Manhattan Project 2
The Genesis Mission

“this time it’s biblical.”

Derek Van Derven
Author of the NMCA AGI Blueprint
November 17, 2026

Contents

1	Essay: Manhattan Project 2 — The Genesis Mission *Corrected and updated, with China’s parallel moves and the visual-thought architecture*	2–15
2	Summary: Genesis, the Manhattan Project of the AI Age? *The record: who said what, and when*	16–19
3	Source document: Executive Order 14363 *The White House page, saved Nov. 17, 2026*	20–27
4	Sources *Numbered list with links*	28–29

The essay is analysis and scenario. The summary and the executive order are the documentary record. Every dated claim is listed with a link in the sources.

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.

The order set a 270-day target (about August 21, 2026) for initial operating capability. In July 2026, the Department of Energy gave the first public demonstration of the Genesis platform, and the White House announced more than $5 billion in federal commitments. No official statement I have found formally declares the 270-day milestone met.

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 1, 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 provides for a classified annex, to be issued within 90 days.

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

Source: NSPM-11, June 5, 2026

China's side of the ledger, 2025–2026

China has been moving in parallel. Its State Council issued the “AI+” plan in August 2025, and in December 2025 it launched an AI agent system on its national supercomputing network to run research tasks with little human oversight. In March 2026 its legislature approved the 15th Five-Year Plan, which stresses AI and robotics breakthroughs. In September 2026, researchers from ByteDance, Tsinghua University and the Shanghai Artificial Intelligence Laboratory published a five-stage path toward recursive self-improvement. Press coverage of that paper quoted experts who see U.S. firms as several months ahead, mainly on computing power.

Sources: China State Council (Aug. 26, 2025); South China Morning Post (Jan. 2 and Sept. 16, 2026); Congressional Research Service (Apr. 2026). Links in the source list.

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 (v5, October 16, 2026; zenodo.org/records/23417247).

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.

What could slow this down

Not everyone agrees this scenario holds. One recent economics study (Whitfill and Wu, 2025) found that whether compute limits can stop AI-driven AI research depends on whether compute and human labor are substitutes, and its own two models gave conflicting answers. Epoch AI's earlier analysis found that training runs far larger than today's are technically feasible by 2030, but that power would be the first constraint, then chip manufacturing, at a cost that could reach hundreds of billions of dollars. Progress could slow, or arrive unevenly across fields.

What today's AI actually is, and what happens when a machine thinks like humans do

What they're built on. Today's leading systems are neural networks trained to predict what comes next, in text, images, audio and video, then refined with human feedback. Their working memory is whatever sits in the context window. Their strength is fluency, and it scales with data and computing power, which is why the race is framed as a contest over chips, power and money.

Where it falls short. The rose page states the problem with one scene. Tell a system to build a blue vase and a red apple on a wooden table. Delete the description. Wait. Then ask where the apple is relative to the vase, change the viewpoint, move an object toward the system, take another entity's perspective and return. A system that only retrieves the stored sentence fails. The object has to still be in a place. Today's systems hold the sentence in context, retrieve it, or update a short-lived latent state, and fluency hides the difference. The same page names the walls: scaling, agents, world models, memory banks, hybrids and robotics may each keep improving while failing this probe.

How humans think. People build an inner scene from a few words and then walk around in it. Einstein chased a beam of light. Tesla assembled his machines in his head. Newton's cannon fires from a mountain. Darwin drew a tree. Humans don't render a photorealistic room. They keep a usable spatial model and bring parts into focus when needed. People with aphantasia, who have little or no visual imagery, and people who report few or no dreams, are exceptions to the rule, and they reason in other ways. The rule across human minds is a spatial scene with symbols attached to its objects.

What NMCA does. NMCA builds the machine the same way. Implicit visual simulation is the substrate of thought, with symbols pegged to the entities in the scene, and nothing is rendered unless someone asks to see it. Embodiment starts inside that scene: an avatar trains on location, reach, hiding and return, then the same operations transfer to a real robot. In the page's example, an agent recalls where the refrigerator is, simulates the route, acts, finds that reality differs from its simulation, and generates a new subgoal. The scientist test asks for a new, unsolved problem and scores only whether the system solves it, using a visual or spatial working model after search and verbal reasoning fail. The full design is the 504-page blueprint, and the case is on the rose page.

