

Nexgen Expert View – AI Governance Series

**GenAI Is Neither Good Nor Evil –
But Executives Who Fail to Decide How It Is
Used Bear the Responsibility**

AI Governance in the Dual-Use Era: The Decisions Companies Face as “Stakeholders,” Not Just Users



Nexgen Japan

May 2026

15 May 2026

GenAI Is Neither Good Nor Evil — But Executives Who Fail to Decide How It Is Used Bear the Responsibility

AI Governance in the Dual-Use Era: The Decisions Companies Face as “Stakeholders,” Not Just Users

Ariki Ono, CEO & Logistics AI Architect

(Specializations: Gen AI Research & Development, Supply Chain Strategy, Geopolitics & Macroeconomics)

Nexgen Japan Inc.

Led AI, DX, and supply chain transformation initiatives at major IT and foreign financial firms. Provides advisory services to corporations combining generative AI and supply chain strategy, while conducting research and lectures on supply network resilience design under geopolitical risk.

This editorial is not predicated on any particular political ideology or partisan stance.

The discussion surrounding generative AI and other emerging technologies encompasses a wide range of positions — development versus regulation, optimism versus caution, and so on. Each carries reasoned arguments and the dialogue itself holds social significance; however, the focus of this article is not such macro-level policy debate.

What this article seeks to examine is a single question: what ethical perspective should each individual who actually wields these technologies — executives, practitioners, researchers, and citizens alike — bring to their engagement with technology? Technology itself bears no inherent good or evil; it is the judgement and responsibility of those who use it that shape the outcome. On this premise, I hope this article will serve as a starting point for readers to reflect on "the ethics of the user."

1. January 2026 — The Fact That “Something Began”

On 9 January 2026, the Secretary of the U.S. Department of War issued a single memorandum to the department. Its title: "Artificial Intelligence Strategy for the Department of War." Yet anyone who grasps what this document truly means will recognise that it marked a turning point in history.

The memo directs, in explicit terms: under "AI Model Parity," to "establish a delivery and integration cadence with AI vendors that enables the latest models to be deployed within 30 days of public release"; to "incorporate standard 'any lawful use' language into any DoW contract through which AI services are procured within 180 days"; and, as Pace-Setting Project #6 (GenAI.mil), to advance "Democratizing AI experimentation and transformation across the Department by putting America's world-leading AI models directly in the hands of our three million civilian and military personnel, at all classification levels."

This is not merely a notice aimed at improving operational efficiency. The U.S. Department of War elevated generative AI from "departmental PoCs (proofs of concept)" to a "department-wide operational foundation," and rapidly advanced an institutional design that positions private AI companies as continuous suppliers of national infrastructure.

On 12 January, Secretary of Defense Pete Hegseth delivered an official speech at SpaceX alongside Elon Musk, declaring that Google (Gemini) and xAI (Grok) would be integrated into GenAI.mil. On 9 February, OpenAI officially announced the integration of ChatGPT into GenAI.mil. As of February, five military branches had adopted it on an enterprise basis, with users reaching 1.1 million.

In March, Google introduced "Agent Designer" — a tool that allows users to generate operational AI agents using natural language alone — into GenAI.mil. Generative AI transformed from a "chat tool" into a "business automation platform." Meanwhile, the Department of War excluded Anthropic as a "supply chain risk" and accelerated procurement of alternative vendors. Anthropic had long maintained a cautious stance on unrestricted military use of its models.

'Governing AI through regulation' is no longer the approach; 'governing AI through procurement conditions and supply chain designation' — this is the structure the United States chose in 2026.

2. The Question Left Behind by the Manhattan Project

Eighty-four years ago, in 1942, the U.S. government launched a single national project — a top-secret research programme to counter Nazi Germany's nuclear development. This was the Manhattan Project. At an annual cost equivalent to roughly USD 100 billion (approximately JPY 15 trillion) in today's money, three secret cities were built at Oak Ridge, Hanford and Los Alamos. More than 500,000 people participated in total, yet the vast majority were not told the ultimate purpose of their work.

Extracting its structure, the Manhattan Project had four essential characteristics: (1) state-led emergency mobilisation, (2) the integration of the private sector (universities and corporations) into a total-war effort, (3) the suppression of public debate through information control, and (4) a collision between ethics and results.

