



Catastrophic Risks: A Plan for Nuclear Weapons, Engineered Pathogens and Advanced AI, Seen from September 2026

Version 1.0 · Current as of September 2026 · specieslevel.com/plans/catastrophic-risks.html

This plan is an initial draft prepared with AI models. It is open for review by subject-matter experts and will be revised as they join the work and as world and national events change.

Introduction

Most of the problems on this site cause harm every day and can be counted in deaths per year. Nuclear war, a deliberately engineered pandemic or a loss of control over very capable artificial intelligence might not happen in any given year, but any one of them could kill hundreds of millions or billions of people in a short time, some irreversibly, and even a modest annual probability of harm on that scale justifies sustained effort.

Each of the three threats depends on a technology that relatively few actors control today and that is governed by rules weaker than the technology itself, and in each case the cost of misuse falls as the tools spread. They also bear directly on the other problems on this site: a nuclear exchange would cause a global famine, and a severe pandemic reverses years of progress against poverty.

This essay sets out a plan for reducing all three risks and tests it against the events of 2025 and 2026. Arms control between the nuclear powers, the expiry of New START and the new weapons in orbit are covered in [Ending War: A Plan Tested Against the World of September 2026](#) and are not repeated here.

Where things stand

Nuclear weapons. At the start of 2026 the nine nuclear-armed states held an estimated 12,187 warheads. About 9,745 were in military stockpiles for potential use, about 4,012 were deployed with missiles and aircraft, and between 2,100 and 2,200 were kept on high operational alert on ballistic missiles, mostly by Russia and the United States ([SIPRI](#)). China has about 620 warheads and is expanding its arsenal faster than any other country ([SIPRI](#)). The world total has generally fallen since the Cold War because Russia and the United States dismantled retired warheads faster than they deployed new ones, but SIPRI

expects that trend to reverse in the coming years as dismantlement slows and new deployments accelerate (SIPRI).

A 2022 study in *Nature Food* modelled the effect of soot from burning cities on crops, fisheries and livestock and found that more than 2 billion people could die after a nuclear war between India and Pakistan and more than 5 billion after one between the United States and Russia, most of them from famine (Xia et al., *Nature Food*, via NASA GISS). A report mandated by the US Congress and released by the National Academies in June 2025 found major uncertainties at each stage of the chain from detonation to climate and social effects, and recommended that US agencies coordinate model intercomparison projects to narrow them (National Academies).

Engineered pathogens. The Biological Weapons Convention, the treaty that bans these weapons, has no verification system. Its Implementation Support Unit was set up in 2006 with three staff and has had four since a temporary post was added for 2023 to 2027 (UNODA). In December 2024, 38 scientists, including some who had once regarded mirror life as a long-term research goal, warned in *Science* that “mirror” bacteria built from molecules of the opposite handedness to natural life would probably evade many immune defences and natural predators, and concluded that such organisms “should not be created” (Adamala et al., *Science*, accepted manuscript). The Global Preparedness Monitoring Board reported in May 2026 that outbreaks are becoming more frequent and more damaging (GPMB).

Advanced AI. The International AI Safety Report 2026, chaired by Yoshua Bengio and written with guidance from more than 100 experts, found that leading systems reached gold-medal performance on International Mathematical Olympiad questions and that AI agents can reliably complete some programming tasks that take a human about half an hour (International AI Safety Report). General-purpose systems can provide information about biological and chemical weapons development, including expert-level laboratory instructions, and in 2025 several developers released models with additional safeguards because they could not exclude that the models would help novices build such weapons. The report judged that current systems cannot yet pose a risk of loss of control, but noted that models increasingly distinguish test settings from real use (International AI Safety Report). Researchers disagree widely about the longer-term danger: in a survey of 2,778 authors of papers at leading AI venues, between 38 and 51 percent gave at least a 10 percent chance to outcomes as bad as human extinction (Grace et al.).

Part I: The plan

Elements 1 and 2 concern nuclear weapons, elements 3 to 5 biology, elements 6 and 7 AI, and element 8 funding.

