

AI Organization Redesign

A book on redesigning work and organizational operations in the age of AI agents.

Kim Kyung-jin

English PDF edition

AI Organization Redesign

Kim Kyung-jin

Start from the table of contents, then move through the preface, six chapters, and closing essay.

Table of Contents

Preface. Why I Came to Write This Book

Chapter 1. The Coming Age of AI Agents: Work Is About to Change Completely

Chapter 2. The 10:20:70 Rule and Work Innovation That Determine the Success or Failure of AI Adoption

Chapter 3. From Rigid Machine-Like Organizations to Living, Breathing Organism-Like Organizations

Chapter 4. New Talent in the Future Workplace and the Rebirth of Jobs

Chapter 5. The Leadership Posture That Drives Success and How to Change Organizational Culture

Chapter 6. A Strategy for Expanding Small Experiments Across the Entire Company

Closing Essay. Will You Hold the Sword, or Design the War?

Preface. The Purpose and Structure of This Book

Preface. Why I Came to Write This Book

In the autumn of 2024, I lost an election for the National Assembly.

I had spent 13 years as a prosecutor, served in the 20th National Assembly, and had a will, in my own way, to try to change the world. But the voters made a different choice. On the night of my defeat, after I had packed up and left the campaign office, it sank in for the first time that I had nothing to do the next morning.

So I packed a backpack. I began a solo journey through 13 countries.

When I stood before the ruins of the Roman Forum, I remembered a passage from Nanami Shiono's *The Story of the Romans*: "Rome was not built in a day, but neither did it fall in a day." Walking among stone columns from 2,000 years ago, I thought: even empires collapse; what, then, is one person's career? You can build again on the ground where something has fallen.

Throughout the trip, AI was beside me. When I needed to find my way through an unfamiliar city, check local laws, or plan the next day's schedule at night in my lodging, I asked AI, and AI answered. Over the course of traveling through 13 countries, AI was my guide, interpreter, secretary, and conversation partner.

That experience changed me.

After I returned, I began writing a book about AI, started teaching AI at a university, and founded a company to enter the AI education business. This book is an interim report on that journey.

The question I wrestled with most while writing this book was: "Who am I writing for?" Books about AI already overflow the shelves. Technical books are for specialists, and books of future prediction often stop at abstraction. I wanted to write a different kind of book. A book for ordinary people working inside companies, for CEOs running small businesses, and for students preparing to enter society. A book these people could read and think, "Ah, this is about me."

When Nanami Shiono wrote Roman history, she wrote with the eyes of a novelist, not a historian. She did not focus on the minutes of the Senate; she imagined what Caesar may have been thinking as he stood before the Rubicon. I have attempted something similar. I did not write about AI's technical specifications, but about how the Monday morning office changes because of AI. I put metaphors before technical terms and stories before statistics.

This book consists of six chapters: what an AI agent is (Chapter 1), why people matter more than technology (Chapter 2), how organizations must change (Chapter 3), which jobs will disappear and which will emerge (Chapter 4), what leaders must do (Chapter 5), and how to expand small experiments across an entire company (Chapter 6). These six chapters form a single story: in the face of a new tool called AI, how must people and organizations change in order to survive?

May 2026, Seoul.

KIMKJ.COM

#AI #AI #AI #AI #AI #AI #AI
#ChatGPT #Claude #GPT #LLM # # # # #GDPR #
#AI # # # #AI #AI #AI #kimkj #kimkjcom

Chapter 1. The Coming Age of AI Agents: Work Is About to Change Completely

1.1 AI That Only Did What It Was Told Is Becoming AI That Thinks and Moves on Its Own

In the days when the Roman Empire ruled the Mediterranean world, the emperor had two kinds of subordinates. One was the messenger. When the emperor ordered, "Deliver this letter to Gaul," the messenger rode off, delivered it exactly as instructed, and returned. Whether the contents of the letter were correct, or what conditions in Gaul were like, was no concern of his. He simply did what he was told.

The other was the governor. A governor was different. If the emperor gave only the broad directive, "Stabilize the province," the governor judged and acted for himself: how to collect taxes, whether there were signs of rebellion, where roads should be built. He read reports through the night, studied local conditions, and, when necessary, moved legions. Even while the emperor slept, the governor's work continued.

The change we are now experiencing in AI looks exactly like this.

For the past 20 years, the technologies we have used were closer to messengers. Search engines were like that. If you typed "Seoul weather" into Naver or Google, the search engine's role ended with showing you the results. Whether to bring an umbrella after seeing those results was something the human had to decide. Every step of gathering information, comparing it, making a decision, and taking action was human work.

The chatbots that came next, conversational AIs such as ChatGPT or Claude, were somewhat smarter messengers. If you said, "Summarize this report," they produced a clean summary. If you said, "Translate this English email into Korean," they translated it. But still, they moved only after a person spoke first. Without an instruction, they did nothing. In other words, they were "passive tools." This kind of AI is called a copilot. Like the copilot in an airplane, it sits beside you and helps, but it cannot get the plane off the ground unless the captain holds the controls.

But beginning in late 2025, an entirely different kind of AI started to appear.

This new AI is closer to a governor than a messenger. If a person simply says, "Please prepare my business trip for next week," this AI searches flight schedules on its own, compares hotels, coordinates meeting times, and organizes the required documents. It judges what it should do next and moves without a person having to issue one instruction after another. This kind of AI is called an "agent," and this mode is called "agentic AI." In English, an agent is a representative: a being that handles work on my behalf even when I do not do it myself, just as a Roman governor ruled a province on behalf of the emperor.

Now, in 2026, agent AI has already left the laboratory and is working in the real world. Gartner, a well-known technology analysis firm, announced in its 2026 report that agentic AI had reached the

"peak of inflated expectations." In April 2026, Salesforce, an enterprise software company, released more than 100 new tools enabling AI agents to operate all of its functions directly. The structure is one in which AI manages customer information and runs marketing behind the scenes without people needing to click through screens. NVIDIA also released a toolset for building AI agents in March of the same year, and fifteen giant companies, including Adobe, SAP, and Cisco, immediately moved to adopt it.

The speed of change is frightening. In March 2026, Cursor, an American coding tool, released a feature that lets AI agents inspect and modify code on their own in the background. AI works even while the developer sleeps. The company's annual revenue surpassed 2 billion dollars, or about 2.8 trillion Korean won, and its revenue doubled in just three months.

From Painkiller to Immune System

This difference is easy to understand if we compare it to the body.

Copilot-style AI is a painkiller. When you have a headache and take medicine, the pain disappears. But if you do not take the medicine, nothing happens. It is a one-off solution that begins working only after I say, "It hurts."

Agent AI is closer to the immune system. The immune cells in our bodies patrol around the clock, watching for strange viruses or bacteria. They detect enemies before I even feel sick and send white blood cells to fight them. They work while I sleep. Nobody orders them to; they act on their own.

Another metaphor is possible. Existing AI was a very good "hammer." It was strong and precise, but it became useful only when a person held it, aimed at a nail, and struck. Agent AI is an "artisan." If you say, "Make me a bookcase," the artisan chooses the wood, measures the dimensions, saws, hammers, and brings you the finished bookcase.

The Engine That Moves Agents: Observe, Think, Act, Learn

The secret that allows agent AI to move this autonomously lies in a four-stage wheel: Observe, Think, Act, Learn. In short, OTAL.

Observe. AI grasps the surrounding situation: whether an email has arrived, whether a schedule has changed, whether there is something strange in the data. Just as people examine the world with eyes and ears, AI reads what is happening now through text, sensors, internet connections, and other channels.

Think. Based on what it has observed, it considers, "What should I do next?" Just as people make judgments by combining experience and logic, AI also builds plans by combining the intuition that detects patterns with rigorous logic.

Act. It puts the plan into execution: booking a flight on a website, entering data into a company system, sending a Slack message. It does not stop at thought; it actually does something.

Learn. This is the most important stage. It looks back at the results of its action and evaluates for itself: "Did I do well or poorly?" If it made a mistake, it tries a different method next time. What makes this possible is AI's memory system: "short-term memory," which briefly stores the task currently being performed; "experiential memory," which records, like a diary, what happened in the past and how it turned out; and "knowledge memory," which contains general knowledge about the world. As these three kinds of memory combine, an agent improves little by little the more it works.

Roman governors were the same. When first appointed, they did not know the province well. But by collecting taxes, experiencing rebellions, and building roads, they gradually became seasoned. As experience accumulated, judgment became faster; as judgment became faster, governance became more stable. Agent AI learns on the same principle.

The Collision of Two Philosophies: Copilot vs. Cowork

There are broadly two ways of looking at this change.

Microsoft chose the "Copilot" path. As the name suggests, it means "assistant pilot." The approach is to embed AI into existing programs such as Word, Excel, and PowerPoint and have it help beside the user while the user works. It is safe and stable. But the fact that the user must issue instructions step by step has not changed.

Anthropic took a different path. Claude Cowork, released in 2025, treats AI like a "newly hired employee." This AI employee directly accesses the user's computer desktop and file folders. It moves across different work tools such as Slack, Salesforce, and Asana, and even while a person is away, it plans on its own and handles complex tasks.

Copilot is the philosophy that "the human is the master, and AI is the assistant."

Cowork is the philosophy that "AI, too, is a member of the team."

No one yet knows which side is right. Just as Rome wavered for a long time between republic and empire, this question will not be settled for some time. But one thing is clear. Both approaches point in the same direction: AI will take on far more work than it does now.

What Agents Need in Order to Work: MCP and Skill Files

For agents to work properly inside a company, they must be able to connect to the company's various systems. They need to check calendars, open customer databases, and send messages in the company messenger. But every system has a different way to open its door. The keys all have different shapes.

This is where MCP, the Model Context Protocol, appears. Put simply, it is a "master key." It is a standard specification that lets a single key open the door to any system. Without breaching a company firewall, AI can safely access calendars or databases within the scope it has been authorized to use. As of 2026, Salesforce alone has newly created more than 60 MCP tools.

There is also something called a "skill file," or Skill.md. This is a work manual for agents, or, more simply, an "employee handbook." It contains rules such as, "In this situation, do this," "If you do not know something, do not make it up; ask a person," and "Write reports in this format." The agent reads this handbook and follows it.

One of the secrets that allowed the Roman Empire to govern vast territory was a unified system of rules called Roman law, along with a road network connecting the entire empire. MCP is the road of the AI world, and skill files are the legal code of the AI world.

From the Era When People Open Apps to the Era When AI Opens Apps

This is where the truly great change occurs.

Until now, people opened programs directly. They logged into CRM, a customer management program, entered customer information, went into ERP, a management program, checked inventory, opened email, and contacted business partners. Human fingers clicked the buttons on the screen.

But when agent AI appears, AI operates these programs on behalf of people. A person only has to say, "Make this month's sales report," and the AI moves across several programs behind the scenes, gathers data, analyzes it, and completes the report.

When that happens, screen design, or the user interface seen by users, loses value. If AI is the one using it anyway, what use is a beautiful button? Instead, API, the channel through which programs exchange data with one another, gains value. AI needs not a pretty screen but a well-organized data pathway.

