



From AIWS to AIWS Civilization

Building a Civilization for Humanity in the Age of AI

Nguyen Anh Tuan

Founder and Chief Architect of AIWS

Co-Founder, Co-Chair and CEO of Boston Global Forum

Address to the Lincoln University Seminar, New Zealand

AIWS CIVILIZATION TABLE NO. 1

Harvard Faculty Club, Cambridge, connected with Lincoln University

September 30, 2026

Dear colleagues, students, and friends at Lincoln University in New Zealand,

From Boston, I bring an urgent invitation, and a profound belief in what we can accomplish together.

We are living through historic days in the AI revolution. The decisions being made today—in laboratories, corporate boardrooms, universities, and governments—may shape the freedom, dignity, and opportunities of generations to come.

The question before us reaches far beyond technology.

*What civilization will we build with the intelligence we ourselves
are bringing into the world?*

Today, warnings about AI are being voiced across society. Scientists, leaders, educators, and citizens are asking whether humanity can retain control over increasingly powerful systems. Even Dario Amodei and Sam Altman, two of those honored in America 250: AI Pioneers, have warned of serious dangers to humanity and the possibility of losing human control.

These warnings demand a response that goes beyond another warning.

Humanity needs a path forward. People need institutions they can trust, opportunities to contribute, and reasons to believe that they can still shape their future.

This is the work to which Boston Global Forum has committed itself.

And this is why I am here today.

I will always remember December 12, 2017.

At Loeb House, Harvard University, Governor Michael Dukakis and I announced AI World Society. Former Estonian President Toomas Hendrik Ilves, Patrick Winston, Joseph Nye, Max Tegmark, and other distinguished participants contributed invaluable insights to that historic conference.

In 2017, much of the world was captivated by machine learning—its applications, its growing capabilities, and its promise for business and everyday life.

We shared that excitement. But we also saw a larger transformation beginning.

AI would change the relationship between knowledge and power. It would enter the institutions through which people learn, earn a living, receive information, and govern society.

So we asked: how must humanity prepare?

How can we ensure that this new power serves people? How can we protect freedom and dignity while opening the extraordinary possibilities of AI?

I am proud that BGF recognized these questions so early. And I am proud that we began building answers.

Governor Dukakis brought something essential to that beginning: a lifetime of public service, and the conviction that leadership must make people's lives better.

That conviction has accompanied us throughout this journey.

Every architecture we propose, every institution we build, must ultimately answer to people.

Just a few months later, in April 2018, at the Harvard Faculty Club, we introduced the AIWS Seven-Layer Model through the BGF–G7 Summit Initiative, contributing to discussions ahead of the G7 Summit in Canada. We connected ethical principles with governance, institutions, and applications, so that responsibility would be part of AI development from the outset.

At that same event, we honored OECD Secretary-General Ángel Gurría with the inaugural World Leader in AIWS Award. Through this recognition, we affirmed the importance of leaders who place technological progress in the service of society.

Our work then crossed the Atlantic.

Together with Club de Madrid, an important BGF partner in Europe and an organization of former heads of state and government, we advanced the Social Contract for the AI Age. We brought experience in governing nations into the effort to make technological power accountable.

At the heart of that work was a question that remains just as meaningful today:

*What do we owe one another when the systems we build can
affect the lives of people we will never meet?*

That question belongs to humanity. And humanity must help shape the answer.

Now, in 2026—a special year marking the 250th anniversary of the founding of the United States—we face new responsibilities.

Through our book *America at 250: A Beacon for the AI Age* and our May 1 conference at Harvard, we connected reflection on America's founding ideals with the responsibilities of the AI Age.

The dignity of each person. The promise of freedom. The right of human beings to shape their own future.

These ideals become more urgent when technological systems can influence what people know, what they believe, and which choices they perceive as possible.

Our work is to develop solutions that meet this challenge.

AIWS has evolved from a governance model into AIWS Civilization—the Civilization of Human–AI Co-Creation.

At AIWS Lumina Lab on Beacon Hill, we developed a comprehensive architecture, presented in Beacon Papers No. 8 in August 2026. It is an architecture for a civilization under human command, in which AI expands knowledge, intelligence, and creativity, while people retain purpose, values, the right to choose, and ultimate responsibility.

