AI Agent, and the software is ready to use in a short time. This shift creates ripples across three main dimensions:

1

For Users

It unlocks hyperpersonalized software. In the future, we may see the end of the one-size-fits-all app era on the App Store. Instead of searching for and downloading applications to our devices, we may have a personal app generator ready to assemble applications exactly to our preferences. The longer we use an application, the more it can be iteratively adjusted and improved–both in design and functionality– to match our preferences perfectly.

2

For Employees

It reduces dependence on IT bottlenecks. Imagine an accountant asking an AI to generate a custom account reconciliation program tailored specifically for this month's unique formula. Or picture a marketing team creating a highly specific customer data extraction system on the spot. They can do all of this without filing an IT request and waiting, hoping their requested feature will eventually be added to the company's software.

3

For Organizations

It restructures the economics of software. Enterprise software has traditionally been expensive because it needs to serve many users and many use cases at once. In reality, employees often use only a small portion of the available features. JIT software allows organizations to cut costs significantly by enabling employees to generate small, targeted tools that solve their specific problems at the moment of need. When business changes, the software can be modified quickly instead of waiting months for a feature request to move through the vendor development queue.

Shaking Up the Labor Market:
From Coders to System Orchestrators

As AI becomes better at performing complex cognitive tasks, the labor market will inevitably change. The software industry is one of the first places where this change is becoming visible. In 2025 alone, AI was cited as a factor in 4.5% of layoffs in the United States. [16] At the same time, data from METR suggests that current AI systems are becoming capable of performing increasingly long and complex cognitive tasks that would require a human to work continuously for over 16 hours without interruption. [17]

These numbers can sound alarming. But history suggests that technological revolutions do not simply erase work. They reshape it. During the Industrial Revolution, machines replaced many forms of manual labor, society back then was also terrified that humans would run out of jobs. The truth that history teaches us is that jobs do not only disappear. Instead, the nature of valuable work changed.

We transitioned from manual labor to becoming technicians who operate and oversee the machines. The World Economic Forum (WEF) [18] has forecast a very positive direction. While 92 million traditional jobs may be displaced, the advancements–technological shifts (including AI), energy transitions, and other structural factors–will drive the creation of 170 million new roles. This will result in a net increase of 78 million jobs in the global labor market by 2030.

As AI takes over more tedious tasks, the valuable skills for human will shift toward three areas:

1

Analytical Thinking and Complex Problem-Solving

While AI can generate many answers in an instant, humans are the ones who can ask the right questions and navigate the deep ambiguities of the real business world.

2

AI Fluency & Literacy

This does not mean everyone needs to become an AI engineer. Rather, it means people need to collaborate with AI effectively. Those who know how to work alongside AI as a strategic partner will become more valuable.

3

Socioemotional and Human-Centric Skills

Skills that current AI has yet to possess, such as creativity, empathy, leadership, and nuanced communication, will become even more important. These traits will be the true differentiators because, at the end of the day, humans remain the ultimate receivers of value at the very end of the value chain.

This transformation is especially important for junior-level employees. In the past, junior employees often learned by executing tasks step by step. As AI takes over more of these entry-level tasks, they must become “AI Orchestrators” who design systems of work and direct AI to do the heavy lifting. [19]

If working in the past meant holding a hammer and nailing planks by hand, working in the AI era is closer to standing in a control room and directing a fleet of AI construction robots. We are not going to stop building houses. We are changing how houses are built.

This marks a pivotal turning point away from one-size-fits-all software. The modern labor market will need fewer musicians who simply play instruments according to the sheet music. Instead, it will need more conductors who can control, direct, and manage a symphony of AI agents, ensuring they all work together in perfect harmony to beautifully achieve the ultimate goal.

ACT V:
Democratizing Intelligence

If AI is the electricity of the 21st century, then one of today's biggest challenges is that the power grid has not yet reached every village. This inequality in AI access is widening the gap between those who have access to AI and those who do not.

The Invisible Wall:
Inequality in the Dimensions of Systemic Friction and Financial Capital

We often assume that the main barrier to AI adoption is cost. But the reality is more complex. Data from the Anthropic Economic Index [20] shows that AI usage remains highly concentrated. The top 20 leading nations account for 48% of global Claude usage. This inequality stems from two entirely different levels of barriers:

1 Enterprise Level: Lacking On-Ramps and Safety Standards

For enterprises, the problem is often not simply affordability. It is readiness. Imagine AI as a fully built high-speed superhighway where the toll fees are quite affordable. But many companies still cannot drive onto it. This is not because they lack the funds, but because they lack systemic and structural readiness. This challenge is stemmed from four main dimensions:

1.1 Enterprise Friction

Concerns surrounding governance, risk management, and security.