What happens when they use it. A system that reasons over a stable scene is grounded, inspectable and cheaper to run, because the scene is working memory, not output. It would also automate research faster, so the country that gets there first holds the advantage this essay describes. The architecture doesn't end the race. It decides which capability wins it.

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.

GENESIS: THE MANHATTAN PROJECT OF THE AI AGE?

The United States government has formally invoked the Manhattan Project to describe its new national AI mission.

On November 24, 2025, President Donald J. Trump signed **Executive Order 14363, "Launching the Genesis Mission."**

This was not simply a speech about America's ambitions in artificial intelligence. It was a presidential executive order establishing a coordinated national effort, led by the Department of Energy, to accelerate scientific discovery through artificial intelligence.

And the historical comparison was deliberate.

1. The President's Executive Order — November 24, 2025

In Section 1, titled *Purpose*, the executive order states:

> *"The challenges we face require a historic national effort, comparable in urgency and ambition to the Manhattan Project..."*

The order establishes Genesis as a national initiative to build an integrated AI platform combining federal scientific datasets, national laboratory computing resources, scientific foundation models and AI agents.

The stated objectives include:

- Automating research workflows.
- Testing scientific hypotheses using AI agents.
- Accelerating scientific breakthroughs.
- Strengthening national security.
- Securing American energy leadership.
- Maintaining American technological dominance.

The order places the Department of Energy in charge of implementation, with White House-level coordination across participating federal agencies.

Who signed it: President Donald J. Trump.

Document: Executive Order 14363, November 24, 2025.

Official source: https://www.whitehouse.gov/presidential-actions/2025/11/launching-the-genesis-mission/

2. The Secretary of Energy — November 24, 2025

On the same day, Energy Secretary Chris Wright publicly reinforced the historical comparison.

He stated:

> *"Throughout history, from the Manhattan Project to the Apollo mission, our nation's brightest minds and industries have answered the call when their nation needed them."*

Wright described Genesis as an effort to mobilize America's national laboratories, supercomputers and scientific resources to establish global AI leadership.

This is significant because the comparison was not confined to the presidential order. The department responsible for executing the mission was also invoking the Manhattan Project as its historical precedent.

Who: Chris Wright, U.S. Secretary of Energy.

When: November 24, 2025.

Official source: https://www.energy.gov/articles/energy-department-launches-genesis-mission-transform-american-science-and-innovation

3. The Genesis Mission Director — December 1, 2025

Perhaps the most revealing statement came one week later.

Darío Gil, Under Secretary for Science at the Department of Energy and newly appointed Director of the Genesis Mission, issued a letter to the Department of Energy and national laboratory community.

Rather than simply repeating the historical analogy, Gil explained why Genesis would need a different organizational structure.

He acknowledged that the original Manhattan Project was conducted through federal government direction and secrecy.

Genesis, he explained, would require something different: an unprecedented partnership involving the federal government, national laboratories, private industry and universities.

He pointed out that approximately two-thirds of American research and development investment comes from the private sector, making private-sector participation essential to the mission.

Then came his statement about urgency:

> *“The one thing we don't have is time.”*

Gil explicitly identified China as America's most formidable scientific and technological competitor and concluded:

> *“This is a race we must and will win.”*

He described the scientific and engineering advances ahead in AI, quantum computing, fusion and biotechnology as technologies that would define the future of the country and the world.

He also suggested that the work could define his generation's legacy.

Who: Darío Gil, Under Secretary for Science and Genesis Mission Director.

When: December 1, 2025.

Official source: https://content.govdelivery.com/accounts/USDOEOS/bulletins/3fdb894

4. What the government actually committed to doing

The executive order goes beyond historical comparisons and establishes concrete organizational and operational requirements.

Among its directives:

- Mobilize the Department of Energy's 17 national laboratories.
- Integrate federal scientific datasets into a unified AI platform.
- Develop scientific foundation models and AI agents.
- Identify at least 20 national science and technology challenges.
- Assess robotic laboratories and facilities capable of AI-directed experimentation.
- Seek to demonstrate an initial operating capability within 270 days, subject to applicable law and available appropriations.

The targeted scientific fields include advanced manufacturing, biotechnology, critical materials, nuclear fission and fusion, quantum information science, semiconductors and microelectronics.

The Department of Energy also announced an ambition to double the productivity and impact of American science and engineering within a decade.