The scale of this change is beyond imagination. In February 2026, about 285 billion dollars in market value, nearly 400 trillion Korean won, evaporated from software company stocks in a single day. ServiceNow fell 7 percent, Salesforce 7 percent, and Intuit 11 percent. This happened because AI had begun working so well that companies no longer needed to buy multiple software subscriptions. Gartner predicted that by 2030, 35 percent of enterprise software would be replaced by AI agents.

The Australian design company Canva is also worth noting. Having achieved annual revenue of 4 billion dollars, or about 5.6 trillion Korean won, in 2025, the company discarded many of the business software tools it had used internally and replaced them with AI-based systems. In April 2026, it went so far as to declare, "We are no longer a company that has attached AI to a design tool; we are an AI company that has attached a design tool."

From the standpoint of existing software companies, this is an event that overturns the market itself. Before, they could sell software licenses in proportion to the number of employees. If a company had 100 employees, they collected payment for 100 seats. But if an AI agent can finish the work of five people on behalf of one person, the company only needs to buy one seat instead of five. Software companies' revenue is bound to fall.

When roads were laid across the Roman Empire, trade on those roads expanded explosively, while local merchants who had done business without roads were eliminated. When the infrastructure of technology changes, the fate of the people who made a living on top of it changes as well. What is happening now is exactly that.

1.2 From People Who Do the Work Directly to People Who Command AI

The New Rules of the Chessboard

What determines victory or defeat in chess?

The number of pieces? No. Both sides receive the same number of pieces at the start. The movement abilities of the pieces? No. The queen moves the same way for both sides. In the end, what determines victory or defeat is strategy: where you place the pieces and in what order you move them. The answer lies not in the pieces themselves, but in the mind of the person holding them.

The same thing will happen in the age of AI agents.

Before long, everyone will have similarly capable AI agents. Everyone will hold powerful pieces in their hands. What makes the difference then will not be the AI itself, but the human judgment that decides how to deploy that AI and toward what goal to make it move.

In the past, people did the work directly. They wrote reports, organized data, and sent emails. The "executor" was the person good at work. But as AI begins to take over this execution, the human role changes: from the "performer" who moves their own hands to the "conductor," or orchestrator, who leads multiple performers.

The conductor of an orchestra does not play the violin. Nor does the conductor play the flute. But the conductor decides when every instrument enters, how strongly it plays, and what emotion it must carry. An orchestra without a conductor creates nothing but noise, no matter how excellent its performers are. This is the role that will be required of people in the future: the ability to conduct performers called AI and create a single result.

The Person Who Designs AI Teams: The Agent Manager

Imagine what work inside a company will look like.

Under the old method, a team leader gave every task to one subordinate: "Do the market research, write the report, and make the presentation materials." No matter how capable that employee was, there was a limit to what one person could handle.

A team leader in the agent era works differently. They assign several AI agents to different roles. One is an "information-gathering agent" that roams the internet collecting the latest news and market information. Another is a "writing agent" that drafts text based on the information gathered. Another is an "editing agent" that refines the writing to fit the company's tone and rules. The team leader simply checks whether this AI team is working well and stamps approval on the final output.

The framework that organizes this is called RPM.

R is Role. What does this agent do? Does it gather information, write text, or proofread?

P is Project. What is the goal to complete this time? It must be specific, such as "complete the press release for the June new product announcement."

M is Method. What rules must be followed? Rules such as "always use pseudonyms for customer names" and "cite sources for figures."

In this way, AI agents divide roles like people, exchange outputs with one another, and complete the work. The human only needs to draw the overall picture and adjust the direction along the way.

The Birth of the Centaur

In Greek mythology there is a being called a centaur. Its upper body is human and its lower body is horse. It possesses both human wisdom and the strength of a horse. The most valuable talent in the AI era resembles precisely this centaur.

Take programmers as an example.

Programmers of the past wrote code by hand from beginning to end. Writing each line with care was proof of skill and a source of pride, like an artisan carving furniture by hand.

But the reality of 2026 is different. After using Anthropic's Claude Code, one Google engineer said, "It produced in one hour what our team spent a year building last year." AI generates repetitive, mechanical code in an instant.

In this situation, people divide into two groups.

One group laughs at the small mistakes in code produced by AI and says, "See, it still can't match the human hand." These people take pride in their manual craft, but they gradually fall behind the times.

The other group entrusts repetitive code writing to AI and spends its own time thinking about how to design the overall system, what users want, and what business value this program will create. This is the centaur that combines the human head with AI hands.

The market wants the latter kind of talent.

Why Humans Must Remain Involved

If AI takes over this much work, will people simply become more comfortable? No. On the contrary, human responsibility becomes heavier.

AI is fast. It processes data and produces results with astonishing speed. But AI sometimes lies. The phenomenon in which it plausibly invents facts that do not exist is called "hallucination." It may fabricate numbers or cite legal provisions that do not exist.

AI also has no conscience. It cannot determine on its own whether a decision harms someone or whether a report is biased. It has the logic of efficiency, but not a sense of ethics.

That is why the principle that a human must remain in the loop becomes important. A person must review the results produced by AI, send them back if there is a problem, and give final approval. If this process is removed, a fast but dangerous machine will come to steer the company.

Roman history offers a similar lesson. The Roman legion was the strongest fighting machine of its age. But the decision about where to send the legion, when to fight, and when to stop belonged to the commander or the Senate. The legion itself never decided on war. No matter how powerful a weapon is, without the hand and mind that hold it, it becomes a disaster.

AI is the same. What AI liberates is the time formerly spent on repetitive tasks. What returns to people instead is a burden that may be even heavier: "responsibility for the result."

New Jobs That Are Emerging

This enormous change is creating new kinds of jobs.

The role of "AI product manager," or AI PM, is a representative example. This person acts like an interpreter. When a company's business division says, "We want to improve customer satisfaction," the AI PM translates that vague wish into concrete work instructions and workflows that AI can understand. The job is to turn human language into machine language, and then turn machine outputs back into a form that people can understand.

"AI ethics engineers" are also emerging. They monitor whether AI has learned biased data and is producing results unfavorable to certain groups, and whether AI's decision-making process is transparent. A role called "AI output evaluator," which assesses the quality of results produced by AI, is also appearing.

These roles have something in common. All of these new jobs require both the ability to use AI well and the judgment that only people possess.

The most valuable person in the future workplace will be neither the person who handles Excel quickly nor the person who codes well. It will be the person who knows precisely what AI can and cannot do, and who knows how to weave into the work process the things machines cannot do: empathy, intuition, ethical judgment, and creative ideation. An age is coming in which such people will be treated best.

It is said that Caesar almost never swung a sword himself in the Gallic Wars. His genius lay in knowing whom to place where, in what order to attack, and when to stop. In 2026, a similar question has been put to us: will we learn how to swing the sword, or will we learn how to design the war?

KIMKJ.COM

#AI #AI #AI #AI #AI #AI #AI #AI
#ChatGPT #Claude #GPT #LLM # # # #GDPR #
#AI # # # #AI #AI #AI #kimj #kimjcom

Chapter 2. The 10:20:70 Rule and Work Innovation That Determine the Success or Failure of AI Adoption

2.1 The Secret of Success Lies Not in Technology but in People

90 Percent Jumped In, but Only 9 Percent Survived

Before Rome gained supremacy over the Mediterranean, hundreds of city-states competed on the same sea. Many countries had warships. Their armor-making technologies were broadly similar. Yet in the end, only Rome took the Mediterranean in its hands. The difference was not ships or armor. It was the people and organizations that used them.

The same thing is happening now in the world of AI.

Ninety percent of companies around the world have already adopted AI. In other words, there are almost no companies left that cannot say in a conference room, "We use AI too." Yet only 9 percent have succeeded in actually making money with AI or changing the way they work. Nine out of ten entered the stadium, but not even one out of ten crossed the finish line.

Why does this happen?

Because most companies make the same mistake: the illusion that "buying expensive AI is enough." They believe that if they spend a large sum to purchase the AI program said to be the best on the market, the company will transform as if by magic. This is like buying the world's most expensive Formula 1 race car and expecting someone without even a driver's license to suddenly become a professional racer. No matter how good the car is, if the driver is not ready, it will smash into the wall at the starting line.

The Meaning of the Numbers 10:20:70

There is a company called Boston Consulting Group, or BCG, one of the most famous management consulting firms in the world. After analyzing many corporate AI adoption cases, BCG proposed one rule: the 10:20:70 rule.

When AI adoption succeeds, 10 percent of the success comes from the performance of the AI program itself.

Twenty percent comes from technical equipment such as computers, servers, and data storage.

The remaining 70 percent comes from changes in people and ways of working.

Ten percent plus 20 percent equals 30 percent. What is this 30 percent? It is what money can buy. A good AI model? You can purchase it. Fast computers? You can order them. Cloud servers? You can pay by card and use them immediately. As of 2026, excellent AI models are being released by companies such as OpenAI, Anthropic, and Google, so anyone can use AI at a similar level. Research from MIT Sloan School of Management reached the same conclusion: AI models are

rapidly becoming a "public good" that anyone can use.

Then where is the real contest decided? In the remaining 70 percent. How people accept AI, how the sequence of work changes, and whether organizational culture permits a new tool: these cannot be bought with money. They require time, effort, and a leader's decision.

This becomes clear when we recall why Rome was strong. Rome's sword, the gladius, was not an especially outstanding weapon. It was short and rough. The long swords of the Celts looked more threatening. But the Roman legions won. Not because the sword was better, but because the legion formation, training system, and command structure that used the sword were superior. If the weapon's performance was 30 percent, the organization and people using it were 70 percent. Whether 2,000 years ago or today, this ratio is surprisingly similar.

If the Company Cannot Keep Up, Employees Run Ahead Alone

Microsoft and LinkedIn produced a report based on a survey of more than 20,000 workers around the world. The result is interesting. In the gap between people who use AI well and those who do not, individual ability accounted for only 32 percent. The remaining 67 percent was determined by the company's culture and support system.

This means that no matter how smart a person is, and no matter how diligently they learn AI on their own, it is of no use if the company blocks them by saying, "Do not use that yet" or "There are security issues."

Yet an interesting phenomenon is taking place. Whether companies block it or not, employees are already using AI. A survey found that 78 percent of workers use AI tools such as ChatGPT or Claude for work personally, without company permission. This is called "shadow AI." It means using AI without the company's knowledge, like a shadow.

As of 2026, the situation has grown even more serious. In 98 percent of companies, employees are found to be using AI without authorization. Forty-seven percent of employees access AI tools through personal accounts to handle company work. Thirty-eight percent say they have entered confidential company information into AI. According to IBM's 2025 report, when information leaks because of this shadow AI, the additional cost averages 670,000 dollars per incident, or about 940 million Korean won.

Are employees doing this because they want to break the rules? No. They do it because they want to work faster and better. Since the company does not provide good AI tools, employees find and use them on their own. In fact, one healthcare organization found that when it provided an official AI tool, unauthorized AI use fell by 89 percent. The answer is not to block AI, but to provide proper tools.

There is one particularly unfortunate group: people who use AI extremely well personally but whose growth is blocked because the company provides no support. These people carve out time on weekends to study how to use AI. They discover ways to handle work ten times faster. But they do not share those methods with colleagues. Why? Because the company does not recognize them for

it, and they fear being glared at for "doing useless things." They are afraid of being criticized for attempting innovation.