The scientists who participated in the project were deeply anguished by this fourth "collision." The physicist Leo Szilard, who had conceived the nuclear chain reaction in 1933 and patented the idea in 1936, came to see the same technology — once envisioned as a transformative new energy source — diverted into a weapon of mass destruction. In July 1945 he drafted a petition to President Truman, asking him to "not resort to the use of atomic bombs in this war unless the terms which will be imposed upon Japan have been made public in detail and Japan knowing these terms has refused to surrender." Seventy Manhattan Project scientists signed the petition. It never made it through the chain of command to President Truman, and was not declassified until 1961.

Robert Oppenheimer, who led the project, later recalled the moment of the Trinity test: "We knew the world would not be the same. A few people laughed, a few people cried. Most people were silent. I remembered the line from the Hindu scripture, the Bhagavad Gita; ... 'Now I am become Death, the destroyer of worlds.' I suppose we all thought that, one way or another." In October 1945 he met

with President Truman and said, "Mr. President, I feel I have blood on my hands." Truman, infuriated, later remarked, "I don't want to see that son of a bitch in this office ever again." Two years later, in a 1947 lecture, Oppenheimer stated, "the physicists have known sin; and this is a knowledge which they cannot lose."



Photo 1. Oppenheimer at the remains of one leg of the test tower. Canvas overshoes kept trinitite off shoes, Source: [Wikipedia](#)

Technology has no inherent good or evil. The same nuclear technology became both power plants and atomic bombs. AI has the same structure.

3. From “Concept” to “Implementation” — A Two-Stage Turning Point

Phase 1 : Concept (2024–2025)

In November 2024, a bipartisan U.S. congressional commission recommended the formulation of an "AI Manhattan Project." According to Reuters, the recommendation argued that "a national project on the scale of the Manhattan Project is needed to prevent the United States from falling behind China in AGI (artificial general intelligence)." OpenAI also requested government funding, and Congress justified large-scale public investment on the grounds that "AI development offers an incredible first-mover advantage."

On 20 January 2025, the announcement of DeepSeek-R1 sent shockwaves through the United States. At a

development cost of only USD 5.6 million, it achieved performance equivalent to GPT-4o. NVIDIA's share price fell by as much as 17%, wiping out approximately JPY 91 trillion in market capitalisation. Just two days later, on 23 January, the Trump administration responded by issuing a presidential executive order repealing AI regulation. The order explicitly called for the "elimination of ideological bias" and directed the formulation of a new AI Action Plan within 180 days.

On 28 February 2025, Secretary of Energy Chris Wright appeared at the "1,000 Scientists AI Jam Session" at Oak Ridge National Laboratory, alongside OpenAI co-founder and president Greg Brockman. The laboratory had served as the central site for uranium enrichment and plutonium production in the Manhattan Project. Speaking before more than 1,000 scientists, Wright declared: "We're at the start of Manhattan Project 2. It is critical, just like Manhattan Project 1, that the United States wins this race." Around the same time, OpenAI and Meta shifted their policies to permit the use of AI for military purposes.

Phase 2 : Implementation (2026)

As 2026 began, the concept transitioned into institutions, procurement and operations.

The Department of War incorporated "deployment of the latest model within 30 days" into its procurement standards, making generative AI companies mandatory continuous-supply contractors. GenAI.mil, once a conceptual idea on paper, was transformed into an operational infrastructure for 3 million personnel.

Deployment into classified environments also became reality. OpenAI announced the introduction of its commercial state-of-the-art models into classified networks, explicitly stipulating contractual red lines such as "no domestic surveillance" and "no independent command of autonomous weapons." However, the government demanded thorough adherence to "any lawful use," and excluded Anthropic — which did not comply — from its procurement list. The structure in which a company's ethical policies are judged through market access via state procurement has become sharply defined.

This is not a story about whether the dual-use era (military × civilian) 'might come'. It has already arrived. What is now being asked is how we will respond to that fact.

4. National Approaches — Not a Difference of Systems, But of Decision-Making Subjects

From the end of 2025 into 2026, the AI governance strategies of the four major powers came into sharp relief. The differences among them go beyond the dimension of "how strict the regulation is" and contain a more fundamental question.

The United States has positioned "winning the AI race" as national strategy, nullifying state-level AI regulations through presidential executive order. Europe enacted the world's first comprehensive AI regulation, the EU AI Act, but — taking into account its impact on industrial competitiveness — is considering postponing the application of "high-risk AI" provisions until the end of 2027. China has continued to strengthen regulation consistently since its Interim Measures for the Management of Generative AI Services (2023), and in November 2025 made algorithmic registration and censorship mandatory, establishing a system in which the state centrally manages both "safety and development." Japan approved its National AI Basic Plan at a cabinet meeting in December 2025 and announced AI-related investment exceeding JPY 1 trillion.

