# PALANTIR: War, Surveillance, Artificial Intelligence

Table of Contents, Preface, 14 Chapters


Attorney Kyungjin Kim

English PDF edition

# PALANTIR: War, Surveillance, Artificial Intelligence

Attorney Kyungjin Kim

PALANTIR: War, Surveillance, Artificial Intelligence is an online AI Library book by Attorney Kyungjin Kim. It covers Palantir, war, surveillance, artificial intelligence, data analytics, national security and is organized as Table of Contents, Preface, 14 Chapters.

## Table of Contents

# Preface

See every chapter

Next Chapter →

In 2022, I paused while watching a video.

I happened to see footage of a Russian tank company being annihilated on the Ukrainian front. The armored forces of the world's second-largest military power. How was this possible? As I read through related materials, I stopped at one word. Palantir. It was a name I had never seen before.

I began searching. A Silicon Valley data analytics company. Founded with investment from the CIA. Rumored to have contributed to tracking Bin Laden. Providing software free of charge to the Ukrainian military. A company led by a philosophy PhD named Alex Karp. A market capitalization of hundreds of trillions of won.

It was strange. A company of this scale, with this level of influence , why had I never heard of it?

Without our knowing, the structure of power in the world had been shifting. An era where algorithms and data wage war, not tanks and fighter jets. An era where engineers in sweaters dominate the battlefield, not generals in uniform. And at the center of it all stood Palantir.

I began digging deeper. I gathered materials, read articles, analyzed reports. The more I learned, the longer this company's shadow stretched. From the U.S. Department of Defense to the British NHS. From counterterrorism operations to COVID vaccine distribution. From Wall Street fraud detection to immigrant tracking.

And then I realized. This is not merely the story of one company. It is the story of an era where data becomes power , of who holds that power and who will watch over it. It is the story of what we are surrendering in the name of efficiency and security.

This book is my journey to find answers to those questions.

Author: Kyungjin Kim

See every chapter

Next Chapter →

Kim Kyung-jin

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com

#KimKyungjin #PALANTIR #Palantir #WarAI #SurveillanceAI #DefenseAI #DataPower #ArtificialIntelligence #AILibrary

# Chapter 1. The PayPal Mafia and 9/11

← Previous Chapter

See every chapter

Next Chapter →

Part 1: The Heretic of Silicon Valley

A. 9/11 and the Lessons of Intelligence Failure

(1) The Tragedy of Disconnected Information (Silos) and the Limits of the CIA

At 8:46 AM on September 11, 2001, American Airlines Flight 11 struck the North Tower of the World Trade Center. Seventeen minutes later, United Airlines Flight 175 tore through the South Tower. In the underground situation room at CIA headquarters in Langley, Virginia, analysts were staring at their screens. None of them had anticipated this. Or, to be precise, it was something they could have anticipated but failed to.

Intelligence agencies began combing through their own files. The CIA knew that two al-Qaeda operatives, Nawaf al-Hazmi and Khalid al-Mihdhar, had entered the United States. They had first identified these two men while monitoring an al-Qaeda meeting in Kuala Lumpur, Malaysia, in January 2000. But the CIA did not pass this information to the FBI. The FBI handled domestic matters; the CIA dealt with foreign intelligence. An invisible wall stood between the two agencies.

Meanwhile, an agent at an FBI field office had submitted a report that summer noting that Middle Eastern men were receiving training at flight schools across the United States. He warned that this could be connected to a terrorist plot. But the report was buried in someone's desk drawer at FBI headquarters. The laptop computer of Zacarias Moussaoui, who had been arrested in Minnesota, contained a flight simulation program. FBI field agents applied for a warrant to search his computer, but headquarters rejected the request. The reason given was that the wording of the warrant application did not meet legal standards.

The NSA, the National Security Agency, had been intercepting communications throughout that summer indicating that "a massive attack was imminent." They reported this information to the President. The title of the Presidential Daily Brief dated August 6 was "Bin Laden Determined to Strike in US." However, this imminent-attack briefing contained no specific dates, locations, or methods. The NSA considered its mission to be collecting intelligence, not analyzing or connecting it.

It was not that there was no intelligence. Intelligence was overflowing. The CIA, FBI, NSA, State Department, Immigration and Naturalization Service , all of these agencies held a piece of the puzzle. The fact that the terrorists' passports were forged, that they had attended flight schools, that they were connected to al-Qaeda, that they were inside the United States. But these pieces were scattered across different rooms.

Each agency had its own database, used its own classification system, and followed its own security regulations. For a CIA analyst to access FBI investigative records, a complex administrative process was required. In most cases, such attempts were never even made.

This is called an "information silo." A silo is a cylindrical storage structure used on farms to store grain. There is a wall between the silo holding wheat and the silo holding corn. The grain on one side does not flow to the other. That was the state of the American intelligence system before 9/11. Intelligence was collected but did not flow. It was not shared. It was not connected.

Two years later, the 9/11 Commission reached this conclusion: "Our principal failure was a failure to connect the dots." The commission pointed out that this failure was not due to individual incompetence but to the structure of the system. The intelligence system, designed during the Cold War, had been built to face a single adversary: the Soviet Union. But the new enemy, al-Qaeda, was different. They crossed borders, moved like a network, and followed no rules. A bureaucratic structure divided into departmental silos could not counter such an enemy.

After the commission's report was released, the U.S. Congress established the Office of the Director of National Intelligence (ODNI). The purpose was to tear down the walls between intelligence agencies. But changing the organizational chart alone was not enough.

The real problem was the data. How to connect hundreds of millions of documents, videos, communications records, and financial transaction records stored in different formats across dozens of agencies. It was impossible for the human eye. Computers were needed. But the computer systems the government had at the time had been designed in the 1980s. They ran different operating systems, were written in different programming languages, and were incompatible with one another.

One man read this report. He was not an intelligence officer. He was not a soldier, nor a politician. He was an entrepreneur who ran an online payments company in Silicon Valley. His name was Peter Thiel.

(2) Peter Thiel's Vision: "Technology that preserves freedom while ensuring security"

One day in 2002, Peter Thiel was sitting in a cafe in Palo Alto, drinking coffee and reading the newspaper. The interim report of the 9/11 Commission was on the front page. "Failed to connect the dots." Thiel could not take his eyes off that sentence. Because he had already solved that exact problem before.

Thiel had just sold PayPal to eBay for $1.5 billion. PayPal was an online payment service. Millions of people sent and received money over the internet every day. The problem was that a significant portion of those transactions were fraudulent. The Russian mafia, Nigerian scam rings, and credit card thieves were using PayPal to siphon off money. At one point, more than 1% of all PayPal transactions were fraudulent. It was a level that could have destroyed the company.

Thiel and PayPal's engineers developed a system called "IGOR" to solve this problem. IGOR analyzed millions of transaction records to find suspicious patterns. For example, if multiple transactions occurred from a single account in a short period of time, it raised a flag. If multiple different accounts were accessed from the same IP address, it sent an alert. It combined dozens of variables , transaction amounts, time zones, locations, card types, shipping addresses , to assign a fraud probability score. However, IGOR did not make the final decision. The last judgment was left to humans. The system filtered out suspicious transactions, and analysts verified them.

Thiel thought about it. Terrorists are no different from fraudsters. They also wire money. They buy plane tickets. They rent cars. They lease apartments. They make phone calls. They send emails. They cross borders. All of these actions leave digital traces. If the government could connect these traces, couldn't it prevent terrorism before it happened?

Thiel had another concern. He was a libertarian. He detested the government surveilling citizens' private lives. After 9/11, American society was gripped by fear. The Patriot Act was passed. The government gained the ability to collect call records without a warrant. The argument that "some freedom must be sacrificed for security" was gaining traction. Thiel believed this was a false dichotomy.

"If technology works properly, you can maintain security without giving up freedom."

Thiel said this in a Forbes interview. Instead of the government indiscriminately collecting everyone's data, the idea was to selectively analyze only data with grounds for suspicion. And to log who viewed what data and when. If an analyst tried to access information beyond their authority, the system would block it. Access records could later be audited. This way, terrorists could be found while still protecting ordinary citizens' privacy.

Thiel began assembling people to bring this idea to life. The first person he recruited was PayPal engineer Nathan Gettings. Gettings was the key architect of the IGOR system. The second and third were undergraduates at Stanford University: Joe Lonsdale and Stephen Cohen. Both were studying computer science and had previously interned during Thiel's PayPal days. They began developing a prototype in 2004.

Thiel needed one more thing: a CEO to lead the company. Thiel was an excellent investor, but he lacked confidence as a manager. He had no talent for managing complex organizations, negotiating with government officials, or preaching a vision to employees. So he thought of a friend from his Stanford Law School days. Alex Karp.

Karp was a strange choice. He was not an engineer. He could not program. He had no experience running a startup. He was a philosophy PhD. He had received his degree from the University of Frankfurt in Germany, in neoclassical social theory. His doctoral thesis was on "aggression and political identity." At the time, he was running a small hedge fund in London, managing the assets of wealthy families.

The reason Thiel chose Karp was precisely because of his background. A company that works with the government, a company that handles classified data, needed legal and ethical sensibility. Karp had a law degree and philosophical training. He was the kind of person who instinctively asked, "What happens if this technology is misused?" Thiel needed someone like that.

On May 6, 2003, a corporation named "Palantir Technologies" was registered in Delaware. The company's mission was "developing software to support counterterrorism operations." Thiel invested $30 million in initial funding. Karp became CEO. They set up shop in a small office in Palo Alto and began writing code.

But other venture investors in Silicon Valley showed no interest. Investing in a startup that worked with the government was considered insane. Government contracts were slow, complex, and low-margin. Michael Moritz of Sequoia Capital spent the entire pitch meeting doodling in his notebook. An executive at Kleiner Perkins told them, "Your company will definitely fail."

But one place , exactly one place , showed interest. It was In-Q-Tel, the CIA's venture investment arm.

B. The CIA's Venture Investment: The Union with In-Q-Tel

(1) The Meeting of Silicon Valley Engineers and Intelligence Analysts

One day in 2004, Alex Karp was standing in front of an ordinary-looking office building in Arlington, Virginia. He wore a loose jacket and disheveled hair. The people waiting inside were former spies and Defense Department officials in crisply pressed suits. Karp was not nervous. Or perhaps the problem was that he was too relaxed.

In-Q-Tel was a nonprofit venture investment organization established in 1999 by then-CIA Director George Tenet. The "Q" in the name was a witty reference to Q, the technology developer in the James Bond films.

In-Q-Tel's mission was clear: to identify cutting-edge technologies in the private sector and introduce them to intelligence agencies. After the Cold War ended, the CIA had been falling behind technologically. While Google was revolutionizing web search, CIA analysts were still organizing intelligence reports with word processors and Excel sheets.

In-Q-Tel's CEO was Gilman Louie. He was a video game designer. The flight simulation game he created was so sophisticated that it was used for Air Force pilot training. Director Tenet had recruited him precisely for that reason. Louie understood what cutting-edge technology was and knew how to apply it to real-world operations.

Usually, Louie did not attend pitch meetings in person. Projects had to be screened by his staff before reaching his desk. But Palantir's case was different. Peter Thiel was involved. Thiel had already earned his reputation through PayPal, and his money and network could not be ignored. Louie decided to attend the meeting personally.

There was something the Karp and Palantir team did not know. At that time, the CIA was in desperate need of new data analytics technology. The program the CIA primarily used at the time was "Analyst's Notebook," a product from the British company i2. The interface was decent, but it had severe limitations in data processing capability. When handling large volumes of data, the system would slow down or crash. Analysts had many complaints.

Karp began his presentation. Instead of PowerPoint slides, he explained verbally. Like a philosopher, he started by identifying the essence of the problem. The failure of 9/11 was not an absence of intelligence but an absence of connection. A system was needed to bring together data from each agency, find hidden patterns, and present them to analysts. But this system must not become a black box. Analysts must be able to understand why a particular result appeared. And it must log who viewed what data and when.

Louie nodded. This was exactly what he had been wanting to hear.

In-Q-Tel decided to invest approximately $2 million in Palantir. By Silicon Valley standards, it was a small amount. But something was more important than the money. In-Q-Tel's investment stamped Palantir as "a government-validated company." And even more important than that was "opening the door." In-Q-Tel helped Palantir's engineers obtain security clearances. And it placed them before actual CIA analysts.

At first, there were clashes. Silicon Valley engineers and Washington intelligence analysts were people who spoke different languages. The engineers wore shorts and slept on the office floor. They disliked authority, ignored rules, and their creed was "fail fast and fix fast." The analysts, on the other hand, were people who had spent their entire lives within bureaucracy. They valued procedure, respected hierarchy, and a mistake meant the end of a career.

Palantir's first prototype failed miserably in front of CIA analysts. The system was flashy, but it couldn't do what the analysts actually wanted. One analyst asked, "If I press this button, will it tell me where Bin Laden is?" The engineers were left speechless. They understood the technology but did not understand the users.

So Palantir adopted a new approach. It was the concept of the "Forward Deployed Engineer." Engineers left the comfortable offices of California and went into the underground bunkers of Washington and the forward bases of Afghanistan. They sat next to the analysts. They observed how the analysts worked. They asked what was inconvenient. And they fixed the code on the spot.

This was completely different from the approach of traditional defense contractors. Large defense companies like Lockheed Martin or Raytheon would receive a requirements document from the government, spend years developing a product, and then deliver the finished item. If the product didn't work in the field, it took more years to fix. Palantir was different. They built prototypes within weeks, tested them in the field, received feedback, fixed them, and redeployed. They built software for the government at Silicon Valley speed.

Slowly, the analysts' attitudes began to change. They realized that with Palantir's software, they could complete analysis that previously took months in just hours. They identified patterns in where improvised explosive devices (IEDs) were being planted in Iraq. They tracked the flow of Taliban funding in Afghanistan. They visualized the networks of drug cartels. Word began to spread among analysts: "Try Palantir. It's different."

(2) Establishing the Identity of "Google for Government"

In 2008, Palantir released its first official product, "Gotham." The name was taken from Batman's "Gotham City" , the image of a hero fighting crime in the darkness. Gotham was a platform for large-scale data analysis, integration, and visualization. It was designed for intelligence agencies and military operations.

The core of Gotham was search. But it was a different kind of search from Google. Google searches the open web. It quickly finds information that anyone can see. Gotham searched locked data. Classified documents, intelligence reports, wiretap records, financial transaction records, airline reservation records, immigration records. It allowed information scattered across different databases to be viewed on a single screen.

Here is an example. An analyst searches for the name of a specific terror suspect. Gotham pulls up all information associated with that name. The phone numbers that person used. The people called from those phone numbers. The bank accounts of those people. The flow of money in and out of those accounts. The airline tickets that person booked. The people on the same flights. All of this information is visualized as a graph. Dots are connected by lines. The analyst can see the entire shape of the network at a glance.

That is why people called Palantir "Google for Government." The nickname was apt yet incomplete. Google only finds information. Judgment is left to the user. Palantir was the same. But there was one critical difference. On Google, anyone can search for anything. On Palantir, they could not.

Strict access controls were built into the Palantir system. Every user had a clearance level. One analyst might only have access to Database A; another might only have access to Database B. If someone tried to view information they were not authorized to see, the system blocked them. And all activity was logged. Who opened what file and when, what search terms were entered, what results were exported , all of it was recorded. Auditors could later review these logs.

This was the meaning of Peter Thiel's "security without surrendering freedom." Not indiscriminate surveillance, but controlled analysis. Of course, this system was not perfect. Access permissions had to be configured by someone, and that someone could make mistakes or abuse their authority. But at least it was better than the previous system. Before, analysts could rummage through file cabinets as they pleased, and no records were left behind.

By 2010, 70% of Palantir's clients were government agencies: the CIA, FBI, NSA, Marine Corps, Air Force, Special Operations Command, National Center for Missing & Exploited Children. The remaining 30% were private financial institutions. Large banks like JP Morgan began using Palantir

for fraud detection. The technology that started at PayPal had expanded to Wall Street.

Mysterious rumors circulated around Palantir. There was a rumor that Palantir had contributed to the 2011 operation that killed Osama bin Laden. The story was that when U.S. Navy SEAL Team 6 raided the compound in Abbottabad, Pakistan, and killed bin Laden, Palantir's software had been used to track him. Palantir neither denied nor confirmed the rumor. The silence fostered mystique. The mystique brought in new clients.

Palantir defined its identity not as a mere IT service provider but as "a critical infrastructure supplier defending democracy." Alex Karp said in an interview: "We believe software can change the nature of warfare." It was not an exaggeration. In Iraq and Afghanistan, U.S. forces used Palantir to predict insurgent movement routes, analyze IED placement patterns, and track enemy funding flows. Faster intelligence means faster decisions, and faster decisions mean faster operations. On the battlefield, speed was life.

The software created by Silicon Valley's heretics was becoming the nervous system of the world's most powerful military.

C. The Lord of the Rings and Palantir

(1) The Meaning of the Mission: "Seeing Stones"

Peter Thiel was a passionate fan of J.R.R. Tolkien. He had read The Lord of the Rings multiple times since his undergraduate days at Stanford. When naming the company, he recalled a mystical object from the novel. The "Palantir" , plural "Palantiri."

Palantir, in Quenya, the Elvish language Tolkien created, means "that which sees from afar." "Palan" means "far," and "tir" means "to watch." In the novel, the Palantiri are indestructible crystal spheres. Those who possess one can see events happening in distant places, transcending space and time. They can glimpse events of the past. They can communicate with those who hold other Palantiri.

According to Tolkien's mythology, the Palantiri were made by the Noldor Elves of Valinor , perhaps the work of the greatest craftsman, Feanor himself. Eight Palantiri existed in total, and seven were in Middle-earth. The Faithful Numenoreans placed them throughout the kingdoms of Gondor and Arnor. These stones formed a network that surveilled the entire realm. The king could know of distant rebellions in advance and detect intruders at the borders.

Thiel believed this name perfectly symbolized the company's mission. To see through the vast data scattered across the world, detect hidden threats within it, and protect the nation. The tragedy of 9/11 happened precisely because of a failure to "see far." Intelligence agencies could not connect fragmented data and failed to predict the coming attack. Palantir's software aspired to become exactly that "far-seeing eye."

But Tolkien's story also contained a warning. In The Lord of the Rings, the Palantir was a double-edged sword. Saruman, the White Wizard, tried to use the Palantir to monitor the enemy's

movements. But he gazed too deeply. The Dark Lord Sauron held another Palantir, and Saruman's mind was dominated by Sauron. He was corrupted and became Sauron's servant.

Denethor, the Steward of Gondor, also used a Palantir. He sought to understand the enemy's forces to defend his kingdom. But Sauron used the Palantir to show Denethor visions of despair. Denethor saw the overwhelming enemy army and lost all hope. In the end, he threw himself into the flames and died.

Tolkien scholar Jane Chance Nitzsche analyzed it this way: Saruman fell because he chose "mere knowledge" and abandoned "true wisdom." The Palantir gave information, but it did not give the wisdom to interpret that information correctly. Instead, it caused the loss of judgment in a flood of information.

Thiel and Karp were fully aware of this literary warning. They knew that the technology they were building could become like Saruman's Palantir. Powerful surveillance capability, when misused, becomes a tool of totalitarianism. That is why they placed access controls and audit logs at the core of the system. The principle was: "We can see everything, but we do not look at what should not be seen."

Palantir's offices also followed Tolkien's world. The Palo Alto headquarters was called "The Shire," the homeland of the Hobbits. The McLean, Virginia, office was "Rivendell," the Elven refuge. The Washington, D.C., office was "Minas Tirith," the capital of the Kingdom of Gondor. Every conference room bore a Tolkien place name. New employees received The Lord of the Rings as required reading.

This was not mere branding. It was Palantir's identity itself. They likened themselves to "the alliance of free peoples defending Middle-earth." The enemy was Sauron , terrorists and totalitarian states. Their weapons were not swords and bows but code and algorithms. And the Palantir they built was the "seeing stone" that protected Western civilization.

Critics saw it differently. They warned that Palantir could itself become the Eye of Sauron. A totalitarian tool for the state to surveil its citizens. The irony of destroying freedom in the name of protecting it. The duality embedded in the name "Palantir" became a debate that ran through the company's entire history.

(2) Seventeen Years of Silence: The Journey from Unicorn to Public Listing

On Wednesday, September 30, 2020, the opening bell rang at the New York Stock Exchange (NYSE). The ticker "PLTR" appeared on the screen for the first time. It was the moment Palantir Technologies' stock began trading on the public market. Starting from a reference price of $7.25, it opened at $10 and closed at $9.50. The market capitalization was approximately $16.5 billion , $21.7 billion on a fully diluted basis.

Palantir was founded in 2003. It took a full 17 years to go public. Considering that Silicon Valley startups typically go public or get acquired within 5 to 10 years, this was an exceptionally long time.

Why did it remain private for so long?

The first reason was the nature of its clients. Palantir's main clients were the CIA, NSA, FBI, and the Department of Defense. Spies and special forces operators. They intensely disliked having the software they used made public. Becoming a public company means disclosing financial statements. Revenue breakdowns must be revealed. Which government agencies paid how much could be exposed. The clients did not want that.

The second reason was Alex Karp's philosophy. Karp despised Wall Street's pressure for short-term results. He refused to change strategy to meet analysts' quarterly expectations. He thought in terms of 10 and 20 years. The goal was to build the technological infrastructure that defends Western civilization. Such a long-term project was difficult to reconcile with quarterly earnings reports.

In 2013, Karp said in a media interview: "Going public would make it very difficult to run a company like ours." That December, Palantir raised $450 million from private investors. The company was valued at $9 billion. Forbes called Palantir "one of the most valuable private technology companies in Silicon Valley."

In 2014, it raised an additional $50 million, and the valuation rose to $15 billion. In 2015, it secured $880 million in funding and reached a $20 billion valuation. The term "unicorn" was no longer sufficient. Palantir had become a "decacorn" , a company worth ten times $1 billion.

Behind the flashy numbers, there was pain. Palantir recorded losses of hundreds of millions of dollars every year. A net loss of $598 million in 2018. A net loss of $588 million in 2019. It was not making a lot of money , it was spending a lot of money. The problem was the business model. For a long time, Palantir operated like a consulting firm. It dispatched engineers to client sites to custom-install software and modify it to meet each client's specific needs. This was an expensive approach. Acquiring a single client required millions of dollars and months of time.

A turning point came in 2016. Palantir launched a new platform called "Foundry." Unlike Gotham, Foundry was designed for private-sector companies. It was a tool that allowed companies in manufacturing, finance, healthcare, and energy to integrate and analyze their own data. Foundry was more standardized than before. Customization was reduced, and deployment time was shortened from months to days. The cost structure began to improve.

In 2018, Morgan Stanley valued Palantir at $6 billion , a sharp drop from $20 billion in 2015. Investors' patience was running thin. How long could one keep betting on a company that had been posting losses for 17 years? Inside Palantir, discussions about going public began.

Then in 2020, the COVID-19 pandemic hit. Ironically, the pandemic was an opportunity for Palantir. The British National Health Service (NHS) used Palantir's Foundry for vaccine distribution and bed management. The U.S. Department of Health and Human Services used Palantir software called "Tiberius" to manage vaccine allocation. Suddenly, Palantir had become a core supplier of national health infrastructure. Revenue surged.

In September 2020, Palantir finally entered the public market. But instead of a traditional IPO, it chose a "Direct Listing." In an IPO, the company issues new shares, investment banks set the price, and the company pays fees to the banks. In a direct listing, existing shareholders sell their shares directly on the market. No new capital is raised, but there are no commissions to pay to Wall Street intermediaries.

Karp chose this method deliberately. He despised Wall Street conventions. The day before the listing, he released a video message filmed in the woods of New Hampshire. He was wearing cross-country ski gear and a helmet. Removing his helmet, he said: "We believe our software makes the world safer. If you don't agree with that belief, don't buy our stock."

On the first day of trading, Karp said in a CNBC interview: "We have built the foundation on which the company carries substantial value. Entering the public market will help our relationships with clients and accelerate growth."

Five years later, by September 2025, Palantir's stock price had risen 1,700% from the first day of trading. The market capitalization surpassed $430 billion. Quarterly revenue broke through $1 billion for the first time. The number of clients grew from 125 in 2020 to 849. Headcount also increased by more than 1,500.

Seventeen years of silence and patience were rewarded. Peter Thiel's original vision , "technology that preserves freedom while ensuring security" , was valued at hundreds of billions of dollars on Wall Street. But whether the future Palantir reveals is truly bright, or whether a dark shadow looms over it as Tolkien warned, that judgment remains in the reader's hands.

← Previous Chapter

See every chapter

Next Chapter →

Kim Kyung-jin

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com



#KimKyungjin #PALANTIR #Palantir #WarAI #SurveillanceAI #DefenseAI #DataPower #ArtificialIntelligence #AILibrary

# Chapter 2. The Philosopher CEO: Alex Karp's Worldview

← Previous Chapter

See every chapter

Next Chapter →

Part 1: The Heretic of Silicon Valley

A. From the Frankfurt School to Defense Company CEO

(1) Why a Neo-Marxist Philosophy PhD Leads a Defense Company

In the late 1990s, winters in Frankfurt, Germany, were cold and damp. In a worn-out cafe near the Goethe University campus, an American graduate student was sipping coffee and reading a thick German-language academic text. His name was Alex Karp.

He could have graduated from Stanford Law School and become a lawyer, but instead he crossed the Atlantic to the birthplace of philosophy. Just a few blocks from where he sat stood the Institute for Social Research, the home base of the Frankfurt School. It was the very building where Theodor Adorno and Max Horkheimer had returned to teach after fleeing the horrors of Nazism into exile.

Karp had his reasons for choosing Frankfurt. He wanted to understand how power works. Studying law, he had grasped the forms of power, but the essence of power remained a mystery. Why do people obey? Why do some societies slide into totalitarianism? Why does a civilization that pursued rationality produce barbarity like Auschwitz? These questions drew him to Frankfurt.

His doctoral advisor, Carola Brede, was an expert in Freudian psychoanalytic theory. Karp had initially intended to study under Jurgen Habermas, but for some reason, the two parted ways. In 2002, he finally received his PhD. The title of his dissertation was "Aggression in the Lebenswelt." It was an abstruse work exploring how human aggression permeates everyday language and culture.

Had the story ended there, Karp would have become a professor teaching critical theory to students in some university lecture hall. But he felt that the world of academia could not answer the questions he harbored.

"Academic discussions are interesting only to those within them." That is what he said in a later interview. The world did not operate according to the theories in books. Power was not an abstract concept but a concrete force. Data, information, and technology were becoming new sources of that force.

During his Stanford days, he had befriended a unique individual named Peter Thiel. Thiel was a former chess prodigy turned law student who, like Karp, saw the world differently. The two debated endlessly about philosophy and politics. Thiel was a libertarian, and Karp called himself a socialist. They appeared to stand at opposite ends of the political spectrum, but the two agreed on one point:

freedom does not preserve itself.

On September 11, 2001, the world changed. Thiel called Karp. He had already become wealthy co-founding PayPal. What Thiel proposed was a new company. The idea was to take the technology that caught credit card fraud at PayPal and use it to build software that tracks terrorists. Karp hesitated at first. Joining hands with intelligence agencies was, from the perspective of the critical theory he had studied, an act of serving state power. It was precisely what the Frankfurt School's teachers had warned against.

But Karp decided to think differently. What he had learned in Frankfurt was not that power itself is evil. It was that power becomes evil when it is uncontrolled. Nazism was born not because power existed but because of the concentration and abuse of power. The question was not how to eliminate power but who holds it. If the power of information and data was going to emerge regardless, should it not be in the hands of democratic nations? That was his conclusion.

In 2003, Palantir was founded. The CEO position was given to Karp. A philosophy PhD who could not code leading a technology company defied Silicon Valley's common sense. But Thiel judged that what Karp possessed was more important than coding ability.

It was the ability to ask questions. The ability to ask not what technology can do, but what technology should do.

Karp had another secret. He had dyslexia. The condition, which makes letters appear to dance, had set him apart since childhood. Because he could not absorb text at face value, he had to think for himself. In a 2025 interview with Fortune, he said: "If you don't have dyslexia, when you read a text, the text becomes you. The more you read, the more the text controls you. If you have dyslexia, it doesn't work that way." This cognitive independence became Palantir's culture , a culture that does not follow consensus but moves according to conviction.

A neo-Marxist philosophy PhD leading a defense company appears paradoxical. But for Karp, this was not a betrayal but a logical conclusion. He had learned the fear of totalitarianism in Frankfurt. His way of putting that fear into practice was to place the most powerful technological weapons in the hands of democratic nations. He left the theory in the lecture hall and chose to transform reality through code.

(2) "Technology Is Not Neutral": Fear of Totalitarianism and Defending Western Democracy

At the 2024 World Economic Forum in Davos, Alex Karp addressed the audience: "Technology is not neutral." He stood on stage in athletic wear, his hair sticking out in every direction. Before the world's financiers and politicians, he directly repudiated Silicon Valley's mythology.

Silicon Valley's engineers had long regarded technology as a value-neutral tool. A knife in a chef's hands is a cooking tool; in a robber's hands, it becomes a weapon. Algorithms are the same , a good tool when used for good purposes, a bad tool when used for bad ones. Therefore, the technologist's responsibility is to make the tool well; how the tool is used is the user's affair. This logic is convenient

because it frees technologists from moral responsibility.

Karp called this a naive delusion. Sophisticated technologies like artificial intelligence and big data are different from a knife. Such technologies restructure society in and of themselves. The way data is collected, the algorithms used for analysis, and the form in which results are presented already embody specific values. For instance, the moment one determines that certain data is "important" and other data "can be ignored," value judgments intervene. Those judgments are etched into code, and code changes reality.

This perspective directly connects to what he learned at the Frankfurt School. In Dialectic of Enlightenment, Adorno and Horkheimer warned that reason and technology could degenerate from tools of human liberation into instruments of domination and control. Nazi Germany was living proof of that warning. Hitler's empire was barbaric yet simultaneously thoroughly rational. Train schedules were precise, and the bureaucracy was efficient. Auschwitz demonstrated the horrifying truth that industrial-scale massacre requires a sophisticated logistics system.

Karp never forgot this history. What he fears is not a science-fiction future in which artificial intelligence dominates humans. It is the present, in which the methods by which humans dominate other humans are becoming more sophisticated. Data and algorithms make that sophistication cheap. China is already surveilling its citizens through its social credit system. It tracks Uyghurs with facial recognition technology. This is the reality of the "non-neutrality of technology" that Karp speaks of. The same technology is used to catch terrorists in democratic nations and to suppress dissidents in authoritarian ones.

In Karp's logic, the crucial next step is this: if technology is not neutral, then who holds that technology is decisive. He puts it this way: "If we don't build it, they will." While the democratic camp hesitates over AI development, authoritarian states are charging ahead unchecked. If one is paralyzed by moral deliberation, one hands the technological advantage to an adversary unencumbered by moral constraints. In Karp's words, "the naive idea that our adversaries are more moral because they are less adept at AI" is dangerous.

In response to criticism that Palantir's software could infringe on civil liberties, Karp offers a rebuttal. His argument is that if terrorism goes unchecked and society is gripped by fear, citizens will ultimately surrender their freedom in exchange for security. A terrified public delegates unlimited power to authoritarian leaders. This is what happened in 1930s Germany. Therefore, to preserve democracy, security must be secured first. And for that security, the most advanced technology is needed.

Karp does not hide the fact that Palantir's technology "is sometimes used to kill people." In a letter to investors at the time of the 2020 IPO, he wrote directly: "Our software is used to target terrorists and keep soldiers safe." In the war in Ukraine, Palantir's technology was used to track Russian tank positions and transmit those coordinates to artillery units. When reports emerged that the Israeli Defense Forces were using Palantir's platform in the Gaza Strip, Karp publicly declared he would stand with Israel.

Such positions naturally provoked fierce criticism. Critics call Palantir "a tool of the surveillance state." Civil liberties organizations denounce its contract with Immigration and Customs Enforcement (ICE). On university campuses, protesters blocked Palantir recruitment events. Some of Karp's own employees resigned in opposition to the company's direction.

Karp does not retreat from this criticism. He appears to regard being criticized as evidence that he is doing the right thing. In a 2024 New York Times interview, he warned that the United States is likely facing "a three-front war against China, Russia, and Iran." The best way to prepare for this war, he argued, is to develop autonomous weapons. His logic was that while the West and the authoritarian axis are approaching technological parity, the West is at a disadvantage because it has moral constraints against using nuclear weapons.

Karp's worldview is straightforward. The world is dangerous, freedom is fragile, and good cannot be protected without strength. This is a philosophical journey that begins in the Frankfurt School's critical theory and lands in cold-eyed realism. That is why his leading a defense company is not a contradiction. Fear of totalitarianism is the engine that drives him, and defending Western democracy is his destination. Technology is the fuel for that journey.

B. The Break with Silicon Valley's Mainstream Culture

(1) Big Tech's Refusal to Cooperate with the Pentagon and Palantir's Choice

On April 4, 2018, tension hung in the air at Google headquarters. A petition signed by more than 3,000 employees was delivered to CEO Sundar Pichai. The petition's title was clear: "Google should not be in the business of war." At issue was "Project Maven," a Department of Defense contract.

Project Maven was a program the Department of Defense launched in April 2017. Its goal was to use artificial intelligence to analyze vast amounts of drone-captured video footage and identify targets. The volume of footage pouring out daily from unmanned aircraft flying over Iraq and Syria exceeded what human analysts could handle. If AI could do this work instead, the time needed for target identification would be dramatically reduced.

Google agreed to provide its TensorFlow AI technology. The contract was not large , under $70 million in the first year. But for the employees, this was not a matter of money. Google's longstanding motto, "Don't be evil," was at stake. They could not accept the fact that their code was being used to improve the accuracy of drone strikes.

The petition read: "Google will be placed on the same line as companies like Palantir, Raytheon, and General Dynamics." At the time, Palantir was an object of contempt among Google employees. It was perceived as a shadowy company that partnered with intelligence agencies to build surveillance technology. The idea of Google being grouped in the same category as such companies was intolerable.

The protests intensified. At least 12 employees resigned. More threatened to strike if the company continued developing military products. Finally, in June, Google's management capitulated. They

announced they would not renew the Project Maven contract. Officially, the contract ended in March 2019.

This incident symbolized Silicon Valley's culture. For California tech companies, the government and military were seen as outdated, inefficient, and sometimes immoral entities. They dreamed of a borderless internet world and wanted to define themselves not as American companies but as global companies. Having the mission of "making the world more connected" placed next to the phrase "building weapons that kill enemies" was a brand catastrophe.

At that very moment, one company made the opposite choice. It was Palantir.

Alex Karp publicly criticized Google's decision. He pointed out that Silicon Valley's elites had forgotten where the freedom and prosperity they enjoyed came from. Companies that grew on the foundation of security maintained by the U.S. military and the rule of law, yet refused to help maintain that very foundation, were guilty of hypocrisy. He said: "While sitting in the safe offices of California enjoying moral superiority, someone else is risking their life fighting on the front lines."

Palantir did not hesitate to fill the vacancy Google left behind. In 2019, Palantir became Project Maven's new partner. In May 2024, the Department of Defense awarded Palantir a sole-source contract worth $480 million for the development of the "Maven Smart System" prototype. This contract was later expanded to nearly $1.3 billion.

Karp's choice was considered "suicide" in the tech industry at the time. At recruitment events, students held protest demonstrations, and investors shook their heads. The company's reputation within Silicon Valley hit rock bottom. But Karp did not back down. Instead, he turned this conflict into an opportunity. He established Palantir's identity as "the antithesis of Silicon Valley."

In the prospectus (S-1) filed at the time of the 2020 IPO, Karp wrote: "Our software is used to target terrorists and keep soldiers safe. If we are asking someone to put themselves in harm's way, we will provide them with what they need to carry out their mission."

We have an obligation. We chose a side." Including such a statement in a securities filing was unprecedented. It was a political declaration in a document sent to investors.

Karp dismissed the attitude of American companies that do business with the Chinese Communist Party yet refuse to work with their own military as "the logic of losers , incomprehensible and unsustainable." He criticized Silicon Valley for wasting its resources "satisfying the fickle demands of consumer culture." While the brightest minds are devoted to "trivial conveniences" like ride-sharing apps, photo-sharing platforms, and advertising optimization algorithms, the truly important issue of national security is being neglected.