These are measurable institutional objectives, not merely aspirational language.

5. Why the Manhattan Project comparison matters

The original Manhattan Project demonstrated what could happen when the federal government concentrated extraordinary scientific, industrial, financial and organizational resources around a single strategic objective.

Genesis is not the same project. Its declared purpose is scientific discovery and technological acceleration, rather than the development of a particular weapon.

But the historical analogy is meaningful.

The government is describing AI-powered scientific discovery as a matter of national strategic importance, placing the Department of Energy and its national laboratories at the center of a coordinated initiative, and explicitly identifying international technological competition as a reason for urgency.

There is also an important difference in scale and structure: Genesis intends to mobilize not just government laboratories, but the private AI industry and the wider scientific ecosystem.

The executive order does not, by itself, guarantee that every proposed activity will be funded or completed. Its implementation is subject to applicable law and available appropriations.

Nevertheless, the combination of a presidential executive order, departmental leadership, national laboratory mobilization, defined deadlines and explicit international competition represents a substantial formal commitment.

The Manhattan Project comparison did not originate with outside commentators trying to dramatize AI. The President put it into the founding document of the mission itself.

And the people appointed to execute that mission subsequently reinforced the comparison, explained its urgency and described the scale of national mobilization they believe is necessary.

That is worth taking seriously.

PRESIDENTIAL ACTIONS

LAUNCHING THE GENESIS MISSION

Executive Orders

November 24, 2025

EXECUTIVE ORDER 14363

By the authority vested in me as President by the Constitution and the laws of the United States of America, it is hereby ordered:

Section 1. Purpose. From the founding of our Republic, scientific discovery and technological innovation have driven American progress and prosperity. Today, America is in a race for global technology dominance in the development of artificial intelligence (AI), an important frontier of scientific discovery and economic growth. To that end, my Administration has taken a number of actions to win that race, including issuing multiple Executive Orders and implementing America's AI Action Plan, which recognizes the need to invest in AI-enabled science to accelerate scientific advancement. In this pivotal moment, the challenges we face require a historic national effort, comparable in urgency and ambition to the Manhattan Project that was instrumental to our victory in World War II and was a critical basis for the foundation of the Department of Energy (DOE) and its national laboratories.

This order launches the "Genesis Mission" as a dedicated, coordinated national effort to unleash a new age of AI-accelerated innovation and discovery that can solve the most challenging problems of this century. The Genesis Mission will build an integrated AI platform to harness Federal scientific datasets — the world's largest collection of such datasets, developed over decades of Federal investments — to train scientific foundation models and create AI agents to test new hypotheses, automate research workflows, and accelerate scientific breakthroughs. The G
bring together our Nation's research and development resources — combining

Click to listen

American scientists, including those at our national laboratories, with pioneering American businesses; world-renowned universities; and existing research infrastructure, data repositories, production plants, and national security sites — to achieve dramatic acceleration in AI development and utilization. We will harness for the benefit of our Nation the revolution underway in computing, and build on decades of innovation in semiconductors and high-performance computing. The Genesis Mission will dramatically accelerate scientific discovery, strengthen national security, secure energy dominance, enhance workforce productivity, and multiply the return on taxpayer investment into research and development, thereby furthering America's technological dominance and global strategic leadership.

Sec. 2. Establishment of the Genesis Mission. (a) There is hereby established the Genesis Mission (Mission), a national effort to accelerate the application of AI for transformative scientific discovery focused on pressing national challenges.

(b) The Secretary of Energy (Secretary) shall be responsible for implementing the Mission within DOE, consistent with the provisions of this order, including, as appropriate and authorized by law, setting priorities and ensuring that all DOE resources used for elements of the Mission are integrated into a secure, unified platform. The Secretary may designate a senior political appointee to oversee day-to-day operations of the Mission.

(c) The Assistant to the President for Science and Technology (APST) shall provide general leadership of the Mission, including coordination of participating executive departments and agencies (agencies) through the National Science and Technology Council (NSTC) and the issuance of guidance to ensure that the Mission is aligned with national objectives.