1. Lengthen nuclear decision time and keep humans in charge

A standing nuclear danger is a launch caused by a false warning or a misread crisis, with leaders given minutes to decide. The plan asks the nuclear powers to take missiles off high alert where they can, to adopt launch procedures that allow more time to verify a warning, and to state formally that no automated system will make or recommend a launch decision without human judgment. The main obstacle is the belief that a slower posture invites a first strike.

In November 2024 the presidents of the United States and China affirmed “the need to maintain human control over the decision to use nuclear weapons” ([White House archive](#)), and section 1638 of the US National Defense Authorization Act for fiscal year 2025 made it US policy that AI should not compromise the integrity of nuclear safeguards ([Lieu and Markey letter](#)). The plan seeks a joint statement by the five permanent members of the Security Council, followed by technical exchanges on how each keeps AI out of launch decisions.

2. Keep the testing moratorium and prepare for the worst case

No country other than North Korea has conducted an explosive nuclear test this century; the United States last tested in 1992 and China in 1996 ([Fortune/AP](#); [OPB/NPR](#)). The informal moratorium limits the development of new warhead designs and removes a visible signal of escalation. The plan asks every nuclear power to state that it will not test, to support the International Monitoring System of the Comprehensive Nuclear-Test-Ban Treaty Organization (CTBTO), and to settle allegations of secret testing with evidence. Frank von Hippel of Princeton has proposed an independent National Academies review of the evidence and reciprocal transparency at test sites, including observers at subcritical experiments ([Bulletin of the Atomic Scientists](#)).

The plan also asks governments to treat a multi-year loss of sunlight and harvests, after a nuclear war or a very large volcanic eruption, as a planning scenario: to fund the modelling the National Academies recommended and to include a global crop failure in national food-security plans, alongside the grain and fertilizer reserves proposed in the [starvation essay](#).

3. Screen DNA orders and govern the riskiest experiments

Anyone trying to build a dangerous pathogen from scratch has to obtain its genetic material, and most synthetic DNA comes from commercial providers. Screening every order against a list of sequences of concern, and checking who the customer is, is among the cheapest controls available. Free, open-source screening software already exists in the Common Mechanism run by the International Biosecurity and Biosafety Initiative for Science ([IBBIS](#)), and the United Kingdom published screening guidance in October 2024 that covers providers, users and benchtop synthesisers ([GOV.UK](#)). The plan makes

screening a legal duty for all providers and makers of benchtop machines in major markets, under a shared international standard so that customers cannot move to an unscreened supplier. The Biosecurity Modernization and Innovation Act, introduced in the US Senate in January 2026 by Tom Cotton and Amy Klobuchar, offers a model: within a year of enactment the Commerce Department would have to require sellers of synthetic DNA and benchtop synthesisers to screen orders against a maintained list, verify customers and detect orders split across providers, with audits and civil penalties of up to \$750,000 for companies ([GovInfo](#)).

Research that deliberately makes pathogens more transmissible or more deadly needs a separate and stricter process: independent review before funding, disclosure of what is being done and where, and equivalent rules for privately funded laboratories. For mirror life the plan supports an international agreement not to create mirror organisms, in line with the *Science* authors' recommendation that research aimed at creating mirror bacteria not be permitted, together with the limit that a UK Government Office for Science meeting suggested on the size of mirror DNA, which would block a self-replicating mirror genome while allowing research on small mirror molecules ([GOV.UK](#)).

4. Give the Biological Weapons Convention a working staff and a science review

The treaty that bans biological weapons is served by four people whose posts cost \$1,229,300 in 2024, paid by states parties with no money from the regular UN budget ([BWC Implementation Support Unit](#)). The plan asks states parties to adopt the three mechanisms their own working group has drafted: an international cooperation and assistance mechanism, a science and technology review mechanism to track developments such as AI-aided design, and an open-ended working group on compliance and verification. The draft would also add three posts to the support unit, bringing it to seven ([UNODA](#)). Full verification is unlikely soon, because biological work is small, dual-use and easily hidden, but a regular scientific review and voluntary transparency visits are achievable now.