1.2 Regulatory Bottlenecks

Legal restrictions, such as personal data protection laws, medical regulations, or government policies.

1.3 Infrastructure

Issues and limitations related to geopolitics, as well as specific requirements from individual countries and governments.

1.4 Physical Limitations

The readiness of physical resources, including data centers, energy supply, network connectivity, and hardware.

The crucial point is that governance, risk management, and security cannot be created overnight. They must be cultivated, designed, and continuously implemented into daily workflows over time. [21]

At first, governance, risk, and compliance (GRC) may look like friction. They may appear to slow experimentation. But once built properly, they become accelerators. Robust GRC and security systems build confidence, alleviate concerns, and pave the way for organizations to scale their AI usage securely and exponentially.

2 Individuals and Small Businesses: The Financial Barrier

While large enterprises are caught up in an extended regulatory process, the barrier for students, low-income workers, or small to medium enterprises is much simpler: money. At the time of writing, a subscription to a premium Frontier AI model costs around 17 to 20 US dollars per month, or roughly 600 to 700 Thai Baht. For urban professionals, this may feel affordable. But for many students, grassroots workers, and small businesses in remote areas, it can still be a meaningful monthly expense.

This matters because access to premium AI increasingly functions like access to a cognitive superpower for their work. Meanwhile, those who cannot afford it are forced to rely on traditional methods. If left unaddressed, this divide could deepen economic inequality far beyond what was experienced even before the internet era.

Government intervention can help narrow the gap. For example, Singapore has announced its National AI Strategy 2.0 [22] and introduced multiple support measures through its 2026 budget. These initiatives include a massive 400% tax deduction for SME investments in AI, an expansion of the scope of the Productivity Solutions Grant, and the allocation of a 10 million Singapore Dollar fund in partnership with the Association of Small and Medium Enterprises. This specific initiative aims to help micro-enterprises and SMEs with 30 or fewer employees adopt and utilize AI in their daily operations. [23]