Sec. 3. Operation of the American Science and Security Platform. (a) The Secretary shall establish and operate the American Science and Security Platform (Platform) to serve as the infrastructure for the Mission with the purpose of providing, in an integrated manner and to the maximum extent practicable and consistent with law:

(i) high-performance computing resources, including DOE national laboratory supercomputers and secure cloud-based AI computing environments, capable of supporting large-scale model training, simulation, and inference;

(ii) AI modeling and analysis frameworks, including AI agents to explore design spaces, evaluate experimental outcomes, and automate workflows;

(iii) computational tools, including AI-enabled predictive models, simulation models, and design optimization tools;

(iv) domain-specific foundation models across the range of scientific domains covered;

(v) secure access to appropriate datasets, including proprietary, federally curated, and open scientific datasets, in addition to synthetic data generated through DOE computing resources, consistent with applicable law; applicable classification, privacy, and intellectual [illegible]ations; and Federal data-access and data-management standards; and

Click to listen

(vi) experimental and production tools to enable autonomous and AI-augmented experimentation and manufacturing in high-impact domains.

(b) The Secretary shall take necessary steps to ensure that the Platform is operated in a manner that meets security requirements consistent with its national security and competitiveness mission, including applicable classification, supply chain security, and Federal cybersecurity standards and best practices.

(c) Within 90 days of the date of this order, the Secretary shall identify Federal computing, storage, and networking resources available to support the Mission, including both DOE on-premises and cloud-based high-performance computing systems, and resources available through industry partners. The Secretary shall also identify any additional partnerships or infrastructure enhancements that could support the computational foundation for the Platform.

(d) Within 120 days of the date of this order, the Secretary shall:

(i) identify a set of initial data and model assets for use in the Mission, including digitization, standardization, metadata, and provenance tracking; and

(ii) develop a plan, with appropriate risk-based cybersecurity measures, for incorporating datasets from federally funded research, other agencies, academic institutions, and approved private-sector partners, as appropriate.

(e) Within 240 days of the date of this order, the Secretary shall review capabilities across the DOE national laboratories and other participating Federal research facilities for robotic laboratories and production facilities with the ability to engage in AI-directed experimentation and manufacturing, including automated and AI-augmented workflows and the related technical and operational standards needed.

(f) Within 270 days of the date of this order, the Secretary shall, consistent with applicable law and subject to available appropriations, seek to demonstrate an initial operating capability of the Platform for at least one of the national science and technology challenges identified pursuant to section 4 of this order.

Sec. 4. Identification of National Science and Technology Challenges. (a) Within 60 days of the date of this order, the Secretary shall identify and submit to the APST a detailed list of at least 20 science and technology challenges of national importance that the Secretary assesses to have potential to be addressed through the Mission and that span priority domains consistent with National Science and Technology Memorandum 2 of September 23, 2025, including:

(i) advanced manufacturing;

(ii) biotechnology;

(iii) critical materials;

(iv) nuclear fission and fusion energy;

(v) quantum information science; and

(vi) semiconductors and microelectronics.

Click to listen

(b) Within 30 days of submission of the list described in subsection (a) of this section, the APST shall review the proposed list and, working with participating agency members of the NSTC, coordinate the development of an expanded list that can serve as the initial set of national science and technology challenges to be addressed by the Mission, including additional challenges proposed by participating agencies through the NSTC, subject to available appropriations.

(c) Following development of the expanded list described in subsection (b) of this section, agencies participating in the Mission shall use the Platform to advance research and development aligned with the national science and technology challenges identified in the expanded list, consistent with applicable law and their respective missions, and subject to available appropriations.

(d) On an annual basis thereafter, the Secretary shall review and update the list of challenges in consultation with the APST and the NSTC to reflect progress achieved, emerging national needs, and alignment with my Administration's research and development priorities.

Sec. 5. Interagency Coordination and External Engagement. (a) The APST, through the NSTC, and with support from the Federal Chief Data Officer Council and the Chief AI Officer Council, shall convene relevant and interested agencies to:

(i) assist participating agencies in aligning, to the extent permitted by law, their AI-related programs, datasets, and research and development activities with the objectives of the Mission in their respective areas of expertise, while avoiding duplication of effort across the Federal Government and promoting interoperability;

(ii) identify data sources that may support the Mission's aim;

(iii) develop a process and resourcing plan in coordination with participating agencies for integrating appropriate and available agency data and infrastructure into the Mission, to the extent permitted by law and subject to available appropriations, including methods under which all agencies contributing to the Mission are encouraged to implement appropriate risk-based security measures that reflect cybersecurity best practices;

(iv) launch coordinated funding opportunities or prize competitions across participating agencies, to the extent permitted by law and subject to available appropriations, to incentivize private-sector participation in AI-driven scientific research aligned with Mission objectives; and

(v) establish mechanisms to coordinate research and development funding opportunities and experimental resources across participating agencies, ensuring agencies can participate effectively in the Mission.