This phenomenon existed in Rome 2,000 years ago as well. Even if a capable general wanted to try new tactics on the battlefield, he could not move if the Senate insisted, "Do it the traditional way." When Hannibal crossed the Alps and invaded Italy, the Roman general Fabius Maximus avoided direct confrontation and chose delaying tactics. The Senate and citizens mocked him as the "Cunctator," the delayer. Anyone who tries a new method always encounters resistance from the existing system.

The AI era is no different. It is not a question of technology. It is a question of people and culture.

Organizations That Permit Failure Win

Many companies leave AI adoption solely to the technology department and pay no attention to employees' feelings. When this happens, "culture debt" accumulates. Just as debt grows with interest if it is not repaid, culture debt grows larger if it is neglected.

Fifty-six percent of companies ignored issues such as employee anxiety and fairness while adopting AI. As a result, more than 60 percent of employees are experiencing "change fatigue." Each time a new tool appears and each time a system changes, they become exhausted and no longer want to keep up.

The key to this problem is "psychological safety." It may sound difficult, but its meaning is simple: an atmosphere where "it is okay to make mistakes." An environment in which people are not scolded or treated as fools if they try a new AI tool and something goes wrong. A workplace where people can honestly say, "I tried it once, and it didn't work."

This is not simply a matter of "employee welfare." It is directly connected to money. Studies show that in teams where employees can try new things without fearing mistakes, new ideas increase by 27 percent, decision-making speed rises by 35 percent, and employee departures fall by as much as 50 percent.

Take SpaceX as an example. Elon Musk's space company does not call it "failure" even when a rocket explodes. It uses the expression "rapid unscheduled disassembly." It sounds like a joke, but there is a philosophy inside the phrase. Failure is not shameful; it is data for doing better next time. When a rocket blows up, the company analyzes "Why did it blow up?" and uses that lesson to make the next rocket better.

Ernesto Pagano, a senior partner at BCG, said the same thing. "Employees accept AI only when they understand it not as a tool that replaces them, but as a tool that helps them." This cannot be forced. Trust must be built.

The Interpreter Between Technologists and Businesspeople: The AI Product Manager

To drive the 70 percent change, a new kind of talent is needed.

Until now, companies have had one great valley. On one side is the technology department: people who handle computers and AI. On the other side is the business department: people who meet customers, sell products, and make money. Although these two departments work in the same company, they have not communicated well.

When the business department says, "We want to improve customer satisfaction," the technology department asks, "What exactly does that mean? Which metric are you asking us to improve?" The business side does not know the language of technology, and the technology side does not know the language of business. Because of this communication failure, the wrong AI gets built, and money and time are wasted.

This is where the "AI product manager," or AI PM, appears. This person is an interpreter. The AI PM translates the business department's vague desire, "Please make customers more satisfied," into a concrete work instruction that the technology department can understand: "Build an AI chatbot that reduces customer inquiry response time from the current 24 hours to 2 hours, while keeping the error rate below 5 percent."

This person does not need to code directly. Instead, they must understand two things at once: how the business works, and what AI can and cannot do. They are, in effect, a bilingual speaker who can use the languages of both worlds.

One reason the Roman Empire was able to integrate dozens of different cultures, including Greece, Egypt, and Gaul, was that it had "interpreter officials" who spoke both Latin and Greek and understood the conditions of each region. Technology and business also need such interpreters. The most valuable talent in the company of 2026 is precisely this AI PM.

The CEO Must Use It First

Finally, what is the most powerful force that brings about the 70 percent change?

The answer from the Microsoft report was unexpectedly simple: "Leaders must show themselves using AI directly." When the CEO personally asks AI questions during a meeting and reviews the results together with others, employees feel, "Ah, this is real." Conversely, if the CEO only receives reports from behind the scenes and says, "Adopt AI well," employees think, "It's just a trend. It will fade away."

BCG partner Pagano made the same observation. If consultants create a 200-page plan and leave, nobody reads it. It is far more effective for the leader to use AI directly and go through trial and error together with team members. He called this kind of leader an "individual contributor CEO," or ICCEO.

In the end, the true contest in AI adoption is not "which AI did we buy?" It depends on "how quickly our company's CEO and employees created a culture of working together with AI."

2.2 The Sequence of Work Must Be Redrawn from the Beginning

52 Projects, and Nothing Changed

There was a large insurance company in the United States worth 7 billion dollars, or about 10 trillion Korean won. This company was serious about AI. It ran no fewer than 52 AI pilot projects. The accuracy of the AI models also reached the target. Technically, it was a success.

But after 18 months, the way the company worked had not changed at all.

Employees reviewing insurance claims were still flipping through documents the old way. AI had been abandoned as a separate screen occasionally consulted in one corner. The money and time spent on 52 projects had almost been scattered into thin air.

Why did this happen?

Because the technology changed, but the sequence of work did not.

Experts call this the "bolt-on" approach. A bolt is a screw. It is the method of forcibly attaching a new part to an old machine with a screw. It is like inserting a cutting-edge 2026 engine into the body of an old car built in the 1980s. No matter how good the engine is, if the body cannot withstand its power, the car breaks apart while driving.

Many companies make this mistake. They leave work procedures from the past 30 years unchanged and simply place AI on top. At first, things seem a little faster. But competitors can make the same degree of improvement. The real contest is between companies that redraw the sequence of work itself from the beginning and companies that merely insert a new tool into an old sequence.

A report BCG released in March 2026 made this point clear. To use AI agents at large scale, companies must redesign the existing work process itself. Well-designed agent-based work processes greatly reduce confirmation steps such as "Did I hear that correctly?" and "Did I miss anything?" Reviews that used to take days can sometimes be completed in minutes.

The question we ask must change. Not "How can we use AI to do existing work faster?" but "In a world where AI exists, how would we build this work from scratch?"

If Humans Must Check Every Time, There Is No Point in Using AI

There is another problem that arises when AI is forcibly inserted into the existing method: approval bottlenecks.

At many companies, when AI creates a report, the person in charge checks it first. Then the team leader checks it. Then the department head checks it. No matter how quickly AI produces the result, if the structure in which people approve it step by step remains intact, the overall speed is tied to human speed. It is AI in name only; in reality, it repeats the old manual process as before.

So real change means changing the approval structure itself. Instead of "procedural safeguards" in which a person checks every time, we need "structural safeguards" embedded from the system design stage.

Here is an example. When AI accesses the company's customer data, a rule can be built into the system from the start: "This AI can view data up to Grade A, but not Grade B." No person needs to confirm each time, "Is it okay for this AI to see this?" The rule is already built into the design.

The goal is clear. It is not to make an approval process that used to have 10 steps 10 times faster with AI. It is to reduce the 10-step process itself to 4 steps. Reduce the number of process steps, while guaranteeing safety through structure.

This is easy to understand if we think of Roman roads. Before Rome laid its roads, wagons carrying goods had to stop in every village and pay tolls. If there were ten villages, they had to stop ten times. Rome changed this structure. It removed checkpoints from major roads and instead standardized road width and specifications so that any wagon could travel without obstruction. It did not make wagons pass through checkpoints faster; it eliminated the checkpoints themselves. Work redesign in the AI era follows the same principle.

Lessons from British Coal Mines in the 1940s

When people talk about redesigning work, it may sound as though this will be only a technology discussion. It is not. Surprisingly, the answer to this problem emerged from British coal mines 80 years ago.

Britain in the 1940s. A very good new mining machine was introduced into coal mines. Judging by the machine's performance alone, coal production should have increased significantly. In reality, production fell. The machine was not bad. The problem lay elsewhere. When the new machine was introduced, the team structure of mine workers changed, and the close teamwork through which they had helped one another broke apart. The technology improved, but because relationships among people were damaged, the overall result became worse.

The scholars who studied this incident reached a conclusion: "The technical system and the human, or social, system must be changed at the same time. If only one is changed, the result can actually fail." This is called "joint optimization."

The same principle applies to AI adoption in 2026. When introducing a new technology called AI, companies must redesign at the same time the way people work, the structure of teams, and the distribution of authority. If only AI is inserted and everything else remains the same, the result will be like the British coal mines of 80 years ago.

What Should Be Entrusted to AI, and What Should Humans Do?

When redesigning work, it helps to think in four boxes.

The first box is "delegate." This includes entering data, summarizing long documents, and making daily reports in the same format. These are tasks people can do, but AI can do much faster and with fewer mistakes. Such work should be handed over to AI entirely.

The second box is "assist." AI writes the first draft, and a person refines it. For example, when AI drafts an email to send to a customer, a person reads it and fixes awkward parts. AI speed and human sensibility combine.

The third box is "amplify." This is when AI multiplies human capability several times over. For example, when a doctor analyzes a patient's test results, AI instantly finds and shows thousands of similar cases. AI's data processing ability is added to the doctor's experience, making it possible to reach judgments the doctor could not have conceived alone.

The fourth box is "the human-only domain." Empathizing with a customer's grief, judging whether something is ethically right, and taking final responsibility. These are things AI cannot do and should not do. They must remain human work.

Putting every task into one of these four boxes is the first step in work redesign.

This Is How a Salesperson's Day Changes

Let us look at a concrete example.

There is a salesperson who sells products at a company. When this person's day was examined, the time actually spent meeting and talking with customers accounted for only 25 percent of the whole. What took up the remaining 75 percent? Entering data into a customer relationship management program, or CRM; drafting proposals; scheduling meetings; and writing reports. The work that actually brings in money, meeting customers, took up only one quarter of the day.

If the work is redesigned, it changes like this. CRM data entry? An AI agent automatically analyzes meeting recordings and enters the data. Proposal drafts? AI analyzes the customer's past transaction history and industry information and prepares a customized draft. Meeting scheduling? AI checks both calendars and finds available times.

Then the salesperson's time meeting customers doubles. This person's role also changes. In the past, the salesperson was a "miscellaneous worker" who even had to enter data. Now the salesperson becomes a "conductor" who reviews strategies created by AI and makes final decisions.

From an Old Machine to a Living Organism

When all of these changes come together, the company itself takes on a different form.

Companies until now have resembled "machines." Walls between departments are high, orders flow from top to bottom, and people move only according to fixed procedures. They are regular like gears, but gears cannot suddenly change direction.

The future company, where AI agents work mixed together with people, becomes closer to a "living organism." When the environment changes, it responds in real time and changes form as needed, as an octopus freely changes the shape of its body when passing through a narrow gap.

Chapter 3. From Rigid Machine-Like Organizations to Living, Breathing Organism-Like Organizations

3.1 The Pyramid Is Collapsing

The Limits of a 100-Year-Old Blueprint

The Roman legion was a pyramid.

At the very top was the commander-in-chief. Beneath him were the legion commanders, then battalion commanders, then centurions, and at the bottom the soldiers. Orders from above were passed down level by level, and information from below was filtered through layer after layer on its way up. This structure functioned for more than 700 years, from the third century BCE to the fifth century CE.

Companies over the past 100 years have had the same shape. At the top sits the CEO; below are vice presidents, then department heads, managers, assistant managers, and employees. This is the pyramid organization. Work is broken down into the smallest possible units and assigned to each department. Middle managers receive reports from below and send them upward, then receive instructions from above and pass them downward. In an era when the world changed slowly, this structure fit well. It was stable and predictable.

But in 2026, this 100-year-old pyramid is shaking.