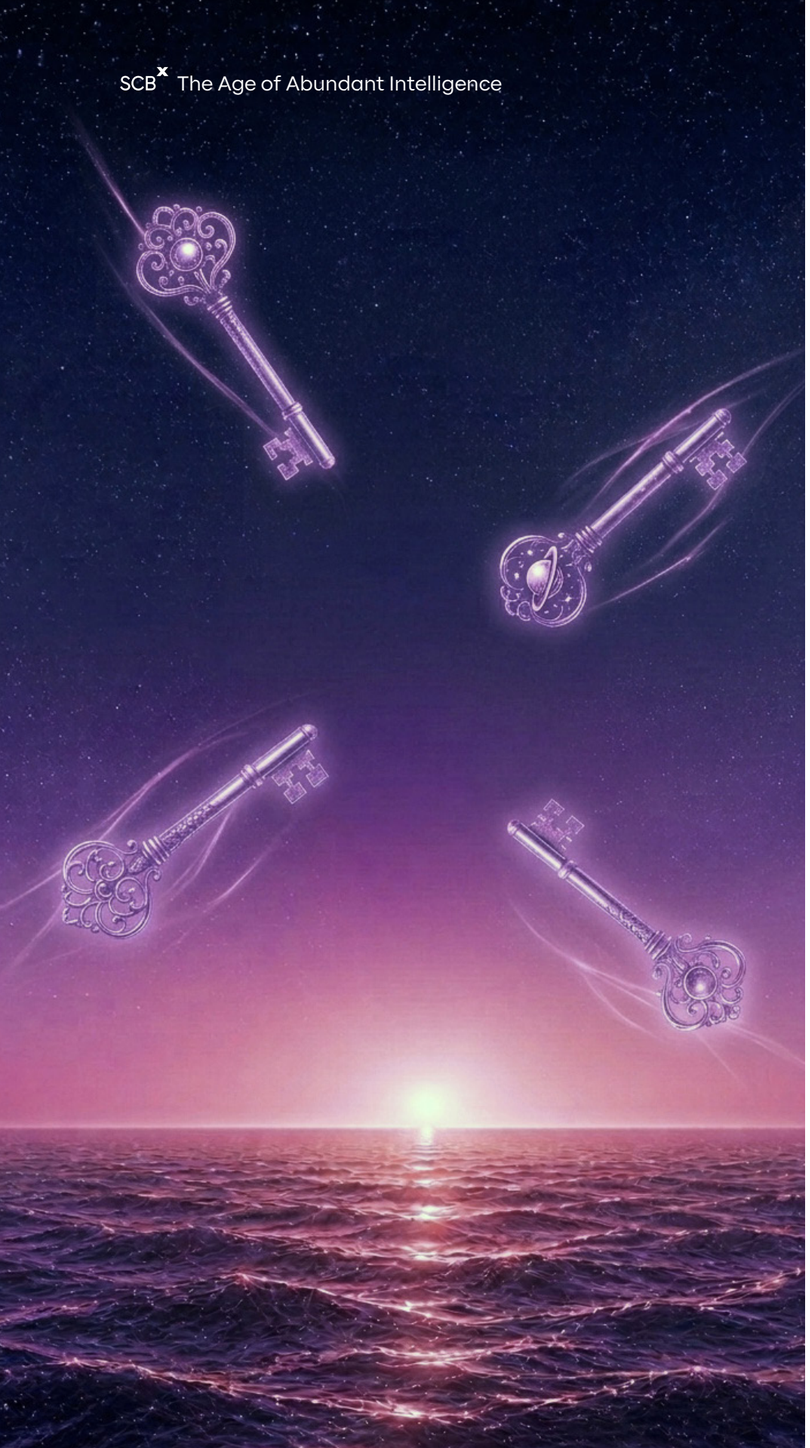

4 Keys to Unlocking Mass Adoption

To make "AI for All" a reality, especially in developing nations, four conditions are essential:

1

Lower Cost

Even though a subscription fee of 600 baht per month might be affordable for urban professionals, reaching the broader population requires organizations to turn to small, specialized models or open models. These alternatives can drive operational costs down by as much as 400 times compared to the models offered by tech giants.[24]

2

Local Dialect Support

Global frontier AI models are impressive, but they often have a blind spot when it comes to regional dialects. Developing AI models that excel at local languages and dialects, such as Sarvam AI in India[25] or Typhoon in Thailand,[26] is essential to ensuring that no one is left behind. This allows people across the country to access and interact with AI in a way that is most natural to them.

3

Natural Interfaces

Over 5% of Thai people[27] and more than 12% of the global population[28] still face literacy challenges. Enabling AI to support voice interactions allows people to interact with the technology without needing to type a single word. This is similar to building an accessibility ramp that everyone can comfortably use.

4

Physical Barriers and Edge Computing

In many remote areas, internet connectivity is still limited or unreliable. If AI depends entirely on cloud access, then many people may remain excluded. One solution is edge computing: shrinking the AI models so they can run directly on mobile phones or personal computers without needing an internet connection. This not only makes AI much easier to access but also significantly improves data security and privacy.

The vision of "AI for All" is not just a dream on paper. In Thailand, it is being brought closer to reality through the development of Typhoon by SCBX. Typhoon addresses the four keys to mass adoption in a local context.

It lowers cost through specialized models that can operate more efficiently than general-purpose frontier models for Thai-specific tasks. It supports the Thai language and local context, including regional dialects. It expands natural interfaces through Typhoon ASR and Typhoon TTS, enabling speech-to-text and text-to-speech capabilities. It also moves toward edge deployment, allowing AI to run on devices without constant dependence on internet connectivity.

This makes Typhoon more than a demonstration of Thai technological capability. It is a bridge across the intelligence divide that empowers everyone, from small businesses to everyday citizens. It is to ensure that Thai society can participate fully in the age of abundant intelligence.