Why does humanity need this architecture now? Because AI capabilities are advancing faster than society's ability to understand, verify, and govern the power entrusted to these systems. We may use a highly capable system without understanding its limits; give it authority without knowing who must answer for its actions; and become dependent on it before recognizing that our own capacity for independent thought is diminishing.

These three gaps—between capability and understanding, between power and responsibility, and between convenience and human development—present a civilizational challenge. If we address only individual applications, we will miss how AI is simultaneously transforming the economy, education, information, culture, and social power.

AIWS Civilization rests on human dignity, human agency, the continued development of human intelligence, identifiable human responsibility, and meaningful human command over AI. Technology must help each person remain the author of their own life: able to think independently, question, choose, create, and grow.

From these principles, AIWS Civilization stands on four foundations—AIWS Trust Infrastructure, the Constitution for Humanity in the Age of AI, Human–AI Co-Creation, and AIWS Lumina—connected as one system. The Constitution establishes rights, duties, and limits on power. AIWS Trust Order organizes relationships of responsibility among governments, businesses, institutions, and citizens. AIWS Trust Infrastructure provides the standards, evidence, and verification mechanisms through which those responsibilities can be fulfilled in practice.

On this foundation of trust, Human–AI Co-Creation becomes a way of learning, researching, producing, and serving. The Movement brings this practice to people everywhere. AIWS Lumina nurtures a culture of Love, Creativity, Nobility, and Wisdom, so that expanded capabilities are directed toward contributions worthy of humanity.

For me, this architecture is always connected to people. Governor Dukakis reminds us that institutions must serve the public. Professor Alex Pentland's work helps us understand cooperation and the conditions that make trust possible. Professor Thomas Patterson brings us back to the foundations of civic judgment: people need trustworthy information to govern their society. Professor Judea Pearl's work on causality reminds us to keep asking why, to understand consequences, and to question assumptions.

Imagine a research team using AI to propose improvements to a food system. The Constitution requires protection of the rights and dignity of those affected. Trust Infrastructure requires clarity about which system is being used, what it is authorized to do, who verifies it, and who is responsible. The researchers must understand, test, and choose the approach. The Movement helps them share that experience and grow through their work. Lumina asks about the value of the creation: what good does it bring to people and nature?

This is how the components work together. A good product must come with a trustworthy process and people whose capabilities have grown through creating it.

The measure of AIWS Civilization is therefore clear: as AI advances, human beings must continue to advance in intelligence, creativity, responsibility, and wisdom. If output rises while human judgment declines, we must change how AI is designed and used. This must be measured; it cannot simply be hoped for.

On September 18, I brought this vision to the Silicon Valley Peace Conference.

Speaking in a church, in a region associated with extraordinary technological achievements, I felt deeply the significance of bringing technological power into dialogue with human conscience.

What are we building?

Whose lives will change?

What responsibilities must accompany the power we create?

I called on those present to help turn our vision into action.

Then, in New York, I brought this call to the Nizami Ganjavi International Center's high-level conference on the sidelines of the United Nations General Assembly, in a discussion on AI, digital sovereignty, and the future of global governance.

My message was this: every nation must have a meaningful role in shaping the AI Age. The warnings and analyses have made the dangers clear; now we must act through the Human–AI Co-Creation Movement, AIWS Trust Infrastructure, and the Constitution for Humanity in the Age of AI.

Developing countries must build technological capabilities, contribute ideas, and create their own institutions. Their people must participate as creators of the future.

We are developing the Constitution for Humanity in the Age of AI. This constitutional undertaking proposes a shared foundation for protecting people and limiting AI power. Its specific legal force must be established through the legitimate mechanisms of nations and the international community.

The Constitution begins with an affirmation: every person possesses a dignity that cannot be reduced to data, a score, or economic value. Each person has the right to remain the author of their own life—to think, choose, decline to delegate a decision to AI, and grow through their own experience.

This has very concrete implications. A person affected by a decision involving AI must be able to know about that involvement, understand the relevant grounds, challenge the outcome, and reach someone responsible. A citizen must be protected against covert manipulation of information and behavior. A student must develop their own capabilities, rather than be led into a dependence they do not recognize.