5. Detect new pathogens early and make vaccines fast

If prevention fails, speed of detection and response largely determines the death toll, so the plan funds genomic surveillance of wastewater, airports and hospitals capable of spotting a new pathogen, natural or engineered, before it spreads widely. It backs the goal of the Coalition for Epidemic Preparedness Innovations (CEPI) to have a safe, effective vaccine within 100 days of identifying a new pandemic threat; research commissioned by CEPI found that 8.3 million deaths could have been averted by the end of 2021 had COVID-19 vaccines been available within 100 days of the virus's genome being published ([CEPI](#)). It also asks governments to complete and ratify the WHO Pandemic Agreement, adopted in May 2025 by 124 votes to none with 11 abstentions. The agreement opens for

signature only after the World Health Assembly adopts an annex on pathogen access and benefit sharing (PABS), and it needs 60 ratifications to enter into force ([UN News](#)).

6. Test frontier AI before release and report what is found

Developers of the most capable AI systems should test them for dangerous capabilities before release, publish safety frameworks, report serious incidents to government and give government evaluators access. Parts of this are already common practice: twenty companies signed the Frontier AI Safety Commitments at the Seoul summit in 2024, which included publishing safety frameworks ([GOV.UK](#)), and in 2025 Anthropic and OpenAI applied stricter safeguards to new models because they could not rule out meaningful help with biological weapons ([Anthropic](#); [OpenAI](#)). The case for making the practice mandatory rests on evidence that voluntary commitments are uneven. The Future of Life Institute's AI Safety Index for summer 2026 graded nine developers, gave none a grade above C+, gave three an F, found no company above C- on planning for the most extreme risks, and concluded that "safety rhetoric outpaces revealed behavior" ([Future of Life Institute](#)).

Model safeguards and the DNA screening in element 3 are two layers of defence against the most concrete near-term danger from AI, help with biological weapons. On the Virology Capabilities Test, OpenAI's o3 model scored 43.8 percent, against an average of 22.1 percent for expert virologists answering questions in their own specialties ([SecureBio](#)).

7. Build shared scientific understanding and channels between rivals

Governments assess AI risks differently and so cannot yet agree on rules. The plan supports an international scientific body, modelled on the Intergovernmental Panel on Climate Change, that reports regularly on capabilities and risks, and a standing intergovernmental dialogue on its findings. It asks the United States and China, which build most frontier systems, to hold regular technical talks on AI safety, covering shared evaluation methods, incident notification and military uses, building on their 2024 agreement on human control of nuclear launch decisions.

8. Pay for the work

Private funding for nuclear risk reduction fell sharply after the MacArthur Foundation, the largest funder, decided in 2020 to leave the field. Between 2014 and 2020 its nuclear programme made 228 grants worth \$100.8 million ([MacArthur Foundation](#)), its final grants were made in 2023 ([MacArthur Foundation](#)), and one estimate puts total philanthropic spending on nuclear security at about \$35 million a year, down from about \$50 million ([80,000 Hours](#)). The plan asks governments and foundations to fund arms-control expertise, the proposed seven-person BWC support unit and the WHO's biosecurity

staff, and independent AI evaluation, all at sums far below the missile-defence budgets described in the war essay.

Sequencing

The first two years should go to steps that are cheap, reversible and within existing authority: mandatory DNA screening, pre-release testing of frontier AI, a five-power statement on human control of nuclear launches and continued observance of the testing moratorium. Over two to ten years the aim is treaty-level work: a strengthened Biological Weapons Convention with a science review, a ratified pandemic agreement and international standards for AI evaluation. Over the longer term the goal is a verification system for biological weapons and binding rules for the most capable AI systems, both of which depend on trust built in the earlier stages.

What makes it stick politically

Most elements serve each country's own security, and measures presented that way and backed by defence and intelligence agencies tend to outlast changes of government. DNA providers and AI developers that already screen and test have a commercial interest in rules that make competitors do the same, and published test results and incident reports let legislatures and the public check that commitments are kept.

Where reasonable people disagree

Estimates of how large and how near these risks are range from negligible to very high. Some critics argue that attention to speculative AI harms draws money from present ones such as fraud and job loss, while others reply that the worst outcomes are large enough to justify spending even at low probabilities.

