

"check against delivery"

ARTIFICIAL INTELLIGENCE — PROMISE AND CHALLENGES FOR SUSTAINABLE DEVELOPMENT AND A PEACEFUL ERA

Statement at AGIC 2026 — Shenzhen International General Artificial Intelligence Conference
Shenzhen World Exhibition & Convention Center, 26–28 August 2026

Daniel Stauffacher, Founder, ICT4Peace Foundation

Distinguished Chair, Excellencies, ladies and gentlemen,

I thank the Shenzhen Artificial Intelligence Industry Association and the organizers of AGIC 2026 for their invitation and their hospitality.

Shenzhen: transformation in one generation

I first came to Shenzhen at the end of the 1980s. It was a border town of some thirty thousand people. Today it is a global center of technology. Shenzhen proves a simple point: with clear policy, investment in people, and openness to the world, technology can transform a society in one generation.

AI — first of all an instrument of progress

We should look at artificial intelligence in that spirit — first as an instrument of progress. It helps farmers predict drought. It brings diagnosis to clinics without specialists. It teaches children in their own languages. It sharpens climate forecasts. It keeps factories running, and brings credit to people banks never reached. Shared fairly, AI can accelerate the Sustainable Development Goals. It can also serve peace.

But the organizers have asked me to speak about the risks. I will be candid.

But where do we stand on AI: the first scientific assessment

Last month, more than 170 countries met in Geneva for the first Global Dialogue on AI Governance. For the first time, independent science was on the table: the Scientific Panel on AI presented its first report. Its essence is this. AI capabilities are advancing fast. Governance is not keeping pace. The Panel calls this the evidence dilemma: by the time harm is proven, it may be too late to act. A few actors dictate the direction of AI. Everyone else depends on their choices. And AI-enabled deception is spreading.

These are not the words of pessimists. They are the first global scientific assessment of AI. The Panel will now report every year. That gives us a yardstick.

Four risks to sustainable development

First, the AI divide. Most of humanity lives without the compute, data, energy or skills for AI. Left alone, AI will deepen inequality — silently, and at scale. The SDGs are already off track. The AI divide could put them out of reach.

Second, the loss of development ladders. For fifty years, nations climbed out of poverty the same way: manufacturing first, then services. Automation may kick that ladder away. We have barely begun to design what replaces it.

Third, the environmental bill. AI data centers already use as much electricity as entire nations. Their hunger for power and water grows — just as the climate demands restraint.

Fourth, trust. Development rests on trust — in institutions, in science, in elections. AI-generated deception corrodes it faster than we can rebuild it. The most fragile societies have the fewest defenses.

The risks to peace

Here I speak from ICT4Peace's two decades in the UN processes on cyberspace. AI is entering military decisions, weapons and critical infrastructure faster than any shared understanding of limits. Decision times shrink exactly where miscalculation is most dangerous — even in the nuclear domain. The drones and robots in this exhibition are dual-use: the platform that delivers medicine can deliver munitions. AI-driven information operations undermine elections and peace processes — cheap, deniable, open to any actor. And in fragile states, unchecked AI amplifies surveillance and repression.

One more risk, from the frontier

Since this is a conference on general AI, let me add one more risk. Systems are learning to plan, to use tools, to write their own code. The Panel was explicit: there is no scientific guarantee that such agents will follow instructions, and our methods to control them lag behind. Control is now an operational question. Safety research and agreed thresholds are not brakes on innovation. They are its license to continue.

How should the international community react? What the ongoing cyberspace peace processes taught us

We have seen capabilities outrun rules before — in cyberspace — and we paid the price. But that taught us the method. Keep the major powers talking. Agree on norms. Build confidence step by step. Keep human control over the use of force. Protect civilians. Require developers and deployers of advanced systems to assess foreseeable risks, conduct rigorous testing and report serious incidents.

And say it plainly: international law applies fully to AI, including humanitarian law and human rights law. We must now apply this method faster, because AI moves faster.