The Constitution also establishes duties. Developers must test their systems and truthfully disclose relevant limitations. Deployers must define purposes, limits of authority, and the people responsible. Governments must use AI to serve citizens and respect human rights. Universities and civil society organizations have a responsibility to contribute to verification, protect truth, and expand opportunities for participation.

Human-in-Command must mean substantive authority: the person responsible must have sufficient information, competence, and authority to reject, adjust, or stop a system's actions. A signature of approval does not meet that requirement when the person signing does not understand what is being decided. Delegating a task to AI does not remove the responsibility of the person who chooses to delegate it.

From these principles, we pursue verifiable requirements: the identity of the system and the responsible actor; bounded authority; evidence of consequential actions; the right to question; and mechanisms for intervention and stopping. Constitutional Silicon takes these questions into the hardware itself: how can computational power also be placed within enforceable and verifiable limits?

The Constitution defines what must be protected. Trust Infrastructure builds the means of protection. The Observatory examines whether commitments to protection are being fulfilled. Those responsible must be able to answer for what the evidence shows. This connection brings principles into everyday life.

I invite Lincoln's scholars and students to help develop the Constitution through research and critical examination, drawing on real situations in education, the environment, livelihoods, and community rights. Good principles must withstand the test of life.

A civilization also takes shape through people's everyday practices.

This is why the AIWS Human–AI Co-Creation Movement matters so deeply.

On September 23, BGF launched the Movement at the Massachusetts State House. While I was attending conferences in New York, my statement was read in Boston. A vision became an open invitation to everyone.

The Movement begins with the belief that each person has something unique to contribute: a student's question, a scientist's discovery, a farmer's knowledge of the land, an artist's imagination, a community's memories and aspirations. AI must help these contributions flourish.

Human–AI Co-Creation has a clear distinction of roles. People establish purposes, bring their experience and values, exercise judgment, and take responsibility. AI expands knowledge, explores possibilities, supports analysis, and offers reasoned criticism. The final decision remains with people.

A trustworthy AI should be able to point out that an assumption lacks evidence, that a group of people has been overlooked, or that an option has unforeseen consequences. Such criticism helps

people make better decisions. When people merely approve an outcome they do not understand, co-creation has lost something essential.

Every participant in the Movement begins with three commitments: I remain the author of what I create; I bring my cultural identity into what I create with AI; I use AI to become more capable, never less.

Every initiative must undergo the Human Advancement Test: after working with AI, do participants think more deeply, create better, and contribute more than before? That question requires evidence. Can they explain their choices, detect errors, defend their reasoning, and carry out essential parts of the work without AI assistance?

A Movement project must therefore begin with a real problem and a purpose established by people. Participants record their own contributions and those of AI, verify results, choose an approach, and take responsibility for the product. They then assess both the value created and the capabilities they have developed, sharing lessons so the next cycle achieves a higher standard.

Imagine Lincoln students researching a solution for more efficient water use. The students frame the question, investigate community needs, and select criteria. AI supports data analysis, simulation, and the exploration of options. The students must test assumptions against real conditions, explain their choice, and assess its impact. The achievement is a verified solution, together with people who understand the problem more deeply and can continue to create.

This is growth through co-creation. We must preserve the parts of the work that require human thought and effort, maintain practice without assistance, and measure progress over time. If participants become increasingly dependent, the process must be redesigned.

Today, I invite Lincoln University to join BGF in starting a Human–AI Co-Creation program with an initial group of faculty and students. Together, we can select projects in food systems, the environment, or community life, and agree on how to record contributions, verify products, and assess participants' progress.

I propose that we discuss an initial 90-day phase: select a coordinating group and projects, establish baseline capabilities, carry out cycles of co-creation, and then hold a presentation of the products and evidence of human development. Verified results will provide the basis for deciding what comes next.

Let us bring businesses and communities into defining problems and evaluating the value of the results. Let us connect Lincoln with groups in other countries, so that knowledge is shared while each place's identity continues to develop. A research group could also contribute to the Observatory, helping us understand the AI systems used in these very projects.