Supporters of rapid AI development argue that regulation will hand the lead to less careful rivals, especially China; supporters of regulation argue that a race without safety testing raises the risk for everyone, including the winner. Open publication of research, pathogen data and model weights aids science and scrutiny as well as potential misusers. Deterrence theorists argue that keeping some nuclear weapons on high alert makes a first strike pointless, while critics regard high alert as the main source of accidental risk.

Part II: The plan in the world of September 2026

Nuclear decision time and testing (elements 1 and 2)

In October 2025 President Trump said he had instructed the military “to start testing our Nuclear Weapons on an equal basis” with other countries. Energy Secretary Chris Wright said days later that the planned tests were “not nuclear explosions” but “noncritical explosions” of weapon systems ([Fortune/AP](#)). President Putin responded by ordering

officials to prepare proposals for possible test preparations, while saying Russia would test only if the United States did ([CSIS](#)).

In February 2026 Under Secretary of State Thomas DiNanno told the Conference on Disarmament in Geneva that China had conducted nuclear explosive tests concealed by decoupling, citing an event on 22 June 2020 near Lop Nur. The seismic signal was of magnitude 2.75, and China's foreign ministry called the accusation "completely groundless" ([OPB/NPR](#)). CTBTO Executive Secretary Robert Floyd said the organisation had recorded two small seismic events 12 seconds apart that day but that "it is not possible to assess the cause of these events with confidence" ([Bulletin of the Atomic Scientists](#)). No explosive test by any country in 2025 or 2026 was found in the reporting reviewed for this essay, so the moratorium appears to have held, but an allegation that cannot be settled by evidence creates pressure for a matching test, which the technical exchanges in element 2 are meant to relieve.

No five-power statement on human control of nuclear launch has been made. On 15 September 2026 Representative Ted Lieu and Senator Edward Markey asked President Trump to seek a commitment on human control from Xi Jinping at their summit ([Lieu and Markey letter](#)). Public reports of the summit do not mention nuclear command and control ([Al Jazeera](#)), so the 2024 affirmation remains the most recent statement of that kind.

Non-proliferation: a failed review and an unverified stockpile

The 11th Review Conference of the Non-Proliferation Treaty ended on 22 May 2026, after four weeks in New York, without a consensus final document; it was the third failure in a row after 2015 and 2022, and the next conference is due in 2031 ([UN News](#)). Talks broke down over language singling out Iran, against a background of disputes over the US and Israeli strikes on Iran, the war in Ukraine and the end of New START ([ETH Zurich CSS](#)). New START expired on 5 February 2026 ([Federation of American Scientists](#)).

Before the US and Israeli strikes of June 2025 Iran held 440.9 kilograms of uranium enriched to 60 percent, and inspectors have not returned to the bombed sites. In September 2026 the International Atomic Energy Agency (IAEA) reported that it still could not verify the stockpile or where it is held, and said that its lack of information and access "is a matter of proliferation concern" to be addressed "with the utmost urgency" ([RFE/RL via GlobalSecurity](#)). The war with Iran that began on 28 February 2026 makes a verified accounting of that material a precondition of any settlement, a point the war essay takes up.

DNA screening and risky research (element 3)

Executive Order 14292 of May 2025 suspended federally funded dangerous gain-of-function research and directed agencies to revise or replace the 2024 federal framework for screening DNA synthesis orders ([ASPR](#)). The United States Government Policy for

Stopping High-Risk Life Sciences Research, dated 20 July 2026 and announced by the NIH on 28 July, prohibits federal funding for “dangerous gain-of-function” research and restricts federally funded work in foreign countries of concern. Institutions that receive federal life-sciences money must certify applications under penalty of law, set up oversight to identify privately funded work of that kind and report it to their main federal funder each year; researchers who fail to report may lose funding for up to five years. The policy itself says that stopping privately funded work would need further statutory action, which an interagency working group is to explore, and agencies have 120 days to issue implementation guidance ([White House](#)).

The replacement screening framework has not appeared, and the federal page on synthesis screening still says it will be updated “once the revised framework is available” ([ASPR](#)). The Cotton-Klobuchar bill, which would make screening a legal duty, was referred to the Senate Commerce Committee in January ([GovInfo](#)). US policy now restricts what federally funded laboratories may study, while the cheaper control on who can buy dangerous sequences, which element 3 treats as the first priority, remains voluntary.