This choice made Palantir not just a software vendor, but a company with an ideological identity. When everyone else was exploiting user data for advertising revenue, Palantir declared it would protect and use data for national security. This was simultaneously a business strategy and a culture war.

(2) Headquarters Relocation (from Palo Alto to Denver) and Patriotic Technologism

One day in August 2020, Palantir quietly packed up. The company had decided to leave Palo Alto, the heart of Silicon Valley where it had been based for 17 years since its founding. The destination was Denver, Colorado , a place where the rugged wilderness of the Rocky Mountains unfolds.

When the news broke, people speculated about the reasons. Economic factors like high rents, taxes, and cost of living were mentioned. Denver's cost of living was far lower than San Francisco's. Colorado's tax structure was also business-friendly. Geographic proximity to major clients like the Department of Defense and Northern Command (NORTHCOM) was likely also a consideration.

But the reason Alex Karp gave was different. He had already foreshadowed it in a May 2020 interview with Axios: "I've grown sick of the increasing intolerance and monoculture of Silicon Valley." He said he would move the headquarters to a location closer to the East Coast than the West Coast. He added that Colorado was a strong candidate.

The term "monoculture" was significant. Silicon Valley pursues diversity technologically, but is remarkably homogeneous politically. Most tech industry workers hold progressive views. There is an atmosphere in which displaying patriotism or supporting the military is treated as unsophisticated, or sometimes immoral. Companies that cooperate with the Pentagon are considered morally inferior.

Karp saw this culture as unhealthy. He criticized Silicon Valley for being detached from the values of America as a whole, trapped in its own ideological bubble. In the IPO prospectus, he wrote: "Our company was founded in Silicon Valley. But we seem to share fewer and fewer of the technology sector's values and commitments."

The move to Denver was the physical manifestation of this cultural separation. The new headquarters was located in the Sugar Cube building at 1555 Blake Street in downtown Denver. Palantir had already been operating a small team there for two years. The official headquarters relocation was announced ahead of the public listing. The timing was deliberate , it was a message to investors: "We are different from Silicon Valley."

Denver was an appropriate choice in many respects. Politically, Colorado is a "purple" state. Democrats and Republicans compete neck and neck. Unlike San Francisco's overwhelmingly progressive orientation, diverse political views coexist. also, Colorado already had an established defense industry ecosystem. Traditional defense companies like Lockheed Martin, Raytheon, and Northrop Grumman had offices there. The headquarters of Space Force and Northern Command were also nearby.

The move was not a Palantir-specific phenomenon. It was part of a larger trend. From the late 2010s, several technology companies began leaving Silicon Valley. Oracle and Hewlett Packard Enterprise went to Texas. Tesla's Elon Musk also left California. High taxes, strict regulations, and soaring costs of living were the surface reasons, but cultural fatigue could not be ignored either.

Karp himself did not stay in one place even after the headquarters move. He is known to work from a barn in the woods of Lyman, New Hampshire. When reports emerged that he owns ten homes around the world, he answered as if joking: "That should be redefined as ten cross-country ski cabins." He sometimes appears at video conferences wearing ski goggles. Such eccentricities demonstrate that he is deliberately distancing himself from the stereotypical Silicon Valley CEO image.

Palantir's relocation symbolizes a new current called "Patriotic Technologism." This term embodies the philosophy that technology is not merely a tool for making money or satisfying engineers' intellectual amusement, but must serve the survival of the nation and society. Karp emphasizes that the values of the American Midwest , diligence, integrity, and devotion to country , are more important than Silicon Valley's elitism.

He argues that the roots of the American software industry lie in the Midwest. Bob Noyce, co-founder of Intel, was from Iowa. Many of the engineers who led the early semiconductor industry came from the Midwest. They had a culture that was open, collaborative, and committed to national goals. From Karp's perspective, Silicon Valley has forgotten these roots. Only borderless capital and profit maximization remain.

This trajectory opened the way for other defense startups. Figures like Palmer Luckey of Anduril and Elon Musk of SpaceX began shaping a culture that takes pride in "building weapons for America." In December 2024, Anduril and Palantir announced a consortium for collaboration on Department of Defense projects. A new generation of defense tech companies is filling the seats that Silicon Valley refused to take.

Palantir in Denver has become a bridge connecting the severed link between Silicon Valley and Washington. They transfused the hot blood of patriotism into the cold logic of technology. To critics, this looks like dangerous militarism; to supporters, it looks like necessary realism. The taste varies depending on the reader's political constitution.

C. "We Chose a Side": The Declaration as Guardians of Western Civilization

(1) "Peace through Strength" and a Realist Security Vision

There is an old unwritten rule in the business world: do not make enemies, because customers could be anywhere. Coca-Cola sells cola to communists, and Apple sells iPhones to citizens of dictatorships. Political neutrality is a survival strategy for global companies.

Palantir broke this unwritten rule spectacularly.

The opening line of the prospectus filed just before the 2020 IPO shocked investors: "We have chosen sides. We support the United States and its allies." It was a political declaration unseen in ordinary corporate filings. Karp hammered in that the company would never do business with adversary nations of the United States, such as China or Russia. The Chinese market with its 1.4 billion people is a vast opportunity, but for Palantir, it is off-limits territory.

Many investors expressed concern. Voluntarily limiting one's market could be seen as undermining shareholder value. But for Karp, this was a non-negotiable principle. He believes that helping regimes incompatible with Western values is a betrayal of corporate ethics.

Karp's security vision is thoroughly realist. He does not believe that peace can be maintained through goodwill or dialogue alone. "Peace through Strength" is his creed. This concept, used by President Ronald Reagan during the Cold War, holds that strong military capability and technological superiority deter war and maintain peace.

Karp reinterprets this logic for the 21st century. Just as the atomic bomb brought about "the long peace," he argues that artificial intelligence will become the new deterrent. He says: "The nuclear age is ending. We are at a similar crossroads in computing science." In the era of large language models, the United States needs "a new Manhattan Project to maintain exclusive control over the most sophisticated forms of artificial intelligence on the battlefield."

Karp frequently quotes historian Samuel Huntington: "The West won the world not by the superiority of its ideas or values or religion, but rather by its superiority in applying organized violence." Karp adds that Westerners often forget this fact, but non-Westerners never do. This is a realist recognition that moral superiority alone is not enough; it requires physical force to back it up.

This belief was put into practice during the Ukraine war. When Russia invaded in 2022, Karp was one of the first Western corporate CEOs to fly to Kyiv. He met President Zelensky and offered to provide Palantir's technology free of charge. This was not a business meeting but a declaration of participation in the war. Palantir's software played a decisive role in enabling the Ukrainian military to defend against Russia's large-scale forces.

After Hamas attacked Israel on October 7, 2023, Karp did not hesitate. He publicly declared, "We will stand with Israel." When reports emerged that the Israeli Defense Forces were using Palantir's platform, some employees pushed back and resigned. Karp acknowledged this but did not change his position.

In a 2024 New York Times interview, he warned that the United States likely faces "a three-front war against China, Russia, and Iran." The best way to prepare for this war, he argued, is to develop autonomous weapons. He also expressed support for military conscription. His logic was that while the West and the authoritarian axis are approaching technological parity, the West is at a disadvantage because it has moral constraints against using nuclear weapons.

Karp's rhetoric is sometimes aggressive. He says Palantir's technology "should instill fear in our adversaries." In February 2025, at an event promoting his book, he said: "I like the idea of sending drones with light fentanyl-laced urine to spray on the analysts who tried to ruin us." Even as a joke, it was a chilling expression.

But his logical structure is consistent. Deterrence is the art of making the adversary believe in your will. To achieve that, you must have actual capability. And you must demonstrate the willingness to use that capability. For Karp, Palantir's very existence is that message. "We hope our software brings

fear to our adversaries and reassurance to our allies."

Critics worry that such a worldview promotes war. But Karp's counterargument is clear. If America loses its technological edge and appears weak, events like Russia's invasion of Ukraine or Hamas's attack on Israel will occur more frequently. Only credible deterrence can maintain order in a chaotic world.

(2) The Vision of the Technological Republic

In February 2025, a book co-authored by Alex Karp and Palantir's head of corporate communications, Nicholas Zamiska, was published. The title was The Technological Republic: Hard Power, Soft Belief, and the Future of the West. The book rose to No. 1 on the New York Times bestseller list immediately after publication. NPR selected it as a Book of the Year.

George Will of The Washington Post wrote in his review: "Since Allan Bloom's The Closing of the American Mind in 1987, there has been no cultural critique this comprehensive." The Wall Street Journal called it "a searing cry against the tech industry's abandonment of its history of helping America and its allies."

The book's core argument is straightforward: Silicon Valley has lost its way. In the 1950s and 60s, Pentagon funding drove Silicon Valley's early growth. The internet, GPS, and semiconductor technologies all originated from military research. The engineers of that era were dedicated to solving national challenges. But at some point, that spirit vanished.

Karp and Zamiska criticize the fact that Silicon Valley's finest minds today are being wasted building "photo-sharing apps and marketing algorithms." "Narrowly pursuing the demands of a late-capitalist economy has become their vocation." The question is whether optimizing chicken finger delivery algorithms is really what the best engineers should be doing.

The book diagnoses this complacency as having spread beyond Silicon Valley to the entire West. "Intellectual weakness" has become pervasive in academia, politics, and corporate boardrooms. Rather than pursuing truth, people seek to avoid uncomfortable conversations. Rather than enduring ideological confrontation, they choose ambiguous neutrality. Karp warns that this is eroding the West's competitiveness.

The core concept of The Technological Republic is the combination of "Hard Power" and "Soft Belief." Hard Power is the physical capacity to build and defend , military strength, industrial base, and technological superiority fall under this category. Soft Belief is the moral clarity that guides what to build and why to defend it. Karp diagnoses the West as losing both.

The book warns that the democratic world faces an existential choice. While authoritarian regimes are integrating artificial intelligence into military, surveillance, and economic systems, democratic nations are dithering. They face "the risk of re-establishing their own national project based on democratic values, institutional strength, and technological excellence, or becoming geopolitically irrelevant."

Karp's prescription is clear. The partnership between government and the technology industry must be rebuilt. The public-private collaboration model that made possible the Manhattan Project, the Moon landing, and the development of the internet in the 20th century must be restored for the 21st century. Through agencies like DARPA (Defense Advanced Research Projects Agency), "asymmetric power" must be created to deter adversary nations.

The book also provides a window into Palantir's internal organizational culture. Karp explains that Palantir's structure draws inspiration from the honeycomb of bees. It is non-hierarchical, relies on collective decision-making, and uses group voting dynamics. It adopted the root cause analysis method learned from Taiichi Ohno, the creator of the Toyota Production System. When problems arise, surface-level assumptions are avoided, and the real reasons for failure (often rooted in interpersonal dynamics) are sought.

Critics express concern about the vision presented in The Technological Republic. The argument for fusing military and technological power and placing software engineers at the center of national defense could be read as a vision for a "militarized technological oligarchy." There are concerns about the paradox of actually strengthening authoritarian control under the guise of defending democratic values.

One reader on Goodreads wrote: "The book suggests that citizenship means conformity, technology means weaponry, innovation means militarization, dissent means disloyalty, and the republic itself is a garrison state built to Palantir's specifications." Another critic assessed the book as "more a manifesto of frustration than evidence-based strategy."

Karp pays no heed to such criticism. His conviction remains unshaken. "The importance of defending Western values and opposing those who would undermine them" is his passion. He asserts that these values are "inherently superior." He sees America as "a unique entity embodying meritocracy and fairness, capable of uniting individuals from diverse backgrounds."

In 2025, Karp was named to Time magazine's list of the 100 Most Influential People in the World. Time described him as "a new kind of Silicon Valley billionaire: the avatar of techno-nationalism, unapologetically projecting Western power." The Economist named him CEO of the Year for 2024. Palantir's market capitalization at one point exceeded $500 billion. Karp's personal wealth reached $18 billion.

Whether the vision of The Technological Republic will be realized remains to be seen. But one thing is clear. Alex Karp has presented the most controversial yet influential vision of the relationship between 21st-century technological capitalism and democracy. His journey from philosophy PhD to billionaire defense company CEO is a living case study of how power and values intersect in the age of technology.

← Previous Chapter

See every chapter

Next Chapter →

Kim Kyung-jin

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com



#KimKyungjin #PALANTIR #Palantir #WarAI #SurveillanceAI #DefenseAI #DataPower #ArtificialIntelligence #AILibrary

# Chapter 3. Ontology: The Magic of Turning Data into Knowledge

← Previous Chapter

See every chapter

Next Chapter →

Part 2: Core Technology, Ontology and the Decision-Making Revolution

A. The Heart of Palantir's Technology: Defining Ontology

(1) Implementing a "Digital Twin" Beyond Simple Data Integration

Late one night in 2008, an emergency meeting was held at the headquarters of a major bank in London's financial district. Thick reports were piled on the table. The CEO asked: how much is our bank exposed to Greek sovereign debt? It was a straightforward question. The chief risk officer opened his laptop and called another department. An hour passed. Wiping the sweat from his brow, he answered: I'm not sure exactly, but it seems to be around $1 billion. The CEO was despairing. One of the world's premier financial institutions didn't even know where its own money was.

The problem was not a lack of data. The bank had complete data. It was just written in different languages. On the lending department's books, Greece was labeled "GR." On the bond trading desk, it was recorded as "Hellas." The derivatives department used a numerical country code. The computers did not know that these three things pointed to the same country. This was the fundamental problem Palantir sought to solve.

Palantir's engineers found inspiration in an unexpected place to solve this problem: philosophy books. The concept they seized upon was Ontology. It is a philosophical term meaning the study of being , the discipline that explores what exists in the world and how those things relate to one another. Palantir took this concept, first established by the ancient Greek philosopher Aristotle, and made it the core principle of its software.

Palantir's official documentation defines ontology as "creating a complete picture of an organization's world." It maps datasets and models into concepts like object types, properties, links, and actions, reconstructing real-world things and events within the platform. It serves as a translator that appears as business language to human eyes and as a computable structure to machine eyes. That is why ontology does not end as a mere catalog or schema design tool. It becomes a resilient foundation for end-user workflows.

This is where the concept of the digital twin enters. Many companies misunderstand the digital twin as merely a flashy 3D display , a pretty rendering of a factory in virtual space. The digital twin Palantir speaks of is something entirely different. It is not decoration. It is a working replica that contains both the semantic system and the behavioral system of an organization. According to Palantir documentation, the ontology serves as the organization's digital twin in many environments,

encompassing not only semantic elements like objects, properties, and links, but also dynamic elements like actions, functions, and dynamic security. There is a reason it bears the name "twin." When reality changes, the screen must change; when the screen changes, the next action must change.

In the spring of 2025, the American fast-food chain Wendy's ran short of syrup. It was a problem affecting 6,450 stores. In the past, it would have taken 15 people working all day to resolve. Pete Surkein, CEO of Wendy's Quality Supply Chain Cooperative, said at a Palantir customer conference: "We solved the problem in five minutes." How was that possible? Because Wendy's digital twin had connected 3,500 trucks, 34 distribution centers, and 6,450 stores on a single screen. The moment the signal that syrup was running short came in, the system instantly calculated which truck from the nearest distribution center should deliver via which route.

The true power of the digital twin lies in its writability. In advanced deployment environments, Palantir's digital twin functions not as a mere mirror but as a control panel. Changes made in Foundry are backpropagated to downstream systems. It can update ERP status, trigger work orders, or reconfigure supply plans. This is the decisive difference that distinguishes a static digital shadow from a true digital twin. Organizations leap from the level of "seeing what is happening" to "continuously re-optimizing what is happening."

In Q1 2025, Palantir's U.S. commercial segment revenue increased 71% year-over-year. In Q2 of the same year, quarterly revenue surpassed $1 billion for the first time. Market capitalization exceeded $430 billion. The Economist noted that Palantir may be the most overvalued company in history, as it was valued at more than 600 times its 2024 earnings. These numbers tell us something: the market is placing astronomical value on the concept of the digital twin. It is betting money not on technology that gathers data, but on a method that defines reality.

(2) Reconstructing the World Through Objects, Properties, and Links

Alex Karp often asked his technologists: do you see the world like an Excel spreadsheet, or like a story? Most data engineers see the world in rows and columns. For them, the world is a giant spreadsheet. But Palantir reconstructs the world in sentences. They use three elements: objects, properties, and links. It is like the subjects, adjectives, and verbs taught in elementary school language class.

First, let us look at objects. Objects are the nouns that exist in the world. In Palantir's documentation, an object type serves as the blueprint for that category of entity. Take military operations as an example. A tank is an object. A platoon leader is an object. An enemy supply truck is an object. In legacy systems, these were scattered across different database tables. The personnel department's database stored the platoon leader's serial number; the logistics department's database stored the tank's specifications. Palantir places all of these on a single map. In private enterprises, sensors, assets, work orders, and employees are each defined as independent object types. In healthcare settings, patients, visits, diagnoses, and medications become objects.

Next come properties. Properties are the adjectives that describe objects. According to Palantir documentation, a property is a schema definition representing a characteristic of some real-world entity or event. The object "tank" has a property called "fuel level." It also has a property called "current location." A property called "last maintenance date" can be added as well. Palantir supports various primitive data types: strings, integers, floating-point numbers, dates, timestamps. Advanced types like geographic coordinates, geometries, media references, structs, and time series are also included. Up to this point, it does not look very different from a conventional database.

The magic happens with the third element: links. Links are verbs. Palantir's system connects objects to each other with lines and names those lines. The platoon leader rides the tank , that is a link. The tank targets the enemy truck , that is also a link. The enemy truck transports ammunition. A link type is the schema definition of the connection between two object types; a link is a single instance of the relationship between two objects. When these simple sentences accumulate, a grand narrative is constructed. Previously, analysts had to imagine these relationships in their heads. They had to find the platoon leader in the personnel file, find the tank in the logistics file, and then cross-reference enemy positions on a map. Palantir's ontology engine predefines these relationships.

Now the system can reason. This is the revolutionary point. Here is an example. The system discovers an object called "enemy truck." That truck has stopped near an object called "ammunition depot." It is then moving toward the front line. The system combines these relationships to infer a new fact: this truck is likely en route to resupply ammunition. On the commander's screen, what appears is not a mere dot but an operational recommendation: "Supply route interdiction needed." This is the foundation of what Palantir calls software-defined warfare.

The German pharmaceutical company Boehringer Ingelheim applied this principle to drug development. They built an ontology connecting cross-team data and relating terms like targets, genes, and diseases. They created an enterprise knowledge graph on top of their data lake. As a result, approximately 90% of R&D data could be served through a unified one-stop space via a semantic layer. When a scientist searches by gene or disease, relevant data from across silos can be instantly reviewed. There is no longer any need to manually combine data from different sources.

Among Palantir's solution architects, there is a common saying: "The ontology is essentially our company's nervous system." There is another one too: "Once you've mapped something into the ontology, you never again need to argue about what a customer is or what an asset is."

One data engineer put it this way: "I spent two years building data models in Databricks, and I threw it all away in two weeks with the ontology." Data is no longer dead numbers. It becomes sentences made of nouns and verbs, and sentences accumulate into context. Translating the complex real world into a language that computers can understand , that is the heart of Palantir's technology.

B. Destroying Data Silos

(1) The Principle of Unifying Fragmented Data into a "Single Source of Truth"

In the early 2000s, on the battlefields of Afghanistan, U.S. forces were fighting an invisible enemy. But their greater adversary was internal: the login screen. For an intelligence officer to determine a terrorist's location, they had to access more than ten different programs. There was a separate program to view drone footage. There was a separate system for accessing wiretap records. Reports from human intelligence sources, called HUMINT, were filed in paper binders. By the time an analyst tried to synthesize all this information, the terrorist had already fled.

This is the tragedy of data silos. A silo originally refers to the cylindrical storage structures used to store grain. The wheat in one silo does not mix with the barley in the silo next to it. The same was true of organizational data. Each department and agency used different systems, and data formats varied. According to one analysis of the pharmaceutical industry, about 48% of drug development executives reported that data silos undermined cross-departmental efficiency. Teams spent valuable time reconciling which numbers were correct. Resources were wasted not on extracting insights but on verifying data integrity.

There is a scene that commonly unfolds in meeting rooms. The sales team's numbers differ from the finance team's numbers. The plant team says "our system is fine." The logistics team says "the field is a nightmare." Everyone is telling the truth , just different truths. A silo is not malice but structure. When systems are separated, languages separate too. And separated languages slow an organization down.

Palantir did not demolish these silos. Instead, it built bridges over them. Many companies, when attempting data integration, try to tear out all existing systems and build one massive new database from scratch. This is called a data warehouse. But this approach takes years and costs astronomical sums. Above all, it faces fierce resistance from the field. Each department does not want to give up the tools they are familiar with. Palantir's approach was different. It chose virtual integration.

The principle works like this: existing databases are left in place. HR continues using Oracle. Finance continues using SAP. Instead, Palantir's Foundry platform connects to all of these systems. It does not physically move the data. It brings over the meaning of the data. This is called mapping. It tells the system that Customer ID in System A and Client Number in System B actually refer to the same person. Users can then access all data through a single window called Foundry.

This is called a Single Source of Truth, or SSOT for short. It is a data management concept in which important data is stored and updated in a central repository. By routing all information through one authoritative source, it ensures everyone works from the same playbook. There is no longer a need to fight over data sources in the meeting room. The revenue figure brought by the sales director will not differ from the one brought by the finance director. On Foundry, everyone sees the same numbers. When someone modifies data, the change is reflected in real time for all users.

The biotechnology company Biogen applied this principle to its manufacturing floor. They built a single source of truth for manufacturing data, treating in-process data as a single source. The results were remarkable. Instead of waiting for a separate end-of-line test silo, continuous verification became possible. Biogen's head of global analytics said: "The combination of time-series data with

manufacturing context enabled advanced applications. Machine learning models could now rapidly assess batch quality."

The economic impact is dramatic. According to research by Veeva Systems, in a siloed environment, a simple protocol amendment required 25 or more manual steps and multiple documents. Using an integrated vault turned it into a single-source, single-step operation. Update times were cut from weeks to minutes. Drug development teams could spend less time searching for data and more time on actual innovation.

It becomes a single-source, single-step operation. Update times were cut from weeks to minutes. This demonstrates how important an integrated data flow across the entire drug lifecycle is for eliminating redundancy and maintaining a single version of the truth for each drug. Palantir extended this integration to all of an organization's data assets. It created an environment where data, logic, and decision-making coexist. Competing versions of data were eliminated, and it was ensured that everyone uses the same correct numbers.

(2) The Design of Metadata, Lineage, and Permissions Systems

Centralizing data is convenient. But it is dangerous. Being able to see all information also means that anyone could steal or manipulate it. Palantir grew up alongside intelligence agencies from its birth. For them, security was not a feature but a matter of survival. That is why they treated data about data , metadata , as more important than the data itself.

Every piece of data in the Palantir system has a tag attached. When was this data created? Who made it? What was the original source? Palantir explains that the ontology can include rich metadata about every field and fine-grained governance over changes. It is a method of recording who defined what, when it was changed, and why it was changed. This record is not an ethical ornament. When an incident occurs in the field, data stands in court too.

Palantir's 2024 Privacy and Governance white paper states: understanding data lineage allows administrators to visualize data flows across the platform. Reading from left to right, one can see how data in the platform flows from ingestion through data transformation to platform applications. A comprehensive view of data lifecycle and interactions provides a clear picture at scale.

Lineage , this concept is like a data resume. It is a mechanism for tracking the path data takes through a system. It visualizes the entire journey from data's origin to its final use. Palantir's data lineage tools provide a holistic view of how data flows through the Foundry platform.

Here is an example. A general is viewing a map on a tablet that shows the predicted movement route of enemy forces. He is curious about how this map was created. With a single click, he can trace the ancestry of this information. He can learn that this route was predicted by AI, and that the AI analyzed it based on drone photos taken 30 minutes ago and wiretapped radio transmissions intercepted an hour ago.

Palantir's 2024 10-K report notes that Foundry identifies the reason data projects fail as the difficulty in understanding and reproducing the steps and methods of building pipelines. To address this, it enables users to track pipelines and follow traces to understand what the rows and columns of a table mean and why they are there. The moment you ask "why is it there," data becomes not numbers but evidence.

This lineage changes politics as well. When someone says "that figure is wrong," the other party can trace the lineage back to find where it went wrong. Arguments become tracing exercises. In the pharmaceutical field, this means a researcher can trace analytical results back to the original instrument files and laboratory notebook entries. Audit trails are maintained at each step. With a proper audit trail in a central clinical data repository, what was submitted to the FDA can be traced back to source data, and post-submission updates are also tracked.

The last element is the permissions system. When silos are broken down, the first fear people think of is: "So will everything be visible?" Palantir says it places fine-grained and flexible security controls within the ontology itself. Permissions can be applied not only to design elements like object types and link types, but also to actual object and link data.

Palantir's data protection and governance documentation lists specific tools. Checkpoint is a Foundry application that enables justification requests before certain sensitive data operations. Sensitive Data Scanner allows administrators to create organization-specific sensitive data definitions and policies for how to handle information when it is identified. Cipher Service obfuscates data through cryptographic operations such as encryption, decryption, and hashing.

This permissions system is not simply a matter of whether you can or cannot open a file. Permissions can be set down to the smallest unit of data. Suppose a CIA analyst and an FBI investigator are looking at the same terrorist file. On the CIA officer's screen, the terrorist's name and overseas hideout information are visible. But the informant's real name is redacted. On the FBI agent's screen, the terrorist's criminal record within the United States is visible, but overseas operational information is not. This is called purpose-based access control. It is not that all data is available just because a user has high rank. Only the data necessary for the specific mission being performed can be viewed.

Palantir tore down the walls of data while simultaneously building the safest vault. It resolved the contradiction between openness and security through technology. Palantir's AIP is built as a traceable and auditable system. It captures a complete audit trail to ensure trust and accountability in responsible human-machine collaboration. It can fully track which data an AI model used, where that data came from, and who had access rights.

C. Human-Machine Teaming

(1) Not Black-Box AI, but a Tool that Amplifies Human Intuition

Many AI companies in Silicon Valley try to exclude humans. Their goal is full automation , because humans make mistakes, get tired, and are emotional. But Palantir's philosophy is the opposite. Alex Karp asks: would you give an algorithm the authority to pull the trigger on the battlefield? Ethics

aside, it is strategically foolish as well. AI is skilled at reading patterns but poor at reading context.

Palantir Technologies was founded in 2003 with a mission born from the national security challenges after 9/11: analyzing vast, heterogeneous intelligence datasets to prevent terrorist attacks without undermining civil liberties. Co-founder Peter Thiel envisioned a mission-driven company that would apply software principles similar to PayPal's fraud recognition system to this complex problem. From inception, the company's philosophy was not about replacing human analysts with omniscient AI. It was about intelligence augmentation , providing human decision-makers with superior tools to navigate complexity.

Palantir rejects black-box AI. A black box is a system where input and output exist, but the process in between is unknowable. Most cutting-edge AI systems based on deep learning are like this. Even the developers do not know why AlphaGo placed its stone in a particular position. With the game of Go, that is fine. But with launching a missile or denying a loan, the story is different. It must be possible to explain why a given decision was made.

Palantir's approach employs two core techniques. The first is chain of thought , Chain-of-Thought prompting. AIP logic functions are configured by default to use this prompting. It causes large language models to respond iteratively according to a structured plan. The prompt instructs the LLM to first outline a plan, then execute the plan step by step, and finally provide a final answer. AIP's LLM Debugger gives users insight into the inner workings of AIP logic functions, allowing them to see intermediate steps and tool calls.

The second is a tool delegation strategy. Instead of processing the entire task within the LLM's black box, AIP delegates specific logical tasks to more interpretable tools. This approach allows reliance on trustworthy and more interpretable logic. If the AI system begins exhibiting unexpected behavior due to tool usage, the specific logic and steps of that tool can be inspected to identify where the logic fails.

An experiment was conducted in the medical field. When explainable AI was used instead of black-box AI , specifically, when heatmap explanations for chest X-ray images were provided , radiologists' work performance improved significantly. Explainable AI went beyond merely providing the AI's prediction score. It offered visual explanations of why a particular judgment was made, enabling medical professionals to critically evaluate the AI's recommendations and combine them with their own expertise.

Palantir's philosophy views AI not as a tool that replaces humans but as one that amplifies human intuition and expertise. It is like the Iron Man suit. When Tony Stark puts on the suit, the suit does not fight on its own. The suit maximizes Tony Stark's capabilities. In 2025, this concept was actually validated in the U.S. Air Force's ShOC-N Capstone experiment. Palantir's Maven Smart System and Maverick AI were tested for integration into the dynamic targeting process to improve speed, scale, and accuracy. The joint integrated experiment improved combatant decision advantage and overall situational awareness.

In the 2024 Forrester Wave AI/ML Platform Q3 report, Palantir was named a Leader. It received the highest ranking in current offering. The report stated: Palantir has one of the strongest offerings in AI/ML and has a vision and roadmap for creating a platform that integrates humans and machines into a co-decision-making model. Palantir's Chief Architect Akshay Krishnaswamy said: Palantir AIP powers the most demanding use cases in both the public and private sectors, and is uniquely designed to connect AI directly to frontline operations. We believe our investment in multimodal guardrails for human-AI collaboration, decision-centric ontology, and the full range of capabilities enterprises need to move from AI prototypes to production has been validated.

(2) Transparent Algorithms for Decision Advantage

Whether in war or business, the essence is the same: a continuous stream of decisions. The side that decides faster and more accurately wins. This is called Decision Advantage. In the past, the side with more information won. Now it is different. Information overflows. The battle now is over who can find meaningful signals in that flood more quickly.

Take the Ukrainian battlefield as an example. A reconnaissance drone spots an enemy tank. In the past, it would have taken dozens of minutes for this information to reach the artillery battalion. Coordinates had to be called out, marked on a map, and radioed in. In the Palantir system, the moment the drone sends its video, AI identifies the tank. It recommends the most suitable friendly artillery within firing range. The time this takes is mere seconds.

But the final button is pressed by a human. The system asks on screen: do you wish to strike? The commander reviews the weapon recommended by the AI and the estimated collateral damage, then approves. The algorithm is transparent. It shows with data why it recommended this artillery unit and why it marked these coordinates. The commander becomes not a slave of the machine but the conductor of the machine.

Palantir's 2024 10-K report describes AIP as providing an interface that enables organizations to operationally use AI and LLMs. It mentions safe handoffs between AI agents and human operators, extensive security and audit controls, and integrated human review checkpoints across workflows. Humans are not designed to be spectators applauding from behind. They are designed as brakes that can stop things at any point along the way.

Palantir's writings on AI governance point in the same direction. They tie critical decisions to human-in-the-loop gates, restrict data and applications visible to each role through access controls, and emphasize data minimization and purpose-based controls. This may look like an ethics statement, but in practical terms, it is a product requirement called "auditable automation." Automate, but leave clear accountability. If accountability is not preserved, automation is banned within the organization.

The core of transparent algorithms is the ability to track and audit AI's decision-making process. Palantir's AIP operates on the principle of building AI systems that are interpretable, understandable, and transparent. AI systems must not become black boxes. Users must understand how the system

works in order to build trust. Palantir's approach ensures that AI cannot independently execute actions such as military targeting operations or major financial transactions. Nothing is executed without explicit human approval.

Transparency is not merely an ethical requirement but a practical necessity. Human-AI collaboration research shows that information asymmetry and capability asymmetry are two key sources of complementary team performance. The best collaboration outcomes occur when humans and AI each possess different information and capabilities. To effectively use this, each agent must know what information the other uses and what strategy it adopts.

The Human-Machine Collaboration Wargame for Decision Advantage, conducted in July 2025, was an important milestone in validating this concept. This series of experiments represents critical advances in AI integration and joint integration in battle management. By fusing vast sensor data with machine learning and human oversight, it improved situational awareness and accelerated military decision-making.

Palantir's ontology is the infrastructure that provides transparency. Data is ingested into Foundry or Gotham, contextualized by the ontology, and activated by AIP, enabling humans and AI to run simulations and propose actions. When decisions are made, those actions and outcomes are recorded back into the ontology. This feedback loop enriches the digital twin over time, making future AI-based recommendations more accurate and enabling organizational leaders to learn from the effects of their decisions.

The long-term value of transparent algorithms is the facilitation of organizational learning. Palantir's closed-loop system creates a compounding advantage where the more an organization uses the platform, the smarter and more indispensable it becomes. This is not simply about collecting more data. It is about systematically capturing the context of decisions, their outcomes, and lessons learned, and reflecting these in future decision-making. Decision advantage does not come from a one-time technological edge but from a culture of continuous learning and improvement. Palantir provides the technological infrastructure that makes such a culture possible.

In 2025, Palantir's annual revenue was approximately $4.4 billion , a 53% increase year-over-year. Operating margin reached 51%. It had no debt and held more than $4.5 billion in cash. As of January 2026, the stock price was trading between $185 and $195. These numbers tell us something: the market is placing enormous value on the concept of decision advantage. Technology computes; humans decide. And the speed of that decision determines victory or defeat. This is what Palantir calls the proper relationship between technology and humans.

← Previous Chapter

See every chapter

Next Chapter →

Kim Kyung-jin

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com



#KimKyungjin #PALANTIR #Palantir #WarAI #SurveillanceAI #DefenseAI #DataPower #ArtificialIntelligence #AILibrary

# Chapter 4. Palantir's Four Major Platforms

← Previous Chapter

See every chapter

Next Chapter →

Part 2: Core Technology, Ontology and the Decision-Making Revolution

A. Gotham: The Watcher in the Dark

(1) A Person-and-Relationship Tracking System for Counterterrorism Operations and Intelligence Agencies

One day in 2008, in an analysis room deep beneath CIA headquarters, a young officer was staring at her screen. Her mission was to track the network of al-Qaeda operatives. In the past, this would have required manually searching dozens of databases. But on the screen, the relationship web between individuals was unfolding in real time like a spider's web. When she clicked on one person, his call records, financial transactions, travel routes, and known associates were connected in a single picture. This was Gotham's first operational deployment.

Gotham was the first platform Palantir launched, in 2008. The name was taken from Gotham City in the Batman franchise. It was no coincidence. In Christopher Nolan's film The Dark Knight, Lucius Fox builds a system that surveils all of Gotham City. In the movie, Batman destroys the system after the mission is complete. The real-world Gotham was not destroyed. Seventeen years later, in 2025, Gotham has become even more powerful.

Gotham's core function is person-centric analysis. The platform integrates data from diverse sources to construct a complete picture of a specific individual or organization. Phone call records. Email metadata. Financial transaction histories. Immigration records. Social media activity. Satellite imagery. Drone footage. All of these fragments are reconstructed into a unified view through ontology technology.

Gotham's interface integrates multiple analysis tools into a single workspace. Geospatial analysis visualizes the movement patterns of specific individuals or vehicles on a map. Network analysis represents relationships between individuals as nodes and links, revealing hidden connections. Call record analysis detects abnormal activity in phone call patterns. AI and machine learning algorithms analyze all of this data in real time, automatically detecting patterns that an analyst might miss.

Gotham also features Mixed Reality capabilities. Commanders can collaborate in virtual operations rooms and visualize battlefield situations in three dimensions. Satellite integration enables acquisition of real-time reconnaissance imagery from anywhere in the world. Algorithms automatically coordinate which sensors should collect which data.

Security is the core of Gotham. The platform is equipped with multilayered permission systems including classification levels, the need-to-know principle, and time-limited access. Every query and analysis activity leaves a complete audit trail. Who saw what and when is recorded. This is a mechanism for ensuring accountability, but paradoxically, it is also a mechanism for surveilling the surveillance system itself.

Gotham's client list is not publicly disclosed. However, certain known user agencies exist: the CIA and NSA, among other intelligence agencies; the U.S. Army and Navy; Europol. The Danish police have been using Gotham since 2017 for a predictive policing project called POL-INTEL. In 2025, Denmark decided to expand Gotham's use to its military and intelligence agencies.