(b) The APST shall coordinate with relevant agencies in establishing, consistent with existing authorizing statutes and subject to available appropriations, competitive programs for research fellowships, internships, and apprenticeships focused on the application of AI to scientific domains identified as national challenges for the Mission, to include placement of program participants at DOE national laboratories and other participating Federal research facilities, with the purpose of providing access to the Platform and training in AI-enabled scientific discovery.

Click to listen

(c) The Secretary, in coordination with the APST and the Special Advisor for AI and Crypto, shall establish mechanisms for agency collaboration with external partners possessing advanced AI, data, or computing capabilities or scientific domain expertise, including through cooperative research and development agreements, user facility partnerships, or other appropriate arrangements with external entities to support and enhance the activities of the Mission, and shall ensure that such partnerships are structured to preserve the security of Federal research assets and maximize public benefit. To facilitate these collaborations, the Secretary shall:

(i) develop standardized partnership frameworks, including cooperative research and development or other appropriate agreements, and data-use and model-sharing agreements;

(ii) establish clear policies for ownership, licensing, trade-secret protections, and commercialization of intellectual property developed under the Mission, including innovations arising from AI-directed experiments;

(iii) implement uniform and stringent data access and management processes and cybersecurity standards for non-Federal collaborators accessing datasets, models, and computing environments, including measures requiring compliance with classification, privacy, and export-control requirements, as well as other applicable laws; and

(iv) establish procedures to ensure the highest standards of vetting and authorization of users and collaborators seeking access to the resources of the Mission and associated research activities, including the Platform and associated Federal research resources.

(d) The APST, through the NSTC, shall, to the extent appropriate, identify opportunities for international scientific collaboration to support activities under the Mission.

Sec. 6. Evaluation and Reporting. (a) Within 1 year of the date of this order, and on an annual basis thereafter, the Secretary shall submit a report to the President, through the APST and the Director of the Office of Management and Budget, describing:

(i) the Platform's operational status and capabilities;

(ii) progress toward integration across DOE national laboratories and other participating Federal research partners, including shared access to computing resources, data infrastructure, and research facilities;

(iii) the status of user engagement, including participation of student researchers and any related training;

(iv) updates on research efforts and outcomes achieved, including measurable scientific advances, publications, and prototype technologies;

(v) the scope and outcomes of public-private partnerships, including collaborative research projects and any technology transitions or commercialization activities; and

(vi) any identified needs or recommendations for authorities or interagency support to achieve the Mission's objectives.

Sec. 7. General Provisions. (a) Nothing in this order shall be construed to in ct:

(i) the authority granted by law to an executive department or agency, or the head thereof, or

Click to listen

(ii) the functions of the Director of the Office of Management and Budget relating to budgetary, administrative, or legislative proposals.

(b) This order shall be implemented consistent with applicable law and subject to the availability of appropriations.

(c) This order is not intended to, and does not, create any right or benefit, substantive or procedural, enforceable at law or in equity by any party against the United States, its departments, agencies, or entities, its officers, employees, or agents, or any other person.

(d) The costs for publication of this order shall be borne by the Department of Energy.

DONALD J. TRUMP

THE WHITE HOUSE,

November 24, 2025.

DOWNLOAD

Related

Fact Sheet: President Donald J. Trump Unveils the Genesis Mission to Accelerate AI for Scientific Discovery

Fact Sheets | November 24, 2025

President Trump Launches the Genesis Mission to Accelerate AI for Scientific Discovery

Releases | November 24, 2025

Trump Administration Announces More Than $5 Billion for the Genesis Mission, a National Mission on AI for Science

Releases | July 22, 2026

Unlocking Cures for Pediatric Cancer with Artificial Intelligence

Presidential Actions, Executive Orders | September 30, 2025

OSTP Director Releases Landmark Report and Recommendations for Renewing American Scientific Discovery

Releases | July 21, 2026

Click to listen



WHWIRE

GET THE FACTS →

Click to listen

SUBSCRIBE TO THE WH NEWSLETTER

Your email SIGN UP

Click here or text 45470 to receive updates

DOWNLOAD THE WHITE HOUSE APP:

Available On The App Store | Get It On Google Play

GET UPDATES STRAIGHT TO YOUR DESKTOP

AVAILABLE ON

Download Browser Extension

THE WHITE HOUSE

JOIN NOW

Click to listen

Sources

Numbered in the order the topics appear in the packet. Where a figure is marked "as reported," it comes from the linked coverage.