According to a survey by PwC, the global accounting firm PricewaterhouseCoopers, 56 percent of companies that adopted AI saw no effect. There is an even more startling figure: one report says that 95 percent of AI pilot projects failed. Why? Because good AI was forcibly inserted into old pyramid organizations. According to one study, when advanced AI is placed inside a rigid top-down organization, productivity actually falls by as much as 30 percent.

The reason the pyramid of the Roman legion endured for 700 years was that the speed of war was slow. It took weeks or months for an enemy to invade. It was acceptable if reports took several days to move upward and orders took several days to come down. But what if the enemy attacked at the speed of light? By the time a report traveled through the battalion commander and legion commander to the commander-in-chief, the war would already be over.

Business in the AI era is precisely that kind of battlefield. A competitor's attack arrives not in days, but in hours. While a company waits for a report to climb to the top of the pyramid and for a decision to come down, a faster company has already eaten the market.

The Flow of Information Has Completely Reversed

Microsoft CEO Satya Nadella identified this change with great precision. In January 2026, at the World Economic Forum in Davos, Switzerland, Nadella said that AI had flattened the entire flow of information. He said organizations must be structurally redesigned.

In the past, things worked like this. Each department in a company gathered information in its own field: the sales team gathered sales information, the finance team financial information, the customer support team customer information. The working-level employees in each department summarized this information and sent it to team leaders. Team leaders organized it again and sent it to department heads. Department heads organized it once more and sent it to executives. This process took weeks. As information moved upward, it was filtered and sometimes distorted. This was precisely the role of middle managers: to gather information from below and send it upward.

But when AI agents appeared, this structure collapsed.

Nadella told a story about preparing for 50 meetings. In the past, each department would have had to spend weeks gathering materials and writing reports to send upward. Now, however, if he tells an AI agent, "Give me a full briefing on this customer," the AI crosses sales data, investment records, and customer support history all at once and creates a complete summary in seconds.

Information no longer climbs slowly from bottom to top. It gathers everywhere at once, in an instant.

When this happens, the role of people who collected and delivered information in the middle disappears: summarizing reports and sending them upward, passing decisions from above downward, stamping approval. AI replaces all of this.

Nadella actually applied this philosophy to his own company. In early 2026, Microsoft carried out a major integration of its Copilot organization. It merged teams that had been divided into consumer and enterprise groups, promoted Jacob Andreou to vice president, and had him report directly to Nadella. It removed the walls between departments.

There was an even more radical measure. Nadella created a weekly "AI acceleration meeting," where senior executives were not allowed to present. Instead, frontline engineers took the stage themselves and delivered raw, unfiltered information directly to the CEO. The voices of the people doing the work could reach the top immediately, bypassing middle managers.

How Amazon Broke the Pyramid

Amazon CEO Andy Jassy moved even more directly.

In September 2024, Jassy issued a company-wide instruction: "Raise the ratio of individual contributors to managers in each department by at least 15 percent." Before the first quarter of 2025 had even ended, Amazon had already achieved this goal.

Jassy's words are striking. "As organizations grow, the number of middle managers increases. As middle managers increase, everyone, with good intentions, tries to put their stamp on everything. Then you get pre-meetings for decisions, pre-meetings for those pre-meetings, and pre-meetings for those pre-meetings." Meetings for meetings for meetings. This is the disease of pyramid organizations.

Jassy also created a "bureaucracy tip mailbox." It is a channel through which any employee can report, "This is a useless procedure." Nearly 1,000 reports came in, and Amazon changed more than 375 procedures based on them.

Jassy's goal was clear: "Move like the world's largest startup." Even as a giant company, make decisions like a fast, light small company.

This is not merely restructuring to save money. It is the act of physically removing from the organizational chart the weakness called "decision delay," which competitors can exploit. In war, an army with a complex chain of command loses to a fast enemy. Business is no different.

Department Walls Are Falling

If the pyramid is being compressed vertically, the walls between departments are collapsing horizontally.

Imagine an old company. The sales team, finance team, and marketing team each sit in their own room. Each uses its own program and tracks its own goals. If they need information from one another, they must send emails or schedule meetings. For information to pass from one room to another, someone has to knock on the door and ask permission. This wall is called a "silo." Like cylindrical storage towers for grain, each department stands apart from the others.

AI agents pass through these walls.

For example, suppose a customer asks an advertising company for a new advertisement. In the past, a salesperson received the customer's request and handed it to the planning department. The planning department organized it and handed it to the production department. The salesperson then brought the finished result back to the customer. The information had to knock on departmental doors three times.

The era of AI agents is different. An agent reads the customer's request in the customer management system, checks the budget in the finance system, identifies available ad inventory in the marketing system, and instantly produces an optimal advertising plan. It does not cross departmental walls; it moves as if the walls do not exist.

This is not because the company wants it to happen. It is because for AI agents to work properly, they must view data from multiple systems at the same time. If departmental walls remain, agents cannot perform their function. Removing the walls is not a choice but a prerequisite.

Lego-Block Companies and Octopus Teams

If the pyramid collapses and departmental walls crumble, what shape will a company take?

Leading companies are becoming organizations like "Lego blocks." Lego blocks can be detached and attached at any time. A payment-function block, a customer-data block, a marketing block: each operates independently, yet when needed, they can be combined instantly to create a new form.

Human organizations are the same. In the past, once someone was assigned to a fixed team with a fixed name, such as "Sales Team 2" or "Planning Team 1," they worked in that team for years. Under the new method, when a project arises, the necessary people gather to form a team, and when the project ends, they disperse and join other teams. This flexible team is called a "modular talent pod."

This is easy to understand if we recall the military organization of the Roman Republic. For each war, Rome selected soldiers from among its citizens, formed legions, and disbanded them after the war ended. It was not a standing army, but a citizen army gathered when needed. This flexibility, gathering and dispersing according to purpose, became a great force in Rome's early expansion.

Companies in 2026 move in a similar way. If a new product must be launched, developers, marketers, designers, and AI engineers gather into a single pod. When the launch is over, the pod disbands, and a new pod forms for another project. Between these pods, AI agents take care of back-office work such as paperwork, schedule coordination, and data organization. People focus on strategy and creative judgment.

Leaders Change from Engineers into Gardeners

In this kind of organization, the leader's role changes completely.

The old leader was a "machine designer." The leader designed how every gear meshed and turned, and if even one gear slipped, fixed it directly. The ability to control everything was the leader's capability.

But in an organization like a living organism, the leader cannot control everything. Just as one cannot individually command the movement of every cell. Instead, the leader becomes a "gardener." A gardener does not command each flower, "You bloom this way, you bloom that way." Instead, the gardener lays good soil, gives water, and creates an environment with good sunlight. The flowers bloom on their own.

When Nadella wrote in an internal message to Microsoft executives, "We must all work and act like individual contributors in our own organizations, constantly learning and unlearning," he was speaking in the same context. The leader must not be someone who merely receives reports from behind the scenes, but a gardener who touches the soil directly.

Nadella added one more warning: "If large organizations cannot keep up with the pace of change, small companies will gain scale thanks to these tools and send you to school." New small companies can design their organizations for AI from the beginning. Large companies, however, must tear down pyramids that have already been built. That difference can decide victory or defeat.

3.2 How People and Multiple AIs Work as One Team

There Is No Single Perfect AI

There is a mistake companies most commonly make when adopting AI: the idea that "one all-purpose AI will solve everything."

Reality does not work that way. No matter how smart a person is, one person cannot do the work of every department in a company. If someone is told to do market research, write reports, handle accounting, and respond to customers, that person will become overloaded, make mistakes, or collapse.

AI is the same. If everything is entrusted to a single AI, the workload becomes too large, and it invents strange answers, a phenomenon called hallucination, or misses important things.

That is why leading companies are dividing multiple AIs by role. This is called a "multi-agent system." The system resembles a human organization to a surprising degree.

According to Databricks' 2026 survey, use of multi-agent systems increased by as much as 327 percent in just four months. This is not a fad; it is a direction.

Let us look at a concrete example of how it works. Suppose a management report worth millions of dollars is being prepared.

There is a supervising agent. This AI plays the role of conductor. It plans the entire task and assigns missions to the agents below it.

There is an information-gathering agent. This AI roams the internet, financial data, and news to collect the latest materials.

There is a structuring agent. It organizes the collected materials logically and lays out the skeleton of the report. It checks whether anything is missing or overlapping.

There is a writing agent. Based on the skeleton, it writes the actual text, draws graphs, adds sources, and completes the report.

There is a reviewing agent. It inspects the completed report from beginning to end. If there is a hole in the logic, it sends the report back to be rewritten. Its role is to filter internally before the report is handed to people.

These five agents exchange outputs and cross-check one another. Even if one agent makes a mistake, another catches it. It is like departments in a company checking one another and improving the quality of a report.

Why was the Roman Senate better than a dictator? Concentrating all authority in one person was fast but dangerous. When multiple people debated and checked one another, the process was slower, but mistakes decreased. A multi-agent system applies this principle to AI. It is a method for securing both speed and accuracy.

Rulebooks and Keys That Control AI Behavior

When multiple AIs work autonomously, a concern arises: "What if they act however they want?" It is a natural concern. Rome also gave governors great authority, but at the same time it had rules called

Roman law and a monitoring system called the Senate. Authority and control always come as a pair.

There are two devices that control the behavior of AI agents.

One is the skill file, Skill.md. As discussed in Chapter 2, this is a work rulebook for AI. It contains rules such as, "Do not make up what you do not know. Ask a person about unverified content. Write reports in this format." The agent reads this rulebook every time and follows it.

The other is MCP, the Model Context Protocol. This is the standard specification agents use to connect to a company's various systems. They do not enter any system they wish. Instead, they move only within an authorized scope, such as, "This agent can view only the calendar and Slack. It cannot access financial data."

To reuse the metaphor of Roman law and roads, the skill file is the law, and MCP is the travel pass.

Where Humans Must Step In

No matter how sophisticated an AI system becomes, there are places where humans must not be removed.

AI is excellent at recognizing patterns. It can find rules in tens of thousands of data points and predict the future faster and more accurately than people. But AI has no conscience. It cannot feel for itself, "Might this decision unfairly harm someone?"

Something like this actually happened in Australia. The government introduced an automated AI system called Robo Debt and sent notices demanding repayment from people it claimed had received welfare payments incorrectly. Eight hundred thousand notices were sent, and many of them were wrong. The system miscalculated or failed to consider circumstances. Vulnerable people who had been receiving welfare suddenly became debtors. This incident was recorded as one of the greatest administrative disasters in Australian history.

This is what happens when we are blinded by speed and efficiency and remove the process of thinking, "How will this decision affect people?"

That is why the principle of "human-in-the-loop" is important.

However, the way humans step in must also be intelligent. Having people stamp approval on every single thing AI does is not the answer. Then the meaning of using AI disappears. As discussed in Chapter 2, instead of "procedural safeguards" that check every time, we need "structural safeguards" built into system design.

Take a bank's loan review as an example. When AI analyzes a loan application, it divides the result into three zones.

Green zone. Cases where AI is at least 95 percent confident that "this can definitely be approved." AI handles these automatically.

Chapter 4. New Talent in the Future Workplace and the Rebirth of Jobs

4.1 Jobs That Disappear, Jobs That Are Born Anew

Where Did the Lamplighters Go?