The most notable recent contract is the $10 billion deal with the U.S. Army concluded in mid-2025. This 10-year contract consolidated 75 individual software contracts into one. Project Maven has also been expanded. This contract, valued at $795 million in 2025, could grow to as much as $1.3 billion by 2029.

Gotham is not a mere analysis tool. As of 2025, the platform supports AI-based kill chains. It automates the process from target identification to strike asset assignment. Of course, the final decision remains with humans. The human-in-the-loop principle is still maintained. But every step before that is handled by algorithms.

This is the essence of Gotham. An eye that sees everything in the darkness. The question is what that eye is watching, and who controls it.

(2) The Truth and Fiction of the Bin Laden Kill Operation (Neptune Spear) Support Claim

In the early hours of May 2, 2011, two U.S. special forces helicopters landed in a residential compound in Abbottabad, Pakistan. Forty minutes later, Osama bin Laden was dead. The code name of the operation, carried out by Navy SEAL Team Six, was "Neptune Spear." After the operation, a question arose: was Palantir involved?

There is no official answer to this question. Neither the CIA nor Palantir has confirmed it. But there is circumstantial evidence.

The tracking of bin Laden was the result of years of intelligence gathering. Between 2002 and 2005, the CIA learned from al-Qaeda prisoners about the existence of a courier named Abu Ahmed al-Kuwaiti. It took another six years to track this courier. In 2010, CIA officers finally located him in Pakistan and followed him to the residential compound in Abbottabad.

The likelihood that Palantir's Gotham was used during this tracking process is high. The CIA had been one of Gotham's earliest clients since 2008. Person tracking, relationship network analysis, and geospatial analysis are precisely what Gotham was designed for , integrating dozens of data sources to construct a unified picture. The bin Laden operation required exactly such capabilities.

Palantir has never directly claimed this connection. However, it has implicitly used it. In the company's marketing materials and investor presentations, the phrase "success in counterterrorism operations" frequently appears. It is a way of stimulating the audience's imagination without mentioning specific operation names.

The truth lies within classified secrets. Immediately after the 2011 operation, Admiral William McRaven, commander of Special Operations Command, ordered all files related to the bin Laden operation deleted from Defense Department computers and transferred to the CIA. Physical evidence has never been released. No photographs, no video, no DNA test results have been made public. All FOIA requests have been denied.

What is certain is this: the bin Laden operation was a case where "connecting intelligence" was the key to victory. The ability to assemble fragmented clues into a single picture , this was exactly what U.S. intelligence agencies desperately needed after 9/11, and what Palantir promised to provide. The operation's success suggests that promise was, to some extent, fulfilled.

Whatever Palantir's role was in the bin Laden kill operation, the operation contributed to strengthening the mythology of the company. The image of "technology that catches terrorists" became one of the company's most powerful weapons. Whether that image is based on fact or carefully cultivated remains shrouded in mystery.

B. Foundry: The Central Nervous System of Enterprise

(1) The Structure of a Data Integration Platform for the Private Sector

In 2016, Palantir stood at a critical turning point. Gotham had achieved success with intelligence agencies and the military, but the government market alone had limits for growth. The company needed to enter the private enterprise market. But companies were reluctant to apply software used by the CIA to their own businesses. The surveillance technology image was the problem. A new brand was needed.

That is how Foundry was born. The name, meaning "a place where raw materials are cast," carries the meaning of transforming raw material (data) into valuable products (insights). The core technology was the same ontology engine as Gotham, but the packaging was completely different. Instead of surveillance, "digital transformation." Instead of tracking, "operational optimization." The language changed.

To understand Foundry's structure, think of three layers.

The first is the data connection layer. Enterprises operate dozens, sometimes hundreds, of different systems , ERP, CRM, production management, logistics, HR. These systems typically do not communicate with each other. Foundry connects to all of these systems and ingests data. Whether through real-time streaming or periodic batch uploads, it ingests regardless of data format or source.

The second is the transformation layer. Ingested data is cleaned and standardized. Throughout this process, data lineage is tracked. Where each piece of data came from and what transformations it underwent are all recorded. This ensures auditability and enables root cause tracing when errors occur.

The third is the ontology layer. This is where the magic happens. Cleaned data is transformed into business objects. For an airline, for example, "aircraft," "route," "passenger," "maintenance record," and "part" each become objects. These objects have properties. An aircraft object has properties such as model, manufacture date, total flight hours, and current location. And objects are connected by links. This aircraft operates on this route, is composed of these parts, and has these maintenance records.

This is the "digital twin" , replicating the complex operations of the real world in the digital world. However, this replica is updated in real time, and simulation and analysis are possible.

Foundry's strength is "closed-loop" operation. It does not end with data analysis. Analysis results are immediately connected to real actions. For example, once a supply chain optimization analysis is completed, the results can be automatically reflected in the ordering system. The time from insight to action is dramatically shortened.

As of 2025, Foundry is recording explosive growth in the U.S. commercial sector. U.S. commercial revenue in Q1 2025 increased 71% year-over-year. The number of clients grew 43% to 849. There were 139 contracts exceeding $1 million, 51 exceeding $5 million, and 31 exceeding $10 million. Customer retention rate reached 98%.

The secret to this growth lies in the "Bootcamp" strategy. Traditional enterprise software sales took months. Palantir compressed this to five days. Potential customers are invited to experience Foundry with their own actual data. A working prototype is built in five days. The customer sees the value firsthand and then decides to purchase. This method has dramatically shortened the sales cycle.

(2) Application Cases in Manufacturing, Finance, and Healthcare

An engineer at Airbus headquarters stared at the screen. Data from thousands of aircraft operating worldwide was streaming in real time. Engine temperatures, fuel consumption, component wear status. An algorithm analyzed the patterns and raised an alert: a specific aircraft model's landing gear component was wearing out faster than expected. A replacement schedule was automatically generated before the part could fail.

This is Skywise , an aviation data platform jointly developed by Airbus and Palantir. Launched in 2017, more than 100 airlines currently participate in this platform. Aircraft operational data, maintenance records, and parts inventory information are all integrated. With predictive maintenance now possible, airlines have been able to reduce delays caused by unplanned maintenance.

In the energy sector, BP is a representative case. BP used Foundry to build a "digital twin" of its global supply chain. The entire process from crude oil production to refining, transportation, and sales

is monitored in real time. Algorithms automatically perform demand forecasting, inventory optimization, and logistics route planning. During the 2022 energy crisis, this system played a critical role in rapidly responding to volatile market conditions.

The healthcare case is controversial. In November 2023, the UK's NHS signed a contract with Palantir worth 330 million pounds over seven years. The purpose is to build a "Federated Data Platform." The goal is to connect patient data from hospitals across England to optimize bed management, surgical scheduling, and resource allocation.

The contract provoked immediate backlash. The British Medical Association (BMA) and the Doctors' Association UK raised concerns about patient data privacy. Cybersecurity experts questioned the transparency of the procurement process. In April 2024, medical professionals protested outside NHS England headquarters. Their demand was for the NHS contract to be canceled because of Palantir's contract with the Israeli Defense Forces (IDF).

As of early 2026, only about 15% of NHS trusts are actually running the platform. The adoption rate is slower than expected. Political and ethical debates, more than technical issues, are holding it back.

In manufacturing, the partnership with Samsung is drawing attention. In 2025, Samsung began a project in collaboration with Palantir to improve semiconductor yield and quality. The goal is to analyze the massive volumes of data generated in semiconductor manufacturing processes, identify causes of defects, and optimize process variables.

In Korea, collaboration with HD Hyundai is also underway. Targeting the digital transformation of shipbuilding, Palantir, HD Hyundai, and Siemens are jointly pursuing a project to automate shipyard operations. KT has taken on the role of partner to spread Foundry and AIP across Korean industry.

Foundry's success cases are impressive. But criticism also exists. Foundry's ontology is a proprietary and closed structure. Once adopted, it is difficult to leave. While data export functionality exists, the business logic embedded in the ontology cannot be taken along. This is another reason explaining the 98% customer retention rate.

C. Apollo: A Revolution in Software Deployment

(1) Seamless Deployment from Cloud to Battlefield Edge

"We bring SaaS to places SaaS has never been."

This is the phrase Palantir uses to describe Apollo. The back seat of a Humvee. Inside a submarine. A satellite. The idea of cloud software operating in such environments was hard to imagine until recently. Apollo made it possible.

To understand why Apollo exists, one must look back at Palantir's history. When Gotham was first deployed in 2008, cloud computing was in its early stages. Most enterprise software was installed on the customer's own servers. Government agencies and the military, in particular, could not use external clouds for security reasons. Software upgrades were rare and painful.

There was a problem. Palantir's software consists of hundreds of individual services. Each service is managed by an independent development team. Deploying and updating these services simultaneously across dozens of different environments , cloud, on-premises, classified networks, and fully isolated air-gapped environments , was a nightmarish task.

Apollo was created to solve this problem. The core idea is a "hub-and-spoke" architecture. There is a central Apollo hub, and spokes (deployment platforms) are installed in each deployment environment. The hub manages all versions of all products. The spokes "pull" updates suited to their own environment.

The "pull" method is important. Traditional software deployment uses a "push" method: developers push out updates. This assumes the developer knows the state and constraints of every environment. As environments multiply, management becomes impossible.

In Apollo's pull method, each environment defines its own constraints. This environment only follows "Release Channel Stable." That environment also accepts "canary" versions. Apollo automatically orchestrates updates while respecting these constraints.

The numbers prove it. Apollo processes more than 41,000 automated updates per week. Lead time for applying changes is under four minutes. Blue-green deployment and automatic rollback capabilities ensure stability. If a problem occurs, the system can immediately revert to the previous version.

Operation in edge environments is Apollo's true differentiator. Environments where internet connectivity is intermittent or nonexistent. Environments where bandwidth is extremely limited. Apollo is designed to function even in such settings. It operates with minimal resources and synchronizes with the central hub when connectivity is restored.

Consider why this matters in military operations. A commander in the field may need to conduct operations for days without connection to headquarters. During that period, the Palantir software installed on his equipment must continue to function. If new threat intelligence becomes available, it must be updated the moment connectivity is restored. Apollo makes this possible.

(2) Lock-in Effects and the SaaS-Based Revenue Model

Apollo is not only a technological achievement but also a business model.

In the traditional enterprise software model, customers purchased and installed software. Upgrades were optional. Many companies used the same version for years. This was unfavorable for software companies. If the customer did not see value in the new version, there was no additional revenue.

The SaaS (Software as a Service) model solved this problem. Customers subscribe instead of purchasing. Software runs in the cloud, and updates are applied automatically. Customers always use the latest version. The company earns continuous subscription revenue.

The problem was that Palantir's main customers , intelligence agencies, military, and heavily regulated companies , could not use the traditional SaaS model. Data could not leave for an external cloud. But reverting to on-premises installation meant losing SaaS's advantages.

Apollo created a new model called "Private SaaS." Software runs in the customer's environment, but updates and management are centrally automated like SaaS. The customer maintains data sovereignty while continuously receiving the latest features. Palantir maintains subscription-based revenue.

The lock-in effect arises here. Palantir software deployed through Apollo becomes deeply integrated into the customer's operational processes. Data pipelines, analytics workflows, and decision-making processes are all built on top of the platform. Over time, this dependency deepens.

What would it take to replace Palantir software with something else? All data integrations would need to be rebuilt. The ontology would need to be redesigned. Users would need to be retrained. The risk of operational disruption would need to be accepted. For most organizations, this is a practically impossible choice.

Average contract duration of 3 to 4 years. A 98% customer retention rate. A 124% net dollar retention rate. These numbers reflect the lock-in effect. A net dollar retention rate exceeding 100% means existing customers are spending more money each year.

Apollo is also available on the AWS Marketplace and Microsoft Azure Government Cloud. It has obtained IL5 and IL6 certifications, enabling operation on classified networks. It also holds FedRAMP certification. These certifications represent barriers to entry that competitors cannot easily overcome.

Apollo is Palantir's unsung hero. It does not receive as much attention as Gotham or Foundry, but it is Apollo that enables these two platforms to operate simultaneously across thousands of environments worldwide. By solving the seemingly mundane problem of software deployment, Palantir has been able to plant its software in places no other company has reached.

D. AIP (Artificial Intelligence Platform): Deploying Generative AI in the Field

(1) Controlling LLM Hallucination and Applying It to Real Work

In April 2023, Palantir announced a new platform: AIP, the Artificial Intelligence Platform. It was a tool for integrating large language models (LLMs) like ChatGPT into enterprise environments. In the announcement video, CEO Alex Karp said: "We are landing LLMs in the real world."

The word "landing" was apt. LLMs have remarkable capabilities, but they also have a serious problem: hallucination , the phenomenon where a model confidently states information that is not true as though it were fact. If an LLM occasionally says something wrong in a casual conversation, you can laugh it off. If an LLM hallucinates during a military operation, medical diagnosis, or financial decision, it becomes a disaster.

Palantir's solution was the ontology. Once again, the ontology.

In AIP, the LLM does not generate answers freely. It is "grounded" in the ontology. When the LLM receives a question, it first searches for relevant data in the ontology. Answers are generated based on verified enterprise data. Sources are cited. Users can verify which data the AI's answer came from.

Here is an example. A logistics manager asks: "What is the current inventory status of the Chicago warehouse?" A typical chatbot would answer based on training data , probably outdated information. An AIP agent retrieves data from the real-time inventory system connected to the ontology and responds. The answer includes the last data update time.

AIP supports multiple LLMs: OpenAI's GPT series, Anthropic's Claude, Google's models, Meta's Llama, and NVIDIA's Nemotron. Customers can select the appropriate model for their use case. For security-sensitive cases, self-hosted models running on-premises can be used.

In August 2024, Palantir announced a partnership with Microsoft. Palantir's entire product suite will be deployed on the Microsoft Azure Government Cloud. This means GPT-4 can be used on classified networks for the first time. Military and intelligence agency analysts can now harness the power of large language models.

Trust is core to AIP's design. The platform provides detailed audit trails for all AI decisions. Who asked what question, what data the AI referenced, and what logic it used to generate the answer , all recorded. There is also a "proposal-based" pattern. The AI does not execute actions directly but only makes proposals. Execution requires human approval.

However, Palantir itself acknowledges the limitations of LLMs. Its documentation warns: "Synthetic reasoning carries the risk of hallucination." Ontology grounding is effective for data retrieval, but errors can still occur in the process where the model interprets and reasons with data. Users must understand this limitation.

(2) Agent-Based Task Automation and Evolution Toward an "AI Operating System"

2023 and 2024 were the years of the chatbot. 2025 is the year of the agent.

Let me explain the difference. A chatbot answers questions. An agent acts. Ask a chatbot "What's the weather tomorrow?" and it tells you the weather. Tell an agent "If it rains tomorrow, change the outdoor event schedule" , it checks the weather, and if rain is forecast, it actually changes the schedule.

AIP Agent Studio makes this possible. Users can build agents. An agent combines an LLM, the ontology, documents, and custom tools. Agents are context-aware. They dynamically execute workflows.

Let us look at real cases. A logistics company can deploy an agent that automatically reroutes deliveries when storms are forecast. A hospital can use an agent that manages nurse schedules in real time. An energy company can run an agent that automatically adjusts the power grid during

demand surges.

In the defense sector, the Maven Smart System has evolved into a fully integrated AI target acquisition and situational awareness tool. It is now used as a standard by multiple NATO member nations. The TITAN (Tactical Intelligence Targeting Access Node) ground station serves as the "brain" of the U.S. Army's AI-defined battlefield.

Let us look at AIP's other components. AIP Logic is a no-code LLM function builder. Even business users who are not developers can create LLM-based functions. AIP Assist is a chatbot available across the entire platform. It can search not only Palantir documentation but also custom documents uploaded by customers. Pipeline Builder can include LLM nodes to automate data classification, sentiment analysis, summarization, and more.

The numbers speak to AIP's impact. In Q1 2025, Palantir signed 139 contracts worth $1 million or more. 51 contracts over $5 million. 31 contracts over $10 million. A significant number of these contracts are related to AIP.

Palantir positions AIP as an "AI operating system." Just as an operating system is the middle layer between hardware and applications, AIP aims to become the middle layer between raw data and AI applications. Any LLM, any data source, any application connects through AIP.

CTO Shyam Sankar led the transition to agentic AI in 2025. According to his vision, future enterprise operations will be a collaboration between humans and AI agents. Agents handle repetitive and data-intensive tasks. Humans focus on decisions requiring judgment and creativity.

Of course, risks exist. Allowing AI agents to act autonomously means relinquishing some control. If an agent makes a wrong decision, who is responsible? How do you ensure an agent's actions are ethical? These questions have not yet been fully answered.

Palantir's answer is, again, transparency and auditability. All agent actions are recorded. Actions requiring approval are distinguished from those that can be performed autonomously. However, while technology advances rapidly, governance and regulation are struggling to keep pace.

AIP is Palantir's present and future. If Gotham and Foundry laid the company's foundation, AIP is the growth engine. Since its 2023 launch, the company's stock has soared. It was added to the S&P 500 at the end of 2024. In November 2025, the stock price reached an all-time high of $207. Market capitalization surpassed $400 billion.

This is what it means to deploy generative AI in the field. Beyond chatbots, AI that acts in the real world. Controlling hallucination, grounding in data, making it auditable. Palantir is presenting one answer to this difficult problem. Whether that answer is the only one, or the best one, is still too early to judge.

← Previous Chapter

See every chapter

Next Chapter →

Kim Kyung-jin

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com



#KimKyungjin #PALANTIR #Palantir #WarAI #SurveillanceAI #DefenseAI #DataPower #ArtificialIntelligence #AILibrary

# Chapter 5. Ukraine: The Testing Ground for AI Warfare

← Previous Chapter

See every chapter

Next Chapter →

Part 3: AI on the Battlefield, Software-Defined Warfare

A. How David Defeats Goliath: Accelerating the Kill Chain

(1) Real-Time Integration of Satellite, Drone, and HUMINT Intelligence

On a morning in October 2022, at an undisclosed location on the outskirts of Kyiv, two Ukrainian officers sat before a laptop. On the screen, a digital map of the Bakhmut front was spread out. The distance between the Russian and Ukrainian trenches was barely a few hundred meters. When a technician clicked the mouse, a thermal image appeared. Another click. This time, a photograph of a Russian tank marked with the letter Z appeared, taken by a ground scout through a fence. Had this been an actual operation, they would have chosen a missile, artillery, or armed drone and issued an attack order within minutes.

Washington Post reporter David Ignatius, who witnessed this scene firsthand, wrote: "This could be a revolution in warfare."

What he saw was Palantir's software. The system integrates multiple eyes on a single screen: images from commercial satellites, real-time video from military drones, photos uploaded by frontline scouts via smartphones, enemy artillery positions detected by counter-battery radar, and even intelligence obtained by hacking Russian surveillance cameras. In the past, each piece of data was trapped in different departments, different systems. It was common for one agency to know something that another did not. It was the exact same problem U.S. intelligence agencies experienced during the 9/11 attacks.

Palantir CEO Alex Karp publicly stated that the company's technology is responsible for a significant portion of Ukrainian targeting intelligence. The company established an office in Kyiv and dispatched engineering teams. They are constantly experimenting and improving on the battlefield.

Modern warfare is a race against entropy. When decisions are delayed, targets move, fire, and hide. When Ukraine received precision weapons like HIMARS from the West, the success of those weapons depended less on range than on the speed and quality of data. Palantir's technology is called "software that compresses time" , because it dramatically shortens the time from target discovery to analysis to strike.

Ukraine also possesses its own domestically developed system called DELTA. Started in 2016 by a volunteer organization called Aerorozvidka, this system was officially deployed across the entire defense and security sector in August 2024. With tens of thousands of users, DELTA integrates

satellite images, drone footage, radar data, cyber intelligence, and even information reported by citizens through Telegram chatbots. In July 2024, DELTA passed a NATO information security audit , the first Ukrainian military system to receive NATO protocol certification.

A 2025 CSIS report assessed that DELTA could rival the U.S. CJADC2 (Combined Joint All-Domain Command and Control) system in terms of design objectives and operational flexibility. Lieutenant Colonel Boyko, who led DELTA's development, said: "We've built modules that have never been used even in NATO."

In April 2025, reports emerged that a NATO member state had begun negotiations to adopt DELTA. If this materializes, it would be the first export of a Ukrainian combat system. Technology born on the battlefield may now be reverse-exported to the world's most powerful military alliance.

(2) The Secret Behind SAKER Drones and the 50% to 80% Improvement in Drone Strike Rate

By late 2023, Ukrainian drone pilots were in despair. Novice operators' hit rates had dropped to as low as 10%. Even veterans struggled to exceed 50%. The reason was electronic warfare. The Russian military had deployed jamming equipment in large quantities to disrupt drone signals. Drones dependent on GPS lost their bearings, and drones that lost their link to the operator crashed to the ground or flew to unintended locations.

That was when SAKER Scout appeared.

SAKER is a reconnaissance drone. It is equipped with Palantir's artificial intelligence software. SAKER identifies targets on its own , tanks, armored vehicles, troop concentrations. The AI determines what they are. It then transmits that information to the command post. The commander decides which weapon to use and when to strike.

SAKER's range is 10 kilometers. It uses inertial navigation instead of GPS. It is resistant to jamming. But the real secret lies not in the hardware but in the software.

Palantir's AI learns. Ukrainian operators fed it Russian military footage , thousands of hours of video. Russian soldiers' uniforms, equipment, weapons. The way they walk, their movement patterns. The AI memorized all of it. Now SAKER can find camouflaged equipment. If a person not wearing a uniform moves like a Russian soldier, the AI recognizes it.

In 2024, according to reports by Time magazine and Britain's Forces News, AI drone hit rates climbed to 80%. From below 50% to 80%. This means nearly twice as many enemy targets can be struck with the same number of drones.

David Kirichenko, writing for the Henry Jackson Society, analyzed: "AI-equipped drones are using technology to counter Russia's numerical advantage in manpower and weapon systems. Ukraine has squarely overturned the conventional wisdom that a much larger great power would overwhelm a smaller country."

But he also expressed concern. If AI can identify Russian soldiers not wearing uniforms based on movement patterns alone, how would it distinguish Russian civilians? Kirichenko said that Ukrainian and Western militaries would halt fire in that situation, but he could not be sure what would happen if the Russian military acquired the same technology.

Forces News concluded: "The wars of the future will be as much about algorithms as armor."

In September 2024, Ukraine's Ministry of Defense officially approved SAKER Scout for military supply. This was thanks to the Ministry of Defense accelerator program established in June. The program ensures that proven innovative technologies can be rapidly deployed to the field without bureaucratic delays. A process that would have taken two years in the past was reduced to a month and a half.

SAKER also integrates with the DELTA system. When a reconnaissance drone discovers a target, that information is immediately displayed on DELTA's digital map. The commander can see available fire assets on the same screen , which battery is within range, which armed drone has completed pre-flight preparation. Once a decision is made, FPV kamikaze drones included in the SAKER system fly toward the target.

Humans make the final decision. For now. But AI's role is growing. Finding, classifying, and recommending targets is now the machine's job. What remains for the human is pressing the approval button.

A Ukrainian Ministry of Defense official said, on condition of anonymity: "First-person-view drones with AI could achieve an 80% hit rate."

The future of warfare is flying over the trenches of Donetsk.

B. Command and Control by Algorithm

(1) Shortening the Time from Target Identification to Strike Approval

In military terminology, there is a concept called the "Kill Chain." It refers to the series of steps: finding a target, tracking it, identifying it, assigning a weapon, attacking, and assessing the result. Each link in this chain consumes time. In traditional methods, this process could take 72 hours. Assuming the target would still be at the location where it was spotted was closer to wishful thinking.

Ukraine has dramatically shortened this time. In some cases, to just two minutes.

The secret was not replacing human judgment but eliminating the time spent waiting for human judgment. In the past, when a reconnaissance unit collected intelligence, it went up to higher command, was reviewed by analysts, decided upon by a commander, and transmitted back down. Reports were written at each step. Meetings were convened. Approval stamps were applied.

DELTA and Palantir systems compress all of this into a single screen.

According to a report published in June 2025 by Ukrainian security expert Kateryna Aniskina, DELTA works as follows: ISTAR (Intelligence, Surveillance, Target Acquisition, Reconnaissance) units collect raw data from all available sources , informant reports, automated collection through chatbots, satellite images, aerial reconnaissance data, radar, cyber data, and more. Then various actors in the security and defense sectors input this information into DELTA. It layers onto the digital map.

What matters here is standardization. The system's tools convert data into a machine-readable standard format. Room for interpretation is removed. An object on the map , say, a specific weapon or piece of military equipment , is exactly that object as entered into the system. It cannot be interpreted differently.

The information processed this way is objectively displayed on the digital map in near real time. And it immediately becomes accessible to all other actors who need battlefield situational awareness.

DELTA is composed of several modules. Monitor is the situational awareness map. Target Hub manages fire missions. Vezha provides real-time video streams. Mission Control handles unmanned vehicle coordination. Element is secure chat. According to Defense Minister Umerov, the value of enemy equipment destroyed through DELTA exceeds $15 billion.

A 2025 Carnegie Endowment for International Peace report noted: "Despite advances in reconnaissance, digitization, AI, and predictive analytics, war remains kinetic, chaotic, and destructive. Technology has shaped decisions but has not eliminated the fog and friction of war."

But the fog has clearly thinned.

There is an AI system called Griselda. This system processes collected data in an average of 28 seconds. The processed information is transmitted to the DELTA battle management system and the GIS Arta fire support system. Twenty-eight seconds. In the time a human analyst reads the first paragraph of a report, the AI has already reached its conclusion.

The true meaning of algorithmic command and control lies here. It is not about improving the quality of decisions. It is about reducing the time to reach a decision. In war, time is life. Literally.

(2) Data Visualization that Clears the Fog of War

Prussian military theorist Carl von Clausewitz coined the concept of the "fog of war." In warfare, commanders must make decisions amid incomplete information, contradictory reports, and unpredictable circumstances. Where the enemy is, how many there are, what they intend to do , almost nothing is known with certainty.

Ukraine is trying to clear this fog. Even if it cannot be completely eliminated, the goal is to make it as thin as possible.

DELTA's digital map is not an ordinary map. It is a living map. When a satellite transmits a new image, the map updates. When a drone sends video, it appears on the map. When a ground scout takes a photo with a smartphone and uploads it, that too appears on the map. Everything in real time.

Everything on one screen.

The Vezha module handles video streaming. A commander can simultaneously watch footage from dozens of drones. In September 2024, at NATO's REPMUS 24 exercise in Portugal, DELTA coordinated more than 50 drones , underwater drones, surface drones, ground robots, aerial drones, and even Rheinmetall's Robotics-L "robot dog." All under a single system.

The power of visualization lies in simplifying complexity. Raw data is meaningless , just a list of numbers and coordinates. But when it is displayed as icons on a map, color-coded, and changes over time are shown as animation, the human brain recognizes patterns. What is important, what is dangerous, what should be dealt with first.

Palantir's system works the same way. In the Washington Post article, the reporter saw a screen with the Bakhmut front's trenches drawn in fine detail. A click brings up thermal imagery. Another click shows a photo taken by a ground informant. All information is layered on the same map. A commander can feel as though viewing the battlefield from a god's-eye perspective.

But there are dangers too: the temptation to believe that what is on the screen is everything. What the sensors missed, what the data misinterpreted, what the AI misidentified , those do not appear on the screen. Clausewitz's fog has not completely disappeared. It has merely changed form.

A 2024 report from the German Marshall Fund warned: "The use of companies like Palantir and the DELTA system has the potential to affect democracy and human rights. The use of advanced data mining tools raises concerns about misuse and privacy violations. The fact that the company has classified contracts with governments underscores the need for ethical use and strict compliance with privacy laws."

Data visualization is a double-edged sword. It gives commanders unprecedented situational awareness. But at the same time, it can breed a new form of ignorance , overconfidence in technology. Believing only what the screen shows and thinking that what the screen does not show does not exist. That is the new fog.

C. GIS Arta and the Maven Smart System

(1) The "Uber-Style" Fire Support System

May 2022. The Siverskyi Donets River in the Luhansk region. Russian forces tried to cross the river. They set up pontoon bridges and attempted to move tanks and armored vehicles to the other side. Days later, satellite photos showed destroyed pontoon bridges and burning vehicles. In two days, Russian forces lost dozens of tanks and armored vehicles. Personnel losses were estimated between 1,000 and 1,500.

Behind this battle was a piece of software called GIS Arta.

GIS Arta , an abbreviation for Geographic Information System for Artillery. Ukrainian developers reportedly began building it around 2015. People call it "Uber for artillery."

Here is how it works. A scout on the front line spots the enemy position. The scout takes a photo with a smartphone and uploads it along with GPS coordinates. There is also footage from drones. There are also enemy artillery positions detected by counter-battery radar. All of this information flows into a central server.

The server's algorithm begins calculating. What type of target is it? Where is it located? Which battery is within range? Which ammunition is appropriate? If infantry is exposed, mortars are effective. If vehicles, howitzers are better. The algorithm finds the optimal combination.

Then it assigns the "job" , just as Uber assigns the nearest driver to a passenger.

Artillery units connect to the network and mark themselves as "available." It is the same as an Uber driver turning on the app and waiting for passengers. When a fire mission comes in, available batteries within range are highlighted. The commander touches the screen. From target acquisition to firing: 30 to 45 seconds.

Another strength of GIS Arta is dispersion. In traditional artillery operations, batteries were clustered in one location for concentrated fire. But concentration is vulnerable to counter-battery fire. If the enemy detects the firing position, return fire pours in.

GIS Arta is different. Multiple guns at different locations can simultaneously attack the same target. The system calculates when the shells fired from each gun will reach the target. It then coordinates firing times so that all shells land on the target at nearly the same moment. A choreography of shells arriving from multiple directions , making it difficult for Russian forces to respond.

Ukrainian batteries practice "shoot and scoot." They fire, then relocate within minutes. In the past, this took hours. It became possible because GIS Arta can assign new fire missions immediately even after repositioning.

On some fronts, Russian artillery outnumbers Ukraine's by a factor of 20. GIS Arta is a "force multiplier" that narrows this gap. Fewer guns strike more targets, more accurately, more quickly, with less risk.

The New America Foundation's analysis put it this way: "GIS Arta is essentially like any other app you can find on your phone or computer. Like Uber or Google Maps, it requires infrastructure and human capital to maintain."

The system requires internet connectivity. In February 2022, Russian forces cyberattacked Tooway, a satellite internet provider. Ukraine's combat systems were paralyzed. That was when Elon Musk's Starlink appeared. Thousands of terminals arrived in Ukraine. GIS Arta began working again.

The Battle of Siverskyi Donets was the result.

(2) The Reality of AI Tactics That Halted Russian Armored Offensives

On March 25, 2025, the NATO Communications and Information Agency (NCIA) signed a contract with Palantir. It was the adoption of Maven Smart System NATO (MSS NATO). From requirements definition to system acquisition in just six months , one of the fastest procurement processes in NATO's history.

Why the urgency?

NATO had been watching what was happening in Ukraine. Software was changing the nature of war. DELTA, GIS Arta, and Palantir's technology had enabled the numerically inferior Ukrainian military to stand and fight against Russian forces. NATO wanted to make these lessons their own.

The Maven Smart System was originally developed for the U.S. Department of Defense. In 2021, it was derived from Project Maven's object recognition AI. Initially, the purpose was to find terrorist targets in surveillance footage. But the system evolved. It now integrates classified and open-source materials into a single searchable database. Recently, generative AI and large language model (LLM) capabilities have been added.

According to the NATO press release, MSS NATO "provides the Alliance with a common data-driven warfighting capability. Through a range of AI applications from large language models to generative AI to machine learning, it improves intelligence fusion and targeting, battlefield awareness and planning, and decision acceleration."

In plain terms, it connects the different databases of different NATO member nations into one. U.S. intelligence, German intelligence, French intelligence all appear on the same screen. This was previously impossible. Each nation's system was incompatible. Formats differed. Security protocols differed.

General Markus Laubenthal, Chief of Staff at NATO's Supreme Headquarters Allied Powers Europe (SHAPE), said: "Innovation is key to our warfighting capability. Maven Smart System NATO enables the Alliance to harness complex data, accelerate decision-making, and thereby add genuine operational value."

Allied Command Operations (ACO) announced it would begin using the new system within 30 days of contract signing, with operations expected from May 2025. According to a SHAPE spokesperson, Maven was already "actively validated, contributing to live operations across nine NATO member states."

Contract sizes in the United States are also surging. In May 2024, the Pentagon signed a five-year contract with Palantir worth $480 million. In May 2025, the contract ceiling was increased by $795 million to nearly $1.3 billion. The reason cited was "increased demand from combatant commands."

According to Breaking Defense, MSS has already become "indispensable" to the Pentagon's Joint Staff and geographic combatant commands, including European Command. A 2024 CDAO (Chief Digital and Artificial Intelligence Office) contract expanded the user base from "hundreds" to "thousands."

According to Janes, the system is designed to link multiple command, control, and communications (C3) systems "to share and exploit data from legacy systems, third-party apps, and even bespoke solutions in ways that were previously impossible."

There are dissenting voices. An analysis by Defense Update notes: "The selection of Palantir, a U.S. technology company, as the central system for the entire Alliance highlights the ongoing tension between the need to use available cutting-edge technology (often originating from the U.S.) for immediate capability improvement and Europe's long-term strategic autonomy aspirations to reduce technological dependence on non-European suppliers."

A 2025 analysis from the Carnegie Endowment for International Peace raises a more fundamental question: "Commercial actors like Palantir are neither neutral suppliers nor patriotic agents. Their primary motivations are commercial, shaped by venture capital, shareholder expectations, and market valuations."

Analysts warn: if future commercial interests are judged to lie elsewhere, companies may shift their loyalties.

Ukraine has become the testing ground for AI warfare. The products of that experiment are now being exported to the world's most powerful military alliance. Algorithms verified on the battlefield are being transplanted into peacetime planning systems. Software is defining the future of war.

The question is who controls that software.

← Previous Chapter

See every chapter

Next Chapter →

Kim Kyung-jin

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com



#KimKyungjin #PALANTIR #Palantir #WarAI #SurveillanceAI #DefenseAI #DataPower #ArtificialIntelligence #AILibrary

# Chapter 6. Pentagon Reform and the Future of Defense

← Previous Chapter

See every chapter

Next Chapter →

Part 3: AI on the Battlefield, Software-Defined Warfare

A. The War with Legacy Prime Contractors

(1) Winning the Lawsuit Against the U.S. Army (2016) and Opening the Door to Commercial Software

In the spring of 2012, soldiers of the 82nd Airborne Division were dying on the outskirts of Kandahar, Afghanistan. In two weeks, six soldiers lost their lives to improvised explosive devices (IEDs). Intelligence analysts struggled to identify enemy patterns, but the Army-supplied software DCGS-A did not work. One analyst recalled that it took a day and a half to create a map of IED placement patterns. The unit commander requested different software from higher command , a product from a Silicon Valley company called Palantir. The request was denied. He asked again. Denied again. A third request met the same fate. The unit returned home with 200 casualties.

In Washington, no one wanted to hear this story.

The Army had already poured 15 years and more than $3 billion into DCGS-A. If there were problems, they could be fixed. It could not be handed over to a startup from Silicon Valley. There was an ecosystem built by giant defense contractors like Lockheed Martin, Raytheon, and Northrop Grumman. They had worked with the Pentagon for decades and employed thousands of lobbyists and retired generals. There was no room for a company like Palantir.

At a 2013 congressional hearing, Republican Representative Duncan Hunter of California clashed with Army Chief of Staff General Ray Odierno. Hunter presented survey results showing 96% of field soldiers preferred Palantir. Odierno grew angry. The scene, beginning around the two-hour mark on YouTube, can still be viewed today. It is a rare spectacle of a general shouting at a congressman.

Palantir waited. Then on June 30, 2016, it filed an 81-page complaint in the U.S. Court of Federal Claims.

The core of the lawsuit was the Federal Acquisition Streamlining Act (FASA), enacted in 1994. This law mandated that government agencies first investigate whether commercial products exist on the market before purchasing. The intent was to avoid wasting taxpayer money redeveloping something that already exists. Palantir's lawyers argued that the Army had violated this law. The Army had not conducted proper market research and had not even reviewed whether Palantir's Gotham platform could meet DCGS-A's requirements.