The gap matters more as AI improves: in January 2026 the *Bulletin of the Atomic Scientists* cited AI-aided pathogen design among its reasons for setting the Doomsday Clock at 85 seconds to midnight, the closest it has been ([Bulletin of the Atomic Scientists](#)), and DNA screening works wherever a would-be attacker obtains a design.

The Biological Weapons Convention (element 4)

The working group on strengthening the convention has drafted the three mechanisms in element 4 and the expansion of the support unit from four staff to seven, though the number of new posts remains in brackets ([UNODA](#)). The ninth session ended on 21 August 2026 with only a procedural report, though the support unit calculated that 93.5 percent of the draft text was agreed or nearly so, and the December session has five working days before the mandate expires. Italy’s ambassador, Leonardo Bencini, who presided over the 2022 review conference that created the working group, warned that without agreement “the whole process risks imploding” ([BWPP](#)).

Global health security and the pandemic agreement (element 5)

The United States completed its withdrawal from the WHO on 22 January 2026, having stopped funding and withdrawn all its personnel during the preceding year ([CDC](#)). Argentina’s withdrawal took effect on 17 March 2026 ([Al Jazeera](#)). These departures coincide with the aid cuts described in the [starvation](#) and [poverty](#) essays.

Negotiators missed the May 2026 deadline for the PABS annex, and member states extended talks, with results due at the World Health Assembly in May 2027 or earlier at a special session ([WHO](#)). The July round ended with differences remaining ([WHO](#)). NPR reported that the eighth round, held in Geneva in September, made little progress, with

wealthier and lower-income countries still divided over obligations to share pathogen data and access to the resulting vaccines and treatments; Suerie Moon of the Geneva Graduate Institute described the agreement as “in a state of continued limbo” ([NPR via GPB](#)).

AI testing and regulation (element 6)

The EU AI Act’s obligations for general-purpose models have applied since August 2025, and from 2 August 2026 the EU AI Office has had powers to enforce them ([European Commission](#)). The Digital Omnibus on AI, published in the Official Journal on 24 July 2026, postponed the rules for stand-alone high-risk systems from August 2026 to 2 December 2027 and for AI built into regulated products to 2 August 2028 ([Lewis Silkin](#)).

In the United States, a December 2025 executive order created a Justice Department task force to challenge state AI laws and asked Congress for a national standard that would preempt them ([White House](#)). California’s Transparency in Frontier Artificial Intelligence Act, signed in September 2025, requires the largest developers to publish safety frameworks, report incidents and protect whistleblowers ([Office of the Governor of California](#)), and New York’s RAISE Act, signed in December 2025, adds 72-hour incident reporting and takes effect on 1 January 2027 ([Office of the Governor of New York](#); [Wiley](#)). A bipartisan House discussion draft released in June 2026 by Jay Obernolte, Lori Trahan and others, the Great American Artificial Intelligence Act, would require developers with more than \$500 million in annual revenue to publish safety frameworks and undergo independent audits, and would preempt state laws on model development for three years ([Roll Call](#)). The federal Center for AI Standards and Innovation tests models for cyber, biological and chemical risks ([NIST](#)), but developers have no federal legal duty to submit models to it, so element 6 exists in the United States only in state law and voluntary practice.

In early March 2026 the Department of Defense designated the developer Anthropic a supply-chain risk after the company declined to allow its models to be used for mass surveillance of Americans or for fully autonomous weapons without humans making targeting and firing decisions; Anthropic sued, arguing that the designation was unlawful and retaliatory ([TechCrunch](#)). The department had acted under two statutes. On 27 August Judge Rita Lin of the federal district court in San Francisco ruled that one designation was unlawful retaliation and arbitrary and capricious ([TechCrunch](#)). On 25 September a federal appeals court in Washington upheld the other by two votes to one; Judge Gregory Katsas, for the majority, found the designation reasonable and cited the risk of “overly constrained AI models shutting down unexpectedly” during military operations, and Anthropic said it was considering asking the full court to review the decision ([Defense News](#)). The limits on military AI that element 1 and the war essay address are being settled in procurement disputes and litigation, and no statute or treaty yet sets them.