On the streets of nineteenth-century London, there was a job called "lamplighter." Around sunset, these were people who walked the streets carrying long poles and lit gas lamps one by one. They lit the London night. At the time, it was a fairly important job.

When electricity was invented, this job disappeared. With a single switch, thousands of lights turned on at once. The skills and experience of lamplighters became useless overnight.

But while electricity eliminated lamplighters, it created new jobs: electricians, power-plant operators, and power-line designers. When some jobs disappeared, other jobs appeared. The same thing happened when the steam engine replaced 60 percent of textile workers. New technicians were needed to operate the machines.

In 2026, a new electricity called AI is spreading through the world. The same pattern is happening again, this time at a much faster speed and across a much wider range.

Dario Amodei, CEO of Anthropic, said, "Within the next one to five years, half of white-collar jobs will disappear." In its January 2025 report, the World Economic Forum, or WEF, predicted that 92 million jobs worldwide would be displaced by AI by 2030.

The numbers alone are frightening. But in the same report, the WEF also said that 170 million new jobs would be created during the same period. Minus 92 million, plus 170 million: on net, that means 78 million jobs will increase.

Things that disappear and things that emerge. We must look at both at the same time. If we see only one side, it is fear; if we see only the other, it is fantasy. Reality lies between them.

How This Technological Revolution Differs from Those of the Past

This time, however, there is one difference from the past.

The steam engine replaced human muscle: lifting, carrying, cutting heavy things. Electricity, too, handed work requiring human physical strength over to machines. But AI is replacing the human mind. It is gathering information, summarizing it, writing reports, writing code, and analyzing data. AI has begun to perform the "thinking work" that until now had been considered something "machines cannot do."

For this reason, some jobs are especially at risk: office jobs with clear rules and repeated daily tasks. Entering data, organizing meeting minutes, answering basic customer inquiries, coordinating simple schedules. AI agents do this kind of work faster than people, without rest, and without mistakes.

The most worrying phenomenon is the "hollowed-out middle." If companies replace the basic tasks done by new employees with AI in order to cut costs, new employees lose the very opportunity to learn work. The process of building skill through making countless mistakes, correcting them, and trying again during the novice period disappears.

This is a serious problem. Ten years from now, experienced people who can supervise and judge these vast AI systems will be needed, but the training grounds that produce such people are disappearing. If the Roman legion eliminated recruit training camps, it would save money right away, but ten years later there would be no centurions left to command wars.

That is why leading companies are changing the work of new employees rather than eliminating it. Instead of having them enter data, they have them use AI to analyze data and propose strategy. Repetitive tasks are entrusted to AI, and new employees are trained early in higher-level thinking work.

New Jobs Pouring Out Like the Cambrian Explosion

Five hundred forty million years ago, an event called the Cambrian explosion occurred on Earth. In a very short period, tens of thousands of new life forms that had not existed before suddenly poured out. Scientists are still studying why this happened.

The job changes AI is causing are similar. Jobs that did not even have names five years ago are suddenly appearing.

Let us examine them one by one.

New Job 1: AI Product Manager, or AI PM

This job appeared in Chapters 2 and 3 as well. Let us look more deeply at why it is so important.

Ninety percent of companies around the world have adopted AI, but only 5 percent have succeeded in actually making money from it. There are also reports that 95 percent of AI pilot projects fail. This is not because the technology is lacking. It is because businesspeople and people who build technology cannot communicate.

The AI product manager is the bridge connecting these two worlds. This person does three things.

First, the AI PM finds the points inside the company that take the most time and would produce the greatest effect if AI were inserted. Metaphorically speaking, this person is a scout who identifies the enemy's weakest point in a war.

Second, the AI PM knows what AI is good at and what it is not good at. This person can distinguish, "This problem needs an AI that writes text, and that problem needs an AI that predicts numbers."

Third, the AI PM converts the business department's vague desire, "Please make customers leave less often," into a concrete blueprint that the technology team can immediately work from.

This person does not need to code directly. Instead, they must be able to speak both the language of business and the language of technology.

As of 2026, the average annual salary for an experienced AI product manager is about 160,000 dollars, more than 200 million Korean won. It has become one of the most valuable jobs in the technology industry.

The Roman Empire had a similar role. Envoys conveyed the will of the Senate to legion commanders, reported the conditions of the legion in the field back to the Senate, and understood the languages of both sides. The outcome of war often depended on the envoy's ability to translate.

New Job 2: AI Ethics Engineer

Until 2025, most AI was chatbot AI that simply conversed on a screen. Even if it gave a wrong answer, it rarely led to a major accident.

But agent AI in 2026 is different. It searches company databases, executes code, sends emails to customers, and even processes payments. It is AI that actually acts. If this kind of AI makes a mistake, the harm does not remain inside the screen. Customer information can leak, or thousands of discriminatory decisions against a particular group can be carried out automatically.

An AI ethics engineer is the person who prevents such accidents. This role is closer to an engineer who designs systems than to a philosopher discussing morality.

This person does two things. One is to restrict in advance the range of information AI can access and build security devices so company secrets do not leak outside. The other is to leave records so that AI's decision-making process can be traced, and to monitor whether AI is making decisions fairly.

Rome had an office called the censor. Its role was to monitor citizens' morals and property and to expel unqualified people from among the senators. The greater the power, the more necessary the watcher becomes. Now that AI's authority is growing, the AI ethics engineer is, in effect, the censor of the digital age.

New Job 3: AI Output Evaluator and Prompt Designer

Can we simply trust reports, code, and analysis results produced by AI? The answer is no.

AI produces plausible outputs that are 95 percent complete. But fatal errors may be hiding in the remaining 5 percent. It may make up facts that do not exist, a hallucination; misunderstand context; or miss subtle nuance. Catching this 5 percent is the human role.

An AI output evaluator carefully inspects what AI has created. This person asks: "Is the source for this number truly correct?" "Is there a hole in this logic?" "Could this expression mislead customers?" The structure is one in which AI creates 95 percent, and the human adds the remaining 5 percent of nuance and intuition to complete 100 percent.

A prompt designer is a person who professionally creates the "instructions" used to give work to AI. This is not merely the level of asking ChatGPT good questions. The prompt designer creates detailed guidelines so that multiple AI agents move precisely according to the company's rules, including skill files. You can think of this role as the person who writes the agent's "work manual."

New Job 4: Multi-Agent Conductor

The "orchestrator" discussed in Chapter 3 is settling in as a profession.

Take a recruiter in the human resources department. In the past, this person's work was to read resumes directly, schedule interviews, and write acceptance notices. Now the work becomes arranging and directing the AI that screens resumes, the AI that schedules interviews, and the AI that sends acceptance notices. The person does not do the work directly; they assign it.

These roles are emerging in every department: marketing, sales, finance, customer support. Everywhere, there is a need for people who combine multiple AI agents and operate them like one team.

All of these new jobs have one thing in common. They are not people who "calculate faster than machines," but people who decide "what to ask machines to do, and why."

4.2 The Rise of the Super Individual: An Age When One Person Does the Work of Dozens

The Era When Hiring More People Increased Revenue Is Over

When the Roman Empire expanded, governing a larger territory required more soldiers and administrators. It had to organize one more legion and send one more governor. The size of territory and the number of people were proportional.

Until the early 2020s, companies grew in the same way. When customers increased, employees also had to increase. They hired more salespeople, recruited more marketers, and appointed more managers. It was taken for granted that revenue and headcount rose together.

AI agents have broken this formula.

Now one person can command several AI agents and do alone the work that several people used to do. A phenomenon is occurring in which company revenue rises while headcount stays the same or even declines.

Such a person is called a "Super IC," or Super Individual Contributor. IC means individual contributor: someone who works directly rather than managing a team. "Super" has been added. It means a person who, with the help of AI, produces the results of ten or twenty people.

The Person Who Built a 50,000-Dollar Team for 500 Dollars

There is a person named Jacob Bank. After working at Google as product lead for Gmail and Calendar, he left and founded Relay.app, an AI agent company. His story came up in an interview

that attracted attention in March 2026.

Bank is the only marketing person at this company. There is hardly a marketing team to speak of, because he is alone. But under his command, about 40 AI agents are working.

One agent tracks competitors' movements 24 hours a day. If a competitor changes its price, a Slack notification arrives immediately. Another agent analyzes industry trends every week and creates a report. Another writes drafts of LinkedIn posts based on those trends. Another automatically sends promotional emails when a webinar, or online seminar, is scheduled.

The cost of operating these 40 agents is 500 dollars per month, about 650,000 Korean won.

If the same work were assigned to people, at least four professional marketers would be needed. Monthly labor costs would be 50,000 dollars, about 65 million Korean won. The cost difference is 100-fold.

With this system, Bank recorded 1.5 million impressions on LinkedIn.

There is one more interesting point. Bank created an agent that neatly organized meeting minutes and sent them to customers, but he noticed that customers did not read them. He immediately "fired" that agent. Firing a human employee requires sensitivity, carries a heavy emotional burden, and involves complicated procedures. If an AI agent is not useful, you turn it off. If strategy changes, one sentence is enough: "Stop what you were doing."

Bank's own words are striking. "For my parents' generation, working at a large company for 40 years may have been the safe path. But now it is the opposite. Staying at one company too long and building only skills that work inside that company is the riskiest path."

A Company Where 7 People Generated 75 Billion Won

In Taiwan, there is a memory semiconductor export company called DAIWA International. Its core personnel number only seven. Yet its annual revenue is 1.8 billion Taiwan dollars, about 75 billion Korean won.

The secret is this. These seven people infused into an AI system the experience and judgment they had accumulated by traveling through global markets for more than 10 years. They built their own AI called "AO Trade Master," enabling marketing documents that previously took specialists 18 hours to write to be completed in 15 minutes. That is 72 times faster.

Instead of going to expensive overseas trade fairs, these people made AI deal with customers around the world 24 hours a day. Even while people sleep, AI emails customers on the other side of the globe, answers inquiries, and prepares proposals.

The reason seven people can generate 75 billion won is that these seven are not simple workers. They inject into AI the judgment and insight that came from 10 years of experience, and AI replicates and executes that insight thousands of times over. Human wisdom is amplified through AI scale.

The coding tool Cursor is a similar case. It has only 60 employees, yet more revenue than a traditional software company of 1,000 people. As of March 2026, its annual revenue surpassed 2 billion dollars, about 2.8 trillion Korean won.

An Era of Investing in AI Usage Instead of Headcount

There is an amusing expression for this change: we have moved from "headcount maxing" to "token maxing."

Headcount means the number of employees. Tokens are the units consumed when using AI, or, more simply, AI usage fees.

In the past, if a company wanted to grow, it had to hire more people. Hiring one employee costs tens of millions of won per year in salary, insurance, office space, welfare, and so on. And that person works only eight hours a day, five days a week.

But AI agents work 24 hours a day, 365 days a year. Their cost is falling sharply every year. So excellent companies, instead of hiring one more person, give the outstanding people they already have a budget to use AI freely. It is more efficient for one capable person to use AI and produce the results of ten people.

From a company's standpoint, a receipt for tens of millions of won in AI usage fees may look frightening. But what that AI is replacing is hundreds of millions of won in labor and management costs. The calculation is clear.