1. Executive Order 14363, "Launching the Genesis Mission," Nov. 24, 2025.
   https://www.whitehouse.gov/presidential-actions/2025/11/launching-the-genesis-mission/

2. U.S. Department of Energy, Secretary Wright statement launching the Genesis Mission, Nov. 24, 2025.
   https://www.energy.gov/articles/energy-department-launches-genesis-mission-transform-american-science-and-innovation

3. Darío Gil, letter to the DOE and national laboratory community, Dec. 1, 2025.
   https://content.govdelivery.com/accounts/USDOEOS/bulletins/3fdb894

4. White House, "Trump Administration Announces More Than $5 Billion for the Genesis Mission," July 22, 2026.
   https://whitehouse.gov/releases/2026/07/45502

5. CDO Magazine, DOE unveils Genesis Mission platform demo at technology summit (updated July 30, 2026).
   https://www.cdomagazine.tech/us-federal-news-bureau/doe-unveils-genesis-mission-platform-demo-at-technology-summit

6. FedScoop, Energy Department Genesis Mission MVP progress, Feb. 23, 2026.
   https://fedscoop.com/energy-department-genesis-mission-mvp-progress/

7. DARPA, "AI Forge: Accelerating AI Breakthroughs for National Security," June 1, 2026.
   https://www.darpa.mil/news/2026/ai-forge-accelerating-ai-breakthroughs-national-security

8. National Security Presidential Memorandum 11 (NSPM-11), June 5, 2026.
   https://www.whitehouse.gov/presidential-actions/2026/06/national-security-presidential-memorandum-nspm-11/

9. State Council of China, guideline on the "AI+" action, Aug. 26, 2025.
   https://english.www.gov.cn/policies/latestreleases/202508/27/content_WS68ae7976c6d0868f4e8f51a0.html

10. South China Morning Post via The Star, China launches AI science system to take on Genesis Mission, Jan. 2, 2026.
    https://www.thestar.com.my/tech/tech-news/2026/01/02/china-launches-super-powered-ai-science-system-to-take-on-donald-trumps-genesis-mission

11. Congressional Research Service, China's 15th Five-Year Plan: S&T and Economic Priorities, Apr. 2026.
    https://www.everycrsreport.com/files/2026-04-16_IF13204_53d4be1a72611630d91a643640c76ba3ff4b4b39.html

12. South China Morning Post via The Star, Chinese researchers chart 5-stage path toward 'last AI built by humans,' Sept. 16, 2026.
    https://www.thestar.com.my/aseanplus/aseanplus-news/2026/09/16/chinese-researchers-chart-5-stage-path-toward-last-ai-built-by-humans

13. Associated Press, Anthropic says its model Claude is helping to build the next version of itself, Sept. 17, 2026 (as syndicated). Figures are as reported.
https://www.local10.com/tech/2026/09/17/anthropic-says-its-model-claude-is-helping-to-build-the-next-version-of-itself/

14. OpenAI, “Research Acceleration”; reporting on the March 2028 automated-researcher target (The Decoder). Figures are as reported.
https://openai.com/index/research-acceleration-view-inside-openai/

15. Whitfill, P. and Wu, C., “Will Compute Bottlenecks Prevent an Intelligence Explosion?” arXiv:2507.23181, 2025.
https://arxiv.org/html/2507.23181v2

16. Epoch AI, “Can AI Scaling Continue Through 2030?”
https://epoch.ai/blog/can-ai-scaling-continue-through-2030

17. Van Derven, D., NMCA: A Neurosymbolic Multimodal Cognitive Architecture, 504-page edition (v5), Oct. 16, 2026.
https://zenodo.org/records/23417247

18. NMCA Visual-3D Architecture summary (the “rose” page).
https://visualthoughtagi.com/rose

19. China’s 2017 New Generation AI Development Plan, translated by Stanford DigiChina (cited in the essay).