The complaint contained one sentence: "Palantir has developed technology that already solves the problems DCGS seeks to address. This technology has been successfully used by Army units and numerous military and intelligence agencies. Field commanders have repeatedly requested this product. Yet the Army issued a solicitation that made it impossible for Palantir to compete. This is unreasonable. And it is a violation of law."

Four months later, on October 31, 2016, Judge Marian Blank Horn of the Court of Federal Claims ruled in Palantir's favor.

The ruling was clear. The Army had violated FASA. It had not adequately investigated whether commercial products existed on the market and had failed to reasonably explain why commercial products could not meet its requirements. The court issued a permanent injunction prohibiting the Army from awarding DCGS-A contracts until it complied with the law. In 2018, the Federal Circuit Court of Appeals upheld the ruling. The appeals court noted that the Army's market research was "a conclusions-only assessment" and that it had not "reasonably used" its research findings.

This ruling became a turning point in the history of U.S. defense procurement.

The Army had to change its approach. In March 2018, it selected two companies , Raytheon and Palantir , and began a competitive evaluation. Soldiers personally used both systems and evaluated them. The Army wanted a system that could run on commercial laptops, be learned in eight hours, and operate in low-bandwidth or disconnected environments. Raytheon submitted a product called FoXTEN. Palantir submitted its Gotham platform.

In March 2019, Palantir was selected as the final winner. The contract was worth more than $800 million.

Doug Philippone, Palantir's head of defense business, assessed that this lawsuit brought transformative change to defense procurement methods. He called it a "try, buy, decide approach" , first test commercial products, then buy them if they work. It was different from the old method of spending more than a decade writing requirements, only to find that requirements had changed while the technology was being developed, yet being unable to stop.

Joe Kasper, chief of staff for Representative Duncan Hunter, said: "It was more than a simple contract dispute. Soldiers in combat were repeatedly requesting a commercial alternative that they claimed saved lives. One division commander said it was a matter of lives and limbs. The fact that this technology was repeatedly denied and actively disparaged is the worst kind of bureaucracy."

Palantir's victory opened a path for Silicon Valley companies. It showed that FASA was a legal weapon. It proved that if commercial products exist and the government ignores them to develop from scratch, one can go to court. This was not just one company's victory. It was the moment the first crack appeared in the fortress walls of the legacy defense cartel.

(2) Breaking the Cost-Plus Contract Practice and Fixed-Price Innovation

In 1917, during World War I, the U.S. government faced a problem.

War materiel was urgently needed, but nobody knew how much it would cost. Tanks, airplanes, warships , these were things being built for the first time. The government could not tell companies "build it at this price." So the government offered a compromise: "Just build it. We'll pay whatever it costs. And we'll add a profit margin on top." This was the beginning of the cost-plus contract.

A century has passed. That stopgap measure became the standard of the American defense industry.

The logic of cost-plus contracts is simple. When developing a new weapons system, uncertainty is large. If companies are forced to accept a fixed price, they either cannot bear the risk and nobody bids, or they set the price absurdly high. So it benefits everyone for the government to reimburse costs and guarantee a reasonable profit.

The problem was the incentive structure this system created.

The more it costs, the bigger the profit. Finishing quickly is a loss. Reducing costs is a loss. Dragging out the project and inflating costs becomes the rational choice. Defense contractors became masters of this game. Projects running over schedule and over budget became not the exception but the rule. Boeing's KC-46 aerial refueling tanker generated more than $7 billion in cost overruns on a $4.6 billion contract. This is not an extreme case. This is the standard.

In 2009, President Barack Obama declared that "the days of giving defense contractors a blank check are over." He issued a policy memo favoring Firm Fixed-Price contracts. Ashton Carter, then the Pentagon's top weapons buyer, also supported expanding fixed-price contracts. Congress passed the Weapon Systems Acquisition Reform Act (WSARA) to promote competition and reduce cost-plus abuses.

But resistance from existing defense contractors was fierce.

In January 2024, Lockheed Martin CEO Jim Taiclet declared at an earnings call: "There are no more must-win programs for us." He meant that the risk of fixed-price development contracts was too great. Northrop Grumman CEO Kathy Warden said something similar. The company had pulled back from fixed-price development contracts after the 2015 B-21 bomber contract. The L3Harris CEO was more direct: "If it's fixed-price development, we won't bid. We don't want to take losses."

This is the worldview of traditional defense contractors. For them, cost-plus is an old friend guaranteeing safe returns. Margins are low but so are risks. Fixed-price contracts are a gamble. If costs exceed projections, the loss falls entirely on the company.

But emerging defense tech companies like Palantir and Anduril operate on a different logic.

In September 2024, the Army released new software procurement guidelines. These included language stating "use cost-plus contracts to the maximum extent." Non-traditional defense companies pushed back. One defense tech executive, speaking anonymously, said: "Cost-plus

favors large traditional defense contractors. They have government-approved cost accounting systems. For smaller companies like ours, building that infrastructure takes enormous time and money."

The tide is turning. Space Force acquisition chief Frank Calvelli has been pushing to expand fixed-price contracts since taking office in 2022. His vision is to purchase small systems using existing technology at fixed prices and deploy them within three years of contract award. He acknowledges exceptions: "I never said to build the next-generation Battlestar Galactica at a fixed price."

The key is balance. For genuine innovation , R&D pioneering unknown territory , cost-plus may be appropriate. But insisting on cost-plus for applying existing commercial technology to the military is a waste of taxpayer money.

Palantir's 2016 lawsuit victory left a decisive precedent in this debate. If a commercial product exists, the government must consider it. Bureaucratic inertia toward developing from scratch can be legally challenged. This was a ruling that shook the foundations of the cost-plus regime. It cracked the "development-first" privilege that traditional defense contractors had enjoyed.

The Trump administration in 2025 is accelerating this trend. Defense Secretary Pete Hegseth released a memo demanding reform of software procurement practices shortly after taking office. The Department of Government Efficiency (DOGE) is targeting wasteful contracting practices. Army Chief Information Officer Leo Garciga said the goal is "to maximize buying power by using enterprise-level discounts."

A wartime stopgap from a century ago became a permanent institution, and that institution became a protective shield for interest groups. Now a new generation of companies is trying to tear down those walls. The fight is not yet over. But Palantir won the first battle.

B. The Shift of Power from Hardware to Software

(1) The Rise of Defense Tech Companies: Anduril, SpaceX, and Others

In June 2014, a young man attended a party on Sonora Island in British Columbia, Canada. It was twenty-one-year-old Palmer Luckey. Two years earlier, he had become a billionaire by selling Oculus, a virtual reality headset company he built in his parents' garage, to Facebook for $2.7 billion. The party was hosted by Peter Thiel's Founders Fund. There, Luckey met thirty-year-old Trae Stephens. Stephens was an investor who had just left Palantir and joined Founders Fund.

The two discovered an unusual commonality. They both thought it was absurd that Silicon Valley venture companies were not working with the Department of Defense. For Stephens, it was incomprehensible that nobody was interested when there was a customer spending tens of billions of dollars. At the time, the only companies receiving venture investment and working closely with the government were Palantir and SpaceX. Stephens was looking for a third company to join that duo.

Three years later, in 2017, Luckey was fired from Facebook. A political donation controversy was the cause. He could have been disheartened, but instead he co-founded a new company with Stephens: Anduril Industries. Named after Aragorn's sword in The Lord of the Rings, it means "Flame of the West" in Quenya.

Anduril's strategy was the exact opposite of traditional defense contractors.

Traditional defense contractors wait for the Pentagon to publish requirements. When requirements come out, they write proposals. If they win the contract, they begin development. If costs overrun, the government bears them. Completion can take 10 or 15 years. Anduril was different. It built products first, using its own funding for R&D. Then it went to the Pentagon and said: "It already exists. It works. We'll give it to you at a fixed price."

This was applying the Silicon Valley startup playbook to the defense sector.

In 2024, Anduril's revenue roughly doubled year-over-year to approximately $1 billion. New contract value exceeded $1.5 billion. In June 2025, the company raised $2.5 billion in a Series G round led by Peter Thiel's Founders Fund. Founders Fund invested $1 billion in this round , the largest check in the fund's history. The company was valued at $30.5 billion, more than doubling from $14 billion eight months prior.

These numbers are significant. Anduril has now entered the territory of traditional defense giants.

The contract list reveals the scale. In February 2025, Anduril acquired the Army's Integrated Visual Augmentation System (IVAS) program from Microsoft , a contract worth up to $22 billion to supply 120,000 augmented reality headsets to the military. It is participating in an autonomous unmanned aircraft program worth up to $9 billion in partnership with General Atomics. It signed a 10-year, $642.2 million anti-drone system contract with the Marine Corps.

In Ohio, a factory called Arsenal-1 is under construction. About $1 billion is being invested in this approximately 5-million-square-foot facility, which is scheduled to begin production in July 2026. More than 4,000 jobs will be created. Luckey calls this factory a "mega-scale autonomous weapons manufacturing facility." Plans call for producing tens of thousands of drones and autonomous systems annually.

Anduril's secret weapon is a software called Lattice. It is a command and control system that collects and analyzes data from various hardware platforms to provide real-time battlefield situational awareness. Drones, sentry towers, and sensor networks are connected through Lattice. When a sentry tower detects a suspicious vehicle, it automatically dispatches a Ghost drone to investigate while simultaneously sending an alert to ground forces.

SpaceX traced a similar trajectory.

Founded by Elon Musk in 2002, the company was initially the subject of ridicule. Rocket manufacturing was Boeing and Lockheed Martin's domain. But SpaceX dramatically lowered launch

costs through the innovation of reusable rockets. It is now a major launch contractor for the U.S. Air Force and Space Force. The Starlink satellite network became critical infrastructure for military communications during the Ukraine war. In 2024, the Pentagon signed a contract with SpaceX to develop a military version called Starshield.

These companies share a common thread: all place software at their core.

Anduril's drones are hardware, but their value lies in the Lattice software. SpaceX's rockets are machines, but what enables reuse is precision control software. Palantir was a pure software company from the start. Together, they prove that power in the defense industry is shifting from hardware to software.

Traditional defense contractors feel this shift too. Northrop Grumman, L3Harris, and Raytheon participate as partners in Palantir's TITAN project. Palantir holds the software leadership. They are in the role of supplying hardware components. A role reversal unimaginable ten years ago.

Luckey said in a CNBC interview: "The idea that we should have the best military in the world is a truly bipartisan one." He stated that Anduril "will certainly become a public company." To win trillion-dollar defense contracts, it needs to be a public entity.

The relationship between Silicon Valley and the Pentagon was once frosty. In 2018, Google employees launched a mass signature campaign protesting participation in Project Maven, and Google abandoned the contract. But now things are different. In December 2024, Anduril announced a partnership with OpenAI , a collaboration to integrate AI into systems that track and neutralize enemy drones. In 2025, it announced joint development of augmented reality headsets with Meta. It was a reconciliation between Luckey and Meta.

The power shift is underway. Anduril, SpaceX, Palantir. What this triad proves is clear: in 21st-century warfare, the winner is not the side with the biggest tank. It is the side with the best software.

(2) Palantir's Role in Building JADC2 (Joint All-Domain Command and Control)

One day in May 2024, a contract was signed in a conference room at the Pentagon.

$480 million. Five years. Maven Smart System. These three terms were the core of the contract. Palantir was tasked with spreading the Department of Defense's AI backbone system to combatant commands worldwide. Central Command, European Command, Indo-Pacific Command, Northern Command/NORAD, Transportation Command. Five combatant commands were the first targets.

Maven's history goes back to 2017.

At the time, the Pentagon faced a problem. Drones captured thousands of hours of footage daily, but there were not enough personnel to analyze it. Human analysts were buried under an endless pile of video. The solution was AI. Project Maven aimed to develop a system that automatically identifies and tracks objects in drone footage using AI and machine learning.

Google participated as the initial contractor but withdrew due to employee backlash. Palantir filled that void. In 2023, the Maven program was transferred to the National Geospatial-Intelligence Agency (NGA) and elevated to an official program of record. Palantir's Maven Smart System became the program's core software platform.

JADC2 , Joint All-Domain Command and Control. What does this acronym mean?

Traditionally, each military branch operated its own command and control system. The Army had its own, the Navy had its own, the Air Force had its own. Intelligence was separate; communications were separate. In joint operations, this became a problem. Even if an Air Force reconnaissance aircraft spotted an enemy tank, it took time for that information to reach Army artillery. Even if a Navy ship detected an enemy missile launch, it was not immediately relayed to Air Force interceptors.

JADC2 is an attempt to break down these barriers. Army, Navy, Air Force, Space Force, Cyber Command , connecting sensors, weapons, and commanders across all domains into a single network. A system where anyone can obtain needed information in real time and respond as quickly as possible with the most suitable weapon.

Palantir's role is to build the software layer of this vision.

Doug Philippone, Palantir's head of defense business, explained: "Our platform is like an operating system for JADC2." More than 2,000 engineers deploy hundreds of updates weekly. The Maven Smart System integrates data from satellites, drones, ground sensors, and human intelligence sources. AI analyzes patterns and identifies threats. Commanders view complex data on visualized screens.

In May 2025, the Department of Defense increased the contract ceiling by $795 million to nearly $1.3 billion. The reason: "increased demand." Vice Admiral Carl Thomas, Vice Chief of Naval Operations, said at the Navy Summit in August 2025: "The number of people using Maven as a data fabric and Foundry for operations centers is growing. Now everyone is using Maven."

TITAN is another key element of JADC2 implementation.

In March 2024, Palantir was awarded a $178.4 million TITAN contract by the Army. TITAN stands for Tactical Intelligence Targeting Access Node. It is a mobile ground station that receives data from space and high-altitude sensors and provides targeting information to ground forces. AI analyzes vast amounts of data to assist human commanders in decision-making.

TITAN comes in two versions. The advanced version is mounted on a large truck and connects directly to space sensors. The basic version is installed on a Joint Light Tactical Vehicle (JLTV) and is more mobile. In March 2025, Palantir delivered the first two TITAN prototypes on schedule. In April 2025, an Army report assessed TITAN as one of the top-performing programs, alongside the Next Generation Squad Weapon and Future Long-Range Assault Aircraft.

In December 2025, another groundbreaking contract was announced: ShipOS.

Secretary of the Navy John Phelan and Palantir CEO Alex Karp made a joint announcement in Washington. This $448 million contract applies AI to shipbuilding optimization. The goal is to surge nuclear submarine production. Palantir's predictive maintenance concepts are applied to the manufacturing floor to reduce costs and increase speed. The contract includes a "shared savings" mechanism tied to performance metrics , meaning if cost savings are realized, both parties share them.

The Pentagon is planning to establish a new program office to unify JADC2 efforts, fusing multiple Department of Defense elements that emerged since the late 2010s: the Air Force's Advanced Battle Management System (ABMS), the Army's Project Convergence, and the Navy's Project Overmatch. The aim is to bind all of these together to realize true joint all-domain command and control.

Palantir is at the center of all these puzzle pieces.

What they are building is not mere software. It is the nervous system of future warfare. The goal is to dramatically shorten the time from sensor to shooter, from detection to strike. The possibility was already proven in the Ukraine war. Now it is being scaled to the entire U.S. military.

Akash Jain, Palantir's President of USG, said when announcing the Maven contract expansion: "The rapid fielding and adoption of key technologies, including the agile, effective, and responsible deployment of AI/ML, is essential for the United States to maintain its competitive advantage."

At the center of that competitive advantage stands Palantir.

C. The Significance of the U.S. Army's $10 Billion Contract in 2025

(1) Consolidation of 75 Individual Software Contracts

On July 31, 2025, the U.S. Army released a brief press statement.

The title read: "Enterprise Agreement Awarded to improve Military Readiness and Operational Efficiency." But the numbers within were not simple. $10 billion. 10 years. Consolidation of 75 contracts. That day, Palantir officially entered the uppermost tier of the American defense industry.

What makes this contract special is not just the dollar amount. The form is new.

Called an Enterprise Agreement (EA), this contract is a type of blanket purchase agreement. It is a system where the Army can buy Palantir products and services as needed, in the quantities needed. No individual bids, individual contracts, or individual approval processes are required. It is similar to an enterprise license agreement that a large corporation signs with a software company.

Why was such a contract needed?

The Army had no fewer than 75 individual contracts with Palantir. Fifteen were contracts where Palantir was the prime contractor, and 60 were contracts where Palantir participated as a subcontractor. The problem was clear: the same software from the same company was being

purchased at different prices, under different terms, through different contracts. Intermediary fees were being added. Administrative costs were wasted. Nobody could see the full picture.

Daniel Moyer, Executive Director of the Army Contracting Command-Aberdeen Proving Ground, explained: "If you have more than one contract with the same vendor, you need to check whether you bought the same thing twice, in a different way or at a different price. This effort is to make sure we're getting the best discount."

This is common sense, but in the Pentagon, it is revolutionary.

Traditionally, defense procurement was decentralized. Each department, each program, each command purchased on its own. There was no coordination. Buying the same product ten times without negotiating volume discounts. Suppliers profited from this inefficiency.

Leo Garciga, Army Chief Information Officer, explained the significance of the enterprise agreement: "This enterprise agreement is a pivotal step in strengthening the Army's modernization capabilities while being fiscally responsible. By streamlining procurement processes and using enterprise-level discounts, we are not only enhancing operational effectiveness but also maximizing our purchasing power."

Examining the contract structure makes the change even clearer.

Most existing contracts were "capability package" based. If the Army said "we want this function," the vendor proposed a complete solution, and the Army purchased the entire package , including features it didn't need. The new enterprise agreement is different. In Moyer's words, it is "an a la carte menu, not an all-you-can-eat buffet." Prices are separated by product and service. The Army selects and purchases only what it needs.

The $10 billion figure also needs clarification.

This is not a commitment by the Army to pay Palantir $10 billion. It is the maximum ceiling that can be purchased through this contract over 10 years. Actual spending could be significantly less. There is a minimum purchase obligation, but it has not been disclosed. The Army can stop purchasing through this contract at any time if a better alternative appears.

Moyer also emphasized competition: "Competition will continue for future programs. If Palantir bids on a program or weapon system and is selected, they would of course use this agreement to get volume discounts and economies of scale. But if Palantir loses a bid, that business goes to another vendor."

This contract is the beginning of a larger trend.

Garciga revealed that the Army is negotiating 10 to 15 additional enterprise agreements. "We're not just focusing on the Palantir contract. This is our big initiative. Essentially, we're looking across the Army at how many times we've purchased the same commercial software and whether it makes sense to consolidate them for maximum discounts."

The implication is clear. Palantir has become the precedent. Other software companies can also sign similar contracts , Microsoft, AWS, Oracle. Anyone willing to offer volume discounts using the Army's scale can sit at the negotiating table.

This aligns with the Trump administration's policy direction. Defense Secretary Pete Hegseth released a memo demanding software procurement practice reform shortly after taking office. The Department of Government Efficiency (DOGE) is targeting wasteful contracting practices. President Trump himself mentioned at a July AI summit that "we buy a lot from Palantir."

Boeing, Northrop Grumman, Lockheed Martin , traditional defense giants doing business with the Pentagon through enterprise agreements is nothing new. But a software company like Palantir joining this club is different. It is a signal that the status of software in the defense industry has fundamentally changed.

The Register commented: "There are no formal criteria for top-tier membership of the U.S. military-industrial complex, but a $10 billion deal that consolidates dozens of contracts into one blanket purchase agreement certainly shows Palantir has earned its place."

From 75 fragments to one massive block. That is what happened on July 31, 2025.

(2) The Trump Administration and the AI-Centric Defense Strategy

On January 20, 2025, Donald Trump began his second term.

His administration declared AI a national priority. The executive orders signed immediately after inauguration had AI development acceleration and deregulation at their core. The defense sector was no exception. Defense Secretary Pete Hegseth released a software procurement reform memo in his first month. The message was clear: faster, more efficiently, with AI at the center.

Palantir is one of the biggest beneficiaries of this policy direction.

CEO Alex Karp's relationship with the Trump administration is close. Karp did not publicly endorse Trump during the 2024 campaign, but his worldview resonates with the administration's "peace through strength" doctrine. Karp has been preaching for years the importance of technological superiority to defend Western democracy. The Trump administration is funding that vision.

The numbers speak.

Just listing the major defense contracts Palantir received in 2025 alone makes a long list. March: first TITAN delivery. May: Maven Smart System contract increased to $1.3 billion. July: Army $10 billion enterprise agreement. December: Navy ShipOS $448 million contract. This is not coincidence. It is the result of policy will.

The Department of Government Efficiency (DOGE) has a role too.

Led by Elon Musk and Vivek Ramaswamy, this organization aims to eliminate federal government inefficiency. Their gaze fell on Department of Defense procurement , duplicate contracts, intermediary fees, administrative waste. The enterprise agreement with Palantir was presented as a model for solving these problems. Consolidating 75 contracts into one, securing volume discounts, eliminating middlemen margins. A story DOGE would like.

But this is not just a cost-cutting issue. It is a matter of geopolitical competition.

China is rapidly expanding its AI capabilities under its Civil-Military Fusion strategy. U.S. intelligence agencies see 2027 or 2028 as a potential crisis point for the Taiwan Strait. Before then, the U.S. military must secure overwhelming superiority in AI capability. JADC2 implementation, kill chain acceleration, autonomous weapons system deployment , everything is a race against time.

Palantir's role is central to this competition.

Axios analyzed the $10 billion contract and wrote: "The deal illustrates three things: Palantir's rise in Washington and especially at the Pentagon. Changes in the way the military tests and buys products, especially software. And a maturing Washington-Silicon Valley relationship."

Wendy Anderson, former Palantir executive and chief of staff to the late Defense Secretary Ashton Carter, commented: "This agreement reflects a broader shift in recognizing software not as a support function but as core to operational readiness. It sets a precedent for how government and industry can partner to deliver real capabilities at speed and scale."

Palantir's stock rose more than 110% in 2025. Market capitalization at one point exceeded $300 billion, surpassing traditional defense giants. Of course, there are bubble concerns. The price-to-earnings ratio runs in the hundreds. But investors are looking at the future , what it means to become the data backbone of the U.S. military.

There is criticism too.

Some analysts worry about excessive dependence on a single company. If Palantir controls the U.S. military's core systems, switching to another vendor later becomes difficult , the so-called "lock-in" effect. The military emphasizes that competition will continue, that if Palantir loses a bid, another vendor gets the business. But realistically, in an enterprise agreement system that consolidated 75 contracts, it will not be easy for competitors to break in.

The Trump administration's AI-centric defense strategy is a bet. A bet that software superiority can overcome hardware disadvantages. Even if China produces more ships and missiles, the logic goes, America can win if it decides and strikes faster with better AI.

Palantir is the key chip in that bet. The $10 billion contract shows the size of that chip.

In 2016, Palantir had to go to court just to earn the right to bid. In 2019, it won an $800 million contract. In 2025, it signed a $10 billion enterprise agreement. From outsider to the center of the military-industrial complex in nine years. That is the speed of the defense industry in the software

age.

← Previous Chapter

See every chapter

Next Chapter →

Kim Kyung-jin

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com



#KimKyungjin #PALANTIR #Palantir #WarAI #SurveillanceAI #DefenseAI #DataPower #ArtificialIntelligence #AILibrary

# Chapter 7. Opportunity in Crisis: Monetizing Chaos

← Previous Chapter

See every chapter

Next Chapter →

Part 4: Digital Transformation of Enterprise and the Public Sector

A. Savior of the Pandemic

(1) Building the NHS Vaccine Distribution and Bed Management System in the UK

In March 2020, during the same week the Nightingale Hospital was being hastily assembled inside London's ExCeL Centre, data managers at NHS England were facing a nightmarish reality. Patient data pouring in from 215 hospital trusts across the country arrived in different formats, from different systems, at different speeds. Which hospital had ventilators to spare, which intensive care unit had an empty bed , nobody could see the full picture.

That's when a phone call came in. On the other end was Palantir's sales team, freshly arrived from the United States. They offered their software for just one pound. There was one condition: give them access to NHS data. Palantir's Foundry platform soon became the nervous system of the NHS vaccine rollout. When the Pfizer vaccine required storage at minus 70 degrees Celsius, the platform connected cold storage locations across the country, transport truck routes, and appointment schedules at each vaccination centre in real time. At the North Tees and Hartlepool Trust, long-stay patients hospitalized for 21 days or more dropped by 36%. In Dorset, 2,500 additional surgeries became possible per year. In North Cumbria, surgical throughput improved by 10%.

The numbers were impressive. But behind them lurked an uncomfortable question: the medical records of 65 million Britons were flowing through an American company's servers.

(2) Establishing Itself as Public Health Infrastructure and the 330-Million-Pound Contract

On November 21, 2023, NHS England signed a seven-year contract with Palantir worth 330 million pounds (approximately 550 billion Korean won). Named the 'Federated Data Platform' (FDP), it was the largest healthcare data contract in British history.

The contract ran to 586 pages. Then something strange happened. When civil society activists opened the contract published on the government register that December, they couldn't hide their bewilderment. 417 pages had been completely blacked out. Seventy-one percent of the entire document was redacted. The 'data protection' sections were almost entirely erased. Financial reporting content had vanished. The roles and responsibilities of key personnel were unreadable.

Ian Brown, a lawyer at the Good Law Project, was furious. "Redacting 70% of the contract is unlawful in itself. It directly violates government transparency guidelines." He immediately launched legal

action, sending a pre-action protocol letter to NHS in February 2024 alongside openDemocracy and Foxglove.

An even more shocking revelation was waiting. Shortly after the contract was signed, emails leaked showing that Palantir had hired PR agency Topham Guerin to pay social media influencers to attack the Good Law Project. The influencers were instructed not to mention Palantir's name.

NHS England launched an investigation. The contract explicitly stated that Palantir needed prior written consent before using the NHS name or brand in marketing or promoting the contract. Barely weeks after signing, allegations of contract breach had already surfaced.

In the spring of 2024, FDP entered full operation after a six-month transition period from the existing national data platform. The goal was to connect 240 healthcare organizations by 2027. In a business case analysis published in October 2025, NHS projected a fivefold return on investment. But reality proved difficult. By the end of 2024, fewer than a quarter of 215 hospital trusts were actually using FDP. Even by May 2025, only 72 trusts , a third of the total , had adopted it. Greater Manchester health authorities wrote in an internal document: "None of the products designed or produced by Palantir as part of the FDP programme currently exceeds NHS Greater Manchester's own capabilities."

Leeds Teaching Hospitals NHS Trust was blunter. In a letter released through a freedom of information request, they wrote to NHS England: "Adopting some of Palantir's platform tools would mean losing functionality, not gaining it."

Eventually, the Department of Health paid KPMG 8 million pounds to 'accelerate' adoption of Palantir's software. A consulting firm having to sell another company's software , that was the state of British public health infrastructure in 2025.

The controversy was amplified by Peter Thiel's remarks. Palantir's co-founder and board chair gave a speech at the Oxford Union in 2024 where he said: "The NHS should be privatized. The British affection for the NHS is a form of Stockholm syndrome." The company distanced itself, calling it a statement made in a 'purely personal capacity.' But the fact that this company's founder was calling for the dismantling of the very public healthcare system whose data he managed was not easily forgotten.

A senior official at Amnesty International UK commented: "Palantir is a deeply concerning choice as an NHS service provider, given the company's reputation surrounding human rights controversies." He was referring to its work with U.S. Immigration and Customs Enforcement (ICE). A 2020 Amnesty International report had warned that Palantir posed a "high risk of contributing to serious human rights violations."

In January 2025, Palantir announced a partnership with Multiverse. The plan was to train NHS staff on using the data platform through an FDP apprenticeship programme. This was no ordinary training initiative. Once staff become accustomed to Palantir's system, the cost of switching to anything else grows exponentially. That is the lock-in effect.

The pandemic crisis opened the door to British public health for Palantir. It took three years for a one-pound contract to grow into a 550-billion-won deal. The alchemy of turning chaos into profit , Palantir had perfected it.

B. Supply Chain Collapse and Recovery

(1) How Global Companies Like Airbus (Skywise) and BP Used Foundry

Paris Air Show, 2017. Airbus's head of digital transformation took the stage to unveil an ambitious vision: connecting all data across the aviation industry into one system. Standing beside him was a Palantir engineer. It was the birth of a platform called 'Skywise.'

Skywise had a distinctive architecture. Palantir's Foundry provided the backend infrastructure, while Airbus layered aviation domain expertise on top.

A single aircraft generates roughly 500 gigabytes of data per day for an A330neo, and up to one terabyte for the latest A350. Engine sensors, flight records, maintenance histories, parts inventories, flight schedules , all of it flowed into a single platform.

Airbus opened this platform to airlines. For free. Skywise Core was included as a standard option with aircraft purchases. Buy a 100-million-dollar airplane, and it comes bundled at no extra charge. Airlines could compare their fleet's operational data against industry averages. But the real revenue came from premium services: Predictive Maintenance, Health Monitoring, Fleet Reliability analysis. Those required a subscription. Once airlines grew comfortable with the free version, they naturally graduated to paid tiers.

By 2023, more than 50% of the global Airbus fleet , over 11,900 aircraft , was connected to Skywise. More than 140 airlines had joined. EasyJet, AirAsia, Cathay Pacific, Philippine Airlines, Malaysia Airlines, and Delta Air Lines were all inside this ecosystem.

Behind the numbers were real stories.

Jetstar, the Australian low-cost carrier, used Skywise to detect system performance degradation in its A320 fleet before problems materialized. The results were striking: engineering-related cancellations dropped by 93%. Predictive maintenance prevented dozens of flight cancellations.

A dramatic case occurred in August 2022. A330neo operators reported high-pressure valve leaks in the engine bleed system. Airbus engineers analyzed operational and sensor data accumulated in Skywise. They discovered that under specific takeoff settings, the bleed monitoring computer's software failed to properly control the high-pressure valve. Excessive stress caused the valve clamping pins to fracture, leading to seal damage and leaks. In a worst-case scenario, high-temperature, high-pressure bleed air could seep into the wing and cause catastrophic structural damage.

Airbus immediately reported to the European Union Aviation Safety Agency (EASA) and recommended issuing an emergency airworthiness directive. They chose to accept the cost of

grounding aircraft rather than risk a catastrophic accident. It was a case of big data preventing disaster.

Airbus's goal was clear: achieve 10 billion dollars in services revenue over the next decade. Skywise was the core engine driving that target. Independent analysts estimated the platform was generating over 850 million dollars in annual revenue opportunities.

The relationship between Palantir and Airbus resembled that between Veeva and Salesforce. Just as Veeva built life-sciences solutions on top of the Salesforce platform and paid 40 to 50 million dollars annually in royalties, Airbus built aviation-specific solutions on Foundry. Of the 26,900 Skywise users, the majority were Airbus employees. Fifteen percent of its 130,000-person workforce used the platform.

Meanwhile, in the middle of the cold North Sea, another digital twin was at work , on BP's offshore oil platforms.

In 2014, BP signed its first contract with Palantir. The wounds from the 2010 Deepwater Horizon explosion had not yet healed. Crude oil gushed for 87 days in the Gulf of Mexico , 7.8 million barrels. It was the worst marine oil spill in history. BP realized that digital transformation was no longer optional but a matter of survival.

Palantir built digital twins across BP's oil and gas production operations. North Sea offshore platforms, Gulf of Mexico deepwater drilling rigs, the Khazzan gas field in Oman. Data flowed in real time from more than two million sensors , pressure, temperature, flow rate, vibration. These numbers were transformed into digital replicas of physical assets.

In 2021, BP extended the contract by five years. Then in September 2024, the two companies agreed on a new level of collaboration: deploying Palantir's AIP (Artificial Intelligence Platform). AIP applied large language models (LLMs) to industrial operations. When a BP engineer asked in natural language, "Analyze the cause of pressure anomalies on Platform 3," AIP synthesized relevant sensor data, maintenance records, and similar past incidents to provide an answer. The critical feature was hallucination control. To prevent LLMs from confidently producing false information, transparency tools were applied to every AI recommendation. Security features constrained the LLM's range of action, and every decision and action was logged as a fully auditable digital record.

Matthew Babin, Palantir's head of energy and natural resources, said: "The goal going forward is the same , driving greater efficiency across BP operations and expanding data integration. Now AIP offers an opportunity to accelerate human decision-making on top of the solid digital twin and deep operational workflows already in place."

BP's subsidiary Castrol built a supply chain digital twin. Azule Energy, a joint venture between BP and Italy's Eni operating in Angola, also signed a multi-year contract with Palantir to manage daily production of 200,000 to 250,000 barrels. In August 2024, a collaboration with NASA was announced: BP's digital modeling technology, developed for drilling in hard-to-reach locations, would be applied to lunar and Mars exploration.

Airbus and BP. What did these two companies share? Both operate complex physical assets. Both depend on safety as a matter of life and death. Both were drowning in oceans of data. Palantir launched a ship on those waters. The price was steep. But without that ship, navigation itself seemed impossible.

(2) 'Warp Speed': A Manufacturing Operating System and American Reindustrialization

On December 11, 2024, Palantir announced a new program called 'Warp Speed.' It was an operating system for manufacturing. The first 'Cohort' included four companies: Anduril Industries, L3Harris, Panasonic Energy of North America, and Shield AI.

Palantir CTO Shyam Sankar declared in the announcement: "When World War II began, we didn't have a 'defense industrial base.' We had an 'American industrial base.' Our future must look the same. America must reindustrialize and mobilize at warp speed."

There was historical context behind that statement. When the United States entered World War II in 1941, automobile factories built tanks, and refrigerator factories made machine guns. Civilian industry converted to military production. Sankar was arguing that the same kind of national mobilization was needed again , only this time, software would make it possible.

The design philosophy was simple: "Software must adapt to the business. Not the other way around." Traditional enterprise software , SAP or Oracle ERP systems, for example , demanded that companies reshape their business processes to fit the software. Years of implementation, hundreds of millions in costs, and results that still didn't match what the floor actually needed. Warp Speed promised the opposite.

Testimonials from the first cohort members were striking.

Anduril Industries is a defense startup building autonomous weapons systems. They were the first company to adopt Warp Speed as part of their 'Arsenal OS' product line. Chief Information Officer Tom Bosco said: "Delivery speed is everything for us. Warp Speed lets us rapidly build manufacturing capacity for our customers. Using this software, we've seen up to a 200x efficiency improvement in our ability to anticipate and respond to supply shortages."

Two hundred times. It sounds like an exaggeration. But in context it makes sense. Dynamic material requirements planning (MRP), efficient phased introduction of new configurations, reduction in material and labor variances , tasks that once took days were now done in minutes.

Shield AI builds unmanned aerial systems. Their flagship product, the V-BAT drone, was facing record demand. CEO Ryan Tseng explained: "Our mission demands rigorous, integrated execution from design to delivery. Production is only as fast as the slowest part or process. Warp Speed will help our teams identify bottlenecks and stay in sync."

L3Harris announced a strategic partnership with Palantir in October 2024. This legacy defense company was using Warp Speed to close the loop between design, configuration, and production.

David Jack, Vice President of Global Operations and Program Excellence, said: "Warp Speed is a key element of our strategy to improve operational efficiency. It enables improvements like process automation and smarter inventory management." Panasonic Energy of North America was transforming American battery manufacturing through a three-year contract. Warp Speed supported the ramp-up at the Nevada Gigafactory and the Kansas De Soto plant. The goals were strengthening U.S. energy self-sufficiency and building resilience against global supply chain disruptions.

Warp Speed offered four core capabilities: Dynamic Production Scheduling, Engineering Change Management, Automated Visual Inspection, and a continuous network planning function called MRP Speed.

Among these, Engineering Change Management applied the CI/CD (Continuous Integration/Continuous Deployment) concept from software development to manufacturing. Design changes instantly reflected on the production line , something hard to imagine in traditional manufacturing.

One line from Sankar's speech stood out: "What we're supporting are companies that manufacture core products underpinning freedom and prosperity." Anduril's drones, Shield AI's unmanned aircraft, L3Harris's communications equipment, Panasonic's batteries , all were essential to American security and economic self-reliance.

Warp Speed was not just manufacturing software. It was an ambitious project aiming to become the software foundation for American reindustrialization , countering China's manufacturing rise, overcoming supply chain vulnerabilities, and accelerating the growth of defense innovation companies. Palantir intended to be the operating system for all of it.