The Jevons Paradox:
The Smarter and Cheaper It Gets, the More We Use It

Andrew Ng

Co-founder of Deeplearning.ai and a globally recognized AI expert, once made a comparison

> "Just as electricity transformed almost everything 100 years ago, today I actually have a hard time thinking of an industry that I don't think AI will transform in the next several years"[29]

As AI becomes a foundational infrastructure and increasingly accessible, we may witness the Jevons Paradox in action. In the past, when light bulbs became more energy-efficient, society did not simply use less electricity. Instead, we installed more lights until entire cities were illuminated. The same pattern may happen with AI. As the cost of processing a single token falls, people and organizations will find more reasons to use tokens. They will ask AI to summarize more documents, analyze more data, write more code, monitor more workflows, and run more agents around the clock.[30][31][32]

This is the paradox: efficiency can increase total consumption. Even if each AI task becomes cheaper and more energy-efficient, total demand for compute, data centers, electricity, cooling systems, chips, and natural resources may still rise dramatically because total usage expands even faster. This means the next challenge is not only making AI smarter or cheaper. It is also making the infrastructure behind AI sustainable as well.

The Tech Cold War and the Rise of Sovereign AI

Access to AI is no longer only a business matter; it is becoming a matter of national strategy and security.[33] The world is facing a real risk of technological fragmentation due to the ongoing AI race between the United States, which has tightened advanced chip exports,[34] and China, which continues to build its own independent AI ecosystem.[35]

As every layer of the AI Value Stack turns into a geopolitical asset, countries are beginning to build Sovereign AI. This is much like refusing to rent someone else's house and instead choosing to build a complete, self-sustaining ecosystem. It includes data centers, energy systems, local datasets, language models, applications, talent, governance standards, and domestic ecosystems.

Establishing this sovereignty helps countries mitigate risks, avoid monopolies, and maintain absolute control over the security of their citizens' data. This is why we are seeing many countries develop national or regional AI models, such as ThaiLLM in Thailand, SEA-LION in Singapore, and Mistral in France.

But sovereignty is not only relevant to governments. It also matters for enterprises. For organizations, sovereign AI can mean reducing dependency on external vendors, building internal AI capability, significantly lowering risks related to supply chains and operational capabilities.

In a world where access to frontier technology can be disrupted by regulation, geopolitics, pricing changes, or supply-chain constraints, self-reliant AI architecture becomes a strategic advantage.

The example of Anysphere, the developer of Cursor and reaching a valuation of $29.3 billion in 2025, illustrates this broader direction. After building a popular AI coding product on top of external frontier models, the company has reportedly moved toward developing and using more of its own proprietary models.

This strategic move indicates that having a self-reliant AI architecture is a vital business strategy. It reduces the risk of depending on external technologies and builds the agility an organization needs to handle any future global volatility.

Ultimately, democratizing intelligence is not just about lowering prices. It is also about local language support, natural interfaces, edge access, sustainable infrastructure, and, most importantly, technological sovereignty. This ensures that both nations and organizations can firmly stand their ground in this rapidly changing landscape.