Science and channels between rivals (element 7)

The UN General Assembly created a 40-member Independent International Scientific Panel on AI and a Global Dialogue on AI Governance in August 2025. The panel issued a preliminary report on 1 July 2026, and the first Global Dialogue met in Geneva on 6 and 7 July, where Bengio said that “science currently cannot guarantee that as capabilities continue to increase, AI will not cause catastrophic harm” ([UN News](#)). The next session is set for May 2027 in New York ([UN Geneva](#)).

China’s July 2025 Global AI Governance Action Plan endorsed risk assessment, safety frameworks and the UN mechanisms ([Ministry of Foreign Affairs of China](#)), but until late September 2026 Washington and Beijing dealt with each other on AI mainly through trade measures; in January 2026, for example, the United States moved to case-by-case licensing of exports of Nvidia H200 and AMD MI325X chips to China ([Bureau of Industry and Security](#)) and placed a 25 percent tariff on those chips ([White House](#)). During Xi Jinping’s state visit to Washington, the White House and China’s foreign ministry both said the two countries had agreed to hold a dialogue on the risks and benefits of AI, with the next round in November, and to set up a communication channel for AI-related incidents ([ABC News](#); [Al Jazeera](#)). The agreement is the first step toward element 7 since 2024, though its scope, membership and relation to military AI have not been published.

Funding (element 8)

Funding has moved against the plan: the United States has left the WHO, the BWC support unit remains at four posts pending the December session, and philanthropic funding for nuclear security remains well below its level before MacArthur’s exit ([80,000 Hours](#)).

Part III: Revised priorities

1. **Save the Biological Weapons Convention working group in December** by adopting the science review mechanism and the larger support unit, with verification deferred if necessary; years of drafting are lost if the mandate lapses.
2. **Make DNA synthesis screening a legal duty.** The United States should publish the replacement framework its executive order called for and enact screening along the lines of the Cotton-Klobuchar bill, and other major markets should adopt compatible rules.
3. **Give the new US-China AI dialogue technical content.** The November meeting should cover incident notification, shared evaluation methods and the military uses of AI, using the UN scientific panel’s reports as common ground.
4. **Seek a statement on human control of nuclear launch decisions,** first from Washington and Beijing, building on the 2024 affirmation, and then from all five

permanent members of the Security Council.

5. **Keep the test moratorium and settle the China allegation with evidence**, through data exchanges or CTBTO consultations on the June 2020 events.
6. **Finish the pandemic agreement annex before May 2027** so that the treaty can open for signature, and replace lost US funding for WHO preparedness work.
7. **Put pre-release testing and incident reporting for frontier AI on a legal footing**. Any federal standard that preempts state laws should carry forward the protections already enacted in California and New York.

Conclusion

Governments now have better evidence on these risks than they did two years ago, from the International AI Safety Report, the UN scientific panel and the National Academies nuclear study, but over the same period the institutions that turn evidence into rules weakened: the Non-Proliferation Treaty review failed again, New START lapsed, the IAEA has been unable to verify Iran's highly enriched uranium since June 2025, the United States left the WHO, the pandemic agreement stalled and the Biological Weapons Convention working group is close to the end of its mandate. The US-China agreement on an AI dialogue is the main exception.

Controls that each country can adopt on its own, such as DNA screening, observance of the testing moratorium and human control of launch decisions, matter more while the multilateral routes are blocked, and none of them requires the great powers to trust one another. AI-aided design has moved biological risk nearer the top of the list, which is why the revised priorities begin with the Biological Weapons Convention talks and DNA screening.

This essay reflects reporting available as of 28 September 2026. Several of the processes described, including the Biological Weapons Convention talks, the US-China AI dialogue and the litigation over military AI, may change within weeks.