The more essential question, however, is that **the difference in national governance is not "what kind of regulation" but "who decides."** The United States regulates private-sector behaviour through procurement conditions. China intervenes directly through the state. Europe draws lines through its legal framework. Japan relies on existing laws, while the private sector exercises its own judgement. It is precisely the difference in "who decides" that constitutes the essential difference among national governance regimes. In the Manhattan Project, scientists were removed from the position of "decision-making subjects," and only the state made decisions. Szilard's petition was buried. Oppenheimer's repentance drew the president's wrath. In today's AI landscape, do private companies function as "decision-making subjects"? Or are they already being consumed merely as "parts of the supply chain"?

5. Companies Are Not “Users” — They Are “Stakeholders”

Corporate leaders should recognise that they are not merely "those who use AI," but stakeholders in the society that AI is shaping. What the Manhattan Project demonstrated is that even the most outstanding scientists can, within the context of a national cause, have their individual consciences silenced.

The same mechanism can be at work in today's AI competition. "Our competitors are doing it," "The government is demanding it," "It is not legally problematic" — these words are structurally identical to those that the scientists of the past once told themselves.

The lesson of the Manhattan Project is universal because it was not "a tragedy caused by evil people," but "the result of a group of well-intentioned experts being structurally moved." That latent structure also exists inside companies.

This is not a call for executives to avoid using generative AI. But there are clear questions that must, at the very least, be debated at the executive level. How should the clean mission statements written on our slide decks be reinterpreted in the age of the AI revolution?

Questions Your Company Should Already Have Answers To

Does your company have answers to the following three questions?

First, how far will our company use AI? Will we use it as a tool for productivity improvement, or place it at the core of our decision-making? Who decides where the line is drawn, and through what process?

Second, what uses will we permit? Against what criteria will we judge cooperation with use cases that could be diverted to military purposes? How will we confront the gap between "not legally problematic" and "ethically right"?

Third, who bears the responsibility? Have we, as an organisation, clearly established where responsibility lies for the consequences caused by AI? Is our structure one in which engineers can say, "I merely wrote the code"?

The "AI Manhattan Project" has completed its conceptual stage in 2025 and entered its implementation phase in 2026. Will we let this fact pass as a mere technology trend, or will we debate it as a core issue of management? That choice, too, is part of the responsibility that corporations bear as "stakeholders."

References and Key Sources


- [1] US Department of War. (2026, January 9). Artificial Intelligence Strategy for the Department of War. media.defense.gov.
- [2] US Department of War. (2026, January 12). War Department Launches AI Acceleration Strategy. war.gov.
- [3] OpenAI. (2026, February 9). Bringing ChatGPT to GenAI.mil. openai.com.
- [4] OpenAI. (2026, February 28 / updated March 2). Our Agreement with the Department of War. openai.com.
- [5] Reuters. (2026, April 9). Pentagon's ouster of Anthropic opens doors to small AI rivals. reuters.com.
- [6] The White House. (2025). America's AI Action Plan. whitehouse.gov.
- [7] National Park Service. (n.d.). Manhattan Project National Historical Park. nps.gov.
- [8] Scientific American. (n.d.). The Manhattan Project Shows Scientists' Moral and Ethical Responsibilities.
- [9] RAND. (2025, April). Beyond a Manhattan Project for Artificial General Intelligence. rand.org.
- [10] Atomic Heritage Foundation. (n.d.). Rotblat Account. ahf.nuclearmuseum.org.
- [11] The Nobel Foundation. (1995). The Nobel Peace Prize 1995 - Rotblat and Pugwash. nobelprize.org.
- [12] Reuters. (2024, November 19). US government commission pushes Manhattan Project-style AI initiative.
- [13] European Parliament. (2025). Defence and Artificial Intelligence. europarl.europa.eu.
- [14] DefenseScoop. (2026, February 2). Military branches adopt GenAI.mil as enterprise AI platform. defensescoop.com.
- [15] DefenseScoop. (2026, March 10-11). DoD GenAI Agent Designer, custom AI assistants, Google Gemini.

The information in this article is current as of the end of April 2026. For the latest developments, please refer to official announcements from relevant government ministries and international organisations.

Ariki Ono, CEO & Logistics AI Architect

Nexgen Japan Inc.

Nexgen Japan is an independent advisory firm that supports corporate transformation across the fields of AI, supply chain, and business strategy. Grounded in research and insights from a global perspective, and with a clear understanding of industrial restructuring and technological evolution, we help organizations make better decisions through structural insights and the design of actionable strategies.

 <https://nexgen.co.jp>

 contact@nexgen.co.jp

NexGen