C. The Bootcamp Strategy

(1) A Unique Sales Approach: Automating AI Workflows in Five Days

In mid-2023, Palantir began an unusual sales strategy: inviting prospective customers to a five-day intensive workshop. They called it the 'AIP Bootcamp.' Participants brought their own real data. Palantir engineers built AI workflows directly on top of it. Five days later, attendees witnessed problems they hadn't solved in years get resolved in a matter of days. CEO Alex Karp explained the strategy's purpose on an earnings call: "The AIP Bootcamp is a challenge we throw down, daring customers to compare against what they've built internally. No matter how many resources they've poured in or how many years it took , run your data on Palantir's platform for just ten hours, and you'll see the difference."

The results showed up in the numbers. At the end of 2022, total bootcamps numbered just 92. By the end of 2023, the count had surged past 500, with over 465 organizations participating. By June 2024, the cumulative total exceeded 1,300.

One case illustrates the power of this approach. Palantir closed a 3-million-dollar contract through outbound sales. Then they ran a bootcamp. Before the same quarter ended, that contract expanded

into an enterprise-wide deal , several times the original size.

CRO (Chief Revenue Officer) Ryan Taylor reported: "By the end of November, we plan to run bootcamps for more than 140 organizations. Nearly half of those will happen this month alone. That's more than all U.S. commercial pilots from last year combined."

The bootcamp's effects appeared across multiple dimensions.

First, it shortened the sales cycle. Traditional enterprise software sales typically take six to eighteen months: requirements analysis, proposal writing, technical validation, pilot projects, negotiations. The bootcamp compressed all of that into five days.

Second, it reduced customer acquisition cost (CAC). Instead of sales teams scheduling dozens of meetings to persuade potential buyers, the five-day intensive experience did that work. Seeing is believing. Executives who watched their own data being transformed were easily convinced.

Third, it created cross-sell opportunities. During bootcamps, participants were exposed to Palantir's various products: Foundry, AIP, Apollo. Many came to solve one problem and ended up purchasing multiple products simultaneously.

Fourth, it built internal capabilities. Palantir's blog explains: "Participants build their own intuition and skills so that in the weeks that follow, they can independently identify and execute AIP use cases." Palantir experts were being created inside customer organizations. These people became internal evangelists driving platform adoption.

Fifth, it strengthened the lock-in effect. Once workflows were built through a bootcamp, the cost of switching to another platform spiked. Data structures, business logic, and user training were all tailored to Palantir. In August 2024, a collaboration with Microsoft gave this strategy new momentum. Palantir's AIP could now be deployed in the Azure Government Secret environment. The two companies planned joint bootcamps for intelligence and military customers , Microsoft's government cloud infrastructure combined with Palantir's AI platform.

(2) The Virtuous Cycle of Customer Acquisition and Lock-In

The bootcamp strategy was directly reflected in financial performance.

In Q3 2025, Palantir's quarterly revenue reached approximately 1.18 billion dollars , a 39% increase year over year. U.S. revenue alone was 628 million dollars, up 55%. The U.S. commercial segment grew 65% year over year, the fastest-growing division.

Behind this growth were bootcamps. Analysts noted that bootcamps had replaced the traditional multi-month sales cycle. Customer acquisition costs plummeted, and profitability surged as a result. In the second half of 2025, Palantir's operating margin hit 51% , an unprecedented figure in the enterprise software industry.

There's a metric called the 'Rule of 40.' It sums revenue growth rate and profit margin , a measure of the balance between growth and profitability. Companies scoring above 40% are generally considered excellent. Palantir's Rule of 40 score was 114%. Management boasted it was the second highest among major tech companies, trailing only Nvidia.

The bootcamp strategy soon became a benchmark for competitors. Snowflake launched its own 'LLM Bootcamps.' But replicating Palantir's approach , building workflows on-site with a customer's actual data , was not easy to copy. C3.ai was assessed as losing market share to Palantir's bootcamp model due to its longer deployment timelines.

CTO Shyam Sankar explained the essence of the bootcamp: "The AIP Bootcamp drives home the point that LLMs alone are not enough. You need tools that provide algorithmic reasoning." He was talking about the ontology , Palantir's core technology that transforms data from mere streams of numbers into 'digital twins' of the real world. The bootcamp was the device that let people experience its power firsthand.

In the second half of 2025, Snowflake and Palantir forged a strategic partnership. Two companies that had once been fierce rivals joined forces. Palantir's AIP could now run natively on Snowflake's data cloud. Snowflake handled data storage; Palantir handled the 'operational logic.'

This partnership opened a new land-and-expand channel for Palantir , the ability to spread through Snowflake's massive customer base.

CEO Karp said commercial success was no surprise. "Given our military experience, it's only natural." He had a point. For 20 years, Palantir had supported decision-making in the fog of war. That experience had simply been transferred to the corporate battlefield.

"AIP lets companies manage LLMs and essentially run penetration tests on their own enterprise." This remark from Karp revealed AIP's true nature. It wasn't about deploying an AI chatbot. It was about deeply integrating AI across an enterprise's data and processes , and ensuring that integration remained safe.

On an earnings call, Karp declared: "My view of what we need to do is to build a product so good that competition stops. Whether in commerce or on the battlefield. That's what we're doing. That's what we're seeing with AIP."

In 2025, Palantir emerged as the biggest beneficiary of the 'AI supercycle.' Companies began demanding real ROI from their AI investments, and Palantir could prove that value in five days.

A simple idea , the bootcamp , became the growth engine powering a company worth hundreds of billions in market cap. Palantir's ability to turn crisis into opportunity didn't come from technology alone. It came from precisely identifying a customer's urgent need and presenting a solution in the fastest way possible. And once that solution took root, pulling it out was nearly impossible.

← Previous Chapter

See every chapter

Next Chapter →

Kim Kyung-jin

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com



#KimKyungjin #PALANTIR #Palantir #WarAI #SurveillanceAI #DefenseAI #DataPower #ArtificialIntelligence #AILibrary

# Chapter 8. Business Model and Financial Structure

← Previous Chapter

See every chapter

Next Chapter →

Part 4: Digital Transformation of Enterprise and the Public Sector

A. Balancing Government and Commercial Business

(1) Long-Term Contracts with the U.S. Department of Defense, Intelligence Agencies, and Allied Governments

On July 31, 2025, the U.S. Army issued a press release just two paragraphs long. But the number contained in that short document sent shockwaves through the entire defense industry. Ten billion dollars. A ten-year contract. Palantir Technologies.

The contract's official name was the Enterprise Service Agreement. Its core purpose was to consolidate 75 separate data and software contracts that the Army had been managing individually. Of those, 15 had Palantir as the prime contractor and 60 had it as a subcontractor. Years of fragmented agreements were being funneled into a single point of contact.

Alex Karp reportedly sat in silence for a long time in his office at the Denver headquarters when he heard the news. The company that had started 17 years earlier with $2 million from In-Q-Tel, the CIA's venture investment arm, had just won the largest defense contract ever awarded to a software company in American history.

There was a catch to the $10 billion figure. It was a ceiling, not a guaranteed amount. The Army had no obligation to spend the full sum. But the contract mattered for a different reason. Palantir had previously delivered its products through middleman resellers. Now those arrangements were converting to direct contracts, eliminating commissions and speeding up delivery. One thick layer of bureaucracy had been peeled away.

Government business had been the backbone of Palantir since its founding. As of Q3 2025, the government segment accounted for roughly 55% of total revenue. U.S. government revenue alone approached 80% of the total government figure. CIA, FBI, NSA, the Department of Defense, Army, Navy, Air Force. Nearly every American intelligence agency and military branch appeared on Palantir's client list.

Allied nations were no exception. Britain's GCHQ, France's DGSI, the Danish police, Australia's Department of Defence. Five Eyes allies and NATO member states analyzed intelligence and conducted operations on Palantir's platform. In December 2025, France's DGSI extended its contract with Palantir by another three years. A partnership spanning nearly a decade would continue.

But heavy dependence on government was a double-edged sword. Government budgets could be cut at any time depending on political decisions. When administrations changed, priorities shifted. By late 2025, a U.S. federal government shutdown had dragged into its fourth week, delaying approval of new contracts. Wall Street analysts worried about this concentration risk.

Karp was well aware of the concerns. So he decided to fire up a second engine.

(2) Commercial Expansion into Finance, Healthcare, and Manufacturing

In early 2024, during Palantir's earnings conference call, an analyst posed a sharp question. "Commercial segment growth, particularly outside the U.S., is slowing down. Could you explain what's happening?"

Karp's answer was brief. "We are focused on the United States. The U.S. is our most important market."

He wasn't bluffing. In Q3 2025, U.S. commercial revenue grew 121% year over year. $397 million. That was a 29% increase from the previous quarter as well. During the same period, international commercial revenue grew only 10%. In Q2 it had actually declined 3%. Palantir was deliberately concentrating its resources on the American market.

The number of commercial customers was surging. From roughly 169 in 2021, the customer count exceeded 1,500 by 2025. More important was the size of the deals. In Q3 2025 alone, 204 contracts exceeded $1 million, and 53 exceeded $10 million. The company was landing big fish.

The client roster was impressive. Airbus used Palantir's Foundry for aircraft maintenance and parts management under the name Skywise. BP adopted it for refinery operations. Merck brought it in for drug development data analysis. Lowe's, the largest home improvement retailer in the United States, used Palantir and Nvidia technology together to build a digital twin of its global supply chain network.

Growth in healthcare was also notable. Multiple hospitals and health systems across the U.S. began adopting Palantir's platform. AI was used to integrate patient data, optimize bed management, and improve staff allocation. The model proven at Britain's NHS had crossed the Atlantic.

In manufacturing, an operating system called Warp Speed drew attention. It was a solution that managed everything from factory production lines to supply chain management and inventory optimization on a single platform. Karp called it the software backbone of American re-industrialization. Finance was another important vertical. Investment banks, hedge funds, and insurance companies used Palantir's platform to analyze market data and manage risk. Ontology-based data integration also proved powerful for regulatory compliance work.

The rapid growth of the commercial segment sent an important message to investors. Palantir was no longer a government-only shop. It was a company running two engines, government and commercial. If one faltered, the other would hold. That was the business model Karp had designed.

As of Q3 2025, the revenue split between government and commercial was 55 to 45. That ratio had held steady for four consecutive quarters. The company had reached an equilibrium. But looking at growth rates alone, the commercial segment was dominant. U.S. commercial at 121% versus U.S. government at 52%. Forecasts emerged that a reversal might not be far off.

B. The Evolution of the Revenue Model

(1) The Mix of Licenses, Subscriptions, and Services

Palantir's early business model was straightforward. Deploy engineers to a client's site, integrate their data, and build a custom solution. These field-based engineers, called Forward Deployed Engineers (FDEs), were practically synonymous with Palantir. They lived in clients' offices, identified problems, and coded solutions.

The model was effective but had scalability limits. People are finite. Every new customer required hiring and training additional engineers. Revenue grew, but margins didn't improve. Wall Street hammered this point relentlessly. "Is Palantir a software company or a consulting firm?"

Karp agonized over this question for years. Then he found the answer. Apollo.

Apollo was a software deployment platform. It allowed Palantir's software to be installed and updated anywhere, from the cloud to a client's on-premise servers, even to edge environments on battlefields with no internet connection. Once Apollo arrived, the FDE role changed. Previously, engineers had to install and maintain software on site. Now Apollo automated much of that work.

The revenue model evolved too. In the early days, many contracts were one-off, project-based deals. Clients paid a large upfront build fee, followed by maintenance charges. But the company gradually shifted to a subscription model. Customers paid a fixed amount annually or monthly to use the platform. It was the standard SaaS approach.

The launch of AIP (Artificial Intelligence Platform) accelerated this transition. AIP was a platform that let enterprises apply generative AI in their operations. It controlled the hallucination problem of large language models through ontology-based data structures and embedded AI into real business processes. Customers paid additional fees on top of their existing Foundry or Gotham contracts to use AIP.

In Q3 2025, Palantir's revenue reached $1.18 billion. That was 63% growth year over year. Even more striking was profitability. Net income was $475.6 million, nearly triple the year-ago figure. A 40% net margin. Adjusted operating margin oscillated between 46% and 51%. Gross margin held between 81% and 84%.

There is a metric called the Rule of 40. It measures the health of a software company by adding revenue growth rate to profit margin; a score above 40 qualifies a company as high-quality. In Q2 2025, Palantir's Rule of 40 score was 94. By Q3 it had soared to 114. Numbers like these were virtually unprecedented in enterprise software.

Adjusted free cash flow was equally strong. In Q3 2025 alone, $540 million in cash flowed in. The trailing twelve-month total exceeded $2 billion. Annual guidance was between $1.9 billion and $2.1 billion. Palantir was no longer a cash-burning startup. It had become a cash-generating machine.

(2) Reaching a $300 Trillion Won Market Cap and Overtaking Legacy Defense Contractors

On November 3, 2025, Palantir's stock hit an all-time high of $207.52. Market capitalization surpassed $490 billion. Roughly 680 trillion Korean won. More than South Korea's entire annual national budget.

What the number meant was clear. Palantir had surpassed Lockheed Martin, Raytheon, and Northrop Grumman, the legacy defense giants. Of course, by revenue alone, there was no comparison. Lockheed Martin's annual revenue exceeded $70 billion.

Palantir's projected 2025 revenue was around $4.4 billion. A 16x gap. Yet the market assigned Palantir a higher valuation.

Why? The answer lay in growth rates. Lockheed Martin's revenue grew 5-6% per year. Palantir grew 50-60%. Investors were buying the future. They bet not on today's revenue but on tomorrow's potential.

Criticism of this rich valuation was fierce. In August 2025, Britain's The Economist called Palantir "the most overvalued company in history." Citron Research, the well-known short seller, calculated Palantir's fair value at $40 per share, less than 75% below the market price at the time. The stock plunged 25% right after Citron's report was published.

Valuation metrics were extreme. As of late 2025, Palantir's price-to-earnings ratio stood at 387x. Price-to-sales was 109x. Legacy defense firms traded at roughly 2x P/S, a gap of more than 50-fold. From a Warren Buffett-style value investing perspective, the numbers were impossible to justify.

Wall Street analysts were divided. Of 26 analysts, 17 rated the stock Hold, and 3 recommended Sell or Strong Sell. The average price target was $155, about 18% below the market price. Major investors were exiting. Legendary hedge fund manager Stanley Druckenmiller sold his entire Palantir position. Cathie Wood's ARK Invest also sold approximately 9 million shares.

Karp responded to the criticism in his characteristic way. In a letter to shareholders after the Q3 2025 earnings release, he wrote: "The critics are confused. We have delivered venture-capital-level returns to retail investors." It was the kind of sentence that invites charges of arrogance. But the numbers backed him up. Since its 2020 IPO, Palantir's market cap had grown from $15.6 billion to over $400 billion. A compound annual growth rate of 86%. That was only possible because of the fervent support of retail investors.

In September 2025, Goldman Sachs data showed that retail investors poured $1.2 billion into Palantir stock in a single month. They ignored the warnings of Wall Street analysts. To them, Palantir was not just a stock. It was a symbol of the AI era.

C. Global Partnership Strategy

(1) Partnerships with Nvidia and Microsoft

October 28, 2025. A convention center in Washington, D.C. On the stage at Nvidia's annual GTC technology conference, Jensen Huang addressed the audience.

"Palantir's Ontology is probably the most important enterprise software stack in the world."

That single sentence from the CEO of the world's largest AI semiconductor company sent ripples through the tech industry. It was not just praise. It was a signal announcing a new partnership between the two companies.

Here was the announcement: Nvidia's GPU-accelerated computing, its CUDA-X data science libraries, and its open-source AI model Nemotron would be integrated with Palantir's Ontology framework. This would allow enterprises and government agencies to build real-time decision intelligence systems.

In more technical detail, it worked like this.

Palantir's AIP (Artificial Intelligence Platform) gained access to Nvidia's Nemotron Super model, an inference model with 49 billion parameters. The 9-billion-parameter Nemotron Nano 2 and the 8-billion-parameter vision model Llama Nemotron Nano VL were also loaded onto the Palantir platform.

A decision-optimization tool called cuOpt was integrated as well. It solved complex problems like supply chain management and logistics optimization using AI. A technology called NeMo Retriever helped companies build AI agents grounded in ontology data.

The first major customer was Lowe's. Lowe's used the combined Palantir-Nvidia technology stack to build a digital twin of its global supply chain network. Where individual nodes had previously been optimized once a week, continuous and dynamic optimization at a global scale was now possible.

Karp said at the announcement: "We are fusing AI-powered decision intelligence systems with the world's best AI infrastructure." In December 2025, the collaboration expanded further. A project called Chain Reaction was announced. It was essentially an operating system to accelerate the buildout of America's AI infrastructure. The goal was to manage the complex supply chains spanning power generation, transmission, and data center construction using Palantir's AIP and Ontology together with Nvidia's AI models. CenterPoint Energy, a major Houston-based energy company, joined as a founding partner.

The collaboration with Microsoft also deepened. The core of a partnership announced in August 2024 was this: Palantir's full product suite, Foundry, Gotham, Apollo, and AIP, would be deployed on Microsoft Azure Government Cloud. This included both Secret and Top Secret classified cloud environments.

More significant was the integration with Azure OpenAI Service. Large language models including GPT-4 became usable within the Palantir platform, and in classified network environments at that. A path opened for analysts at the Department of Defense and intelligence agencies to use generative AI in their work with security guaranteed.

Palantir CTO Shyam Sankar described the partnership this way: "Bringing the combined capabilities of Palantir and Microsoft to national security organizations is a fundamental shift in how we support the defense and intelligence community."

The two partnerships, Nvidia and Microsoft, carried a clear message. Palantir was no longer fighting alone. It was joining hands with the key infrastructure companies of the AI era, positioning itself at the center of the ecosystem.

(2) Entering the Korean Market: Alliances with HD Hyundai, KT, and Samsung

The morning of October 14, 2025. KT's headquarters near Gwanghwamun, Seoul. Alex Karp walked into the building. His visit to Korea had been kept strictly confidential. It was not on any public schedule.

KT CEO Kim Young-shub greeted him. It was their first meeting on Korean soil since signing a strategic partnership in the U.S. in March. In the conference room, they reviewed the results of seven months of collaboration and discussed strategy for spreading Palantir's technology across Korean industry.

That same afternoon, KT hosted the AX Leader Summit. AX stood for AI Transformation. On the surface, it was a forum for sharing enterprise innovation strategies using AI. But the event's real purpose lay elsewhere.

Four executives at the chairman or vice-chairman level of major Korean conglomerates attended, each personally selected by Karp. Korean Air Vice Chairman Woo Ki-hong, Meritz Financial Group Vice Chairman Kim Yong-beom, LS Electric Chairman Koo Ja-kyun, and POSCO Holdings CEO Lee Ju-tae. Each had a 25-minute one-on-one meeting with Karp. Each session took place in a separate room, and attendees did not know the contents of each other's meetings. A person familiar with the matter said: "Every meeting was kept secret. The goal was to prevent each company's AI transformation strategy from being exposed."

This secretive approach was typical Palantir. They handled clients' data and strategies. Trust and confidentiality were central to the business.

The partnership with KT was already producing tangible results. At the time of the March signing, KT became the first Korean company to officially join Palantir's Worldwide Partner Ecosystem. Since then, KT had built its own cloud-based work environment and piloted Foundry and AIP in select business units. The results were good. KT now planned to offer the platform as a service to financial and public sector clients.

Karp spoke about the KT collaboration: "We value the partnership with KT highly. This cooperation is an important step toward secure cloud-based data use and industry-specific innovation."

Palantir's moves in Korea did not stop at KT.

In April 2024, a memorandum of understanding was signed with HD Hyundai Heavy Industries at Palantir's Washington, D.C. office. The two companies agreed to jointly develop an unmanned surface vessel (USV).

The USV was named Tenebris, Latin for "darkness." A mid-size vessel 17 meters long and weighing 14 tons, it targeted a speed of 50 knots and an operating range of over 1,000 nautical miles. It was designed to operate in rough seas up to Sea State 6.

Autonomous navigation technology developed by Avikus, an HD Hyundai subsidiary, would be combined with Palantir's AI-based mission autonomy system. The plan called for developing a reconnaissance USV first by 2026, then expanding to combat variants. In May, a scale model of Tenebris was unveiled at an AI expo in Washington. South Korea's Ambassador to the U.S., Cho Hyun-dong, visited Palantir's booth for a briefing.

HD Hyundai Heavy Industries CEO Joo Won-ho said of the collaboration: "The unmanned vessel market is a new blue ocean where advanced technology determines the winner. Building on the achievements and trust both companies have accumulated, we will pioneer this field." Cooperation with Samsung was also underway. Applications of Palantir's platform in supply chain management and manufacturing were being discussed. Partnerships with other Korean companies, including LIG Nex1, DL E&C, Kolon Benit, and Waiker, were expanding as well. Palantir's interest in the Korean market was serious. During his October visit, Karp told a reporter: "Korea is the most exciting market outside the United States."

A Palantir pop-up store even opened in Seoul's Seongsu-dong neighborhood. It was a marketing effort to raise brand awareness. A Silicon Valley defense-AI company running a pop-up store in one of Seoul's trendiest districts. Ten years ago, that scene would have been hard to imagine.

Korea was important to Palantir in several ways. World-class IT infrastructure, a corporate culture eager to use data, and a security environment defined by the North Korean threat. The conditions for Palantir's technology were in place in both the government and private sectors. The global rise of K-defense also broadened opportunities for collaboration.

Karp's Seoul visit lasted just one day. But the seeds he planted had the potential to reshape Korea's AI landscape. If major Korean conglomerates begin making decisions on Palantir's platform, it would mean not just adopting new software but a fundamental change in how they run their businesses.

← Previous Chapter

See every chapter

Next Chapter →

Kim Kyung-jin

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com



#KimKyungjin #PALANTIR #Palantir #WarAI #SurveillanceAI #DefenseAI #DataPower #ArtificialIntelligence #AILibrary

# Chapter 9. Big Brother's Eye: Surveillance and Predictive Policing

← Previous Chapter

See every chapter

Next Chapter →

Part 5: Controversy and Regulation, The Shadow of the Surveillance State

A. The Reality of Minority Report

(1) The Achievements and Controversies of LAPD's Predictive Policing Program (Operation Laser)

One day in 2011, a police officer in the Newton Division of the Los Angeles Police Department (LAPD) was staring at a tablet inside his patrol car. On the screen was a photograph of a young Black man. Name, address, tattoo locations, vehicle license plate number, and a number written beside them. 15 points. This young man had been classified as a "Chronic Offender." The officer's mission was to find this person and make contact. What crime he had committed did not matter. The algorithm had designated him as a dangerous individual.

This was how Operation LASER worked. LASER stood for "Los Angeles Strategic Extraction and Restoration." Translated literally, it means "Los Angeles Strategic Extraction and Restoration." The people who designed the program embedded a medical metaphor in the name. The official documents from the Bureau of Justice Assistance under the federal Department of Justice stated the following: "Operation LASER can be likened to laser surgery. Just as a trained physician uses modern technology to remove a tumor or improve eyesight."

Tumor. They called certain citizens tumors.

Palantir's Gotham platform served as the brain of this program. The system integrated dozens of data sources, including arrest records, field interview reports, automatic license plate reader data, gang databases, and social media activity. And it assigned a score to each individual. Five points if they had previously possessed a firearm. One point if they were on probation. One point was added each time they had contact with police. Once a score exceeded a certain threshold, the person's name was placed on the "Chronic Offender Bulletin." Like a wanted poster, photographs and personal information were distributed to patrol cars.

The problem lay in the operating logic of this system. If the score went up with every police contact, then people living in areas where police frequently patrolled automatically became high-risk individuals. And where did police patrol frequently? Areas where many crimes had been reported in the past. This created a self-fulfilling prophecy loop. Police went more often to poor Black and Hispanic neighborhoods, residents of those areas were stopped and searched more frequently, so they received higher scores, and so they became surveillance targets more often.

In March 2019, the audit results from the LAPD Inspector General's office were released. The results were shocking. Of the 233 people classified as chronic offenders, 84 percent were Latino or Black. These two groups accounted for less than 60 percent of Los Angeles's total population. There were even more surprising facts. Nearly half of the chronic offenders had never been arrested for a violent crime, or had been arrested only once. And about 20 percent, one in five, had no record whatsoever of "meaningful contact" with police.

The audit report pointed to another problem. It was impossible to demonstrate the program's effectiveness. Of the 13 areas designated as LASER zones, six areas actually saw increases in violent crime or showed no difference from non-LASER zones. Data collection methods were inconsistent. Some officers diligently entered their contact records, while others did not. There was no way to verify the accuracy of the system itself.

Hamid Khan, co-founder of the Stop LAPD Spying Coalition, had fought against this program for years. His organization secured internal documents through public records request lawsuits and organized protests with community residents. Khan said in an interview with MIT Technology Review: "We officially launched the campaign in 2016. The goal was to understand the program, inform the community, and dismantle it."

In April 2019, LAPD finally officially terminated Operation LASER. But the story did not end there.

(2) Algorithmic Bias and Racial Discrimination Issues

The very month Operation LASER was abolished, LAPD announced a new program. It was called "Data-Informed Community-Focused Policing" (DICFP). Only the name had changed.

In 2021, new documents obtained by the Stop LAPD Spying Coalition were made public. The documents showed how LASER and another predictive policing program called PredPol had automated and reinforced existing racial biases. PredPol was a location-based system that predicted 500-foot by 500-foot areas with a high probability of crime based on past crime data. It was an algorithm inspired by earthquake aftershock prediction models.

Sarah Brayne, a sociologist at the University of Texas at Austin, participated as an embedded researcher at LAPD for several months and observed how predictive policing actually worked. Her book, Predict and Surveil, contains shocking scenes. One sergeant explained to her: "There are a lot of minor violations to stop someone for. Yesterday, this person may have been stopped for jaywalking."

Brayne also witnessed how police used automatic license plate readers installed in front of emergency rooms. Family members or friends of shooting victims often drove the patient to the hospital and left quickly. Police tracked those license plates to build the victim's social network. Even without evidence of a crime, people who entered the network became surveillance targets.

Andrew Ferguson, a law professor at American University and author of The Rise of Big Data Policing, defined the essence of the Palantir system this way: "Palantir is less a data-driven policing

system than a data-driven surveillance system. Palantir helps police surveil people, not keep people safe."

In 2023, the nonprofit investigative outlet The Markup analyzed the actual performance of PredPol (later rebranded as Geolitica). After reviewing more than 23,000 predictions over the course of 2018, the success rate was less than 0.5 percent. Fewer than 100 cases saw a predicted category of crime actually occur at the predicted location and time and be reported to police. The Plainfield police chief told The Markup that the money spent on Geolitica software "could have been better spent elsewhere."

In April 2020, LAPD announced that it was terminating its PredPol contract, citing budget constraints due to COVID-19. But critics pointed out that this was not the real reason. Years of accumulated evidence showing the program was ineffective and discriminatory had led to this result. The Palo Alto Police Department had already terminated its contract after three years of using PredPol, saying it "didn't get any value."

The core of the problem is not that algorithms create bias, but that they amplify existing bias. Shakeer Rahman, an activist with the Stop LAPD Spying Coalition, explained it this way: "Data-driven policing automates existing police logic. That includes targeting poor people, targeting unhoused people, targeting Black and brown and disabled people. Now this is helping automate those practices, automate harm, automate deportation, automate the exclusion that police have always been in the business of."

In February 2024, the city of Chicago refused to renew its contract with ShotSpotter, a gunshot detection and predictive policing system. Mayor Brandon Johnson stated in a press release that "the City of Chicago will deploy resources toward strategies and tactics proven to be effective in accelerating the current trend of declining violent crime." Skepticism toward predictive policing was spreading.

Yet Palantir was still thriving. Through COVID-19 tracking and vaccine safety analysis contracts, it had gained entry to the National Institutes of Health (NIH) and the Food and Drug Administration (FDA). Since its stock market listing in September 2020, its share price had more than tripled. The failure in Los Angeles seemed to have no impact on the company's growth.

B. ICE and the Deportation Machine

(1) Supporting Immigration and Customs Enforcement's Tracking of Undocumented Immigrants

One day in April 2025, an agent in the Enforcement and Removal Operations division of U.S. Immigration and Customs Enforcement (ICE) logged into a new system. On the screen appeared the name "ImmigrationOS." This platform, developed by Palantir, gave the agent remarkable capabilities. With a few clicks, they could view a specific individual's passport records, Social Security files, IRS tax data, driver's license information, and even automatic license plate reader data at a glance. The system could track voluntary departures "in near real-time," identify visa overstays, and designate individuals associated with transnational criminal organizations like MS-13 or Tren de Aragua as

targets.

This was the deportation machine of the Trump second-term administration.

The relationship between Palantir and ICE dates back to 2014. At the time, Palantir built a tool called Falcon for ICE's Homeland Security Investigations (HSI). Falcon enabled agents to track specific individuals' locations over time by cell phone number, analyze air travel records, and share field interview information in real time. This contract, signed during the Obama administration, continued through the Biden administration.

But the ImmigrationOS of 2025 was on a different level. In mid-April 2025, ICE signed a $30 million contract with Palantir. According to federal contract records, a prototype was to be delivered by September 25, 2025. In the contract justification documents, ICE stated there was an "urgent and compelling need." And it declared Palantir was the "sole source."

"Palantir possesses deep institutional knowledge of ICE operations," the documents stated. "The company is already collecting and processing data from multiple federal agencies, placing it in a unique position to develop a prototype within six months." It warned that any delay would undermine the enforcement mission specified in President Trump's executive orders.

ImmigrationOS was designed to provide three core functions. First, target selection and enforcement prioritization. This function helps determine who to deport first, including "violent criminals," gang members, and visa overstays. What data or criteria are used to identify "violent criminals" was not disclosed. Second, voluntary departure tracking. It monitors whether individuals leave the United States voluntarily. Third, immigration lifecycle management. It simplifys the entire process from identification to deportation, minimizing time and resource waste.

In September 2025, ICE signed an additional contract worth $29.9 million. It was for software license renewal, operations and maintenance, and adaptive maintenance support. The contract was valid through September 2027. According to public records obtained by the New York Times, the total value of federal contracts awarded to Palantir since Trump's inauguration exceeded $900 million. This was nearly double the $541.2 million from 2024.

The system's power lay in its data integration capability. ImmigrationOS pulled in not only government databases but also information from the private sector. Credit information, utility records, Department of Motor Vehicles data, and airline passenger information were all connected. According to analysis by the Migration Policy Institute, ICE is believed to be among the largest government purchasers of commercial credit, utility, and vehicle registration information.

In April 2025, Palantir's Chief Technology Officer (CTO) and President of Palantir USG, Akash Jain, wrote in an internal Slack message: "A quick update on our work with ICE." According to the leaked message, the company had developed "a new set of data integrations and workflows with ICE" in response to "the new administration's focus on using data to drive enforcement operations."

According to CNN reporting, ImmigrationOS would use artificial intelligence systems to allow immigration agents to "authorize raids, book arrests, generate legal documents, and route transfers to deportation flights or detention facilities from a single interface."

(2) Internal Employees' Ethical Backlash and Exodus of Resignations

In May 2025, thirteen former Palantir employees published an open letter. The title was "The Scouring of the Shire." It referenced the scene in The Lord of the Rings where the hobbits' homeland is devastated even after Sauron's defeat. Considering that the company's name derived from Tolkien's world, there could be no sharper criticism.

"Companies are kowtowing to the Trump administration, suppressing dissent, and aligning with his xenophobic, sexist, plutocratic agenda," the letter stated. They warned that the company's ethical guardrails, originally designed to prevent discrimination, disinformation, and abuse of power, were being dismantled.

This backlash was not new. CEO Alex Karp frankly acknowledged it in a March 2024 CNBC interview. "We have lost employees. We will lose more." He continued: "If you have a position that doesn't cost you employees, it's not a position."

Brianna Katherine Martin was one of the employees who left the company over ethical concerns. She said projects like ImmigrationOS constituted "complicity in human rights and constitutional violations."

Criticism from human rights organizations was even fiercer. Mariana Olaizola Rosenblat, a policy advisor at NYU Stern's Center for Business and Human Rights, said in an interview with Axios: "By supporting the Trump administration's deportation apparatus, Palantir is complicit in those human rights and constitutional violations."

The American Immigration Council pointed out fundamental problems with the system. ImmigrationOS pulls information from government databases regardless of their accuracy or reliability. Outdated or incorrect information can lead to life-altering consequences. Detention, loss of status, or wrongful deportation.

One of the most uncomfortable facts was the conflict of interest issue. Stephen Miller, the chief architect of the Trump administration's immigration policy, held significant financial stakes in Palantir. The person designing the government's immigration policy was a shareholder in the company providing the technology to enforce that policy. Was this a coincidence, or a designed outcome?

In December 2025, the Washington Post reported on how Palantir had pivoted. The article's lead read: "For years, Alex Karp declared that his data management firm was 'engaged in supporting progressive values' and 'repeatedly walked away' from contracts it deemed targeting minorities or unethical."

But now the company had become the technological backbone of the most aggressive mass deportation campaign in American history. Palantir argues that it merely provides tools. It is not the

company that decides who to target, deport, and surveil, they say.

But critics push back. Design choices shape real-world outcomes. AI architecture becomes policy. When an algorithm classifies certain groups as higher risk, that is not a technical decision but a political one.

Likhita Banerji, deputy director of Amnesty Tech at Amnesty International, told El País: "These technologies can systematically promote racism, discrimination, and oppression, and are routinely used to advance racist and xenophobic agendas."

C. Lavender and the Gaza Strip

(1) Official Declaration of Support in the Israel-Hamas War

One day in January 2024, an unusual gathering took place in a conference room in Tel Aviv. It was the first time the Palantir board of directors held a meeting outside the United States, and in Israel no less. Peter Thiel and Alex Karp sat alongside Israeli Ministry of Defense officials for a photograph. Danny Gold, director of the Defense Research and Development Directorate (DDR&D), known for leading the development of the Iron Dome missile defense system, was also present.

The photograph did not appear in the Israeli press. But Bloomberg reported it. The news was that Israel had agreed to purchase an AI-based system called AIP from Palantir. This system could support decision-making based on intelligence, analyze enemy targets, and suggest combat actions. The contract was expected to be worth tens of millions of dollars.

Alex Karp did not hide his position. In a February 2024 earnings conference call, he said: "I am incredibly proud that within weeks of October 7th, we were on the ground and involved in operationally significant activities in Israel."

Karp became even more candid during a conversation with legendary Israeli entrepreneur Yossi Vardi at Tel Aviv University. "There is nobody who doesn't have a position on Israel, if you're a big company. There are plenty of people who are not as pro-Israel as I am, but they think of Israel as a very special place and they are more sympathetic to Israel's position and appreciate what Israel has accomplished in building a state in the desert."

The company took out a full-page advertisement in the New York Times. "Palantir stands with Israel."

The relationship between Palantir and the Israeli military establishment had begun much earlier. In 2015, the company opened an office overlooking Rothschild Boulevard and Yehuda Halevi Street in Tel Aviv. This was a strategic move to deeply integrate into the Israeli tech ecosystem. Karp himself acknowledged it in a December 2023 interview with Fox Business: "We are very well known in Israel. Israel thinks highly of our products."

But the timing of the January 2024 agreement, which Bloomberg described as "a strategic partnership for battlefield technology," raised questions. This agreement was concluded at a point when the carnage in Gaza had already been underway for three months. The Business and Human

Rights Resource Centre asked: "Was it a calculated move by Palantir to use an escalating conflict as an opportunity to test its AI models against civilians? Did it turn Gaza into a horrific proving ground for technology?"

(2) AI Target Identification Systems and the "War Crimes by Algorithm" Debate

In April 2024, the Israel-Palestine specialist outlet +972 Magazine and the Hebrew-language outlet Local Call published a shocking investigative report. The headline read: "'Lavender': The AI Machine Directing Israel's Bombing Spree in Gaza." Six Israeli military intelligence officers testified on condition of anonymity.