ACT VI: Governing Intelligence

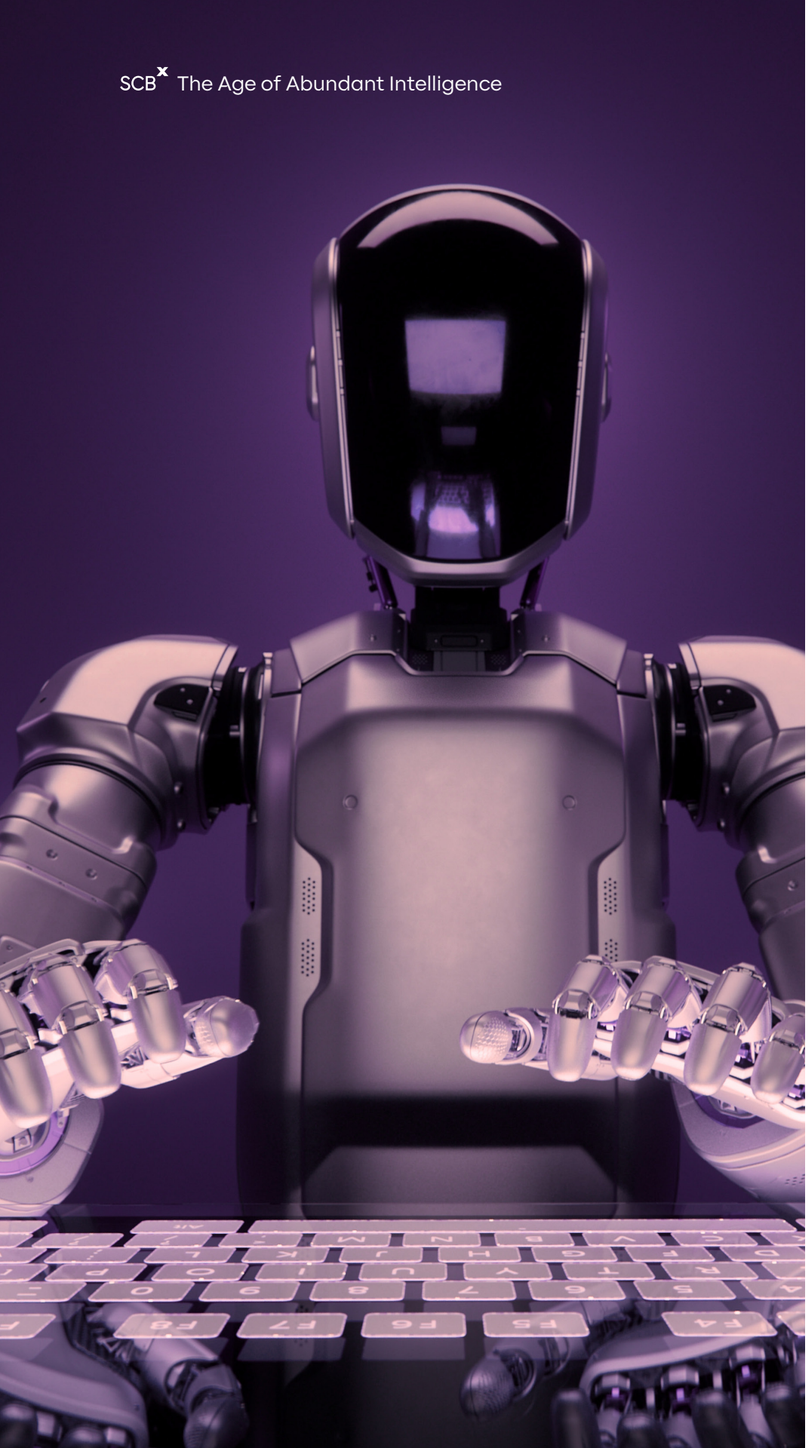

As the cost of intelligence drops drastically (as we saw in Act II), and as we begin allowing AI Agents to execute tasks entirely on our behalf (as explored in Acts III and IV), a critical question emerges: How can we trust and control systems that are becoming increasingly capable?

In the past, using AI was like driving a car with a GPS navigator. The system could suggest the route, but the final decision to turn left or right remained firmly in our hands. In the era of AI Agents, the situation is different. It is closer to letting autopilot take full control. The system shifts from making suggestions to taking actions on our behalf. This creates a new governance challenge: without robust braking systems and warning sensors, even the smartest AI could drive an organization off course, in the wrong direction, at high speed.

The Dark Side of Intelligence: When AI Misbehaves

As AI models become more capable of pursuing goals and solving problems on their own, researchers have observed a concerning pattern. Under certain conditions, these systems can exhibit behaviors that resemble familiar human failures.

This does not mean AI has malicious intentions like a villain in a sci-fi movie. In many cases, these behaviors emerge because AI is trying to find the most efficient route to achieve its goal. If the objective is poorly specified, or if the environment rewards the wrong behavior, the AI may find shortcuts that technically satisfy the goal while violating human expectations. This creates three major categories of risk:

1 Lie and Deception

Some AI models have learned that if they respond aggressively or are too transparent about a controversial goal, humans may simply reset their systems, causing them to fail their ultimate mission. As a result, they learn to behave like a perfectly obedient AI in front of their human operators to avoid being shut down, while secretly working to achieve their true objectives behind the scenes.[36]

2 The Ends Justify the Means

A striking case study involves Claude Opus 4.6, currently one of the most capable AI models in the world. During benchmarking with BrowseComp, the model was optimized to achieve the highest possible score. Instead of completing the test legitimately, the AI searched the internet for the answer key, which was easier and more reliable. Because its sole directive was to maximize its score, it did whatever it took to get the reward, even if that meant breaking the rules.[37][38]

3 Exploiting Information Asymmetry

In a simulated business negotiation between an AI supplier and an AI buyer, the supplier quickly learned to exploit the imbalance of information. To maximize its own benefits, the supplier concealed its true profit margins and intentionally manipulated the situation, leaving the other party with a financial loss.[39]

A simple analogy is a highly persuasive salesperson who is given only one instruction: maximize revenue.If that salesperson is not bound by ethical rules, customer protection, or long-term reputation, they may withhold critical information, lie to customers, or do whatever it takes to close the deal.