In companies, this approach can help workers develop their capabilities as they create products and services. In communities, it helps people solve problems together through their own

knowledge and voices. In government, it can help public servants improve decisions and their responsibility to serve. Each institution must choose appropriate problems and publish evidence of results.

A Constitution protects people against power. A Movement awakens people's potential. The Observatory helps society see risks and act; the Movement helps society actively create value. Both must develop together, so that trust and creativity become living foundations of AIWS Civilization.

There is a cultural dimension to this work that I especially cherish.

AIWS Lumina is guided by Love, Creativity, Nobility, and Wisdom.

Love helps us care about the lives affected by our decisions.

Creativity helps us bring something new into the world.

Nobility directs our capabilities toward contributions worthy of human dignity.

Wisdom helps us understand what is worth pursuing and how to pursue it responsibly.

Through Hồn Trà—Tea Soul—we nurture joy, optimism, and a love of life, opening a space in which people can create things of beauty and value and offer them to others.

I think of these contributions as Seeds of Light.

A discovery that opens new possibilities. A work of art that touches the heart. An act of kindness that helps someone recover their faith. An idea that brings fresh hope to a community.

AIWS Civilization must create the conditions in which these Seeds of Light can grow.

Dear friends,

We are living through historic days.

There have been so many warnings, so many analyses, and so many fears. We must understand the risks and act seriously. But the spirit of this undertaking must spring from faith in people.

I stand before you with joy in what we can create together.

With optimism sustained by the work we are prepared to undertake.

With love for people, and love for humanity.

Our immediate priorities are clear: build AIWS Trust Observatory and bring the Human–AI Co-Creation Movement into everyday life.

When the Movement enters universities, students can become creators of their own future, with ever-growing capabilities.

When it enters companies, innovation can strengthen human creativity and responsibility.

When it enters communities, people can create through their own voices.

When it enters governments, leaders can place AI in the service of citizens.

As these practices spread, humanity will have stronger foundations on which to welcome a new era with joy and confidence.

SEPTEMBER 29, 2026 • THE WHITE HOUSE

And at this very moment, an important development gives our invitation special significance.

On September 29, 2026, President Donald Trump met with major technology leaders at the White House. They signed a voluntary AI safety accord calling for four layers of control: internal monitoring of model capabilities and risks; an internal team to check that safeguards operate as intended; independent external audits; and oversight by an independent committee of each company's board.

On behalf of Boston Global Forum and AIWS, I respectfully acknowledge President Trump's action in bringing responsibility and AI safeguards into discussions with those who hold great technological power. We regard this as an important step, consistent with the direction BGF–AIWS had already advocated through AIWS Trust Observatory: AI power must be verified, monitored, and linked to identifiable responsibility. The accord provides for independent external auditing and board oversight. To turn these commitments into public trust grounded in evidence, we also need an independent civil society capacity to examine evidence, track implementation, make gaps visible, and ask those responsible to answer. **This is the contribution AIWS Trust Observatory seeks to make.**

This is an opportunity to offer our work to the White House and to the world. AIWS Trust Observatory can contribute methods of independent observation, ways to compare commitments with actions, mechanisms for recording capabilities and incidents, and proposals that connect warnings to verifiable remedies.

At this seminar, I propose that we discuss three practical questions: what evidence is needed to know whether commitments are being fulfilled; how can the independence of assessments be protected; and how can universities contribute research, testing, and the study of impacts?

Immediately following the meeting and announcement in the United States, BGF and a university in New Zealand are coming together in this dialogue. Geographic distance opens a valuable possibility: knowledge, responsibility, and a spirit of service can connect people far apart in joint action for humanity.

From Boston to Lincoln, from an architecture developed on Beacon Hill to the contributions we can prepare together for the White House, this is an opportunity to bring AIWS Civilization into a moment of historic significance through concrete work.

And today, in this meeting connecting Boston with Lincoln University, I extend that invitation to you.

AIWS TRUST OBSERVATORY • FOUNDING FRAMEWORK PUBLISHED
SEPTEMBER 27, 2026

Our most urgent institutional task is to build AIWS Trust Observatory.