When a Super Individual Arrives at Work in the Morning

Let us draw a concrete picture of how a super individual's day unfolds.

They arrive at work in the morning. Overnight, the AI information-gathering agent has summarized the movements of 50 competitors. Reading it, the person sees that one competitor has announced a new product. The super individual instructs the AI writing agent, "Write a draft of our company's response strategy to this competitor's new product." Ten minutes later, the draft arrives. The person reads it, personally adds the 5 percent the AI missed, such as customers' emotional reactions or internal circumstances known only to the company, and presses the approval button.

This person's role has three parts: setting direction, filling in the human parts AI cannot see, and taking responsibility for the final output.

Do you remember the RPM framework? It appeared in Chapter 1.

R, Role. Agent A is an "intelligence officer" that gathers information. Agent B is a "writer" that writes text. Agent C is an "editor-in-chief" that refines the text according to company rules.

P, Project. "Complete the competitive analysis report by this week."

M, Method. "Do not use unsourced content. Write according to the company template. If you do not know, do not make it up; tell me."

You do not need to know how to code. The core is knowing the logic and goals of the business clearly. The ability to decide "what must be done, and why." This is the weapon of the super individual.

The Person Who Designs the Game Record Wins

In Chapter 1, I used the metaphor of chess. Let us bring it out again.

Before long, everyone will hold the same powerful AI pieces. What decides the outcome then is not the piece itself, but the ability to design the game record: in what order, for what purpose, and where to place each piece.

The person who compresses with AI the time that dozens of people used to waste exchanging emails and Excel files, and pours energy only into judgment and creativity that humans alone can provide: that person is the super individual.

If you read Caesar's Commentaries on the Gallic War, scenes in which he personally swings a sword in battle almost never appear. Instead, you repeatedly see him designing which legion to send where, how to maintain supply lines, and how to exploit conflicts among Gallic tribes. Direct fighting was entrusted to the soldiers; drawing the picture of the war was Caesar's work.

The workplace of 2026 is moving toward the same structure. Execution is entrusted to AI, and designing strategy becomes the work of people.

Lamplighters disappeared, but the city's nights became brighter. There is no need to mourn disappearing jobs. But we must begin preparing now for the abilities required by the jobs newly being created. The core of those abilities is simple: knowing how to give work to machines, and knowing what machines cannot do.

KIMKJ.COM

#AI #AI #AI #AI #AI #AI #AI #AI
#ChatGPT #Claude #GPT #LLM # # # #GDPR #
#AI # # # #AI #AI #AI #kimkj #kimkjcom

Chapter 5. The Leadership Posture That Drives Success and How to Change Organizational Culture

5.1 The CEO Must Take the Wheel First

CEOs Who Buy Supercars and Hire Substitute Drivers

When Julius Caesar set out on his Gallic campaign, he did not sit in Rome waiting for reports. He mounted his horse himself and walked the battlefield himself. He confirmed with his own eyes what the land of Gaul smelled like, how high the soldiers' morale was, and what the enemy formation actually looked like. The Gaul written in reports and the Gaul he stepped on with his own feet were entirely different lands.

In 2026, on the new battlefield called AI, most corporate leaders are behaving in the exact opposite way from Caesar.

Dr. Chien Lee-feng, former head of Google Taiwan, offered a painfully sharp analogy. Companies spend billions of won buying a supercar called cutting-edge AI, but the CEO does not take the wheel. The reason is fear or inconvenience. Instead, the CEO hands the driving over to an assistant who has ridden only bicycles all their life, then sits in the back seat and scolds, "Why aren't we going faster?" A survey found that 80 percent of Taiwanese SME executives do not use AI directly and push it off onto assistants. This is not only Taiwan's problem.

If the leader sits in the back seat, the entire company moves at bicycle speed. No matter how good an engine the company buys, it is useless.

According to PwC's survey, 56 percent of companies that adopted AI saw no results. There are also reports that 95 percent of AI pilot projects fail. This is not a technology problem. It is because leaders have not touched the technology themselves, and therefore do not know what it truly can and cannot do.

The Hands-On CEO: A New Leader Called the ICCEO

Old management textbooks said this: "A good CEO is someone who steps away from detailed operations, delegates authority, and looks only at the big picture from above." In the AI era, however, this formula is wrong.

Howie Liu, CEO of Airtable, a company valued at about 12 billion dollars, or about 17 trillion Korean won, shows a completely different kind of leader. He calls himself an "ICCEO," an Individual Contributor CEO, which in Korean would mean a "hands-on chief executive."

Howie Liu is the person at his company who spends the most on AI usage, or inference costs. He ranks first. He writes code himself and uses AI to analyze tens of thousands of sales meetings recorded at the company over the past year. He willingly spends hundreds of dollars on API calls. As a result, in just a few days he directly discovers things that previously would have required paying an

outside consulting firm millions of dollars and waiting months: what customers truly want, the most effective sales methods, and the company's position in the market.

In early 2026, Howie Liu made an even more unconventional move. He canceled most regularly scheduled executive meetings and instead urged team members, "Cancel every meeting for one week and personally try every AI product on the market." He meant: touch it yourself, collide with it, and fail with it.

He says the CEO must become the company's "Chief Taste Officer." When a chef obtains a new ingredient from outer space, the chef cannot know the ingredient's true value by merely reading the manual and telling the cooks below to use it. The head chef must touch it, smell it, and chew it personally.

Indeed, Howie Liu publicly stated that he is a heavy user of Claude, using it directly every day. In January 2026, Airtable released a new AI product called Superagent, and in April it unveiled Hyperagent as well. These were the first new products outside the company's main business in its 13-year history. This was possible because the CEO personally felt AI's potential in his own body.

Four Stages of a Leader's Self-Revolution

"That's a good story. But I don't know technology well. How am I supposed to start?" Many leaders will ask this question.

Dr. Chien Lee-feng proposes four stages.

Stage 1. Experience pain.

It sounds cruel, but the effect is certain. Send the leader's assistant on vacation for one or two weeks. Schedule coordination, email sorting, meeting material preparation, data organization: have the leader do all of this "miscellaneous work" bare-handed. It will be painfully difficult. In that pain, the leader will for the first time feel with the whole body "where exactly our company's work is getting stuck."

The Roman emperor Hadrian is famous for personally walking across and inspecting the entire empire. If he had sat in Rome merely reading reports, he could not have judged that Britannia needed a wall. Because he walked it himself, felt the cold himself, and saw the enemy's invasion routes with his own eyes, he could make the historic decision known as Hadrian's Wall.

Stage 2. Use the tools yourself.

Once you experience pain, you develop a desperate desire to find a solution. A leader struggling without an assistant turns on ChatGPT or Claude directly in order to survive. They ask AI to draft presentations, write replies to tangled emails, and organize messy meeting minutes. It will be awkward, but that awkwardness is learning.

Stage 3. Discover the possibility.

As you use AI yourself, a moment arrives that feels like being struck on the head. "The draft AI produced in five minutes is better than the report a 10-year assistant used to take several days to bring me?" At that moment, you realize AI can become not a chore processor, but the best strategic partner.

Stage 4. Lead by example.

A leader who reaches this point no longer relies on the words of outside consultants. They know for themselves which department and which process in the company should receive AI. They have experienced it directly. The change this leader pushes is real. It is not change read in a report, but change felt in the body.

In SMEs, the CEO Is the Strategy

It would be a mistake to think this story applies only to large companies. In fact, this issue is more urgent for small and medium-sized enterprises.

In a large company of 3,000 people, AI can become a tool for reducing costs. If 3,000 employees each save only 5 percent of their time, the combined labor cost savings are enormous.

But a small business with 20 or 30 employees is different. Even if employees finish Excel work 30 percent faster and leave at 4 p.m., company revenue does not increase. That is because the revenue limit of an SME is tied not to employee speed, but to the CEO's capability.

The real goal of AI in an SME is not to reduce an employee's report time by five minutes. It is to liberate the CEO's time. The CEO must hand over to AI the time stolen by chores such as information gathering, email responses, and schedule coordination, and use that time to meet key customers, imagine new businesses, and win decisive contracts.

In Chapter 3, I discussed how Nadella personally used AI while preparing for more than 50 bilateral meetings in Davos. He instructed AI, "Give me a 360-degree briefing on the relationship with the CEO I will meet tomorrow, investment status, and current issues," and received a perfect summary in seconds. If the head of the world's largest IT company uses AI directly, is there any reason for an SME CEO to sit back and leave it to an assistant?

5.2 Invisible Debt: Culture Debt Must Be Repaid

Debt That Cannot Be Seen but Is Accumulating

When the Roman Empire began to decline after its peak, there seemed to be no problem from the outside. The army was still strong, roads were still laid, and law still functioned. But cracks had begun in places that could not be seen. Citizens' loyalty to Rome wavered, soldiers' morale fell, and senators began putting their own interests ahead of the public good. These invisible cracks eventually brought down the empire.

Companies also have invisible cracks like these. This is called "culture debt."

Just as debt accrues interest if it is not repaid, culture debt grows larger if neglected. Then, at some point, it collapses the company from within.

Culture debt accumulates like this. A company adopts AI while caring only about efficiency and performance and ignoring employees' anxiety. A survey found that 56 percent of companies completely ignored employees' emotions and fairness when adopting AI. According to Deloitte, today's employees must experience as many as 15 major organizational changes a year. New systems, new tools, and new procedures constantly pour in. Seventy percent of employees complain of "change fatigue." They are exhausted and no longer want to keep up.

Why Employees Hide Their AI Skills

In Chapter 3, I discussed "shadow AI": employees using personal AI tools without the company's knowledge. Let us go a little deeper here.

According to the latest Microsoft and LinkedIn report, individual employees are learning AI at frightening speed. They carve out weekends and after-work hours, teach themselves how to use AI, and discover ways to process work ten times faster. Sixty-five percent of employees feel anxious that "if I do not use AI, I will fall behind."

Yet 55 percent of employees hide from the company the fact that they have become good at using AI.

Why? The reason is bitter. Even if they produce astonishing results with AI, only 13 percent of employees expect to receive praise or reward from the company. Eighty-seven percent think, "They won't recognize it anyway." Some even fear, "If I tell the company how I finish my work quickly with AI, won't the company fire me?" If I hand my know-how over to AI, won't I become useless?

This fear is understandable. But the result is destructive. Seventy-eight percent of employees are using personal AI tools without company approval, and 38 percent have entered company secrets into AI. The company does not know who is doing what where. The risk of information leakage is rising, but nobody sees it.

It resembles the phenomenon in the late Roman Empire when frontier legions began moving on their own without following instructions from the Roman center. Because the center did not understand the reality on the ground, the field began ignoring the center.

Companies Where Mistakes Are Allowed Win

The key to this problem is psychological safety. It appeared in Chapter 2 as well, but because it is such an important concept, let us examine it more deeply.

Psychological safety is "an atmosphere where you will not be scolded for making mistakes." It is "an environment where foolish questions are not dismissed." It is "a culture where even if you fail while trying something new, the failure itself is recognized as valuable."

This is not a soft story about "the conditions of a good workplace." It is directly connected to money.

In teams where employees can try new things without fear, new ideas increase by 27 percent. Decision-making speed becomes 35 percent faster. The rate at which core talent leaves the company falls by as much as 50 percent.

