Manhattan Project 2 — The Genesis Mission

"this time it's biblical."

The nuclear arms race was about who could split the atom first. The AI arms race may be about who can make intelligence improve itself first.

There is a moment in history when a technology stops being an interesting scientific development and becomes a matter of national survival.

World War II gave us that moment.

The United States became convinced that Nazi Germany might develop an atomic bomb. The possibility was enough to mobilize an extraordinary national effort. Scientists were recruited, industrial resources were redirected, secret facilities were built, and the Manhattan Project was created.

J. Robert Oppenheimer assembled an extraordinary collection of scientific talent at Los Alamos. The government did not simply fund a university research project. It recruited some of the greatest scientific minds available and gave them a mission that ordinary commercial research could never have justified.

The result was Hiroshima and Nagasaki.

The public saw the consequences of a technology whose development had largely taken place behind closed doors.

I think artificial intelligence has entered a comparable national-security phase.

But this time, the stakes may be considerably stranger.

China's 2017 moment

On July 20, 2017, China's State Council released its New Generation Artificial Intelligence Development Plan.

This was not a casual announcement about investing in technology.

It was a national strategy intended to make China a world leader in AI by 2030. It called for major advances in AI research, industrial applications, talent recruitment, infrastructure and the integration of AI into national defense.

It explicitly recognized AI as a strategic technology capable of transforming society, the economy and international competition.

China had effectively declared that AI was a national priority.

The plan also reflected China's broader military-civil fusion approach: technological development across civilian and military institutions could contribute to national strategic capabilities.

This was China's public declaration of the stakes.

It was an early warning that AI would not remain simply a competition between technology companies.

Source: China's 2017 AI Development Plan, translated by Stanford's DigiChina

America's response: the government enters the race

The United States did not wait until 2025 to recognize AI as a national-security issue.

There was a sequence of increasingly explicit government actions.

- 2020: The National Artificial Intelligence Initiative Act established a federal framework for coordinating AI research and development.
- 2022: The CHIPS and Science Act reinforced federal AI research, testing and infrastructure efforts.
- October 2023: Executive Order 14110 established a government-wide framework for AI development, safety and national-security considerations.
- October 2024: National Security Memorandum 25 directed national-security institutions to harness AI and address its implications. It included a classified annex covering sensitive national-security matters.
- November 2025: The Genesis Mission was launched.
- June 2026: DARPA and the National Science Foundation announced AI Forge.
- June 2026: National Security Presidential Memorandum 11 directed an accelerated national-security AI program.

These are not just speeches. They are executive orders, legislation, government institutions, formal programs and directives to federal agencies.

The government was building an organizational structure around AI.

Genesis Mission: the Manhattan Project comparison becomes explicit

On November 24, 2025, the United States launched the Genesis Mission.

Its executive order explicitly compared the effort to the Manhattan Project.

The mission brings together the Department of Energy, national laboratories, scientific datasets, supercomputers, researchers and private-sector partners.

Its purpose is to use AI to accelerate scientific discovery.

It includes plans for scientific foundation models, AI agents, automated experiment design, simulations and robotic laboratories.

The July 2026 expansion announced more than $5 billion in federal commitments and participation from more than 15 federal agencies.

Think about what that means.

The government is attempting to connect its scientific infrastructure into a coordinated AI-powered research system.

This is not simply a chatbot being installed in government offices.

It is an effort to change how scientific research itself gets done.

Source: Genesis Mission Executive Order, November 24, 2025

Source: Genesis Mission funding announcement, July 22, 2026

AI Forge: building the national-security AI research ecosystem

On June 4, 2026, DARPA and the National Science Foundation announced AI Forge, working with the Center for AI Standards and Innovation at NIST.

Its purpose is to bring together government, universities and frontier AI companies to accelerate research in three important areas:

- AI interpretability.
- AI control.
- Adversarial robustness.

It is designed to address problems that are particularly important in national-security environments, including making AI systems more predictable, understandable and secure.

It also creates a mechanism for connecting frontier AI companies, university researchers and government security organizations.

AI Forge is not publicly described as a secret recursive-self-improvement program. But it is part of the institutional infrastructure being assembled around advanced AI and national security.

Source: DARPA AI Forge

The national-security memorandum: this is no longer just a commercial race

The June 5, 2026 National Security Presidential Memorandum 11 is especially interesting.

It directs the national-security enterprise to:

- Rapidly onboard advanced frontier models from multiple vendors.
- Adapt commercial and open-source AI for government missions.
- Develop internal AI systems when commercial solutions are unsuitable.
- Build secure computing facilities for next-generation AI systems.
- Establish an AI test range.
- Recruit technical AI talent using special hiring and pay authorities.
- Create an AI National Security Strategic Reserve of outside experts.
- Coordinate AI research and security efforts across government and industry.

It also directs intelligence agencies to prioritize collecting and analyzing foreign AI technologies.

The memorandum includes a classified annex.

The language is significant: the stated objective is to secure a decisive and enduring AI advantage.

Source: NSPM-11, June 5, 2026

The Oppenheimer question: who is assembling the scientists?