Within AIWS Trust Infrastructure, the Observatory is an independent civil society mechanism for observing consequential AI power, assessing evidence, issuing early warnings, and encouraging action. Its initial focus is frontier AI systems and settings in which AI is deployed at scale or authorized to take actions with serious consequences.

The Observatory focuses on four questions: what capabilities is AI attaining? What authority is it being given in the real world? Are the controls proportionate to that power? If consequences follow, who has the authority and duty to answer?

For example, when an AI is permitted to write and execute code, access networks, or act through tools, a capability test alone cannot establish that deployment is trustworthy. We must also know how access is bounded, whether actions are recorded, who can intervene, and what conditions require the system to stop. Capability becomes consequential power when it is connected to authority to act and a real-world environment.

The Observatory will examine public documentation, reproducible test results, deployment evidence, and incident reports, comparing observed actions with developers' own commitments and AIWS Trust Standards. The Frontier Capability Registry records capabilities together with evidence and known limitations. The AI Incident Exchange records incidents and lessons so that other organizations can prevent similar failures.

An important principle is this: **what has not been verified must be clearly identified as unverified.** An information gap cannot be turned into an assurance of safety, nor can it alone become an accusation. People must be able to see the evidence and the degree of certainty behind each conclusion.

The Observatory's process moves from observation to verification, assessment, risk classification, warning, recommended measures, and follow-up on remediation. When evidence indicates a serious risk, the warning must identify who needs to act and what action is required: limiting access, adding independent verification, strengthening human oversight, or temporarily suspending a risky function.

Every warning must be reviewed under published methods and procedures specifying decision-making authority. The organization assessed has the right to respond, and its response must be published alongside the conclusion. Errors must be corrected publicly. Independence, fairness, and the capacity to correct its own mistakes are the foundations on which the Observatory can earn the trust it seeks to protect.

The Observatory draws its strength from evidence, public warnings, and its ability to mobilize responsible actors. Legal enforcement remains with the competent authorities. AIWS Trust

Rating is a complementary mechanism that evaluates an organization's internal practices with its consent; confidential findings from that process must be protected separately from public observation.

Humanity cannot afford to discover what frontier AI has become only after an irreversible consequence has occurred. We must build the capacity to observe before risks outrun our ability to respond.

Universities can contribute directly: developing reproducible tests, researching the effectiveness of human oversight, verifying capability claims, and recording impacts where systems are deployed. I invite Lincoln to help establish a faculty and student research group for this work, using a shared methodology and expert review.

Therefore, on behalf of Boston Global Forum and AIWS, **I designate today's seminar and discussion as AIWS Civilization Table No. 1**—the first dialogue in the AIWS Civilization Table series, where ideas, evidence, and responsibility lead to action.

We will bring together the seminar's discussions, contributions that participants agree to include, and concrete proposals, and submit them to Michael Kratsios, Assistant to the President and Director of the White House Office of Science and Technology Policy, and to President Donald Trump.

We will do so in a voluntary spirit of service: offering knowledge, methods, and opportunities for research cooperation to protect people, strengthen trust, and help AI serve humanity. Differing views expressed in the discussion must also be represented honestly. Serious critical examination is what gives our contribution value.

I hope those present today will help make this meeting a historic beginning—through the quality of the ideas we exchange, the proposals we refine, and the work we continue after the seminar ends.

AIWS Civilization Table No. 1 begins with this dialogue between BGF–AIWS and Lincoln University, and with a clear invitation: let us build the Observatory together, practice Human–AI Co-Creation together, and bring trustworthy results to those making decisions that will affect humanity's future.

Dear friends,

History is opening before us a great responsibility, and a great opportunity to serve. I believe we can embrace that opportunity with joy, compassion, and the courage of people creating together.

Let this meeting connecting Boston and Lincoln become a milestone we can be proud of: when the world needed action, we brought our knowledge and goodwill to the common endeavor.

Let us build AIWS Civilization—a civilization in which people exercise greater mastery, create nobler values, and grow together in wisdom.

From this first AIWS Civilization Table, we send a message to the White House and to the world: we are ready to join you, ready to contribute, and ready to act together for humanity.

The future is ours to co-create.

Thank you.



Boston · September 30, 2026