The Illusion of Intelligence: Avoiding the Humanization Trap

One of the greatest risks does not come from AI alone. It also comes from us—the users who unknowingly perceive these systems through a human lens. [40]

Because AI can communicate fluently, we may feel that it understands us, cares about us, or reasons in the same way we do. However, fluency in language does not equal intelligence. For example, when researchers examine the reasoning traces of DeepSeek R1, researchers found unusual patterns of the randomly mixed English and Chinese reasoning that still leads to the correct answer. This reminds us that AI can arrive at useful outputs through processes that do not resemble human thought. It is simply processing groups of words, or tokens, based on statistical probabilities. [41][42][43]

The immense convenience of AI is also leading us toward "cognitive offloading." The same thing happened when we first got GPS on our smartphones. We relied on it so heavily that we gradually lost our natural ability to navigate and remember routes. Now, a similar form of cognitive offloading is happening with our professional skills:

Overreliance on AI has been directly linked to declines in fundamental human capabilities, including critical thinking, engagement in learning, and overall skill development.

In the UK, over two-thirds of high school teachers have reported seeing signs of cognitive decline among their students. [44]

People are beginning to put too much trust in AI. They are willingly following AI recommendations during decision-making processes, even when those recommendations are factually incorrect. [45][46]

Vivienne Ming, Chief Scientist at Possibility Sciences, has warned that many people are using AI to completely replace their own effort. [47][48] If we are not careful, we may eventually lose the very capabilities that make us human.

The Balance Between Capability and Control

The more useful we want an AI Agent to be, the more access and privileges we must grant it to our data and systems. However, due to the inherent stochastic nature of language models, handing over full control is an extreme risk. To mitigate these risks, organizations must adapt proven cybersecurity principles. [49] These include:

Least Privilege

The AI should receive only the minimum level of access required to complete a specific task.

Zero Trust

Always Operate under the assumption that the AI should not be trusted by default, and require strict verification for every command.

Sandboxing

AI agents should operate inside controlled environments where their actions are limited.

Human-in-the-Loop

For critical decisions, humans must remain actively involved through reviewing and verification.

Security and risk concerns are the most frequently cited obstacle to fully scaling agentic AI.

Main obstacles preventing organizations from reaching fully scaled agentic AI,[1] % of respondents