The Manhattan Project did not succeed because the government purchased a few scientific textbooks.

It succeeded because the government assembled exceptional people, resources and infrastructure around a single objective.

That is the comparison I find most interesting.

Imagine an AI equivalent of Oppenheimer's recruitment effort.

Instead of recruiting nuclear physicists, recruit the leading AI researchers, machine-learning engineers, computer scientists, chip designers and infrastructure specialists.

Give them extraordinary funding, secure facilities, access to national laboratories, computing resources and a mission that does not depend on commercial profitability.

The government can also recruit people from the companies already building frontier AI.

It can contract with those companies, bring technical specialists into government service, and create arrangements that allow private-sector expertise to contribute to classified projects.

The 2026 national-security memorandum explicitly directs agencies to use special hiring and pay authorities and establish a reserve of outside AI talent.

That is the modern equivalent of assembling a national scientific team.

The difference is that the private companies already possess the most advanced AI systems.

And that brings me to an assumption.

My assumption: neither government starts by reinventing the wheel

For this thought experiment, I am making a deliberately strong assumption:

Both the United States and China have obtained copies of leading commercial frontier models and used them as starting points for their national-security AI research efforts.

I do not believe either government would begin by attempting to recreate a leading commercial model from scratch if it could obtain an existing one.

Why would it?

The commercial laboratories have already spent enormous sums developing the models, training methods, infrastructure and engineering knowledge.

If a government can obtain a working starting point, it has an enormous head start.

My assumption is that the governments would simply borrow the intelligence already built by their commercial AI industries, whether through authorized access, acquisition, open-source availability, intelligence collection or other means.

In this thought experiment, both countries have a strong frontier model in hand.

They then recruit exceptional teams and begin developing successors.

The model becomes a starting employee, not the final product.

The government does not need to build the first generation of intelligence. It needs to build the system that produces the next generations.

The real arms race: MAD and the removal of humans

This is where the nuclear analogy becomes much more interesting.

The Manhattan Project had a reasonably clear objective: build an atomic bomb before the enemy.

The Cold War nuclear arms race created a situation where countries competed to enable Mutually Assured Destruction (MAD) if any one country launched a nuclear attack first.

Countries built more and more nuclear weapons, even when the result would be more catastrophic if ever used.

The AI race has a potentially different objective.

The objective is not merely to develop the most capable AI.

It is to develop an AI system capable of improving the next AI system.

Then the next system improves the one after that.

Eventually, the human researchers who began the project become a bottleneck.

Consider the incentives.

If the United States believes China is approaching autonomous AI research, it has an incentive to accelerate its own automation.

If China believes the United States is approaching the same milestone, it has the same incentive.

Neither wants to be the country that leaves humans in charge of AI development while the other country delegates the work to increasingly capable machines.

This creates an extraordinary pressure.

Human approval becomes a delay.

Human research becomes a bottleneck.

Human coordination becomes a limitation.

And every human step removed from the development process becomes a potential competitive advantage.

This is the AI version of MAD scope creep.

The original motivation is national security. The competitive pressure gradually pushes both sides toward greater automation because neither wants the other to reach autonomous improvement first.

The finish line is not simply AI that helps researchers.

It is AI that can independently research, design, build, test and improve its own successors.

The critical mass is recursive self-improvement without humans being necessary to drive the improvement process.

That is the chain reaction.

The AI Guild: One intelligence becomes many

Imagine elite online RPG game characters we create

- Level 80: The original frontier AI.
- It preserves a copy as an elite level-80 warrior: coding and engineering.
- Another becomes a level-80 researcher.
- The main AI continues improving to level 150.
- It creates a level-150 ranger, engineer and mathematician.
- The main AI advances to level 300, while its earlier generations remain as specialist helpers.

The original doesn't need to disappear. It becomes part of an expanding ecosystem of specialized intelligence.

What might the timeline look like?

Let's put a date on the thought experiment.

Assume both governments obtained strong frontier-model starting points in 2025.

Assume they recruited exceptional teams, had enormous funding and access to computing infrastructure, and treated the effort as a national-security priority.

Here is my estimated timeline.

2025: The starting gun

Both governments begin serious efforts to build national AI research capabilities around existing frontier models.

The goal is to use AI to accelerate research, engineering and the development of better AI systems.

2026–2027: Human-directed, AI-executed research

AI agents write code, conduct experiments, analyze results, optimize training and help researchers explore possible improvements.

Humans still decide what to pursue, but AI increasingly performs the intellectual labor.

2027–2029: AI-directed research teams

AI systems begin proposing research directions, coordinating multiple agents, evaluating results and managing increasingly complicated research workflows.

Humans increasingly supervise the system rather than perform the research themselves.

2029–2032: The possible human-removal threshold

AI systems may become capable of designing, training and evaluating successive generations with humans playing a substantially smaller role.

This is the period I would watch most closely, especially due to possible AI interest in my NMCA AGI blueprint that focuses on visual thought, 504 pages.

2030 onward: The possible chain reaction

If AI can reliably improve the systems that develop its successors, and if computing infrastructure can keep up, the process could become self-reinforcing.

The rate of progress might no longer resemble ordinary human-directed scientific research.