Lavender was an AI-based database designed to mark all suspected members of Hamas and Palestinian Islamic Jihad (PIJ) military organizations, including low-ranking operatives, as potential bombing targets. The system analyzed virtually everyone in the Gaza Strip and assigned a score from 1 to 100. That score indicated the likelihood that the person was actively involved with an armed group.

According to the sources, during the first few weeks of the war, the Israeli military relied almost entirely on Lavender. The system marked 37,000 Palestinians as suspected combatants, and their homes as potential airstrike targets.

Even more shocking was the human review process. The intelligence officers who testified said they spent an average of 20 seconds approving a target generated by Lavender. In many cases, the verification amounted to checking whether the target was male. Under the assumption that women were not combatants.

Everyone knew the system had an error rate of about 10 percent. But it was considered an acceptable level. One source told +972: "We were not interested in killing Hamas operatives only when they were in a military building or engaged in military activity. On the contrary, the IDF bombed them in homes without hesitation, as a first option. It's much easier to bomb a family's home. The system is built to look for them in those situations."

There was another system linked to Lavender. It was called "Where's Daddy?" This system tracked individuals marked by Lavender until they entered their homes. It deliberately targeted nighttime, when families would be together.

The acceptable civilian casualty thresholds were also revealed. For killing a single low-ranking Hamas operative, the deaths of 15 to 20 civilians were permitted. For senior commanders, that number rose to between 100 and 300.

One source testified: "The targets are endless." Once the bombing started, commanders knew "there are another 36,000 waiting."

In May 2025, evidence recorded in classified Israeli military databases was made public. Of the 53,000 Palestinians killed in Gaza, only 17 percent were combatants. This means 83 percent were

civilians. The UN verified at the end of 2024 that approximately 70 percent of the dead were women and children.

UN Secretary-General António Guterres said he was "deeply concerned." He pointed out that Israel's use of artificial intelligence in military operations endangered civilians and obscured accountability.

Ben Saul, the UN Special Rapporteur on human rights and counter-terrorism, was more direct: "If the reports about Israel's use of AI are true, many Israeli airstrikes in Gaza would constitute war crimes of disproportionate attacks."

Marc Owen Jones, a professor at Hamad bin Khalifa University, said of the Lavender system: "Let's be clear. This is AI-assisted genocide, and there should be calls for a moratorium on the use of AI in warfare going forward."

Whether Palantir developed Lavender has not been confirmed. It is reported to have been developed by Unit 8200, the Israeli military's intelligence unit. But the timing of the January 2024 strategic partnership, and the fact that Palantir's technological capabilities could be applied to similar systems, raises questions.

What is certain is that Palantir actively supported this war. In a statement submitted to the Business and Human Rights Resource Centre, Palantir acknowledged providing "expanded, critical target identification capabilities" to the Israeli military. Those capabilities were likely used to carry out devastating strikes in Gaza, and possibly beyond.

Karp appeared visibly shaken when questioned about Lavender at a Cambridge University debate in May 2024. Investor Catherine Austin Fitts observed: "Peter Thiel may be legally protected, but I think he's beginning to realize that a reckoning could come in the court of public opinion."

But Thiel's position remained unshaken. "I broadly believe that the IDF gets to decide what it wants, and that they are broadly right."

Protests continued in front of Palantir headquarters, at the Hill and Valley Forum in Washington, D.C., and on university campuses worldwide. Palestinian-American Sumer Mobarak rushed the stage during a Karp speech, shouting: "You are making money off of killing Palestinian people. Palantir kills Palestinian people with their AI and their technology. You are killing my family in Palestine."

Karp's response was this: "The primary reason people are dying in Palestine is because Hamas figured out it has millions of 'useful idiots.'"

The video of this scene was viewed millions of times. But Palantir's stock price continued to climb. Market capitalization exceeded $200 billion. Amid the debate over war crimes by algorithm, investors were making money.

Who watches the algorithms. That is the question this era poses.

← Previous Chapter

See every chapter

Next Chapter →

Kim Kyung-jin

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com



#KimKyungjin #PALANTIR #Palantir #WarAI #SurveillanceAI #DefenseAI #DataPower #ArtificialIntelligence #AILibrary

# Chapter 10. Legal Checks and Europe's Counterattack

← Previous Chapter

See every chapter

Next Chapter →

Part 5: Controversy and Regulation, The Shadow of the Surveillance State

A. The German Federal Constitutional Court's Unconstitutionality Ruling (February 2023)

(1) The Ruling Declaring Palantir Use by Hesse and Hamburg State Police Unconstitutional

On the morning of February 16, 2023, a ruling was read aloud in the chambers of the German Federal Constitutional Court in Karlsruhe. Presiding Judge Stephan Harbarth spoke in a calm but firm voice: "Section 25a, paragraph 1 of the Hesse Public Safety and Order Act, and Section 49, paragraph 1 of the Hamburg Police Data Processing Act, are unconstitutional."

The eleven plaintiffs gathered in the courtroom held their breath. With the help of the civil liberties organization Gesellschaft für Freiheitsrechte (GFF), they had carried on this fight for six years. The target was software developed by Palantir, a data analytics company from Silicon Valley.

The Hesse state police had been operating Palantir's Gotham platform under the name "hessenDATA" since 2017. Some 2,000 police officers had access to the system. The official purpose was terrorism and organized crime investigations. But reality was different. Its use rapidly expanded to include residential burglaries, ATM bombings, and traffic accident witness tracking. According to Hesse Ministry of Interior data, police used the software thousands of times per year.

Hamburg had not yet activated the system. It had only established the legal basis and was preparing for deployment. The Constitutional Court ruled on both states' statutory provisions.

The core of the ruling was clear. The court held that automated data analysis violated the "right to informational self-determination" guaranteed by Articles 2(1) and 1(1) of the Basic Law. The ruling stated: "Particularly in light of the broadly defined text of the authorization, the basis for interference, with regard to both data and methods, falls far short of the constitutionally required threshold of identifiable danger."

Put simply, this means: for police to analyze citizens' data with Palantir software, a specific and clear danger must exist. A vague notion of crime prevention is insufficient. The abstract goal of preventing terrorism cannot serve as a legitimate basis for violating individual privacy.

The court ordered Hesse to amend its law by September 30, 2023. In Hamburg's case, since the law had not yet been implemented, it was immediately invalidated.

(2) Violation of the 'Right to Informational Self-Determination'

The "right to informational self-determination" invoked by the German Federal Constitutional Court is a landmark concept in German constitutional history.

In 1983, the Constitutional Court issued the Census Judgment (Volkszählungsurteil). At the time, the German government was planning a detailed census. Citizens objected. In this case, the court derived a new fundamental right: every individual has the right to determine who uses information about them, for what purpose, and in what manner.

The birth of this right was rooted in Germany's dark history. The Nazi Gestapo collected and categorized citizens' information to identify Jews, homosexuals, and political opponents. After the war, East Germany's Ministry for State Security, the Stasi, built an even more sophisticated surveillance apparatus. By some estimates, out of East Germany's population of 16 million, approximately one million served as the Stasi's informal collaborators. It was a society where neighbors spied on neighbors and family members informed on family members.

In the 2023 ruling, the court recalled this historical context. The ruling explained that automated data analysis constitutes a qualitatively different new form of interference compared to traditional data collection. Connecting individually collected pieces of information generates "new knowledge." Combining police records, communications metadata, social media activity, and vehicle location data can reconstruct an individual's behavioral patterns, social network, and movement history. This goes beyond the original purpose of information collection.

The court emphasized the principle of purpose limitation and the principle of proportionality. Data collected for a specific purpose must be used only for that purpose. Using it for other purposes requires a separate legitimate basis. also, government interference must be proportionate to the objective pursued. Extensive surveillance cannot be permitted to prevent minor crimes.

(3) Warning About "Profiling with a Single Click"

Franziska Görlitz, a lawyer with the GFF, spoke at a press conference immediately after the ruling: "A person who filed a complaint, a crime victim, or someone who was simply in the wrong place at the wrong time can all become subjects of this software's analysis."

Her statement was no exaggeration. According to internal documents from the Hesse state police, hessenDATA could integrate police records systems, criminal investigation databases, communications metadata, and social media information. In December 2022, Hesse Interior Minister Peter Beuth told the state parliament that the software had been used in the investigation of the far-right underground organization "Reichsbürger," which had plotted to overthrow the German government.

Constanze Kurz, a spokesperson for the Chaos Computer Club (CCC), was more biting: "This is Palantir-style dragnet investigation. Police are combining different datasets for purposes other than the original ones. Automated mass analysis must not become routine in law enforcement."

The core of the problem was the opacity of the algorithms. Palantir's software is written in proprietary code. Neither outside experts nor the agencies using the system fully understand how the algorithms create their connections. In 2023, the state of Bavaria commissioned the Fraunhofer Institute for Secure Information Technology to review the source code. The institute concluded there were no "secret backdoors." However, this evaluation report was classified as confidential, citing "security concerns" and "Palantir trade secrets."

The ruling pinpointed this issue precisely. "The fact that unlimited data analysis technology is not currently available does not resolve the problem. Constitutional requirements must be based on the legally possible scope of interference." As technology advances, more extensive analysis becomes possible. If the legal scope of permission is broadly set, surveillance can expand to fill that scope.

(4) The Order to Amend Hesse's Law by September 2023 and the Invalidation of Hamburg's Law

After announcing the ruling, Presiding Judge Harbarth emphasized one point to reporters: "Each state retains the option of constitutionally shaping the legal basis for further processing of stored data files."

This statement was double-edged. The Constitutional Court did not ban Palantir software itself. It determined that the laws supporting its use did not meet constitutional requirements. This also meant that continued use was possible if laws specifying stricter conditions were enacted.

Hesse amended its law by the deadline. The new law specified the requirements for data analysis. But critics argued the amendment was insufficient. As of 2025, the GFF has filed a new lawsuit before the Hesse State Constitutional Court.

Hamburg chose a different path. After its law was invalidated, it did not pursue new legislation. Fundamental skepticism toward Palantir adoption had emerged within the Hamburg state parliament.

The ruling's aftershocks spread throughout Germany. Police in Bavaria and North Rhine-Westphalia were also using Palantir software under the names VeRA and DAR, respectively. The GFF and the Chaos Computer Club filed constitutional complaints against these two states as well. The lawsuit against Bavaria is pending before the Federal Constitutional Court as of 2025.

B. GDPR and the Era of Data Sovereignty

(1) Setbacks and Recalibration of the European Expansion Strategy

Alex Karp took the stage at a corporate investment conference in February 2024. It was the FII Institute event in Miami. In his characteristically rapid-fire speaking style, he appealed to the audience.

"We stopped countless terror attacks in Europe. And frankly, if those things hadn't been stopped, you'd be looking at a very different political reality in the West right now. It's the truth." He paused briefly, then added: "I love that I get yelled at in European cities. Keep yelling. The only reason nobody in between me and you is doing the goose step is our products. Say thank you."

The goose step. It was a reference to the marching style of 20th-century German and Italian fascists. Karp was claiming that Palantir had saved Europe from a revival of fascism.

This remark caused a major stir in Germany. It starkly revealed the barriers Palantir's European expansion strategy faced. Ironically, the very history Palantir claimed to have prevented became the source of resistance against Palantir's adoption.

Distrust of state surveillance runs deep in Germany. The experiences of the Nazi Gestapo and the East German Stasi remain as collective trauma. This history was what GFF lawyers repeatedly invoked in their constitutional complaints.

The European Union's General Data Protection Regulation (GDPR) was born out of this historical context. Enacted in 2018, GDPR is regarded as the world's strictest personal data protection framework. It has several core principles. Data minimization: only the minimum data necessary to achieve the purpose may be collected. Purpose limitation: data cannot be used for purposes beyond the original collection purpose. Right to be forgotten: individuals can request the deletion of their data.

Palantir's business model conflicts with these principles. The essence of Palantir software is to integrate and connect data from diverse sources to discover new patterns. It does not minimize data; it maximizes it. It uses data for purposes beyond the original collection intent.

The European Court of Justice's 2020 Schrems II ruling made the situation even more complicated. The court determined that U.S. surveillance laws, particularly Section 702 of the Foreign Intelligence Surveillance Act (FISA), did not adequately protect EU citizens' personal data. The EU-U.S. Privacy Shield agreement was invalidated. Transferring EU citizens' data to U.S. companies itself became a legal problem.

(2) The Polizei 20/20 Project and the Merz Government's Re-Push

Palantir's European strategy has not experienced only setbacks. Attempts to turn crisis into opportunity are underway.

The structural characteristics of Germany's police system provide the backdrop. Germany is a federal state composed of 16 Länder (states). Police forces belong to each state government. At the federal level, there is the Federal Criminal Police Office (BKA) and the Federal Police, but day-to-day policing is handled by state police. The problem is that each state police force's IT systems are different. Information collected in one state is not shared with others. Federal-level databases and state-level databases are not connected.

In 2019, the federal and state governments launched a project called "Polizei 20/20" to solve this problem. A budget of 300 million euros was allocated. The goal was to integrate the decentralized police databases of the 16 states.

Conservative politicians advocated introducing Palantir software to this project. Friedrich Merz and Alexander Dobrindt of the CDU/CSU were prominent voices. They emphasized that Palantir had

already been proven in intelligence agencies and militaries around the world.

In 2023, then-Interior Minister Nancy Faeser of the Social Democratic Party blocked this plan. She prevented the Federal Criminal Police Office, the Federal Police, and the Customs Investigation Office from using Palantir software. The decision came shortly after the Federal Constitutional Court ruling.

But the situation changed when the CDU/CSU won the early 2025 general election. A new coalition government led by Chancellor Merz was formed. The Interior Ministry portfolio went to the CSU's Dobrindt, a longtime supporter of Palantir adoption.

In March 2025, the Bundesrat (upper house) passed a resolution calling for the rapid introduction of a joint "automated data analysis platform" for federal and state police. The resolution did not explicitly mention Palantir. However, under the framework agreement Bavaria had already signed with Palantir, other states and the federal government could join without a new procurement process.

Not all states agreed. According to an investigation by Bayerischer Rundfunk (BR), Hamburg and Mecklenburg-Western Pomerania explicitly demanded a European alternative that would "exclude the use of products from the market-dominant American supplier Palantir." Bremen, Lower Saxony, Saarland, Schleswig-Holstein, and Thuringia supported this position. But the motion failed to secure the required majority in the Bundesrat.

(3) The Dilemma of Digital Sovereignty

One of the reasons several state governments oppose Palantir is the issue of "digital sovereignty."

Peter Thiel, Palantir's co-founder and chairman of the board, is a prominent supporter and political donor of President Donald Trump and Vice President JD Vance. At least ten officials in the current Trump administration are reportedly holders of Palantir stock. In April 2025, Palantir signed a $30 million contract with U.S. Immigration and Customs Enforcement (ICE). The contract was to develop an operating system called "ImmigrationOS" for identifying undocumented immigrants and tracking voluntary departures.

Social Democrat parliamentarian Johannes Schätzl told German newsweekly Stern: "Palantir is not a neutral IT service provider. It is a company with deep connections to U.S. intelligence agencies that pursues clear geopolitical interests."

There are also legal problems. The U.S. CLOUD Act, enacted in 2018, allows U.S. authorities to compel U.S. companies to produce data stored anywhere in the world. In theory, even if Palantir software processes only German citizens' data on German police servers, the possibility that the U.S. government could demand access cannot be entirely ruled out.

Palantir dismisses these concerns. Jan Hiesserich, Senior Vice President of European Strategy, told German business newspaper Handelsblatt: "Palantir brings software, not data. Which data is relevant to an investigation is determined solely by the customer in accordance with the relevant regulations."

Palantir works with Deutsche Telekom subsidiary T-Systems in Germany. The software runs on police-owned servers. The official position is that no data is transferred to the United States.

But critics argue this is not enough. Green Party federal parliamentarian Konstantin von Notz said: "A nationwide rollout should be off the table, especially in an era where the U.S. government is becoming increasingly unreliable."

The new coalition's agreement contains an interesting phrase: "Digital policy is power politics." A goal of reducing dependence on foreign technology is also stated. But at the same time, they are re-pushing Palantir adoption. It is a contradiction to emphasize digital sovereignty while introducing American surveillance software.

C. The Crisis of Democratic Control

(1) When an Unelected Technology Company Takes Over Core State Functions

At the end of July 2025, the Baden-Württemberg state parliament passed legislation authorizing the use of Palantir Gotham software. It was a Green-CDU coalition. Even the environmental party, the Greens, did not oppose it.

According to reports, Baden-Württemberg had already signed a 25-million-euro contract with Palantir in March 2025. This was before the legal basis had been established. The order was reversed. Contract first, law second. The inertia of technology adoption outpaced democratic deliberation.

This is the "vendor lock-in" phenomenon critics warn about. Once the Palantir system becomes deeply integrated into police workflows, switching to another system becomes increasingly difficult. Specialized training, data migration, and workflow redesign require enormous costs. The promise to switch to a future European alternative becomes realistically unachievable.

The case of Bavaria's VeRA system illustrates this. The framework agreement Bavaria signed with Palantir in 2022 includes a provision allowing other federal states and the federal government to join without a new procurement process. A decision made by an individual state becomes a nationwide precedent. De facto standards are created without sufficient deliberation by elected representatives.

A more fundamental problem exists. Nobody knows exactly how Palantir's algorithms work. The source code is a trade secret. Even the Fraunhofer Institute's review results were classified as confidential. When police authorities accept the algorithm's recommendations, there is no way to verify the logic behind those recommendations.

Researchers call this "automation bias." Humans tend to over-trust computer system outputs. When an algorithm classifies a particular individual as "high risk," investigators are more likely to follow that judgment than question it. When the algorithm is wrong, it is difficult to detect and correct that error.

(2) Questions of Sovereignty and Accountability: Who Watches the Algorithms

Peter Thiel wrote a provocative sentence in a 2009 essay: "I no longer believe that freedom and democracy are compatible."

Alex Karp has said at multiple public events: "Our products are sometimes used to kill people."

Palantir's founders do not hide their values. They claim to be defending Western democracy. But who decides the manner and scope of that "defense"? A democratically elected government, or Palantir?

Palantir's governance structure makes this question even sharper. The company employs a differential voting rights structure. The three founders (Thiel, Karp, and Stephen Cohen) exercise approximately 80% of voting rights through Class F shares. Control by shareholders and oversight by the board of directors are both limited. The founders determine the direction of the company.

This company handles core national security functions of democratic states. Palantir's algorithms determine who is a terrorist suspect, who is a surveillance target, and how data is connected. An unelected technology company has entered a core domain of national sovereignty.

The German Federal Constitutional Court's 2023 ruling was one answer to this problem. But it was not a definitive answer. The court did not ban the use of Palantir itself. It merely demanded that the legal basis be made stricter.

The GFF and the Chaos Computer Club continue to fight. Constitutional complaints against Bavaria and North Rhine-Westphalia are underway. But the political tide is moving in the opposite direction. Under the Merz government, Palantir's proliferation is accelerating.

No fundamental solution is in sight at the EU level either. GDPR proclaimed strong data protection principles, but grants each member state discretion regarding the use of surveillance technology in the national security domain. The EU AI Act introduced regulations for high-risk AI systems, but many exceptions apply to the law enforcement domain.

In the end, the problem is not technology. It is a choice. What balance to strike between efficiency and privacy, between security and freedom, between convenience and democracy. The answer to this question cannot be provided by an algorithm. Citizens must decide for themselves.

The 1983 Census Judgment that the German Federal Constitutional Court recalled, and the history of the Nazis and the Stasi that it evoked, serve as a warning. Technology is neither inherently good nor evil. But when it becomes a tool of power, when it goes unchecked, when it lacks transparency, technology can threaten democracy.

Palantir's name comes from the "palantír" in Tolkien's The Lord of the Rings. The seeing stone. But in Tolkien's story, many of those who used the palantír were captivated by its gaze and met their doom. The watcher became the watched. The tool dominated the user.

The question facing Germany and Europe is this: Who will control the palantír? Will we use it, or will it use us?

← Previous Chapter

See every chapter

Next Chapter →

Kim Kyung-jin

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com



#KimKyungjin #PALANTIR #Palantir #WarAI #SurveillanceAI #DefenseAI #DataPower #ArtificialIntelligence #AILibrary

# Chapter 11. The Black Box of Power: Internal Culture and Governance

← Previous Chapter

See every chapter

Next Chapter →

Part 5: Controversy and Regulation, The Shadow of the Surveillance State

A. "Our Technology Sometimes Kills People"

(1) CEO Alex Karp's Statements and Company Culture

One afternoon in May 2020, CNN Business cameras were focused on Alex Karp. This CEO, who holds a doctorate in philosophy, sat deep in a chair wearing a worn sweater. When the reporter asked about the essential nature of Palantir's business, Karp paused for a moment. Then he delivered a single sentence: "Our products are sometimes used to kill people."

It was a statement no other Silicon Valley CEO would ever have made. Google executives say they "organize information." Facebook says it "connects the world." Amazon says it "serves customers." Technology companies know how to wrap their business in abstract, positive language. Karp broke that rule.

After the remark went public, Karp did not back down. He repeated it. In shareholder letters, investor meetings, and media interviews, he echoed the same message. "We make the enemies of the West afraid. Sometimes we eliminate them." In a 2024 New York Times interview, Karp went a step further. "The West is likely to fight on three fronts simultaneously against China, Russia, and Iran. To prepare for that war, we must develop autonomous weapons." A philosopher was talking about weapons.

This manner of speaking is no accident. It is strategy. He refuses to package Palantir as a "neutral technology platform." Instead, he declares, "We have chosen a side." That side is "Western democracy," "a free way of life," and "peace through strength." In a May 2025 earnings call, Karp said, "We are a warrior culture." Warrior culture. The message that word sends to the company's 4,000 employees is clear: you are not software engineers; you are soldiers on the front line.

This culture begins with the hiring process. Palantir operates a unique job category called "Forward Deployed Engineer." These individuals are not posted to air-conditioned offices in Palo Alto but are dispatched to client sites. It could be a military operations center, a police command post, or an intelligence agency's analysis room. There, engineers work alongside soldiers, analysts, and investigators. They write code while seeing with their own eyes what that code actually does.

Karp believes this experience transforms employees. In one interview, he mentioned personally visiting new engineers to talk about "philosophical responsibility." Not about coding techniques or

business objectives, but about "the ethical implications of building tools used by governments." He asked: "Just because our software can do something, should it?" The question was posed, but the company's answer was already decided. It should.

Palantir's internal culture has a distinctive vocabulary. Employees call themselves "Palantirians." Those who leave the company become "veterans." Client demands are reinterpreted as "battlefield needs," and overtime becomes "mission execution." This language transforms work. Designing databases becomes "defending Western civilization." Optimizing algorithms becomes "tracking terrorists."

In November 2025, according to Anadolu Agency reporting, Karp confronted protesters at a Palantir demonstration site, telling them, "Our technology primarily kills terrorists." The word "primarily" was key. "Primarily" does not erase the remainder. The remainder stays in the shadows. Misjudgments, mistaken bombings, excessive enforcement, bad data, biased models. They are pushed into the statistical shadows.

Karp's worldview emerged from his academic background. He earned his doctorate at the University of Frankfurt under the supervision of Jürgen Habermas. His dissertation topic was the cultural origins of aggression. The idea that language shapes social structures and sometimes injects poison. This philosophy became Palantir's DNA. Data is not neutral. Technology is not neutral. Who uses it, how, and why determines everything.

But there is a gap in that logic. Karp says "technology is not neutral," while simultaneously saying "we only provide the tools and the final decisions are made by humans." The first sentence acknowledges responsibility. The second evades it. When both sentences come from the same person's mouth, the audience does not know which sentence to believe. Perhaps both are true. That is precisely the essence of the company called Palantir.

(2) The Coexistence of Cult-like Loyalty and Ethical Conflict

In the summer of 2019, something strange was happening at Palantir's Palo Alto headquarters. Letters were circulating among employees. There were two kinds. One letter criticized the company's contract with Immigration and Customs Enforcement (ICE). The other supported it. In the same hallway, next to the same coffee machine, employees signed different letters. According to Business Insider reporting, more than 200 employees conveyed their concerns to the CEO.

This incident exposed a rift inside Palantir. On the surface, the company appears perfectly united. "We have chosen a side." "We defend the West." "We are a warrior culture." These slogans hang on walls, echo in conference rooms, and are embedded in email signatures. But beneath them, some employees were quietly asking: "Are we really doing the right thing?"

The ICE contract was a particularly hot potato. During the Trump administration, ICE used Palantir's software to track undocumented immigrants and conduct family separation operations. Photos of wailing children appeared in the news. For some Palantir employees, those photos were not just news. They were the consequences of code they had written.

Karp's response was firm. He acknowledged family separation as "a really hard, complex, and shocking moral issue." But at the same time, he said he supported "a fair but rigorous immigration policy." And he presented a core argument: "Does a technology company have the right to censor the policies of an elected government?" His answer was clear. No. Supporting operations lawfully conducted by the government is a corporate duty.

This logic persuaded some employees. Others resigned. The Washington Post reported that departures increased during this period alongside internal disputes. Karp did not try to retain those who left. He opened the door for them. "If you don't like what our company does, leave." His message was blunt.

Through this process, Palantir was refined. Those who harbored ethical doubts left, and only those who agreed with the company's mission remained. The result was greater internal cohesion, but also greater homogeneity. The space for critical thinking narrowed. Karp himself once jokingly described Palantir's culture as "like a cult." It may not have been a joke.

Cults have several characteristics. There is a powerful leader. There is disconnection from the outside world. There is a shared narrative. Those who question that narrative are excluded. Palantir possesses all these characteristics to some degree. Karp is a charismatic leader. The company avoided media and public scrutiny for 17 years. The grand narrative of "defending Western democracy" binds all employees together.

In December 2025, at a DealBook event, Karp delivered another artful sentence: "We are a very ethical company. But don't believe it." The sentence is paradoxical. He claims to be ethical while telling you not to believe the claim. It appears humble, but in truth, it shifts the burden of verification to the audience. Saying "don't believe it" actually means "believe it." It is a rhetorical device that does not persuade critics but reinforces the solidarity of those who already believe.

When the Gaza war broke out in 2024, Palantir declared "full-throated support" for Israel. Karp called pro-Palestinian campus protests a "pagan religion" and "an infection within our society." These remarks would have made some employees uncomfortable. But within Palantir today, there is little space to express that discomfort.

In 2025, Time magazine named Karp one of the world's 100 most influential people. His net worth exceeded $18 billion. He meditates at a farmhouse in New Hampshire, practices tai chi, and goes cross-country skiing with bodyguards who are former Norwegian special forces members. He owns ten houses but jokes that they are "ten cross-country ski cabins." The philosopher became a billionaire, and the billionaire philosophizes about war.

Palantir's internal culture ultimately comes down to a single question: "Do you agree with this mission?" If you agree, you stay. If you disagree, you leave. There is no middle ground. This is how cult-like loyalty and ethical conflict coexist. The conflict does not disappear. It is simply pushed outside.

B. Governance Issues

(1) Class B/F Shares and the Three Founders' 80% Voting Control

On September 30, 2020, an unusual listing took place on the New York Stock Exchange. Palantir chose a direct listing instead of a traditional initial public offering (IPO). It was a method where existing shareholders directly placed their shares on the market without investment bank intermediation. Wall Street bankers missed out on fees, and Palantir saved tens of millions of dollars. But the truly unusual thing lay elsewhere.

When TechCrunch analyzed the S-1 filing leaked before the listing, editor Danny Crichton said: "Wow, this is a really complex ownership structure." Palantir created three classes of stock: Class A, Class B, and Class F.

Class A shares available to ordinary investors carry one vote per share. Class B shares held by insiders carry ten votes per share. Up to this point, it is a dual-class structure common in Silicon Valley. Facebook, Google, and Snap all use similar methods.

The problem is Class F. F stands for Founder. Peter Thiel, Alex Karp, and Stephen Cohen. A special class of stock created for these three people alone. Class F shares perform an astonishing bit of magic. Under certain conditions, the voting rights of these shares automatically adjust to constitute 49.999999% of total voting power. The reason for the precision to six decimal places is that exceeding 50% would classify them as "controlling shareholders," subjecting them to additional regulations. Palantir sits right at the regulatory boundary.

Consider what this structure means. Even if the three founders hold only 2% of the company's shares, they can still exercise nearly 50% of voting rights. Add the 10x voting power of their Class B holdings, and effective control can reach 70 to 80%. According to a 2025 AInvest analysis, the founders maintain 49.99% voting control, influencing major decisions from board elections to mergers and acquisitions.

In a typical company, when executives sell shares, they lose power. Not at Palantir. No matter how many shares a founder sells, the magic of Class F ensures management control is preserved. Economic stake and control are separated. TechCrunch's Crichton described this as a structure where "in the extreme case, founders could hold just 2-3% of the company's shares while controlling nearly 70% of the voting power."

Why did Palantir create such a structure? The official explanation goes like this: "To enable the founders to pursue the company's long-term vision without capitulating to short-term market pressures." Karp puts it more bluntly: "We need to be able to take on unpopular projects. We cannot be swayed by Wall Street's quarterly earnings pressure." Government contracts and defense software require long time horizons. Five years, ten years, sometimes twenty. The logic is that such long-term strategy must not be whipsawed by the demands of hedge fund managers obsessed with short-term stock prices.

There is another reason as well. Palantir handles the most sensitive data of U.S. intelligence agencies and the Department of Defense. A hostile state or force purchasing a stake and interfering

with management could pose a national security risk. In this context, founder control functions as a kind of safety mechanism.

But critics see it differently. Richard Windsor, founder of Radio Free Mobile, gave this structure an "F grade." "Until the founders die, the remaining founders will have absolute control of this company, and no other shareholders will have any say whatsoever." He argued that at least a 30% discount should be applied to Palantir stock as a risk premium for terrible governance.

In practice, some institutional investors shy away from Palantir because of this structure. In its 2024 10-K report, Palantir itself included a warning to the effect that "governance may negatively influence institutional investors' purchase decisions." The company is aware of this fact. It chose this path anyway.

In 2022, shareholders even filed a lawsuit over these "magic shares." According to Law360 reporting, Palantir paid a settlement including $5.5 million in attorney fees and closed the case. But the governance structure itself remained unchanged.

The essence of the problem is this. A publicly listed company is, in principle, owned by its shareholders. Shareholders have the authority to elect, supervise, and if necessary, dismiss management. At Palantir, this principle does not function. No matter how many ordinary shareholders come together, they cannot overcome the voting rights of the three founders. It is formally a public company, but in substance, it is the founders' private kingdom. This is what "the black box of power" means.

(2) Excessive Stock-Based Compensation and Frequent Executive Selling

In February 2025, Bloomberg reported a shocking number. Palantir insiders had cashed out more than $4 billion in stock during 2024. CEO Alex Karp sold approximately $2 billion, and co-founder Peter Thiel sold one-third of his stake for $1.5 billion. And they did not stop in 2025. Karp announced an additional stock sale plan worth roughly $1 billion.

The numbers alone look simple. But what those numbers signify is complex.

Like a typical software company, Palantir compensates employees with stock rather than cash. This is called Stock-Based Compensation (SBC). It is a common practice at startups. When cash is tight, promising future wealth is necessary to attract talented people. "You earn less now, but if the company succeeds, you become wealthy too." This is the logic of stock options.

The problem is that Palantir is no longer a startup. Revenue in 2024 exceeded $2.9 billion. Market capitalization once surpassed $300 billion. Cash flow is healthy. Yet the scale of stock compensation remains enormous. According to the 2024 10-K filing, a significant portion of research and development (R&D) and selling, general, and administrative (SG&A) costs are composed of SBC. These costs dilute existing shareholders' stakes. Each time new shares are issued, each slice of the pie grows smaller.

What is even more contentious is the pattern of executive selling. According to Sherwood News reporting, searching FactSet data reveals virtually no record of Palantir insiders purchasing their own company's stock on the open market. Records of selling, on the other hand, overflow.

Over the course of 2024, Karp sold 40.7 million shares at an average price of $47.99. That is $2 billion. Co-founder Stephen Cohen sold 3 million shares for $152 million. CTO Shyam Sankar also cashed out more than $380 million. In February and March 2025, Karp executed an additional $45 million in sales, and Sankar sold $38 million more.

All of these sales were made through 10b5-1 plans. The 10b5-1 is a mechanism designed to avoid insider trading allegations. When executives set up a sale plan in advance, it executes automatically. Because the plan is established when executives do not possess material non-public information, it is legal.

But investors see a different scene. According to Fortune reporting, $1.4 billion of Karp's 2024 stock sales were concentrated in the weeks surrounding the presidential election. He sold when the stock price surged. Legal, but bitter from the perspective of ordinary investors.

Jefferies analyst Brent Thill pointed out in a March 2025 report: "Karp has sold 21% of his total Palantir stake." He added: "There are many reasons insiders sell stock. Tuition, buying a house, paying off debt, portfolio diversification. But there is only one reason insiders buy stock. They think the price will go up."

Palantir's governance structure makes this problem even more complex. In a typical company, when executives sell in large quantities, two things happen. First, the stock price drops. Second, executives' control weakens. At Palantir, the second does not happen. Thanks to the magic of Class F shares.

This is why investors feel uneasy. Executives cash out by selling shares while maintaining management control. The company's economic performance and executive control are decoupled. Imagine the worst-case scenario. Executives make a bad decision. The stock price crashes. Ordinary shareholders suffer losses. But executives have already sold at high prices and still control the company. This is a "structure that keeps the power but sells off the risk."

According to 2025 Forbes analysis, Karp's 2024 compensation reached $6.8 billion. Most of it was stock option gains from the rising share price. This scale ranks among the very top of U.S. corporate CEO compensation. Of course, this is the story when the company is doing well. Palantir's stock rose 340% in 2024. But in 2022, it had also crashed more than 70%.

A Motley Fool analyst summarized it this way: "Palantir is clearly a great company. But at the current valuation, it may be wise to follow management's lead and take some profits." Management is selling, so why are you holding? That is the question.

C. Diversity and Discrimination Controversies

(1) Asian-American Hiring Discrimination Lawsuit and $1.7 Million Settlement

On certain days in 2010 and 2011, interviews were being conducted at Palantir's Palo Alto office. They were hiring software engineers. Applicants demonstrated their coding abilities, presented their academic credentials, and explained their project experience. Then they went home and waited for the results. Many received rejection notices. What happened during the interviews is not on record. But who was hired and who was rejected was recorded.

In 2016, the Office of Federal Contract Compliance Programs (OFCCP) under the U.S. Department of Labor analyzed those records. And it filed a lawsuit. The allegation was that Palantir had systematically discriminated against Asian applicants.

The numbers tell the story. There were more than 1,160 qualified applicants for software engineer positions. About 85% of them were Asian. But the final hires were 14 non-Asians and 11 Asians. For QA engineer positions, the disparity was even more dramatic. Of more than 730 qualified applicants, about 77% were Asian, but of the 7 people hired, only 1 was Asian. In intern hiring too, Asian applicants were the majority, but non-Asians were hired four times more often than Asians.

The Labor Department statisticians did the math. What is the probability that these results occurred by chance? For the QA engineer hiring, the probability was 1 in 741. This meant statistically significant discrimination.

Asian applicants were "systematically eliminated" at the resume screening stage and the phone interview stage, the Labor Department alleged. Even though they had qualifications equal to or better than white applicants. The Labor Department also suspected that Palantir's internal referral system was designed to exclude Asians.

Palantir denied the charges vehemently. The company submitted a 15-page response. It pointed out that the Labor Department had analyzed only 3 of 44 positions. It countered that 36% of all hires during the same period were Asian, which was higher than the share of Asians in the external labor market. It also criticized the Labor Department's statistical analysis for "ignoring applicant qualifications and looking only at racial composition."

But the dispute did not last long. In April 2017, Palantir paid $1.7 million and settled. Specifically, the terms were $1,659,434 in back wages and stock option value, plus re-employment opportunities for eight victims.

Palantir did not admit wrongdoing. The company's statement read: "We disagree with the allegations made by the Department of Labor. We settled this matter, without any admission of liability, in order to focus on our business. We continue to stand behind our hiring record."

Why did they settle? The reason is simple. Palantir is heavily reliant on federal government contracts. Since 2010, government contracts alone had totaled more than $340 million. If they had lost the Labor Department suit, they could have lost their eligibility for federal contracts. The $1.7 million was a rational cost to avoid that risk.