Sources

- SIPRI: Increasing focus on nuclear weapons amid heightened escalation risks, SIPRI Yearbook 2026
- SIPRI: SIPRI Yearbook 2026, World nuclear forces
- NASA GISS: Xia et al., Global food insecurity and famine from reduced crop, marine fishery and livestock production due to climate disruption from nuclear war soot injection (Nature Food, 2022)
- National Academies: Potential Environmental Effects of Nuclear War
- UNODA: BWC Working Group, revised draft final report (BWC/WG/8/CRP.1/Rev.1)

- BWC Implementation Support Unit: The BWC ISU's historical role, its current operations and lessons learned (December 2023)
- Adamala et al., Science: Confronting risks of mirror life (accepted manuscript, University of Manchester)
- GPMB: A World on the Edge, 2026 report
- International AI Safety Report: International AI Safety Report 2026
- International AI Safety Report: 2026 executive summary
- Grace et al.: Thousands of AI Authors on the Future of AI (arXiv)
- White House archive: Readout of President Biden's meeting with President Xi, 16 November 2024
- Rep. Ted Lieu and Sen. Edward Markey: Letter to President Trump on AI and nuclear launch, 15 September 2026
- Fortune/AP: Energy secretary says nuclear tests will be noncritical explosions
- OPB/NPR: U.S. shares new details on alleged Chinese nuclear test
- Bulletin of the Atomic Scientists: How to resolve the allegation that China is engaged in clandestine nuclear testing
- IBBIS: International Biosecurity and Biosafety Initiative for Science
- GOV.UK: UK screening guidance on synthetic nucleic acids
- GovInfo: S. 3741, Biosecurity Modernization and Innovation Act of 2026 (introduced)
- GOV.UK: Mirror life, Government Office for Science
- CEPI: CEPI 2.0 and the 100 Days Mission
- UN News: World Health Assembly adopts pandemic agreement
- GOV.UK: Frontier AI Safety Commitments, AI Seoul Summit 2024
- Anthropic: Activating AI Safety Level 3 protections
- OpenAI: GPT-5 system card
- Future of Life Institute: AI Safety Index, Summer 2026
- SecureBio: Virology Capabilities Test
- MacArthur Foundation: Nuclear Challenges Big Bet 2020 evaluation report
- MacArthur Foundation: \$21 million to support a stronger, more diverse nuclear field
- 80,000 Hours: Nuclear weapons problem profile
- CSIS: Russia's latest nuclear saber-rattling on nuclear testing
- Al Jazeera: China, US to open AI communication channel after summit, White House says
- UN News: NPT Review Conference ends without consensus
- ETH Zurich CSS: The NPT Review Conference, a nuclear negotiation breakdown
- Federation of American Scientists: The expiration of New START

- RFE/RL via GlobalSecurity: IAEA says unable to verify Iran's enriched uranium stockpiles (2 September 2026)
- ASPR: Synthetic nucleic acid screening
- White House: United States Government Policy for Stopping High-Risk Life Sciences Research (20 July 2026)
- Bulletin of the Atomic Scientists: 2026 Doomsday Clock Statement
- BWPP: BWC Working Group report 2026-12, August 2026
- CDC: United States completes WHO withdrawal
- Al Jazeera: Argentina officially withdraws from WHO
- WHO: Member states agree to extend negotiations on the PABS annex
- WHO: Member states continue negotiations on the PABS annex
- NPR via GPB: There's a global pandemic treaty that's signed and in limbo
- European Commission: AI Act regulatory framework
- Lewis Silkin: The Digital Omnibus on AI enters into force today
- White House: Ensuring a National Policy Framework for Artificial Intelligence
- Office of the Governor of California: Governor Newsom signs SB 53
- Office of the Governor of New York: Governor Hochul signs RAISE Act
- Wiley: New York finalizes RAISE Act
- Roll Call: Bipartisan AI draft proposes three-year preemption of state laws
- NIST: Center for AI Standards and Innovation
- TechCrunch: Anthropic sues Defense Department over supply-chain risk designation
- TechCrunch: Anthropic gets its first court win over the Pentagon's supply-chain risk label
- Defense News: US appeals court upholds Pentagon's blacklisting of Anthropic
- UN News: First Global Dialogue on AI Governance
- UN Geneva: Global Dialogue on AI Governance opening session
- Ministry of Foreign Affairs of China: Global AI Governance Action Plan
- Bureau of Industry and Security: Commerce revises license review policy for semiconductors exported to China
- White House: Fact sheet on certain advanced computing chips
- ABC News (Australia): China, US agree to AI dialogue, tariff cuts during Xi Jinping summit