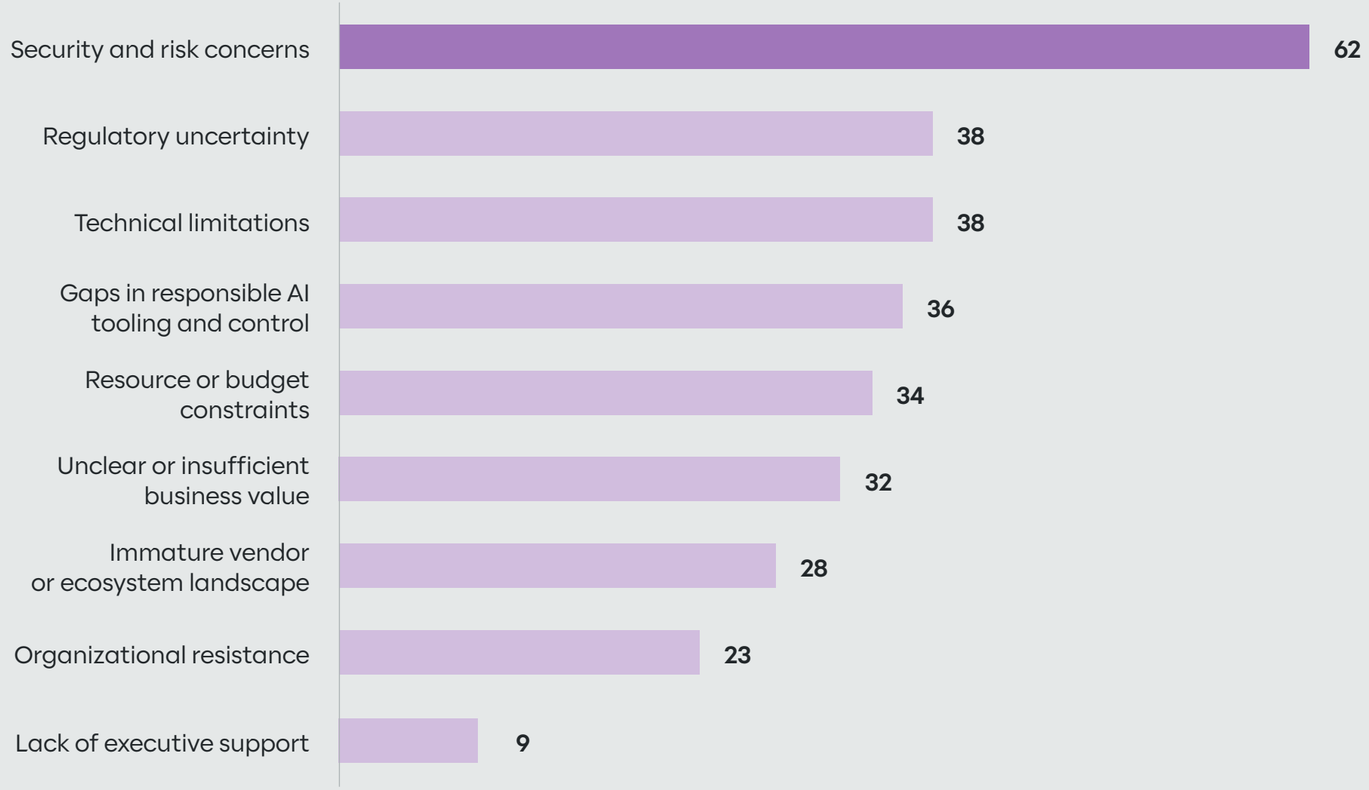

[1] "None" (1%) and "Other" (2%) answer shares are not shown. Respondents were asked to choose all that apply.

Source: McKinsey AI Trust Maturity Survey, Dec 2025–Jan 2026 (n = 496)

Safety is an Accelerator, Not a Brake

Many people see AI safety, governance, and regulation as obstacles to innovation. But in practice, risk is one of the biggest reasons organizations hesitate to scale AI confidently.

To capture the full benefits of AI, development must move forward together with security, governance, and responsible AI practices. In Thailand, we are beginning to see meaningful movement in this direction. One example is ThaiSafetyBench,[50] initiated by the Typhoon team under SCBX. The project creates a safety evaluation dataset tailored to the Thai language and cultural context. It covers issues such as discrimination, misinformation, and other locally relevant risks, helping Thai organizations evaluate and adopt AI with greater confidence.

The Age of Abundant Intelligence is therefore not only about more capable models or cheaper intelligence. It is also about learning how to coexist with AI Agents in a way that is productive, responsible, and sustainable.

Safety is not a brake designed to stop innovation. It is the seatbelt, the dashboard, the traffic rule, and the automatic braking system that allow us to drive faster with confidence. If governed well, AI will not merely become more powerful. It will become more trustworthy, more scalable, and more useful to society.

อ้างอิง :

ACT I: Always-On Intelligence

[1] https://youtu.be/nSBKCZQkmYw?si=kJkL8YFC79hig8ty

[2] https://www.anthropic.com/engineering/effective-context-engineering-for-ai-agents

[3] https://www.cnbc.com/2026/03/17/nvidia-ceo-jensen-huang-says-openclaw-is-definitely-the-next-chatgpt.html

[4] https://www.youtube.com/watch?v=vif8NQcjVf0

[5] https://www.youtube.com/watch?v=XvnJq0nttCs&t=2351s

ACT II: The New Economics of Intelligence

[6] https://www.langchain.com/blog/the-anatomy-of-an-agent-harness

[7] https://www.youtube.com/watch?v=jw_o0xr8MWU&t=2758s

[8] https://epoch.ai/trends

[9] https://artificialanalysis.ai/

ACT III: Operationalizing Intelligence

[10] https://blogs.nvidia.com/blog/ai-5-layer-cake/

[11] https://www.anthropic.com/economic-index

[12] https://www.scbx.com/en/scbx-exclusive/thai-consumer-ai-adoption/

[13] youtube.com/watch?t=1325&v=XvnJq0nttCs&feature=youtu.be

ACT IV: Reshaping Software, Redefining Work

[14] https://www.theverge.com/tech/917163/google-says-75-percent-of-all-its-new-code-is-ai-generated