This case conflicts with Palantir's brand. Palantir had positioned itself as an embodiment of "meritocracy." "We don't look at where you come from. We only look at ability." But the Labor Department's data told a different story. When ability was equal, race determined the outcome.

There is something even more ironic. Palantir is a data analytics company. Finding biases in the world through data is its business. Yet it failed to find bias in its own hiring data. Or rather, it chose not to find it. The problem came to light only when an external audit was conducted.

It was a period when all of Silicon Valley was being criticized for diversity issues. Google, Facebook, and Twitter all faced criticism for being "too white and male-centered." But Palantir's case was unusual. Asians form the largest technical talent pool in Silicon Valley. Indian, Chinese, and Korean engineers are present in large numbers. Yet Palantir actually excluded from this pool.

Why? Only speculation is possible. The vague criterion of "cultural fit" may have been at work. Palantir has a distinctive culture. "Defending Western democracy," "warrior culture," "choosing a side." When searching for people who fit this identity, one may unconsciously favor certain backgrounds.

Palantir did not abandon its "meritocracy" rhetoric after this incident. It intensified it. The "Meritocracy Fellowship" launched in 2025 is evidence of this. But the 2017 settlement remains on the record. A record of a company that preaches meritocracy but excluded people regardless of merit.

(2) The 'Meritocracy Fellowship' and the Meaning of Expanding Non-College Hiring

In April 2025, Palantir plastered posters across elite university campuses nationwide. "Instead of college. Instead of debt. Instead of brainwashing. Earn the Palantir degree." It was a provocative message. A few weeks later, the company officially announced the "Meritocracy Fellowship."

This program targets high school graduates or soon-to-be graduates. Only those not enrolled in college can apply. Fellows work at the New York office for four months, receiving a salary of $5,400 per month. Qualifications? SAT score of 1460 or above, or ACT score of 33 or above. These scores fall in the top 1 to 2 percent. The program demands Ivy League-level academic ability while telling applicants not to go to the Ivy League.

According to Fortune reporting, more than 500 applied and 22 were selected. An acceptance rate of 4.4%. Lower than Harvard's acceptance rate. Among those selected, some felt college was unappealing, while others had been rejected from their preferred schools. They spent the fall 2025 semester at Palantir.

Karp connected this program to his critique of education. He has long attacked the American university system. In the Q2 2025 earnings call, he said: "The opaque admissions criteria at many American universities have pushed out meritocracy and excellence." He called campuses "a hotbed of extremism and chaos." The remark targeted the 2024 pro-Palestinian protests.

The program's content differs from a typical internship. Fellows attend seminars on American history and "the foundations of the West." They hear lectures from historians, philosophers, scientists, and writers. There are communication classes as well. And they are deployed to actual projects with Palantir clients. After four months, they receive a "Palantir degree." Those who perform well are offered regular employment interviews.

Karp does not hide the program's intent. "When you come to Palantir, whether you went to which school or didn't go to any, it doesn't matter. The moment you become a Palantirian, nothing else matters. This is far and away the best credential in tech." There is irony in Karp, a Stanford Law graduate, saying college is unnecessary.

The Wall Street Journal described the program as an "anti-credential experiment." Korean outlet Chosun Biz also covered it with interest. On the surface, this represents an expansion of opportunity. It opens a new path for young people who cannot afford college. It saves four years of tuition and living expenses. They start working and building careers right away.

But critical perspectives exist as well. First, Palantir wants to secure talent early and educate them in its own way. College spends four years teaching critical thinking, diverse perspectives, and independent judgment. Palantir skips those four years and injects the company's values first. A young person exposed at age 18 to the narratives of "defending Western democracy" and "warrior culture" may find it difficult to question those narratives.

Second, the definition of "meritocracy" is problematic. Is an SAT score of 1460 the sole measure of ability? Getting that score requires good schools, good tutoring, and a good environment. Students from low-income families have a lower probability of achieving that score than those from wealthy families. While touting "meritocracy," the structure may actually favor a particular class.

Third, there is the connection to the 2017 Asian discrimination settlement. On one hand, they say "we only look at ability," while on the other, government investigation revealed discrimination. The Meritocracy Fellowship may serve as an image strategy to cover that blemish. Diluting the past record with the message "We're an open company that even hires high school graduates."

The "meritocracy" Karp espouses rejects diversity policies (DEI) based on race or gender. He advocates a "color-blind meritocracy." Looking purely at ability without regard to background. But critics point out that this approach ignores structural inequality. Applying the same standards to people with different starting lines produces unequal results.

Palantir's Meritocracy Fellowship has another dimension. Peter Thiel has operated the "Thiel Fellowship" since 2010. Each year, he gives $100,000 to 20 to 30 young people under 23 and has them drop out of college. Thiel has long called college a "bubble." Palantir's program is an extension of this philosophy. An anti-establishment streak shared by the founders is spreading throughout the company.

In November 2025, the first 22 fellows completed the program. How many received full-time offers was not disclosed. But Palantir has already begun recruiting the 2026 cohort. The program continues.

In the diversity debate, Palantir occupies a unique position. It rejects traditional DEI while emphasizing "diversity of thought." What matters is not racial or gender diversity, but diversity of perspectives and beliefs. But critics ask: in an organization where only those who agree with the "warrior culture" remain, does diversity of thought truly exist?

Palantir's diversity controversies always follow the same pattern. The company says "we hire on merit." Critics ask "does the word merit mask discrimination?" The 2017 settlement put that question into a legal document. The 2025 fellowship elevated that question into the language of the culture wars. Palantir does not avoid controversy. It uses controversy as fuel.

← Previous Chapter

See every chapter

Next Chapter →

Kim Kyung-jin

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com



#KimKyungjin #PALANTIR #Palantir #WarAI #SurveillanceAI #DefenseAI #DataPower #ArtificialIntelligence #AILibrary

# Chapter 12. The U.S.-China AI Hegemony Competition and Geopolitical Risks

← Previous Chapter

See every chapter

Next Chapter →

Part 6: The Future and Implications, Pax Technica

A. The West's Response to China's Technological Rise

(1) Civil-Military Fusion and China's Palantir: Deepexi

In the fall of 2024, a research team at Georgetown University in Washington, D.C. was poring over a massive Excel file. On the screen were 2,857 contracts. They were AI-related procurement contracts signed by the Chinese People's Liberation Army from January 2023 to December 2024. Researchers at the Center for Security and Emerging Technology (CSET) analyzed this data and discovered a pattern. Of the 1,560 contracting companies, a significant number were private firms. Three-quarters of them were not state-owned enterprises. What was even more striking was that two-thirds of these companies had been founded after 2010.

What this data conveyed was clear: China's civil-military fusion strategy was actually working.

The term civil-military fusion first appeared in China in the late 1990s. But it transformed from a mere slogan into a tangible structure only under Xi Jinping. In 2017, Xi established the Central Commission for Integrated Military and Civilian Development. The Chinese Communist Party's highest-level body would directly oversee this strategy. That same year, the Ministry of Science and Technology issued a five-year civil-military fusion plan, and the State Council released an artificial intelligence development plan. Three documents were aligned toward a single goal: converting civilian technological innovation into military capability.

There is a reason the West finds this strategy difficult to understand. In the United States and Europe, the civilian and military sectors are separated. Google employees revolt against drone projects, and Microsoft engineers send protest letters over Pentagon contracts. But in China, no such boundary exists. Alibaba's cloud computing capabilities can be directly linked to the PLA's command-and-control systems. Tencent's facial recognition technology is applied to surveillance cameras in the Uyghur region. Hikvision's security monitors are already being supplied to the Chinese military's surveillance and video processing systems.

Former U.S. Assistant Secretary of State Christopher Ashley Ford summarized the situation this way: "Every technology that is accessible in China, and that Beijing deems useful for military and national security purposes, will be used for those purposes. Come what may."

In this context, the name Deepexi emerges. This company revealed itself to the world when it pursued a listing on the Hong Kong stock exchange in October 2025. According to its prospectus, Deepexi defines itself as a "provider of enterprise-scale large AI application solutions." It explains that it bundles data integration, real-time analytics, and predictive modeling into a single package to support corporate decision-making. It is remarkably similar to what Palantir started doing in the United States twenty years ago.

This is why Western intelligence analysts call Deepexi "China's Palantir." Both companies integrate data, visualize it, and use AI to accelerate decision-making. But there is a critical difference. Palantir grew under the demanding constraints of Western legal restrictions and privacy protections. Deepexi faces no such constraints. It enjoys the Chinese government's full support and unlimited data access. Under the civil-military fusion system, data collected by civilian companies can be immediately repurposed for national security.

Georgetown's CSET analysis shows how effective this strategy actually is. Among the top 15 companies in AI-related military procurement contracts, eleven were still state-owned enterprises or national defense research institutions. Names like China Electronics Technology Group Corporation (CETC), China Aerospace Science and Technology Corporation (CASC), and NORINCO. But the rest were private companies. iFlytek, a leader in voice recognition and natural language processing, won 20 military contracts over two years. They included translation software and AI-based decision-support systems. PIESAT, a satellite and geospatial analytics company, supplied combat simulation platforms and automatic target recognition technology.

This is why the West is on edge. China's defense industry ecosystem is expanding. Young, agile private companies are entering domains that were once monopolized by massive state-owned enterprises. These companies are supplying cutting-edge AI technology to the military, rapidly boosting the PLA's technological capabilities.

Under Xi Jinping's direction, China is now seeking to make civil-military fusion a core pillar of the 15th Five-Year Plan (2026-2030). This initiative, called "AI Plus," aims to deeply integrate AI across the economy and society. In March 2024, Premier Li Qiang declared at the National People's Congress that he would "develop new quality productive forces at a faster pace." This signifies the completion of a system where civilian AI advances automatically translate into military strengthening.

The West's response is proceeding on multiple levels.

The first is narrowing the gateway for semiconductors and computing power. The U.S. Bureau of Industry and Security (BIS) tightened advanced semiconductor export controls in December 2024. In January 2025, a global "AI Diffusion" rule was announced, an attempt to block circumvention through third countries.

The second is the realignment of alliances. AUKUS Pillar 2 institutionalizes joint development among the United States, the United Kingdom, and Australia in AI, autonomous systems, cyber, electronic warfare, and quantum technology. This is not a simple arms purchasing alliance. It is a technology

alliance that shares software and data.

The third is the classification and sanctioning of Chinese companies. The U.S. Department of Defense maintains a "Chinese Military Company" (CMC) list under the FY2021 National Defense Authorization Act. Section 805 of the FY2024 NDAA goes a step further. Starting June 30, 2026, the Department of Defense will be unable to contract not only with companies on this list but also with any other companies that supply goods or services produced or developed by those listed companies.

But export controls and sanctions alone are not sufficient. In October 2025, China expanded export controls on rare earth elements including holmium, erbium, thulium, europium, and ytterbium. These elements are essential for F-35 fighter jets, Tomahawk missiles, radar systems, and AI semiconductors. China dominates 60% of global rare earth mining and 90% of processing. It is a strategy of converting economic use into a tool of defense power.

Palantir's Alex Karp says of this situation: "China and we are in an AI arms race. We are already in a state of war." His CTO Shyam Sankar shares the same view. He analyzes that China is stealing Western intellectual property, economically infiltrating through the Belt and Road Initiative, and pressuring the West with hybrid tactics.

Ultimately, the West's response to China's technological rise goes beyond a simple technology competition. It is a contest between systems. A system in which the state integrates and mobilizes everything versus a system that pursues innovation while preserving democratic values and market principles. Palantir stands at the center of that contest, asking the West: Can you secure technological advantage while still protecting freedom? The answer to that question has not yet been written.

(2) Taiwan Strait Crisis Simulations and AI as Deterrence

On July 31, 2025, retired military officers and security experts gathered in the auditorium of the Center for Strategic and International Studies (CSIS) in Washington, D.C. On the screen appeared the title "Lights Out?" The lights go out. The title was a metaphor for what would happen in the event of a Taiwan blockade. The research team presented results from 23 wargames. The conclusion was that a Chinese maritime blockade of Taiwan could paralyze global trade and trigger a large-scale conflict in the Western Pacific.

Retired Marine Colonel Mark Cancian looked at the audience and said: "This would be a large-scale convoy battle reminiscent of Operation Pedestal in the Mediterranean in 1942, and ONS-5 in the North Atlantic in 1943." It was a warning that the largest naval battle since World War II could occur.

The Taiwan Strait is no longer a distant sea. It is the front line where the U.S.-China AI hegemony competition could physically collide. And on that front line, AI is emerging as a new form of deterrence.

The CSIS wargame series began in 2023. The first report, "The First Battle of the Next War," simulated a Chinese amphibious invasion of Taiwan. Across 24 games, in most scenarios, the combined U.S., Taiwanese, and Japanese forces repelled the landing. But the cost was devastating. Two U.S. aircraft carriers were sunk, and hundreds of fighter jets were shot down. Japanese bases were devastated by missile strikes.

The second report in December 2024 posed an even more uncomfortable question. What is the probability that nuclear weapons would be used? In 15 simulations, nuclear catastrophe occurred in three games. These were cases where Chinese forces, on the verge of defeat, decided to use nuclear weapons to prevent regime collapse.

The researchers' conclusion was bleak: "Complete victory is impossible. The United States must wage a high-intensity conventional war while simultaneously offering the adversary an honorable exit. Otherwise, nuclear catastrophe ensues."

The third report in July 2025 addressed the blockade scenario. Across 26 games, a Chinese maritime blockade could cut off Taiwan's energy supply and paralyze its economy. But if the United States and allies sent convoy escorts, a large-scale naval battle was inevitable. The winner of that battle was unclear.

The lesson these simulations show is singular: traditional military power alone cannot make deterrence work. If nuclear deterrence was based on the terror of "assured destruction," digital deterrence operates on a different principle. It plants the conviction that "your plan has already been detected."

The way AI becomes a deterrent in the Taiwan Strait operates on multiple layers.

The first is early warning. Preparing for a Taiwan invasion would require months of military mobilization by China. Ammunition stockpile increases, troop movements, spikes in communications volume, civilian ships being converted for military use. All of these signs leave data trails. If AI integrates and analyzes satellite imagery, communications intercepts, and open-source intelligence (OSINT), it can detect invasion preparations weeks to months in advance.

Palantir's Gotham system performs exactly this role. It fuses information streaming from thousands of data sources in real time to find patterns. "Fuel output from missile bases in Fujian Province was three times the usual amount last week." "Communications volume from the Eastern Theater Command increased 40% compared to last month." These subtle signals come together to form a picture. The United States can pre-position aircraft carriers or send warnings through diplomatic channels based on this information.

The second is the speed of the battlefield. In modern warfare, the kill chain is becoming ever faster. The time from target detection to strike has been reduced to minutes. In the Ukraine war, Palantir's system shortened this process from hours to minutes. The same principle applies in the Taiwan Strait. The moment Chinese landing craft approach the shore, AI analyzes satellite data, transmits coordinates, and Taiwanese missiles are launched.

This is why the U.S. Department of Defense increased the Maven Smart System contract ceiling with Palantir to $1.3 billion in May 2025. This system is deployed across five combatant commands, supporting operations in real time. Pacific Command, European Command, Central Command, Northern Command, and Transportation Command. In a Taiwan crisis, the Pacific Command commander would grasp the battlefield situation in near-real-time through this system.

The third is masses of autonomous weapons. The U.S. Department of Defense's "Replicator" initiative aims to deploy thousands of low-cost autonomous drones to counter China's numerical superiority. But humans cannot individually control thousands of drones. AI must become the brain of those drones. Palantir's AIP (Artificial Intelligence Platform) applies large language models (LLMs) to military operations, enabling command and control of autonomous weapons systems.

According to the Department of Defense annual report, China also invested in AI for various military applications in 2024, including unmanned systems, intelligence, surveillance, and reconnaissance (ISR) analysis, decision support, cyber operations, and information warfare. Civil-military fusion is broadening the supply lines for that investment. Both sides are strengthening their AI capabilities.

Here a paradox emerges.

The way AI strengthens deterrence simultaneously accelerates the speed of crises. When both sides gain "faster decision-making," both sides reduce the other's reaction time.

This is precisely the risk participants at the 2025 UN Military AI Dialogue pointed out: "AI tools increase the speed of interpreting complex information, but rapid decision-making cycles reduce human deliberation time and weaken commanders' comprehension. This raises the possibility of errors, escalation, and miscalculation in stressful crisis situations."

In August 2025, 25 international experts gathered at Syracuse University to conduct an unusual wargame. They designed not America's defense, but China's attack. It was an attempt to enter Beijing's strategic mindset. The results were uncomfortable. The insight was that the basic assumptions of U.S. military planning , large-scale amphibious operations and preemptive strikes on U.S. bases , might not align with Beijing's actual calculations. Some participants argued China had to act within ten years; others predicted that Xi Jinping, at 72, would seek to achieve unification before his death.

Taiwan's government is also preparing. In December 2024, Taiwan's Presidential Office conducted its first-ever tabletop wargame. Vice President Hsiao Bi-khim and National Security Council Secretary-General Wu Zhao-xie led the exercise, which tested the government's response capabilities in a scenario of rising military tensions with China. Civilian movements are underway as well. Taiwan's Mizo Games released a board game called "2045." Players become military commanders, secret agents, and civilian resistance fighters battling a hypothetical Chinese invasion. It raised NT$4 million (about $120,000) through crowdfunding.

AI as deterrence ultimately delivers a single message: "Your attack will not succeed." The moment China decides to invade Taiwan, it shows data proving that decision is wrong. What percentage of

landing craft will be sunk before reaching the shore, how long it will take to achieve air superiority, how long supply lines can hold. AI can push all these calculations in the adversary's face.

But this is a double-edged sword. A Department of Defense report warns that China may use generative AI for information operations. In May and October 2024, China timed military exercises around Taiwan to combine official accounts with proxy accounts impersonating Taiwanese citizens, exaggerating PLA capabilities and spreading disinformation about U.S. willingness to defend Taiwan. Evidence that AI can serve not only as a tool of deterrence but also of deception.

Whether AI in the Taiwan Strait becomes a shield preserving peace or a spark igniting war remains unknown. The only certainty is this: an invisible war of data is already underway beneath those waters.

B. Digital Deterrence Theory

(1) Gaining Superiority in the AI Arms Race Prevents War

"If you want peace, have smarter software than your enemy."

Alex Karp said this at the 2024 Munich Security Conference. He was offering a modern reinterpretation of the Roman strategist Vegetius's maxim. The original was "If you want peace, prepare for war (Si vis pacem, para bellum)." In Karp's version, war preparation is not tanks and missiles but algorithms and data.

This is the core proposition of digital deterrence theory. The logic is that achieving overwhelming superiority in the AI arms race paradoxically prevents war from occurring.

Traditional nuclear deterrence was based on "Mutually Assured Destruction (MAD)." If you attack me, I will destroy you. Because both sides perish, nobody strikes first. This logic has prevented nuclear war for 70 years. But deterrence in the AI era operates on a different principle.

Digital deterrence is based on "information asymmetry." If one side can track all of the opponent's movements in real time, predict their intentions, and instantly execute optimal responses, the opponent will not even attempt an attack. They know the probability of victory approaches zero.

The case supporting this logic is the Ukraine war. Russia possessed thousands of tanks and overwhelming artillery power. But when the war began, that massive armored force was nothing more than large, expensive targets before Ukrainian drones and precision strikes. Palantir's system identified Russian forces' positions in real time and transmitted that information to Ukrainian artillery units. The results are what we witnessed.

The U.S. Department of Defense embraced this lesson. In August 2025, the U.S. Army signed a $10 billion Enterprise Software Agreement (ESA) with Palantir. It consolidated 75 individual software contracts into one. Daniel Moyer, director of the Army Contracting Command, explained: "It's not a bulk purchase, but rather an à la carte menu. You can pick and choose the software you need." This was a declaration that software had become a core national defense asset.

That same May, the Department of Defense increased the Maven Smart System contract ceiling with Palantir to $1.3 billion. According to officials, this was because "combatant commands have increased their use of MSS for commanding and controlling dynamic operations and activities in their respective theaters." Military commanders in the Pacific, Europe, and the Middle East had begun to rely on AI-based systems.

NATO is moving in the same direction. In August 2025, NATO's Joint Warfare Centre (JWC) integrated the Maven Smart System NATO. More than 100 participants joined the exercise. JWC Commander Major General Ruprecht von Butler said: "Our warfighters demand a faster, more agile, and more realistic training environment punctuated by technological innovation." MSS NATO is NATO's first AI-based battle command-and-control system, deployed across Allied Command Operations (ACO).

The logic behind all these moves is the same: lead in AI, and the enemy will not dare to provoke. Former Google Chairman Eric Schmidt has warned that the slope of AI development is so steep that even a six-month lead can create a permanent gap. If the democratic bloc wins the AI development race, it can perpetuate economic and military hegemony. Conversely, if China gets there first, digital totalitarianism will expand worldwide.

But this logic harbors dangerous traps.

First, if AI superiority is opaque to both sides, the adversary may become more dangerous. A 2024 RAND report warned that AI military systems could increase the risk of unintended conflict without human judgment. What AI provides is not prophecy but calculation. If that calculation is wrong, it becomes catastrophe.

Second, as the AI arms race accelerates, decision-making speed exceeds human cognitive capacity. If the time to determine whether an enemy missile has been launched and issue an intercept order shrinks to within seconds, humans may ultimately have to delegate trigger authority to AI. This gives rise to the danger of "Flash War." Just as algorithmic collisions in stock markets cause instantaneous crashes known as "flash crashes," an accidental data error could lead to an unwanted war.

Third, AI technology spreads rapidly. In November 2024, Reuters reported that Chinese military-affiliated researchers had developed military tools based on the open-source Llama model. The open-source ecosystem makes it difficult to monopolize deterrence. When one side locks the technology, the other enters through a different door.

Participants at the 2025 UN Military AI Dialogue pointed out this dilemma: "Arms race dynamics are an additional risk factor. Perceptions of strategic competition and the psychological drive to gain an edge over an adversary create pressure to rapidly deploy capabilities, and this can outpace assurance, testing, and governance processes."

Nevertheless, proponents of digital deterrence argue there is no alternative. If the West sits idle while China strengthens its AI capabilities, the balance of power collapses. When balance collapses, miscalculation occurs. Miscalculation invites war. Therefore, "Responsible Acceleration" is the only

path. They argue that slowing development out of concern over AI risks amounts to "unilateral disarmament," ceding the initiative to adversaries who do not follow regulations.

Alex Karp says: "Our technology is sometimes used to kill people." An uncomfortable truth. But according to his logic, the possibility of that killing paradoxically prevents greater death , war. The cold conclusion of digital deterrence.

(2) The Future of Defense Determined by Software

In March 2024, the U.S. Army awarded Palantir a $178.4 million contract. It was to develop a next-generation targeting system called TITAN (Tactical Intelligence Targeting Access Node). This system connects data collected from high-altitude and space-based sensors to ground forces, supporting beyond-line-of-sight targeting. Bryant Cheung, Palantir's senior vice president of defense solutions, said: "TITAN greatly improves capabilities, but we want software and AI to reduce complexity."

That sentence captures the essence of future defense capability. Hardware provides the ability, but software makes that ability real.

Past arms races were numbers games about who had more tanks and fighter jets. In World War II, the Soviet Union produced 80,000 T-34 tanks to overwhelm Germany's armored forces. During the Cold War, the U.S. and the Soviet Union stockpiled tens of thousands of nuclear warheads to maintain a balance of terror. But on the 21st-century battlefield, that equation has broken down.

The lesson the Ukraine war revealed is clear. Russia's massive armored forces were helpless before low-cost drones and precision-guided weapons. The era has arrived when a $100 million fighter jet can be suppressed by a swarm of $5,000 drones. Palantir CTO Shyam Sankar expressed this by saying, "Hardware has become the accessory of software."

This change is fundamentally reshaping the defense industry's structure. In the past, giant manufacturers like Lockheed Martin, Boeing, and Raytheon were the masters of defense. They built aircraft carriers and fighters over decades, and the government paid the costs. It was a practice called "Cost-Plus" contracting. When the manufacturer submitted costs, the government paid those costs plus a margin. Even if costs exceeded estimates, it did not matter. The government bore the additional expense.

But now engineers who write code have taken center stage in defense. Palantir filed suit against the U.S. Army in 2016 and won. It opened the door for commercial software to participate in defense procurement. Since then, "defense tech" companies like Palantir, Anduril, and SpaceX have begun encroaching on traditional defense companies' territory.

The $10 billion contract in August 2025 was the pinnacle of this shift. The U.S. Army consolidated 75 individual software contracts into one. In 60 of those contracts, Palantir had been a subcontractor, but now it became the prime. Contracting Command Director Daniel Moyer said: "This is our big initiative. It's not just about the Palantir contract. We looked at how many times the same commercial software

was purchased across the Army and whether it makes sense to bundle it all for maximum discounts."

The advantages of software-centric defense are multiple. First, upgrade speed. Upgrading hardware requires taking equipment to a factory and replacing parts. It takes months or years. But software patches can be deployed remotely. Just as Tesla updates its cars, tanks and missiles can become smarter overnight.

Second, interoperability. JADC2 (Joint All-Domain Command and Control) is an initiative to link land, sea, air, space, and cyber into a single network. Palantir's software serves as the operating system of this network. Different weapon systems share the same data and communicate in the same language. Interoperability between allied nations works the same way. This is why NATO adopted the Maven Smart System.

Third, cost efficiency. Traditional defense programs are notorious for cost overruns and schedule delays. The F-35 fighter program has exceeded $1 trillion. But software-based systems can be developed faster and cheaper. Palantir's TITAN prototype was delivered within months.

But as software determines defense capability, new vulnerabilities also emerge: cyberattacks. If code, models, and data are hacked or manipulated, the physical equipment itself is rendered useless. In January 2025, CISA, DARPA, the Office of the Under Secretary of Defense for Research and Engineering (OUSD R&E), and the NSA issued a joint report identifying the "software understanding gap." The situation where AI systems are so complex that humans cannot understand how they operate. If you do not understand it, you cannot find the vulnerabilities either.

Another risk is dependence. Excessive reliance on a specific country's or company's proprietary software and cloud deepens technological and political subordination. This is the issue of digital sovereignty. It is one of the reasons Europe is wary of Palantir.

Nevertheless, the trend cannot be reversed. In December 2025, U.S. Secretary of the Navy John Phelan and Palantir CEO Alex Karp unveiled "ShipOS," a software platform for ship operations. That same year, NATO first operationally tested the Maven Smart System in its largest-ever allied exercise, "Steadfast Dart 2025." Exercise planners said: "Maven will improve integration of cross-domain data sources, improve decision-making capability, and drive digital transformation. We will get better and faster."

The future of defense capability is no longer measured by "who has the bigger ships." It is measured by "who can run a faster OODA loop (Observe-Orient-Decide-Act)" and "who can clear the fog of war with more sophisticated algorithms." In a future where software determines defense capability, code is the weapon, and data is the ammunition.

The era of Pax Technica is approaching. If Rome's peace came from the mobility of Roman roads, and Britain's peace came from naval supremacy, then a technological peace comes from superiority in data and algorithms. Palantir positions itself as the guardian of that peace. Their logic is clear: only if the Western democratic bloc secures AI , the most powerful spear and shield , can it deter provocations from authoritarian states and maintain world order.

Whether this is an uncomfortable truth or a necessary reality is for the reader to judge. One thing is certain: 21st-century power is etched in code on silicon wafers.

← Previous Chapter

See every chapter

Next Chapter →

Kim Kyung-jin

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com



#KimKyungjin #PALANTIR #Palantir #WarAI #SurveillanceAI #DefenseAI #DataPower #ArtificialIntelligence #AILibrary

# Chapter 13. Korea's Choice: Proposals for an AI G3 Leap

← Previous Chapter

See every chapter

Next Chapter →

Part 6: The Future and Implications, Pax Technica

A. The Need for Korean Defense AI

(1) Responding to North Korean Threats and Regional Arms Races

In the fall of 2024, a massive A-shaped structure rose inside the KINTEX exhibition hall in Goyang, Gyeonggi Province. It was a Hanwha Group exhibition pavilion. Inside were AI-piloted unmanned combat aircraft models, hybrid weapons combining drones and loitering munitions, and smart warships that halved crew requirements. The theme of the Seoul International Aerospace and Defense Exhibition ADEX 2025 was "The Future Battlefield Created by AI." A Hanwha representative explained to visitors: "Now it is data, not pilots, that commands the battle."

At the same time, on the other side of the Demilitarized Zone, a different kind of innovation was underway.

North Korea launched ballistic missiles dozens of times in 2024 alone. Among them were weapons claiming to be hypersonic glide vehicles. Nuclear warhead miniaturization technology had become more sophisticated. Simultaneously, North Korea's cyberattack capabilities had reached world-class levels. According to National Intelligence Service estimates, North Korea steals hundreds of millions of dollars in cryptocurrency annually, with a significant portion of that funding going to nuclear and missile programs. It is a strategy of offsetting conventional force inferiority with asymmetric capabilities.

The problem is not North Korea alone.

Northeast Asia is the region where the arms race is accelerating fastest in the world. China has openly declared the goal of acquiring the capability to forcefully unify Taiwan by 2027. The PLA Navy has already surpassed the United States in ship count. Japan reinterpreted its constitution to raise defense spending to 2% of GDP. Russia's invasion of Ukraine left a dangerous precedent that changing the status quo by force is possible in Northeast Asia as well.

The military personnel available to Korea for responding to all these threats is shrinking. As of 2025, annual births number approximately 250,000. This figure has declined to one-quarter of the one million recorded in the 1970s. According to Ministry of National Defense estimates, by 2040 the pool of military service-eligible personnel will bottom out. The demographic cliff is not merely a social issue; it is the essence of a security crisis.

The solution the Ministry of National Defense has produced is "Defense Innovation 4.0."

The core of this plan, announced in 2023, is replacing manpower with technology. Manned-unmanned teaming combat systems where AI-based drones and robots conduct operations alongside soldiers; intelligent C4I systems where AI analyzes battlefield situations and supports commanders' decisions; Joint All-Domain Command and Control (JADC2) systems where satellites and drones collect real-time intelligence linked to strike systems. All of this is contained in the roadmap.

In April 2024, the Defense AI Center was established under the Ministry of National Defense. In January 2025, this organization was expanded and reorganized into the Defense AI Technology Research Institute. Its budget is approximately 15 billion won. Each service branch , Army, Navy, and Air Force , is pursuing more than 30 AI-based force transformation projects. Automated surveillance video analysis systems, cloud systems using AI semiconductors, and supercomputers for combat staff are being introduced.

But an enormous gap exists between reality and goals.

Korea's defense R&D budget is about 5% of total defense spending. The United States invests over 15%. An even more serious problem is data. AI requires large volumes of high-quality data to learn. But the Korean military has not established clear governance on what data is, where it exists, or how to use it. The Army, Navy, and Air Force each operate their own networks, and those networks are not standardized. The "silo" problem U.S. intelligence agencies experienced during 9/11 is being repeated inside the Korean military 20 years later.

On the Ukraine battlefield, Palantir's software integrated satellite imagery, drone reconnaissance, and human intelligence in real time, completing a kill chain in 20 minutes. According to Ministry of National Defense plans, it will take Korea 10 to 15 years to implement such a system. In the meantime, threats from North Korea and neighboring countries will not wait.

One defense expert assessed it this way: "We are still drawing the future on paper. But other countries are already testing that future on the battlefield."

(2) Public-Private Partnership (PPP) Lessons from the Palantir Model

In 2016, Palantir sued the U.S. Army.

Under federal procurement regulations, the government must first assess whether commercial software can meet requirements. But the Army was skipping this procedure and awarding contracts only to traditional defense contractors. Palantir's argument was straightforward: "If our product is better and cheaper, why aren't we even given a chance to compete?" The court ruled in Palantir's favor. The ruling included this statement: "The Department of Defense's failure to seriously consider commercial solutions is unlawful."

This ruling changed the landscape of the U.S. defense industry.

It opened the door for Silicon Valley startups to enter a market monopolized for more than 70 years by traditional defense contractors like Lockheed Martin, Raytheon, and Northrop Grumman. After that, companies like Anduril, SpaceX, and Shield AI won defense contracts. The $10 billion contract Palantir signed with the U.S. Army in 2025 was the pinnacle of this trend , a project to integrate 75 individual software systems into a single platform.

What can Korea learn from this model?

First, institutional mechanisms are needed to channel civilian technology capabilities into defense. Palantir's core technology, ontology, was originally developed for fraud detection at PayPal. Airbnb's data processing technology, Tesla's autonomous driving algorithms, and Starlink's satellite communications network were all born in the civilian sector and expanded into the military domain. Korea also has world-class semiconductor companies, AI startups, and software talent. The problem is that access points for these players to enter the defense sector are blocked.

Korea's defense procurement system remains closed. The Defense Acquisition Program Administration's complex regulations, lengthy project cycles, and cost estimation methods are unsuitable for fast-moving technology companies. One industry representative at ADEX 2025 said: "We create technology where generations change every three years. But government project cycles are 10 years. If we don't close this gap, we cannot maintain technological superiority."

Second, the "cost-plus" contracting practice must be broken. Cost-plus means paying the contractor their costs plus a set profit margin. Under this structure, contractors have no incentive to reduce costs. In fact, inflating costs increases profits. Palantir directly challenged this practice. "We will deliver results at a fixed price. Even if costs exceed the estimate, we bear the burden." This fixed-price contracting approach improved the Defense Department's budget efficiency.

Third, a shift in mindset regarding software capability is needed. Traditionally, defense was the domain of steel, ammunition, and engines. How many tanks and fighters one had was the measure of military power. But the Ukraine war upended this equation. Russia's military, once rated the world's second-strongest, suffered massive losses from Ukraine's software-based kill chain. Hardware superiority could not offset software inferiority.

Korea's defense industry senses this change. At ADEX 2025, Hanwha Aerospace unveiled the "Chunmoo 3.0" system, where AI enables real-time collaboration between reconnaissance drones and loitering munitions for target engagement. LIG Nex1 introduced an AI-based Battle Management System (BMS 2.0) integrating the entire battlefield. Hanwha Systems' Smart Battleship uses AI to reduce crew from 60 to the mid-30s. The four major defense companies' order backlog for Q3 2025 reached 91 trillion won, up 21.6% year-over-year, and an annual backlog exceeding 100 trillion won is all but certain.

But this alone is insufficient.

The true core of the Palantir model is not technology but governance. Trust between the private sector and government, free flow of data, a culture that tolerates failure, and a performance-based

reward system. These elements must combine for innovation to be possible. The Presidential Committee on National AI Strategy, launched in September 2024, selected "Defense and Security" as one of its 12 strategic areas. The committee's defense and security division head is the director of the AI and Informatization Research Division at the Korea Institute for Defense Analyses. The government has begun addressing defense AI at the national strategy level.

One defense analyst offered this forecast: "Korea's defense exports are evolving from hardware-centric to system-centric. The next step is software. If we can export a data platform alongside weapons systems, as Palantir does, K-defense will enter an entirely new dimension."

Defense exports of $20 billion by 2030 , that is the government's target. To achieve it, tanks and self-propelled guns alone are not enough. What is needed is the software connecting them into one, and the public-private partnership ecosystem to create that software.

B. Data Sovereignty and National Security Strategy

(1) Risks of Dependence on Foreign Platforms and Building Indigenous Capabilities

In early 2025, the Personal Information Protection Commission released the results of its review of the Chinese AI service DeepSeek.

The results were alarming. DeepSeek was storing Korean users' data on servers in China. Under China's National Intelligence Law, Chinese companies are obligated to cooperate with the government's intelligence collection. In theory, the Chinese government could access Korean user data held by DeepSeek. The Personal Information Protection Commission issued a corrective recommendation.

This case posed a single question: when data crosses borders, where does sovereignty reside?

What Palantir sells to the U.S. government is not just software. It is control over the data that software processes. Intelligence agency data entered into the Gotham platform, corporate data integrated into Foundry, real-time battlefield data analyzed by AIP , all of it passes through Palantir's system. Who operates the system is the question of who controls the information.