There is a reason this safety has become even more important in the AI era. AI is not perfect. It sometimes hallucinates and produces biased results. People must be able to say honestly, "I tried using AI, but it produced a strange result." In an atmosphere where "you will be scolded unless you do it perfectly," nobody tries new tools. They are afraid of failure.

The CEO Must Show Failure First

What is the most powerful way to create psychological safety? Research from Microsoft and LinkedIn gives the answer: "When leaders show themselves using AI directly, employees' trust in AI and use of AI increase explosively."

Even more effective is for leaders to disclose failure.

Imagine a CEO saying this during a meeting: "I also tried making a presentation with AI yesterday, and AI invented completely wrong numbers. It was a total mess. But this part was truly astonishing. It finished in 10 minutes what would have taken me three hours."

This one sentence sends a powerful message to employees. "The CEO fails too. Then it must be okay for me to fail." This honesty from the leader gives employees permission: permission to try AI.

Nadella's internal message to Microsoft executives was in the same context. "We must all work like individual contributors, constantly learning and unlearning." And: "Every time I receive a note about how differently a friend at an AI startup works, I laugh. In fact, that kind of innovation is happening right here at Microsoft too. Our job is to find it, nurture it, and learn from it."

Facing the Fear of "Will I Be Fired?"

The deepest root eating away at psychological safety is the fear of job loss. "If I teach AI all of my work methods, won't the company cut me?"

Most leaders avoid this question. They offer vague consolation: "AI will not replace you; it will only assist you." But employees are not fools. They do not believe that comfort.

To build real trust, leaders must be honest. IBM's approach is worth noting. "Certain repetitive tasks will disappear because of automation. This is true." First, state the painful truth. Then make a promise at the same time: "But we will create new roles that only people can perform, and we will support training so that you can move into those positions."

Painful truth and concrete promises build more trust than vague comfort.

Find the AI Geniuses Hidden Inside the Company

Finally, let me tell the most interesting story.

When companies try to adopt AI, many leaders attempt to hire expensive outside AI experts. But the real treasure may already be inside the company.

Someone in purchasing who has managed complex ordering processes for 10 years; someone in finance who has combined hundreds of Excel files to produce reports; someone in HR who understands the entire process from recruitment to retirement. These people may not know how to code. But they know the flow of work better than anyone else. They know where things get stuck, where time is wasted, and in what order work is efficient.

In the age of AI agents, it is possible to combine AI and automate work without coding. When these people hold AI tools in their hands, they can create systems that solve the headaches of their own departments. Because they are the ones who understand the problem best.

The leader's role is to create a "safe playground" for these hidden geniuses. A space where security rules are followed, but within those rules experimentation is free. An atmosphere where results are not immediately judged by ROI, or return on investment, but where people ask, "Did you discover something interesting?" When these people solve problems in their own departments, the effect naturally spreads to colleagues around them. "That person did it. Maybe I should try too?" This is the strongest form of internal diffusion.

One reason the Roman Empire was able to expand far was the reward known as "Roman citizenship." Even frontier soldiers or residents of provinces could receive citizenship if they were loyal to Rome and contributed. This citizenship meant new opportunities and rights. Because there was a clear reward for contribution, people moved on their own for the empire. Companies are the same. When employees who create innovation using AI receive recognition and rewards, more people begin moving on their own.

The Way We Learn Must Also Change

According to global talent analyst Josh Bersin, companies around the world spend 400 billion dollars a year, about 560 trillion Korean won, on employee training. Yet 74 percent of companies admit, "Our training cannot keep up with the pace of change."

Why? Old-style training is "push." The company sets the training schedule, seats employees in a classroom, and has them listen to fixed content. They begin getting sleepy during the lecture, and after it ends, they forget 90 percent.

Education in the AI era must change to "pull." The moment employees get stuck while doing their work, AI finds and shows the necessary information in real time. When a salesperson gets stuck while writing a proposal, AI automatically displays industry trends for that customer and past success cases. They do not learn by sitting in a classroom; they learn while working.

Learning what is needed, at the moment it is needed, immediately. This is education in the AI era.

The Conclusion Is Culture, Not Technology

Chapter 6. A Strategy for Expanding Small Experiments Across the Entire Company

6.1 How to Escape the Purgatory of Pilot Projects

Why Successful Experiments Fail to Change Companies

When Rome conquered Gaul, winning the first battle was only the beginning. Capturing one village and turning all of Gaul into a Roman province were matters on completely different levels. Tactics that worked in one village sometimes failed in another. Local tribal resistance was fiercer than expected, and as supply lines lengthened, unexpected problems burst forth.

The process of adopting AI inside a company is exactly the same.

A small team runs an AI pilot project. The result is good. When demonstrated in front of executives, the reaction is, "Wow, that's impressive." Yet even after 18 months, the way the company works has not changed at all. The pilot project succeeded, but it did not spread across the company.

This state is called "pilot purgatory." Purgatory is neither heaven nor hell. A pilot project wanders around in a state that is neither success nor failure.

More than 70 percent of companies worldwide experience this purgatory while adopting AI. It is not because the technology does not work. The AI models work well. The problem is that people and organizations do not change. When a new technology collides with old ways of working, the organization pushes it away instead of absorbing its power.

So to expand a small experiment across the entire company, AI must be approached not as a "technology project," but as "change in people and organizations." A useful framework here is the five-stage method called SHIFT.

S: Story and Strategy; Explain Why First

Most companies begin with this when introducing AI: "This new chatbot summarizes documents 10 times faster." They start with a feature explanation.

But what employees are curious about is not the feature. "Why are we doing this?" "What does this mean for me?" "Could my position disappear?"

The leader must answer these questions first. "This is why we are doing it. This is how it connects to the company's survival. And this is how it will help each of you personally."

Without a clear Story, fear fills the empty space.

When Caesar began his Gallic campaign, he clearly explained to the soldiers "why we are fighting this war." It was not only the promise of spoils. It was for Rome's safety, for Rome's glory. An army with a clear purpose is strong. An army that fights without knowing its purpose scatters at the first

defeat.

H: Understand Employees' Emotions

AI adoption is different from installing one new program. It touches employees' "identity."

When you tell a veteran employee with 10 years of experience, "From now on, AI will write the reports you used to write," it sounds like, "The skills you have honed for 10 years will now be replaced by a machine." This is not a demand to learn a new tool. It is a threat to one's sense of worth.

William Bridges, a scholar of change, explained the process by which people accept change in three stages.

The first is a sense of loss. It is the sadness of losing what was familiar: "The old way was better."

The second is confusion. The old way has been abandoned, but the new way is not yet familiar. It is an anxious middle zone: "What am I supposed to do, and how?"

The third is acceptance of a new beginning. People adapt to the new method and find their place within it.

When employees are in the first and second stages, leaders must not ignore them or rush them. Instead of empty words like, "Don't worry, everything will be fine," leaders must say, "I know you are anxious. That is natural. Let us cross this together."

When Rome incorporated new provinces, it could govern them stably when it respected local people's culture and gods. But when it ignored local emotions and unilaterally forced Roman methods on them, as in the famous rebellion of Boudica in Britannia, it encountered fierce resistance. Ignoring human emotion always carries a price.

I: Build Structural Supports

Emotional understanding alone is not enough. Employees need concrete supports that allow them to actually learn the new way.

Showing training videos once or twice is not enough. Employees must be given time to touch and experiment for themselves. Cases tailored to the actual work of each department must be shown. And people trusted by colleagues in the field should be designated as "AI helpers," or AI champions, so they can help from beside them.

Once the new method takes root, the standards for evaluating performance, or KPIs, and the reward system must also change. If people are told to work in the new way while performance is still measured by the old method, nobody moves. People move in the direction that rewards point.

F: Create an Atmosphere Where Mistakes Are Okay

Psychological safety, discussed deeply in Chapter 5, is central here as well.

The moment someone is too afraid to ask, "How do I use ChatGPT?" because they worry they will look foolish, innovation stops. The leader must first say honestly, "I do not know either. Let us learn together." The leader must disclose personal failure: "I completely messed up yesterday while writing an AI prompt."

The effect of this one sentence is measurable. Studies show that in teams that allow failure, 27 percent more new ideas emerge, decision-making speed increases by 35 percent, and turnover among core talent falls by 50 percent.

T: Use Resistance as Data, Not as an Enemy

Let me tell the story of a large European insurance company. When this company tried to introduce AI into insurance review, experienced claims examiners strongly resisted. It was not because they disliked technology. It was because their professional knowledge had been completely ignored in the process of designing the AI system. "You are going to ignore my 20 years of judgment and leave it to a machine?"

Leadership changed its method. Instead of forcing AI down from above, it began designing the AI workflow together with the examiners. It changed the structure so that AI would not "replace" the examiners' judgment, but would "assist" it by summarizing vast documents. Once the examiners' expertise was respected, participation rose explosively. This department's success became a model case that spread across the whole company.

There is a core lesson. Employee resistance is not an enemy; it is a signal sent by the system. It is feedback saying, "This part is wrong." Ignoring this signal is like ignoring a reconnaissance report in war. Caesar did not ignore the grievances of Gallic chieftains. He identified the causes of discontent, resolved them when possible, and detoured around them when impossible. Pressing down resistance by force was the last resort.

Prove It with Numbers: Abandon Vanity Metrics and Measure Real Performance

Along with changing organizational culture, companies must also change how they measure performance.

Many companies count "how many people logged into the AI tool." This is a vanity metric. People who merely log in and do nothing are included in the count.

What really needs to be watched has two layers.

Leading indicators. These show whether the AI system is operating healthily: number of users, amount of work automatically processed by AI, the rate at which AI produces strange answers, or hallucination rate, and how confidently employees use AI.

Lagging indicators. These show whether AI use actually led to business performance: Did customer satisfaction rise? Did processing time fall? Did mistakes decrease? Did operating profit increase?

For example, a company must prove that developers receiving AI assistance while writing code, a leading indicator, leads to shorter product development cycles, a lagging indicator. If this connection cannot be shown, the AI program loses its budget and disappears.

6.2 How to Use AI Safely While Protecting the Company's Core Assets

Speed and Safety Must Both Be Secured

Roman roads were laid straight and wide so that people could travel quickly. But roads were not all Rome had. Guard posts stood at key points along the roads, and patrols moved through sections where bandits appeared. Roads for fast movement and public order for safe movement: because Rome possessed both, goods and people could move freely throughout the empire.

The same principle applies in the AI era. Companies must be able to use AI quickly, while at the same time controlling it safely.

As of 2026, 42 percent of enterprise AI projects fail, and one major reason is that security and management systems were not reflected in the design from the start. The problem of "shadow AI," in which 78 percent of employees secretly use personal AI, remains. Nearly half of these employees are entering company secrets into outside AI systems.

At the end of 2025, the data transfer agreement between the EU and the United States was invalidated by a European court. More than 34 countries are strengthening laws that regulate where data is processed. Under the EU AI Act, requirements for transparency and data management for high-risk AI systems begin full-scale enforcement in August 2026. We have entered a world where the old attitude of "just put it in the cloud" no longer works.

Having Humans Check Every Time Is Not the Answer

To use AI safely, many companies require a human approval stamp on everything AI does. It looks safe from the outside. But this makes the use of AI meaningless. Humans become the bottleneck, and the whole speed is tied to human speed.