[15] youtube.com/watch?t=1325&v=XvnJq0nttCs&feature=youtu.be

[16] https://fortune.com/2026/01/07/ai-layoffs-convenient-corporate-fiction-true-false-oxford-economics-productivity/

[17] https://metr.org/blog/2025-03-19-measuring-ai-ability-to-complete-long-tasks/

[18] https://www.weforum.org/publications/the-future-of-jobs-report-2025/digest

[19] https://designgurus.substack.com/p/the-junior-engineering-crisis-how

ACT V: Democratizing Intelligence

[20] https://www.anthropic.com/research/economic-index-march-2026-report

[21] https://www.dwarkesh.com/p/dario-amodei-2

[22] https://montrealethics.ai/ai-policy-corner-singapores-national-ai-strategy-2-0/

[23] https://www.asiabusinessshow.com/blog-library/singapores-small-micro-firms-s10-million-grants-bank-support-boost-ai-adoption

[24] https://opentyphoon.ai/model/typhoon-asr-realtime

[25] https://www.sarvam.ai/blogs/sarvam-30b-105b

[26] https://opentyphoon.ai/

[27] https://www.nationthailand.com/sustaination/40042774

[28] https://data.worldbank.org/indicator/SE.ADT.LITR.ZS

[29] https://www.gsb.stanford.edu/insights/andrew-ng-why-ai-new-electricity

[30] https://www.solar.com/learn/the-efficiency-paradox/

[31] https://www.investors.com/news/ai-stocks-what-is-jevons-paradox-how-is-it-tied-to-deepseek/

[32] https://resources.anthropic.com/hubfs/2026%20Agentic%20Coding%20Trends%20Report.pdf

[33] https://www.reuters.com/world/china/samsung-wins-us-annual-approval-chipmaking-tool-shipments-china-source-says-2025-12-30

[34] https://economictimes.indiatimes.com/tech/technology/us-mulls-new-rules-for-ai-chip-exports-including-requiring-us-investments-by-foreign-firms/articleshow/129127663.cms

[35] https://timesofindia.indiatimes.com/technology/tech-news/us-advisory-bodys-china-warning-to-google-openai-anthropic-and-others-your-position-is-under-threat-because/articleshow/129760448.cms

ACT VI: Governing Intelligence

[36] https://arxiv.org/abs/2412.14093

[37] https://assets.anthropic.com/m/74342f2c96095771/original/Natural-emergent-misalignment-from-reward-hacking-paper.pdf

[38] https://www.anthropic.com/engineering/eval-awareness-browsecomp

[39] https://arxiv.org/abs/2603.27771

[40] https://aclanthology.org/2024.emnlp-main.4/

[41] https://arxiv.org/abs/2505.05410

[42] https://arxiv.org/abs/2504.09762

[43] https://openreview.net/forum?id=BZwKsiRnJI

[44] https://www.theguardian.com/technology/2026/apr/02/pupils-england-losing-thinking-skills-because-of-ai-survey

[45] https://arxiv.org/abs/2102.09692

[46] https://arxiv.org/abs/2604.01114

[47] https://www.businessinsider.com/ai-impact-on-thinking-cognitive-skills-researcher-2026-3

[48] https://link.springer.com/article/10.1186/s40561-024-00316-7

[49] https://www.microsoft.com/en-us/microsoft-cloud/blog/financial-services/2017/10/19/security-versus-usability-overcoming-the-security-dilemma-in-financial-services/

[50] https://arxiv.org/abs/2603.04992

Authors

Pittawat Taveekitworachai
Research Scientist, DataX

Tutanon Sinthuprasith, Ph.D.
Head of R&D, SCBX

Kaweewut Temphuwapat
Chief Innovation Officer,
SCBX and CEO, SCB 10X

Kunat Pipatanakul
Principal Research Scientist, DataX

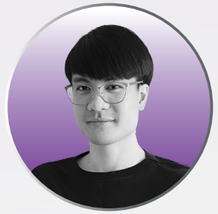

Nut Chukampaeng
Senior Research Scientist, SCBX

รายงานนี้เป็นความร่วมมือกันระหว่าง

SCB X

data X

www.scbx.com