Korean government agencies and companies are increasingly dependent on foreign cloud platforms. Amazon Web Services (AWS), Microsoft Azure, and Google Cloud dominate the domestic market. As of 2024, foreign companies' share of the domestic cloud market exceeds 70%. Government ministries and public institutions face the same challenge as they pursue cloud migration.

Dependence on foreign platforms creates three categories of risk.

First, "kill switch" risk. Platform access can be cut off in the event of diplomatic conflicts or sanctions. When Russia invaded Ukraine in 2022, Western companies immediately suspended services to Russia. Microsoft, Apple, Google, and Cisco all withdrew from the Russian market. Palantir officially sided with Ukraine and terminated all cooperation with Russia. If Korea enters conflict with a particular bloc, what happens when core infrastructure depends on foreign platforms?

Second, technology lock-in risk. Once data is uploaded and systems built on a specific platform, migration to another is difficult. Transition costs are enormous, compatibility issues arise, and trained AI models are not easily transferred. This is why Palantir's Apollo platform exerts a powerful lock-in effect on customers. Once you enter the Palantir ecosystem, it is hard to leave.

Third, weakened bargaining power. Without alternatives, there is no room to negotiate prices or terms. When Britain's National Health Service (NHS) signed a 330-million-pound contract with Palantir in 2023, fierce opposition erupted domestically. The debate was whether it was appropriate to entrust NHS data to an American company. But the NHS was already in a state where it could not operate without Palantir's system. In a negotiation without alternatives, you cannot avoid the subordinate position.

Recognizing these risks, the Korean government has raised the discourse of "sovereign AI."

The National AI Strategy Committee has presented three principles: sovereignty over national goal-setting, sovereignty over infrastructure operation, and sovereignty over talent acquisition. The core is that Korea must possess the key capabilities needed for the AI era indigenously. The government announced plans to secure 10,000 GPUs by 2025, introduce a top-six-ranked supercomputer by the first half of 2026, and build a National AI Computing Center by 2027.

Infrastructure alone is insufficient. True sovereignty comes from the software capabilities running on that infrastructure.

Palantir's value does not reside in servers or GPUs. It lies in ontology technology that connects data and extracts meaning, the engineering capability to implement it, and the Forward Deployed Engineer (FTE) workforce that understands clients' problems and designs solutions. Even if Korea can make world-class semiconductors, if the software platform running on them comes from abroad, sovereignty remains incomplete.

The AI Basic Act passed in December 2024 is significant in this context. Its official title is the "Basic Act on Advancement of Artificial Intelligence and Establishment of Trust." Of 264 members present, 260 voted in favor. This law takes effect in January 2026. It contains transparency obligations for high-impact AI, safety assurance measures, and requirements for appointing domestic representatives. Foreign big tech companies must comply with domestic regulations to provide AI services in Korea.

But regulations alone cannot win the competition.

The ultimate answer is to build a Korean Palantir, a Korean data platform. A sovereign platform capable of securely integrating and analyzing government data, defense data, and public data. Without it, Korea will forever remain a consumer of foreign platforms.

(2) The Fusion Potential of K-Defense and AI Technology

The exhibit that drew the most attention at the ADEX 2025 show floor in October 2025 was LIG Nex1's electronic attack aircraft (EA aircraft) model.

The researcher giving the presentation told visitors: "This aircraft detects and analyzes radio signals across North Korea to jam communications or protect friendly operations. Domestic technology has reached the level of operations at 40,000 feet." On the screen, a demonstration showed electronic jamming signals spreading across a virtual battlefield map.

The evolution of K-defense is remarkable.

The negotiation volume for orders at ADEX grew from $29.4 billion in 2023 to $44.9 billion in 2025, a 52.7% increase. Hanwha Aerospace exported K2 tanks to Poland and K9 self-propelled howitzers to Norway. Seventy-two percent of Hyundai Rotem's defense solutions revenue comes from overseas. The government has set a target of $20 billion in defense exports by 2030.

But hardware alone cannot sustain a competitive edge.

The fastest-growing segment of the global defense market is software. It is symbolic that Palantir's market capitalization surpassed that of traditional defense giant Northrop Grumman. The U.S. Department of Defense's $10 billion investment in Palantir in 2025 was not for aircraft carriers or stealth fighters but for a data integration platform. The era has arrived where it is not the weapons but the systems connecting them that determine the outcome of war.

For K-defense to leap to the next level, AI software capabilities must be fused with weapons systems.

That movement has already begun. Hanwha Aerospace's "Chunmoo 3.0" is a system where reconnaissance drones and loitering munitions share intelligence in real time through AI collaboration. When the launcher opens, drones take off first to search for targets, and based on the acquired information, loitering munitions execute precision strikes. A Hanwha representative explained: "Each round collaborates in real time."

Hanwha Ocean's "Smart Battleship" uses AI to integrate the Combat Management System (CMS), Engine Control System (ECS), and Integrated Bridge System (IBS). Automatic target recognition is built in. A ship that typically requires 60 crew members can be operated with mid-30s personnel. A critical solution for the demographic cliff era.

LIG Nex1 developed an AI-based Battle Management System (BMS 2.0). It connects all weapons, sensors, and command systems on the battlefield into a single network to support real-time decision-making. Long-range air-to-surface and air-to-air guided missiles for the KF-21 fighter were also unveiled. The company has partnered with KAIST to begin developing AESA radar and SAR semiconductors.

Korea Aerospace Industries (KAI) has set a goal of completing a Manned-Unmanned Teaming (MUM-T) system by 2030. An AI pilot simulator and the development of the KUS-FS, which will serve as an unmanned wingman for the KF-21, are underway. It has also partnered with Samsung

Electronics for AI semiconductor development.

What is noteworthy is that these technologies are being connected to export products.

Hanwha Systems signed a contract to supply Boeing with the ELAD (Electronic Laser Attack Demonstrator) for F-15 fighters. This is the first case of a Korean defense company entering the U.S. defense market. LIG Nex1's guided rocket "Bigong" passed the U.S. Foreign Comparative Testing (FCT), earning eligibility for U.S. export. KAI is pursuing the U.S. Navy trainer aircraft program, a large-scale contract worth some 10 trillion won.

But the true game changer is not individual weapons but system integration capability.

What Palantir demonstrated in Ukraine was not the performance of individual weapons. It was the shortening of the kill chain by connecting satellites, drones, ground forces, artillery, and missile units into a single data network. There is a world of difference between Korea exporting K2 tanks, K9 self-propelled howitzers, Chunmoo multiple rocket launchers, and Cheongung missiles individually, and exporting them together with an integrated platform connecting them all.

The "K-Air Defense Belt" vision presented by LIG Nex1 moves in this direction. It is a plan to export the Korean Air and Missile Defense (KAMD) lineup, composed of L-SAM, Cheongung-II, Haegung, and Singung, as a package together with AI-based command-and-control systems. The concept is to expand customized multi-layer air defense solutions to the Middle East, Asia, North Africa, and Europe, with technology transfer and local production models.

This is the potential for a Korean Palantir.

In hardware, Korea has already reached world-class levels. The remaining tasks are the software connecting that hardware, the platform integrating the data, and the talent to build it. It is a good sign that the four major K-defense companies saw operating profits in Q3 2025 surge 133.8% year-over-year to 3.5 trillion won. If these profits are reinvested in AI and software capabilities, K-defense can leap from a hardware exporter to a battlefield systems exporter.

C. Implications for Education, Administration, and Public Policy

(1) Public Data Governance and Applying the Palantir Model

In the spring of 2020, Britain's National Health Service (NHS) was in chaos.

COVID-19 patients were surging, but the NHS could not even determine how many beds were available and where. Thousands of healthcare institutions operated their own systems, and those systems did not talk to each other. Patient data, equipment inventories, staffing allocations, and vaccine supplies were scattered across separate databases. It was an information silo.

The NHS called in Palantir.

After the Foundry platform was introduced, things changed. Oxygen supply volumes, mask inventories, ventilator locations, and bed occupancy rates were displayed on a single dashboard. When vaccines arrived, AI calculated how many should be allocated to which region. The appointment system matched supply and demand in real time. An NHS official assessed: "For the first time, we could see the entire system at a glance."

What can Korea's public sector learn from this example?

The Korean government possesses one of the largest public data repositories in the world. Resident registration, health insurance, tax, real estate, vehicle, financial, and education data are meticulously accumulated. The problem is that these datasets are not connected. Each ministry operates its own system, and those systems are not interoperable. The databases of the Ministry of the Interior, the Ministry of Health and Welfare, the Ministry of Land, and the Ministry of Education each exist independently.

Under this structure, it is difficult to measure the effectiveness of policies. How long do unemployment benefit recipients stay employed after finding a job? How do the economic circumstances of families moving into public housing change? What impact does educational investment have on future income? Answering these questions requires linking data from multiple ministries. But under the current system, such linkages are extremely limited.

The core of the Palantir model is a "Single Source of Truth."

It integrates scattered data into a unified ontology so that the whole picture can be seen. Just as Hospital A's patient data was connected to Hospital B's equipment data in the NHS case, if Korea's welfare data were linked with employment data, policy blind spots could be identified and resources allocated more efficiently.

The National AI Strategy Committee, launched in September 2024, presented "nationwide AI transformation" as a core policy pillar. Applying AI to the public sector is an important agenda. AI-native government platforms, integrated citizen service platforms, and redesigning public systems using private-sector capabilities are under discussion. The Lee Jae-myung government has announced plans to build high-quality data cluster centers for AI learning and expand the opening of public data.

But data integration requires preconditions. Specifically, governance.

Who can access the data, for what purpose, how long it can be retained, and who is accountable. Without clear rules, data integration only creates chaos. This is why Palantir emphasizes "metadata and lineage systems." All data must be tracked , where it came from, who accessed it when, and how it was transformed , so that accountability is clear.

The "Guide to Personal Information Processing for Generative AI Development and Use" released by the Personal Information Protection Commission in August 2025 is important in this context. This guide divides the AI lifecycle into four stages , purpose setting, strategy formulation,

learning/development, and application/management , and presents considerations for personal information protection at each stage. The core is establishing a governance framework centered on the CPO (Chief Privacy Officer).

To apply the Palantir model in the public sector, it is necessary to go beyond simply adopting a platform and redesign the governance framework. Inter-ministry data sharing principles, access rights management systems, audit trail mechanisms, and accountability systems for misuse , these elements must be in place to reap the benefits of data integration while controlling the risks.

(2) Finding the Balance Between Efficiency and Privacy Protection

In February 2023, the German Federal Constitutional Court delivered a historic ruling.

The Palantir-based systems used by police in Hesse and Hamburg were unconstitutional, the court decided. The court declared: "This system enables profiling of individuals with a single click. This violates the right to informational self-determination." Hesse was required to amend its law by September 2023, and Hamburg's relevant law was invalidated.

This ruling reveals the dual nature of Palantir technology.

In terms of efficiency, Palantir is overwhelming. In criminal investigations, it can connect hundreds of databases to track suspects. In counter-terrorism, it can detect anomalies in advance. In public health crises, it can optimize resource allocation. All of this is possible because it allows you to "see the whole picture."

But that very capability becomes a threat to privacy.

Seeing the whole picture also means scrutinizing every detail of an individual's life. The LAPD's predictive policing program "Operation LASER" predicted the locations and people where crime would occur and deployed police accordingly. It was effective. But controversy erupted when the algorithm concentrated police in neighborhoods of particular racial groups, reproducing racial discrimination. Prediction reproduced bias.

Korea is not free from this dilemma.

Korea is one of the countries with the most tightly woven personal information infrastructure in the world. Resident registration numbers track individuals from the moment of birth for a lifetime. Health insurance data contains all medical records. Tax data details income and assets. Communications data reveals location and behavioral patterns. What becomes possible if AI connects all of this?

The AI Basic Act passed in December 2024 foregrounded "trust."

This is why the law's official title is the "Basic Act on Advancement of Artificial Intelligence and Establishment of Trust." The law presents three fundamental principles: the principle of human dignity, the principle of public good, and the principle of technological fitness for purpose. It also specifies ten core requirements: human rights protection, privacy protection, respect for diversity,

prohibition of infringement, public interest, solidarity, data management, accountability, safety, and transparency.

Special obligations are imposed on "high-impact AI." High-impact AI is AI that can have significant effects on human life, safety, or fundamental rights. Autonomous driving systems, medical diagnostic AI, and financial credit scoring AI fall into this category. Operators of such AI bear transparency obligations, safety assurance obligations, and AI impact assessment obligations. Foreign big tech companies providing services in Korea must designate a domestic representative.

But law alone is not enough to achieve balance.

Technical mechanisms are needed. The "Human-Machine Teaming" concept emphasized by Palantir offers useful reference. AI must not operate as a black box; humans must be able to understand and review why a particular decision was made. Explainable AI (XAI), audit trail systems, and bias detection tools are essential.

Institutional mechanisms are also necessary. The role of the Personal Information Protection Commission must be strengthened. It must be able to monitor what data was used and how when AI is employed in public decision-making. Procedures must be established for citizens affected by algorithmic decisions to file objections. Just as the EU AI Act specifies the right to "human oversight," Korea must build similar institutions.

Above all, social consensus is needed.

Where to set the trade-off between efficiency and privacy. How much surveillance to permit for crime prevention. How much personal data linkage to allow for improving public services. There are no correct answers to these questions. Every society will arrive at different answers.

Germany is sensitive to surveillance because of historical experience. It remembers how the Nazi Gestapo and the East German Stasi abused personal information. That is why the Constitutional Court put the brakes on the Palantir system. The United States chose to prioritize security after 9/11. That is why Palantir became a core tool of intelligence agencies.

What choice will Korea make?

The North Korean threat, the demographic cliff, and the demand for administrative efficiency strengthen the temptation of data integration. But memories of authoritarian history, experiences with personal information breaches, and growing citizen awareness of privacy push in the opposite direction. The AI Basic Act seeks this balance with the expression "trust-based." But how the principles written in the law will be implemented in reality remains unknown.

Alex Karp of Palantir has said: "Our technology sometimes kills people." He did not hide this fact. He acknowledged both the power and danger of technology simultaneously. What Korea should learn from the Palantir model is not just technology. It is honest dialogue about the impact that technology has on society, awareness of the need for democratic control, and the wisdom to manage the

ceaseless tension between efficiency and freedom.

Harnessing the power of data without being consumed by it. That is Korea's true challenge on the path to the AI G3 leap.

← Previous Chapter

See every chapter

Next Chapter →

Kim Kyung-jin

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com



#KimKyungjin #PALANTIR #Palantir #WarAI #SurveillanceAI #DefenseAI #DataPower #ArtificialIntelligence #AILibrary

# Chapter 14. The Future at a Crossroads

← Previous Chapter

See every chapter

Part 6: The Future and Implications, Pax Technica

A. Savior or Ruler

One early morning in March 2020, a doctor in the emergency room of Newham General Hospital in East London picked up her phone. On the screen was a text message from hospital headquarters: "Oxygen ventilators assigned today: 0. No additional stock expected." She looked down the hallway. Thirty-two patients lay there, and eleven of them would not survive until the next morning without ventilators. At that moment, the same scene was being repeated across all of Britain. Someone was drinking coffee in front of a warehouse overflowing with ventilators, while someone else was holding a patient's hand with empty hands. The problem was not a shortage of ventilators. The problem was that nobody knew where anything was.

This is the background for how Palantir entered Britain's National Health Service (NHS). The British government was floundering in a sea of data chaos. Each hospital managed inventory with its own system, and those systems did not communicate with one another. There was nowhere that could show at a glance how many beds were available at which hospital, how much vaccine remained in which region, or which medical staff were infected with COVID and in isolation. Palantir's engineers connected the nationwide databases within days. Through the system they built, the British government could see national medical resource status in real time for the first time. Equipment was moved from places with surpluses to places with shortages, and vaccine distribution priorities were determined based on data. By 2024, the value of this contract reached 330 million pounds. Palantir undeniably saved lives.

(1) The Benefits of Efficiency and Security That AI and Data Integration Bring

The word efficiency gives a cold impression. But for someone dying in a hospital corridor, efficiency is another name for life. The first benefit of data integration is straightforward: it gathers scattered puzzle pieces in one place so the full picture becomes visible. As the UK example shows, the tragedy that occurs when information exists in the right places but fails to reach where it is needed is not a technical failure but a systemic one. Palantir's software served as the glue filling that failure.

On the battlefield, this effect becomes even more dramatic. We have already examined what happened in Ukraine. When hundreds of Russian tanks rolled forward, what enabled Ukrainian forces to hold was not an abundance of ammunition. It was a system that calculated in real time where the enemy was, where friendly artillery was positioned, and which weapon to strike with at what moment. Palantir's software wove satellite imagery, drone footage, field scouts' reports, and intercepted communications into one and displayed it on the commander's screen. The target identification and strike authorization process that had previously taken hours was shortened to minutes. In military

terms, this is called "kill chain acceleration." Put coldly, the side that shoots faster wins. Palantir enabled Ukraine to shoot faster.

On July 31, 2025, the U.S. Army signed its largest-ever contract with Palantir. This contract, worth up to $10 billion over 10 years, consolidated 75 previously scattered individual software contracts into one. Army Chief Information Officer (CIO) Leo Garciga said in the announcement: "This enterprise agreement represents the Army's commitment to modernizing our capabilities while acting in a fiscally responsible manner. By streamlining procurement processes and using enterprise-level discounts, we are maximizing purchasing power while enhancing operational effectiveness." There is a key point here. The Army did not simply buy a new weapon. It bought the "integration" of binding 75 different systems into one. This integration is Palantir's true product.

The same logic operates in the corporate world. Airbus sources parts from around the globe to build aircraft. The number of parts in a single aircraft runs into the millions. Tracking in real time where those parts' inventories stand, when they will arrive from which supplier, and which production line is at risk of stopping is a task beyond human capacity. The Skywise system Airbus built integrates all this information and presents it to managers. When global supply chains collapsed during the COVID pandemic, the difference in recovery speed between companies with this system and those without was stark.

BP offers a similar case. Energy companies operate refineries, pipelines, storage tanks, and gas station networks. The data generated at each facility is enormous , pressure gauge readings, temperature sensor records, truck routes, fuel expiration dates. If this data is not integrated, management is effectively blind. They cannot know in advance which equipment at which refinery is about to fail, or where fuel demand is surging. Palantir's Foundry platform aggregates all this data onto a single screen. Not remediating after a problem occurs, but predicting before it happens. This is the second benefit of data integration.

Alex Karp repeats: the problems facing Western society are too complex to solve with human cognition alone. There is merit to his point. Climate change, pandemics, supply chain disruptions, terrorism, cyberattacks , all these threats are enormous, complex, and interconnected. It is impossible for a single genius to consider every variable and make the optimal decision. But machines are different. Machines simultaneously process millions of data points, find patterns invisible to humans, and calculate the optimal choices. Humans need only make the final decision.

In May 2025, the U.S. Department of Defense increased the Maven Smart System contract with Palantir by $795 million. Maven is a system that uses AI to analyze battlefield video and imagery. In the past, analysts had to manually review thousands of hours of drone footage by eye. Now AI does that work. The machine does not get tired, does not lose focus, and does not miss patterns. Analysts only need to review the critical scenes AI has filtered out. This is what "human-machine symbiosis" actually looks like. The machine does not replace the human; it amplifies human capabilities.

The third benefit of data integration is transparency. It may sound paradoxical. A company criticized as a surveillance tool touting transparency as an advantage? But Palantir's system has a built-in

"audit trail" function. It records who accessed which data, when, and how it was used. The "blame avoidance" common in bureaucratic organizations becomes difficult. Because the data underlying every decision is fully preserved. Of course, the point that this transparency is a double-edged sword is something we will examine shortly.

The UK's pandemic response, Ukraine's war effort, the U.S. Army's procurement innovation, and global corporate supply chain management. All these examples converge on a single conclusion: in a complex world, the ability to integrate and analyze information is power. Palantir has made that power into a product. That their technology saves lives, leads wars to victory, and improves corporate efficiency is an undeniable fact. But as with all power, this power can produce entirely different outcomes depending on the intentions of the person wielding it. Now it is time to look at the other side.

(2) The Dystopia of Privacy Loss and a Controlled Society

In March 2025, reports emerged that the U.S. State Department had launched a program called "Catch and Revoke." The program's purpose was to use AI to analyze the social media of student visa holders, searching for statements sympathizing with terrorism. Tens of thousands of international students had their tweets, Instagram posts, and Facebook comments scanned by algorithms without their knowledge. Suppose a student posted something critical about the situation in the Gaza Strip. Could that be classified as "terrorist sympathizing"? If so, who decides? The algorithm, or a person?

This is where Palantir enters. According to reporting by the Washington Post and the New York Times, the Trump administration was expanding Palantir's technology across the federal government. At least four federal agencies , the Department of Homeland Security, the Department of Health and Human Services, the Social Security Administration, and the Internal Revenue Service , had adopted or were discussing adoption of Palantir's Foundry platform. If these platforms are connected, what happens? A single person's tax records, Social Security information, medical claims history, immigration records, and bank account information could be integrated into a single profile. A bureaucrat could peer into a citizen's entire life with a few clicks.

This is technically possible. Palantir's "ontology" is designed to do exactly this. Connecting scattered data to create a "digital twin." If this technology was used at Britain's NHS to track the location of ventilators, at the U.S. Department of Homeland Security, it could be used to track the location of people. The technology itself is neutral. But the context in which technology is used is never neutral.

On February 16, 2023, the German Federal Constitutional Court delivered a historic ruling. The Palantir-based data analysis systems used by police in Hesse and Hamburg violated the constitution. The court's logic went like this: even data that appears harmless individually, if combined without limit to create a "profile" of a person, constitutes a serious violation of the "right to informational self-determination." The judges warned: "If profiling becomes possible with a single click, citizens will not even know which of their actions are being surveilled."

In Los Angeles, this had already happened. LAPD's "Operation LASER" was a predictive policing program using Palantir's technology. The algorithm analyzed past crime data to predict areas and people where crime was likely to occur in the future. The problem was that the past crime data itself was biased. If police patrolled a particular area more frequently, more crimes were detected there. The data recorded that area as having "high crime rates." The algorithm learned from that data and classified the same area as "high risk." Police again concentrated on that area. A vicious cycle. As a result, specific racial groups and low-income neighborhoods were excessively surveilled, arrest rates climbed, and the data once again classified them as "dangerous groups." This is the "automation of discrimination by algorithm."

The relationship between ICE and Palantir is even more contentious. ICE uses Palantir's system to track undocumented immigrants. The system does not merely look at immigration records. It connects vehicle registration information, phone numbers, addresses, family relationships, employment records, and even school enrollment information to build a "network" around a person. It identifies who the undocumented immigrant's relatives are, where those relatives work, and whether other undocumented immigrants are at that workplace. This mapped information is used in large-scale enforcement operations. According to 2025 reporting, the Trump administration is pursuing plans to further expand this system by integrating federal databases.

Even within Palantir, voices of concern are emerging. The New York Times quoted anonymous Palantir employees: "There are worries about bringing this much sensitive information together in one place. The company's security practices are only as good as the people using them." Another employee worried about the impact of cooperation with the Trump administration on the company's reputation. But their concerns did not change management's decisions.

What happened in the Gaza Strip shows these concerns in their most extreme form. The "Lavender" system reportedly used by the Israeli military operated on a structure where human commanders approved AI-generated target lists. According to reports, approval took an average of just 20 seconds. Despite AI having a 10% error rate, human deliberation was skipped in favor of speed and efficiency. Whether Palantir was directly involved in this system has not been confirmed. But it is a fact that Palantir declared its official support for Israel in the Israel-Hamas war. Alex Karp wrote in his 2024 shareholder letter: "We stand with Israel. In this moment, and going forward."

The core of the problem is this. Palantir's technology offers choices. The same ontology can track the location of ventilators or track the network of an immigrant. The same kill chain acceleration can eliminate a terrorist or kill a civilian. Technology does not know purpose. Purpose is set by people. But when technology becomes too powerful, the situation can arise where technology narrows people's choices rather than people controlling technology. The logic of "we have this data, so we must use it" pushes out the question of "should we collect this data in the first place?"

As of 2025, Palantir's stock is at all-time highs. Market capitalization has exceeded $250 billion, surpassing traditional defense giants Lockheed Martin and Raytheon. President Trump said at an AI summit: "We buy a lot from Palantir." The entanglement of government and technology companies has never been deeper. When an unelected technology company assumes core national functions,

where should democratic control operate?

Amnesty International warned in a 2025 report: "Palantir is failing to live up to its responsibility to protect human rights when its software is being used to facilitate serious human rights violations." Internal U.S. Army memos warned of "fundamental security" issues and vulnerabilities in joint battlefield communication systems built by Palantir and its partners. But these warnings did not prevent the signing of a $10 billion contract.

In Tolkien's novels, the palantír is "the seeing stone." Those who possess this crystal ball can see events taking place far away. But in the novels, this stone becomes not a blessing but a curse. What it shows is only a portion of the truth, and that portion can mislead the viewer. Saruman and Denethor were bewitched by what they saw through the palantír and met their doom. The warning Tolkien wrote in the 1950s is, uncannily, becoming reality 70 years later.

B. A Question for the Reader

One ordinary afternoon, a woman in a Seoul apartment was looking at her smartphone. On the screen was a notification from the Health Insurance Review and Assessment Service: "Your medical record has been updated." She cleared the notification without a second thought. At that same moment, her medical record was being entered into an insurance company's underwriting system. Also at that same moment, the keyword "cause of headaches" that she had searched for yesterday was being added to an advertising platform's profile. She knew none of this. She did not think she needed to know. But years later, when she tries to buy insurance, or goes to a job interview, or applies for a loan, these data points could influence her fate. Without her ever knowing.

This is the world we live in.

(1) In the Era of Data Power, What Must We Watch?

Palantir is a mirror. Looking into this company, you can see how the data power of our era operates. But we cannot just stare at the mirror. The question we must now ask is this: in the era of data power, what must we as citizens watch?

The first thing to watch is "the flow of decisions." Not the data itself. Data is merely raw material. What matters is how that data is combined, what conclusions it produces, and what actions those conclusions lead to. When LAPD's predictive policing system classified a particular area as "high risk," on what basis was that classification made? Was there a procedure to challenge that classification? When the classification was wrong, who was held accountable?

Bruce Schneier and Nathan Sanders wrote in their 2025 book, Rewiring Democracy: "AI is fundamentally a power-amplifying technology. It executes one person's command at greater speed, larger scale, broader scope, and more sophisticated manner than any human individual could." The question is to whom this amplified power goes. To citizens, or to the bureaucrats and corporations managing citizens?

The second thing to watch is "the fiction of consent." Every day we "consent" to dozens of apps and services. But nobody reads the terms. The terms are written to be difficult to read, and even if you do read them, there is often no option to refuse. Research frequently finds that data is still collected even when users have declined tracking. Consent has become a mere formality. Real control does not rest with the user.

Research on Facebook (Meta) reveals an interesting phenomenon. On the front stage, the platform manages user trust through visible measures like "privacy settings" and "real-name policies." But backstage, much more significant data flows move almost invisibly through recommendation algorithms and ad targeting systems. In theatrical terms, it is a separation of "front stage" and "back stage." What we need to watch is not the show on the front stage, but what is actually happening backstage.

The third thing to watch is "the gaps in regulation." Blockchain, the metaverse, generative AI, cross-border data flows , in these new technological environments, gaps emerge that existing laws and institutions fail to adequately address. Comparative law research on metaverse platforms shows that countries still take fragmented approaches to consumer protection, jurisdiction, privacy, and liability governance. These gaps lead to the expansion of "private regulation," where companies set their own rules. Who makes the rules, and who watches the rule-makers? This is the essential question.

The fourth thing to watch is "the use of data on vulnerable populations." Research on children's app marketplaces reveals a shocking reality. Platforms constitute child users as "commodified audiences." Children's clicks, viewing times, and preferences are collected as data and sold to advertisers. As data in education, welfare, and healthcare becomes increasingly quantified and serves as material for algorithmic evaluation, the risk grows that the lives and possibilities of vulnerable populations will be narrowed by data models.

According to a 2025 Carnegie Foundation survey, 55% of California residents responded that they were "very concerned" that AI-generated content would deepen political violence and polarization. At the same time, citizens were divided on whether AI could make them "more informed citizens." This data suggests an important point: anxiety about technology is high, but how to translate that anxiety into action has not yet been agreed upon.

The fifth thing to watch is "apathy itself." Many people press the "agree" button without reading privacy policies. This is not individual laziness. It is the result of structural fatigue caused by a data environment that is difficult to understand and designed unilaterally. If this situation persists, data power expands without meaningful resistance. Citizens' democratic agency gradually shrinks. That is why efforts are needed to reposition data issues, through education, media literacy, and public discourse, not as the domain of experts but as part of citizens' common sense and political imagination.

Ultimately, the central question of the data power era goes beyond "Are we being surveilled?" The more important question is this: "What, and how, shall we watch together?"

(2) Beyond Palantir: Imagining the Next Generation of Data Governance

Moving beyond Palantir does not mean eliminating Palantir. It means telling a bigger story. Palantir is one case showing the extremes of the current paradigm. How quickly data integration can make a state's decisions, how a platform becomes infrastructure, and what ethical burdens that infrastructure carries , all of this can be seen through the lens of Palantir.

What we must now imagine is the next generation of data governance.

First, "democratic data agency" must be institutionalized. This means a structure where users are not mere consent providers but subjects who participate in discussing and deciding data policies and usage directions. New institutional imagination is needed: data cooperatives, citizen panels, participatory data governance bodies. In February 2025, California announced a new deliberative democracy program , an experiment using AI to analyze and integrate the opinions of thousands of citizens and reflect them in policy decisions. It is still in its early stages, but it shows the possibility that technology can be used to expand citizen participation.

Second, "responsible governance" principles must be built into the design phase of technological infrastructure. The European food and nutrition data platform case offers a good example. This platform applied design principles that holistically considered data types and technologies, ownership and intellectual property, privacy and security, and ethical governance structures. It demonstrates that a data platform can be not merely a means of increasing efficiency, but a space where individual and collective privacy, power relationships, and accountability are designed together. Data governance is not only a matter of law and policy. It is also a matter of "design ethics" embedded in platform architecture and operating norms.

Third, moving beyond closed, proprietary models, we must explore alternative platform ecosystems that strengthen public interest. California's proposed "CalCompute" initiative envisions a structure where open-source AI systems operated on public servers are governed by public institutions. A public interest data archive model openly accumulates and analyzes materials for monitoring and holding platforms accountable , contracts, court documents, impact assessments, patents, and civil society reports , providing a foundation for researchers, journalists, and activists to demand platform accountability.

Fourth, the framework for international data governance must be reconfigured around human rights and the public good, not only competition and security. Currently, platform governance centered on the United States and China emphasizes algorithmic sovereignty and data localization in the name of security and economic competitiveness. But there are also significant gaps in governance and accountability for data flowing across borders. The next generation of data governance requires redefining cyberspace not as a "deregulated, deterritorialized zone" but as a space where international human rights standards and accountability norms apply.

Fifth, data sensitivity and imagination must be cultivated at the educational and cultural level. Studies analyzing surveillance capitalism warn that the current data regime is permeating daily life with little

social resistance and forming a new social order. To change this structure, data and algorithmic issues must be made topics that citizens, students, and children can discuss together , not the exclusive domain of technologists or policy experts. When learning about and discussing data rights, privacy, algorithmic bias, and platform accountability becomes routine in schools, homes, and communities, the next generation will be able to imagine the post-Palantir data order more critically and creatively.

In a September 2025 report, the OECD recommended: "Government AI adoption is growing, but many initiatives remain at the pilot stage. Successful AI adoption requires seven key elements: governance, data, digital infrastructure, skills, investment, procurement, and partnerships with non-government actors. A whole-of-government, coordinated approach to AI transformation that is user-centric, meaningfully engages stakeholders and the public, and establishes clear and proportionate guardrails will enable managing risks while remaining agile and adaptable to future changes."

The final question returns to the reader. Whether to view the era of AI and data integration solely as a narrative of "salvation" or solely as a narrative of "domination" is not a predetermined fate. The boundary between the two will shift according to what governance structures we design, what power we monitor, and what language of participation and resistance we create.

Alex Karp says: "Technology is not neutral." He is right. Technology reflects the worldview of those who created it. If so, we must demand technology that reflects not the worldview designed by a handful of genius engineers or philosopher-CEOs, but democratic values agreed upon by ourselves.

If Palantir's software is used to catch terrorists, it must also be usable to monitor police abuse of power. Just as technology serves those in power, it must be designed to serve as a shield protecting the vulnerable.

In Tolkien's novels, the palantíri are ultimately thrown into the sea. Because it was the only way to escape their temptation and danger. But the real-world Palantir cannot be thrown into the sea. The era of data integration and AI has already arrived, and there is no going back. What we can do is participate in deciding who uses this powerful tool, how, and for what purpose.

The Technological Republic has arrived. But whether the sovereign of that republic is technology or citizens has not yet been decided. It is not Palantir standing at the crossroads. It is us. Remember that even the most powerful software in the world cannot execute a single line of code without human will. The moment our vigilance and participation ceases, on the other side, someone's finger will be poised over the Enter key.

The Night Outside the Window

As I write the final sentences of this book, Seoul's night is quiet. Lights twinkle outside the window. The people sleeping beneath those lights have likely never heard the name Palantir tonight. They will not know where their data flows or who is watching it.

Throughout the writing of this book, one question never left my mind. Are we spectators in this game, or pawns?

Palantir's technology was born in the embrace of the United States, a military superpower. They created this tool for their own national security, and they selectively permit it to allies. Korea is one of those allies.

But behind the comfort that the word "ally" provides lies an uncomfortable truth. We did not create this technology. We are in the position of borrowers.

There is a term called data sovereignty. The principle that a nation's data should be controlled by that nation. But in reality, this principle wavers before the logic of power. Wherever Palantir's servers are located, whatever the people who designed those algorithms are thinking, we cannot know. There is no way to verify how our information is processed inside the black box.

This is the source of regret.

As a lawyer, as a former member of the National Assembly, as someone who researches artificial intelligence, I believe in the power of technology. But when that power is concentrated in the hands of a few, and when those few are people from another country, I have not found the answer to how democracy should function.

Palantir says it defends Western democracy. Whether we are included within that defense fence, or whether we are on the outside looking in, I am not sure.

In January 2026, President Trump declared he would "get" Greenland, an autonomous territory of Denmark, "one way or another."

If not the nice way, then the hard way. To use his expression, it was "the hard way." Denmark is a founding member of NATO. It has been America's ally for over 70 years. Yet the president of the United States was publicly threatening the territory of that ally. Vice President Vance visited the U.S. military base in Greenland without an invitation and gave a speech saying "Denmark has let down the people of Greenland." The wife of a White House aide posted a map of Greenland covered in American flags on social media and wrote "SOON."

Danish Prime Minister Mette Frederiksen's response was urgent: "If the United States takes military action on Greenland, it will be the end of NATO." Denmark's intelligence service, for the first time in history, listed the United States alongside Russia and China on its national security threat list. In Copenhagen and Nuuk, citizens took to the streets carrying placards reading "Yankee go home."

A country that styles itself the defender of Western democracy is threatening a core ally of Western democracy. A country that invoked international law to criticize China and Russia is now ignoring the same international law and coveting an ally's territory. We are witnessing the moment when the slogan "peace through strength" is being distorted into "submission through strength."

One cannot help but ask: if even a founding NATO member receives this treatment, what about Korea?

We do not belong to the "West" they speak of. When our data passes through their servers, when our military intelligence is analyzed by their algorithms, are we partners, or customers? Or perhaps, products?

No matter how superb Palantir's technology may be, we cannot control the intentions of those who own that technology. The Greenland crisis showed with brutal clarity that today's ally can become tomorrow's coercer. That data sovereignty is not an empty slogan but a matter of survival , I only now feel this in my bones.

I want to ask the reader who closes this book: What are we giving up in exchange for convenience? And who should judge whether that trade is fair?

← Previous Chapter

See every chapter

Kim Kyung-jin

Attorney · Former Member of the National Assembly · AI Policy Researcher

kimkj.com



#KimKyungjin #PALANTIR #Palantir #WarAI #SurveillanceAI #DefenseAI #DataPower #ArtificialIntelligence #AILibrary

# 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.