This is a point repeated in Chapters 2 and 3. We need "structural safeguards" instead of "procedural safeguards."

Structural safeguards mean embedding safety rules into the design from the beginning when building the system. There are three methods.

First, permission inheritance. If an employee cannot view financial data at the company, that employee cannot access that data through AI either. The person's permissions pass directly to AI.

Second, information masking. When an employee asks an external AI a question, a security device intervenes in the middle and automatically masks secret information such as customer names or exact amounts. Instead of being delivered to AI as "If we propose Y amount to customer X," it is changed to "If we propose virtual amount B to virtual customer A." We borrow AI's brain, but do not

send out secrets.

Third, classification by confidence. Cases where AI is at least 95 percent confident are processed automatically, while only the ambiguous 5 percent with low confidence are handed to people. People can focus only on difficult cases that truly require human judgment.

IBM's Five Pillars of AI Management

IBM proposed five pillars companies must have in order to safely manage autonomous agent AI.

First, alignment. This means confirming whether AI is acting in line with the company's values and purposes. It monitors in real time the "drift" phenomenon in which AI departs from its original goal.

Second, control. Strictly limit the range of data and tools AI can access. In case AI runs out of control, an emergency stop button, or kill switch, must be included in the design.

Third, traceability. Leave records of every decision AI makes. When a problem occurs, it must be possible to trace, "Why did it make this decision?"

Fourth, security. Protect the system from attacks in which outside hackers insert wrong instructions into AI, known as prompt injection, or inject malicious data. AI must run in a secure zone, or sandbox, separated from the company firewall.

Fifth, legal compliance. Clearly define who bears legal responsibility for AI's decisions. A leading approach is also to include devices that allow AI itself to review the latest laws, such as the EU AI Act.

"Your Questions Are Your Most Expensive Asset"

Let us go one level deeper.

NVIDIA CEO Jensen Huang said this at the World Economic Forum in Davos in January 2026: "Every country must build its own AI infrastructure. It must create its own AI by using the natural resources of its own language and culture, and cultivate its own national intelligence."

This statement does not apply only to countries. It applies equally to companies.

For the past 10 years, the common sense of the IT industry was, "Put everything in the cloud." Having one's own servers was considered old-fashioned. But in the AI era, this common sense is shaking.

Let us see why with an example. Suppose the CEO of a manufacturing company preparing for global expansion asks an external AI this question: "We are planning to enter the Indian market next quarter. With a 200 million dollar budget to build a local factory, what are the biggest barriers? How should we position our new materials technology in order to take market share from competitor A?"

The answer AI gives may be good. But the problem lies elsewhere. The question itself contains the company's most secret strategy: "We will invest 200 million dollars in India," "We are targeting competitor A," "New materials technology is our weapon." This information has moved to an external

server.

In the AI era, answers have become something anyone can obtain. What is truly expensive is the question. A question is a map showing what the company is thinking about and where it is weak. If a competitor gets this map, it is more fatal than having code stolen.

Jensen Huang went one step further at a ServiceNow event in May 2026. He said that the amount of computing required for agent AI would increase by at least 1,000 percent, or 10 times, compared with existing generative AI. His vision is a world where "10 billion digital AI agents work together with human employees." Imagine AI at this scale running on servers outside the company and processing all of the company's questions and data. The risk grows geometrically.

Keep the Core in Our Own House and Borrow the Rest

That is why leading companies are choosing a "hybrid strategy."

Running all of a company's AI on its own servers costs too much. Conversely, entrusting everything to the external cloud creates a high risk that secrets will leak. Both are extremes. The answer lies in the middle.

Company work is divided into three grades.

Grade A. Core secrets directly tied to the company's survival: proprietary technology, key customer data, strategic algorithms. No matter how much it costs, these are processed on the company's own servers, on-premises. They are kept in the vault.

Grade B. Work that is sensitive but may be processed externally after secret information is masked. Customer names and amounts are hidden before cloud AI is asked to analyze.

Grade C. General work: email translation, drafting marketing copy, summarizing meeting minutes. For this, the latest AI in the external cloud can be used freely. No secrets are included.

Raise the core in your own house, and borrow outside strength for surrounding work.

The Roman Empire's defense system was similar. The defense of the city of Rome was directly handled by elite legions. Auxiliary defense on the frontiers was entrusted to allied forces or mercenaries. Guard the core directly, delegate the periphery. If every place is filled with elite legions, the cost becomes unbearable. If every place is entrusted to mercenaries, the risk is too great. Balance is the answer.

Data Is the Real Moat

Finally, let us discuss the ultimate purpose of all these strategies.

In the early AI era, AI models themselves, such as GPT-4 or Claude, appeared to be the source of competitiveness. But over time, performance differences between models are shrinking. Everyone can use AI at a similar level. Algorithms become standardized.

Then what remains? Data.

Decades of market data accumulated by BlackRock, the global asset manager, while managing more than 11 trillion dollars in assets. Clinical records accumulated by a large hospital while treating millions of patients. Such data cannot be scraped from the internet. It cannot be copied. This is the real moat. A moat is the waterway surrounding a castle, a defensive line enemies cannot cross.

In Chapter 2, we discussed BCG's 10:20:70 rule: 10 percent algorithm, 20 percent technology infrastructure, 70 percent people and process. One more thing must be added. Beneath all of this lies the ground called data. No matter how good a building, or algorithm, is built and how good the people, or process, gathered, if the ground, or data, is weak, the building tilts.

Companies must accumulate their own data in a safe place and train AI on that data so it fits their own company. This is the only competitiveness competitors cannot buy with money.

You Can Drive Fast Only When the Brakes Are Good

Many CEOs regard AI security and management systems as "annoying regulations that slow innovation."

The opposite is true.

Why can an F1 race car tear around the track at 350 kilometers per hour? Because the engine is good? Of course. But the real reason is that the car has the most powerful and reliable brakes. Because the driver trusts the brakes, the driver can press the accelerator all the way down on the straight. A car without brakes cannot go fast. It is too frightening.

AI governance is the same. Employees at companies with solid security and management systems can experiment with and use AI most boldly, without fear of mistakes or information leakage. They do not have to wonder every time, "Is it okay to do this?" The system has already been designed safely.

Safety devices do not slow speed. They make speed possible.

Roman roads enabled rapid movement not only because the roads themselves were wide. People and goods could move quickly and safely because the public order system that protected the roads, relay stations placed at regular intervals, and standardized road specifications all operated together. Roads and security were not separate; they were one system.

In 2026, companies trying to escape pilot purgatory and expand AI across the entire company must remember this: technology matters, people's emotions matter, and safety systems matter. These three must not play separately; they must be woven together as one system, just as Rome's roads, laws, and legions formed one empire.

KIMKJ.COM

#AI #AI #AI #AI #AI #AI #AI #AI
#ChatGPT #Claude #GPT #LLM # # # #GDPR #
#AI # # # #AI #AI #AI #kimkj #kimkjcom

Closing Essay. Will You Hold the Sword, or Design the War?

While I was writing this book, the world changed faster than I could write.

An AI tool I introduced while writing Chapter 1 had already become an older version by the time I was writing Chapter 6. We live in a world where Cursor's revenue doubles in three months, Salesforce releases more than 100 AI tools all at once, and 400 trillion won in software company market value evaporates in a single day. No book can remain up to date in the face of this speed.

That is why, in this book, I tried to focus on unchanging principles rather than the technical details.

After reading all 15 volumes of *The Story of the Romans*, what remains in the reader's mind is not the tactics of individual battles. It is a feel for the fundamental question: "Why do some countries rise, while others fall?" Nanami Shiono found the reason Rome was strong not in military power, but in culture: openness that granted citizenship even to defeated enemies, pragmatism that did not feel ashamed of failure, and the principle that everyone is equal before the law.

What I have repeated throughout this book is ultimately the same point.

Technology can be bought with money. Buying a good AI model and installing fast servers is a matter of budget. But people and culture capable of using technology properly cannot be bought with money. An atmosphere that permits failure, the courage of leaders to jump in first, the honesty to answer employees' fears directly, and the flexibility to break down departmental walls: these can be created only with time and will.

BCG's 10:20:70 rule says the same thing. AI algorithms are 10 percent, technical equipment 20 percent, and people and ways of working 70 percent. Seventy percent of the contest lies outside technology.

In Chapter 1, I used the metaphor of chess. Before long, everyone will hold equally powerful AI pieces in their hands. What decides the outcome then is not the piece itself, but the ability to design the game record. Even now, as I finish Chapter 6, I repeat the same point. In the end, what matters is not the machine, but the person.

If you read Caesar's *Commentaries on the Gallic War*, scenes in which he swings a sword himself almost never appear. Instead, you see him constantly designing which legion to send where, how to maintain supply lines, and how to exploit divisions among the enemy. Holding the sword was the soldier's work. Designing the war was Caesar's work.

In the AI era, I ask: will you hold the sword, or will you design the war?

Execution increasingly belongs to AI: entering data, writing reports, writing code. Machines perform this kind of "sword-swinging" faster and more accurately. What remains for people is deciding "why this war is being fought." It is the work of giving purpose, setting direction, filling in the human parts machines miss, and taking responsibility for the result.

What this book wants to give readers is not a technical manual. It is a feel for finding one's place in the face of change. Rome was not built in a day. Your change will not be either. But taking the first step is something you can do today.

After reading this book, I hope you will try using AI directly tomorrow morning. It is all right if you are awkward. Whether you are a leader or a new employee, the first step is awkward for everyone. Caesar, too, was once an unknown politician. What matters is beginning.

KIMKJ.COM

#AI #AI #AI #AI #AI #AI #AI #AI
#ChatGPT #Claude #GPT #LLM # # # #GDPR #
#AI # # # #AI #AI #AI #kimkj #kimkjcom

Preface. The Purpose and Structure of This Book

Next Chapter

This text was written to explain what organizations must change in the age of AI agents.

AI adoption is not simply the purchase of technology. It is a matter of organizational redesign, requiring changes in work methods, decision-making structures, the role of talent, leadership, organizational culture, and scaling strategy.

This book organizes the issue into six connected flows. Chapter 1 explains what AI agents are and how they change the way work is done. Chapter 2 addresses the 10:20:70 rule and process innovation that determine the success or failure of AI adoption. Chapter 3 explains the structural shift from machine-like organizations to intelligent organizations. Chapter 4 examines new talent models and changes in jobs in the future workplace. Chapter 5 addresses leadership and organizational culture. Chapter 6 presents a strategy for moving beyond pilot experiments toward enterprise-wide scaling.

The central question is simple. In an era when AI takes on more work, what must people and organizations redesign?

This text is not written only for technical experts. It is structured so that people who run organizations, people who work in the field, and people preparing to enter the world of work can understand AI-driven change as their own issue.

May 2026.

Next Chapter

KIMKJ.COM

#AI #AI #AI #AI #AI #AI #AI #AI
#ChatGPT #Claude #GPT #LLM # # # #GDPR #
#AI # # # #AI #AI #AI #kimki #kimki.com

Support This AI Library Book

If this book was useful, you can support future translations, public editions, and open AI knowledge projects through PayPal.



PayPal

<https://paypal.me/kimkjkj>

Scan the QR code or open the link above.