The United Nations: these institutions now exist

This year, for the first time, the institutions to do so exist. The Global Digital Compact, adopted by the General Assembly in 2024, created the Global Dialogue on AI Governance and the Scientific Panel: every state now has a seat, and a shared evidence base. On security, the General Assembly endorsed the consensus report of the Open-Ended Working Group and established a permanent Global Mechanism on ICT security. It held its first substantive session in New York last month. It is the natural home for the cyber dimension of AI. On development, WSIS+20 made the Internet Governance Forum permanent and linked the WSIS Action Lines to both the SDGs and the Global Digital Compact. On military AI, states now meet under a General Assembly mandate, alongside the Geneva talks on autonomous weapons.

A word of caution: AI cooperation must not harden into exclusive blocs that force countries to choose sides. What the world needs is inclusive cooperation — open to all countries, anchored in the United Nations and the multistakeholder community — not competing coalitions demanding exclusive loyalty.

But Government institutions alone are not enough

Governments cannot govern AI without the companies that build and deploy it, the scientists and technical experts who test it, and the civil-society and humanitarian organizations that see its effects on people. Communities affected by these systems — particularly in developing countries and fragile contexts — must have a meaningful voice in setting the rules, not merely be consulted after deployment.

This matters even more as AI agents begin to interact with each other. Risks then emerge between systems — in the gaps between providers, deployers, evaluators and regulators — and no single actor can close those gaps alone. The whole lifecycle must work as one.

Shanghai: China's initiative for global AI cooperation

Ten days after the Global Dialogue on AI Governance in Geneva, President Xi Jinping opened the World AI Conference. He set out four principles — openness, security and human control, inclusivity, global solidarity — and backed them with commitments: the World AI Cooperation Organization, with twenty-nine founding states; five thousand training places for developing countries; regional centers with ASEAN, the Arab League and the African Union; and early-warning systems for thirty nations. ‘A single string cannot make music,’ he said, ‘and a single tree does not make a forest.’

Special focus on extreme risks: a convergence by major powers is emerging

On a single day this July — 16 July — twenty-nine states in Shanghai founded the World AI Cooperation Organization, and two hundred experts and Nobel laureates in Rome called for effective human control over nuclear decisions. Two very different voices, one conclusion: AI must never become a threat to peace and security.

And national laws are already moving in parallel: Washington's executive order on pre-release evaluation, Beijing's first framework for AI agents, binding laws in Brazil, the EU and South Korea, legislation advancing in India. Convergence is achievable because it has already begun.

In July, ICT4Peace joined the IFRC and the Geneva Centre for Security Policy to launch the Strategic Foresight Group's proposal for a Global Compact on Extreme AI Risks. It asks states

to converge on four risks no country can escape — AI-enabled cyber attacks, including on nuclear command; biological and chemical weapons; mass manipulation; and loss of human control — and on a few practical instruments: common evaluation of frontier models, incident reporting, and a two-key release for the most dangerous systems. Light, achievable, and open to both AI superpowers.

Sovereignty and openness: two answers, one direction

Behind all of this lies the sovereignty question. Nearly all leading AI comes from a handful of firms in two countries. So every state now wants sovereignty — over compute, over data, over the models it depends on. The choice is simple. Does sovereignty divide the world into blocs? Or does it connect us? There is a way to connect: openness. China releases its leading models with open weights. The world can use them. Switzerland has built Apertus — a fully open public model. Its data and recipes are published. It speaks over a thousand languages. Any state can inspect it, adapt it, run it. Openness makes sovereignty shareable. And safety rules should follow what systems can do — not who built them.

The road ahead: the Geneva AI Summit

In June 2027, Switzerland hosts the Geneva AI Summit — after Bletchley, Seoul, Paris and Delhi. That is where these strands — the UN Dialogue, the security and development tracks, the Shanghai commitments, the work on extreme risks — should be woven into one architecture. What is still missing is connection. Conferences like AGIC do something valuable: they keep builders, governments and civil society in the same room. The harder task is to carry that mix into the places where it is still missing — into evaluation, into release decisions, into the standards that will govern these systems long after we have all gone home.

A simple standard

Excellencies, let me close with a simple standard. Judge technology by its impact on people's lives, not by its power. The risks I have described are real. But they are not destiny. They are choices — about investment, restraint and cooperation. Those choices are still in our hands.

Trust, once lost, is the hardest infrastructure to rebuild. Let us build it now, together. Shenzhen transformed itself in a generation. Let it also be remembered as the place where the world chose to govern AI wisely — for sustainable development, and for a peaceful era.

Thank you.