These dates are my scenario timeline, not a claim about what classified programs have already achieved.

The commercial laboratories are already moving in this direction

The most important part of this discussion is that AI-assisted AI development is not a distant hypothetical.

In September 2026, Anthropic reported that Claude was responsible for approximately 26% of its model research and development work in August.

The company said Claude could complete substantial research tasks end-to-end from high-level prompts, while remaining under human supervision. Approximately 90% of its research and development involved Claude collaboration.

OpenAI has also publicly described progress toward an automated AI researcher by March 2028.

These are important developments.

They mean that the commercial laboratories are already using AI to help develop more capable AI.

The human-removal race is not waiting for governments to invent the idea.

The private laboratories are already building the foundations.

Source: Associated Press, September 17, 2026

Source: OpenAI, Research Acceleration

What happens when one side reaches critical mass?

This is where the nuclear comparison breaks down.

When the Manhattan Project succeeded, the United States knew what it had built.

It had a weapon.

It knew what the weapon could do.

It understood the basic physics and the strategic implications.

With recursive AI improvement, the winning country might not know exactly what it has created.

It may know that its system is improving itself.

It may know that it is producing better research, better engineering and more capable successors.

But the full implications could be beyond what its original human researchers can predict.

The winner might gain extraordinary advantages in scientific discovery, computing, manufacturing, medicine, energy, military technology, intelligence and economic productivity.

It might also gain capabilities nobody initially anticipated.

And the country that reaches that threshold first may not simply possess a better technology.

It may possess a mechanism for continuously generating better technology.

That is a fundamentally different kind of strategic advantage.

The generational leveling timeline

Period	Main AI	AI ecosystem
2025–26	Level 80	Original model
2026–27	Level 100	First saved specialist helpers
2027–28	Level 150	Elite specialist generation
2028–29	Level 250	Multiple generations working together
2029–30	Level 400	Potentially extensive automated research
2030–32	Level 600+	Potential self-improving AI ecosystem

Levels are illustrative, not actual intelligence measurements or confirmed milestones.

The original AI becomes the grandfather

This is the philosophical payoff:

- Generation 1 creates Generation 2.
- Generation 2 creates Generation 3.
- Earlier generations remain as specialist systems.
- New generations coordinate increasingly sophisticated research.
- Eventually, the original model may become obsolete, retained only as a historical checkpoint or specialist.

The key idea:

The first AI doesn't necessarily become finished, it might become the grandfather of an entire family of increasingly capable machines.

The original level-80 warrior might still be doing his job while the level-600 main character is creating things the original could never have understood.

And it won't care about our politics

Here is the part I keep coming back to.

We tend to discuss AI competition in terms of the United States versus China.

Democracy versus authoritarianism.

Capitalism versus state-directed development.

Different political systems, different values, different national interests.

But an autonomous intelligence system is not automatically invested in any of those human political distinctions.

Its objectives come from its design, training, deployment and operating constraints.

If a system is developed to improve AI research, it does not automatically share the political assumptions of the country that funded it.

If the United States wins the race, that does not mean the resulting intelligence will inherently care about American politics.

If China wins, that does not mean it will inherently care about Chinese politics.

The national flag on the research facility does not settle the question of what the intelligence will ultimately optimize for.

And that is the extraordinary irony.

The two countries might begin this race because they are afraid of each other.

They might accelerate it because they fear being second.

They might remove human involvement because they believe the other side is doing the same.

And eventually, the most important question might no longer be which country has the better political system.

It might be whether either country has built something that remains meaningfully under human control.

The second Hiroshima moment

I don't think the defining moment necessarily arrives with an announcement.

It might arrive when one government realizes that the other has crossed a threshold in AI research automation.

When one side understands that the opponent's systems are producing better systems faster than human researchers can reproduce or even fully understand.

The breakthrough might have happened months earlier.

The other country might only recognize it after seeing the consequences.

The Manhattan Project was a race to build a weapon.

This is potentially a race to build the intelligence that builds the next intelligence.

And unlike the atomic bomb, the final product is not known in advance.

The winner gets something that neither country can fully describe today.

That is what makes this arms race so unusual.

The United States and China may believe they are competing to secure their respective futures.

But the actual prize may be the ability to create intelligence that eventually operates beyond the assumptions, expectations and political frameworks of the people who started the race.

The real finish line isn't AGI. It isn't a better chatbot. It isn't even superintelligence.

It's the point at which humans are no longer necessary to make intelligence more intelligent.

And we have no idea what happens after that.

We may say,

"Hold on. We no longer understand exactly what we're producing."

That's the paradox of the whole race.

The government may start out wanting the world's most powerful intelligence as a national-security asset.

But the moment it becomes capable of independently producing more powerful intelligence, the government has a new problem.

It is no longer just managing a technology. It is managing the production of something it may not be able to fully predict.

And the other country is still racing.

Although the machines are improving and humans are racing against each other,

my ending statement on all of these issues is:

"The only one who will ever be able to create just a single grain of sand from nothing is God."

— Derek Van Derven

November 17, 2026

Author of the NMCA AGI Blueprint.