



American Chamber of Commerce **in Italy**

Milan and the future of global governance of artificial intelligence

Functions for a new UN pillar for AI governance

Updated edition following the high-level week of the 81st United Nations General Assembly. It continues the series launched on 27 July 2026 with the paper *Milano e il futuro della governance globale dell'intelligenza artificiale. Perché candidare Milano a ospitare un Hub ONU per la governance etica dell'AI.*





American Chamber of Commerce **in Italy**

Milan and the future of global governance of artificial intelligence

Functions for a new UN pillar for AI governance

September 2026

Foreword	5
Executive Summary	7
1. From location to functions	9
1.1 Nature of the paper	9
1.2 Why Milan	9
1.3 Why Italy	10
1.4 Why Europe	11
2. Summer 2026: what has changed	12
2.1 The G20 and the Carolina Principles	12
2.2 The United States' position	12
2.3 The demand coming from the labs	13
2.4 The United Nations process	14
3. New York, September 2026	14
3.1 President Trump's speech	14
3.2 The Secretary-General and the call of 21 September	15
3.3 Altman and Amodei at the Security Council	16
3.4 The positions of governments	17
3.5 From divergence to a common language	18
4. Criteria and scope of the mandate	19
4.1 Where the tasks come from	19
4.2 Seven rules and four exclusions	20
5. The twelve functions of the pillar	20
6. What the pillar produces and how it is organised	27
6.1 Un'architettura leggera	27
6.2 Il finanziamento	28
6.3 Il raccordo con il sistema	28
7. The implementation sequence	28
8. What it means for Italy and for Milan	29
Conclusion	31
Methodological note	31
Sources	32

Foreword

On 27 July 2026 AmCham Italy published a policy paper devoted to a precise question: why Milan should put itself forward to host a United Nations Hub for the ethical governance of artificial intelligence.⁵ That paper opened a debate which, in the weeks that followed, produced tangible results: a cross-party promoting committee, the support of economic, academic and parliamentary institutions, and a public discussion that reached well beyond our country's borders.

Since then, the question we are asked has changed. Our interlocutors, in Italy and abroad, no longer ask whether Milan is a suitable location. They ask what this body should do: which functions, with which tools, within which limits and in what relationship with the organisations already operating in the multilateral system. It is a demanding question, and it is the right one. This second paper attempts to answer it.

The answer comes at a particular moment. During the high-level week of the 81st General Assembly in New York, the governance of artificial intelligence became, for the first time, a front-rank issue for heads of State and government. The President of the United States rejected "any attempt to construct a globalist scheme" to control it.⁴⁰ The day before, twenty-two leaders, later thirty, had called for exploring an international institution able to set standards, enable verification and convene States when the capabilities of AI systems cross critical thresholds.⁴³ At the Security Council, Sam Altman and Dario Amodei called for common standards, verification systems and rapid channels for incident reporting.^{46, 48}


At first sight these positions seem far apart. On a closer reading, they leave open the same space: that of a body which does not regulate, authorise or sanction, but which makes shared methods, comparable measurements, early warning and technical expertise available to all. This is the space that this paper proposes to entrust to a new United Nations pillar, and that Milan can host. Measuring does not mean prohibiting. It means making every subsequent choice possible, including the choice not to regulate.

At the heart of the paper are twelve functions. They are not an invention: they translate into operational tasks the principles of the UNESCO Recommendation on the Ethics of Artificial Intelligence, adopted in 2021 by all Member States, within the framework opened by Resolution A/RES/79/325. A single criterion orders them: technology must remain at the service of the human person.

A bid is sustained by the quality of what it proposes, not by the insistence with which it is presented. Milan can offer the United Nations a place; Italy, together with the place, a clear institutional project; Europe, the legal framework in which that project acquires universal meaning. These pages are our contribution to that project and an invitation to discuss it on its merits.

Managing Director

Simone Crolla



Executive Summary

The international community is no longer debating whether permanent institutions for artificial intelligence are needed. It is debating what they should do and, above all, what they should not do. This paper, which follows the AmCham paper of 27 July 2026, answers with twelve functions for a new United Nations pillar and proposes Milan as its seat.^{1, 5}

The twelve functions of the pillar. The core of the proposal, developed in chapter 5, is a set of twelve technical functions. They translate into operational tasks the principles of the 2021 UNESCO Recommendation and are ordered by a single criterion: technology at the service of the human person. Each has a declared limit and a performance indicator, and is exercised together with the organisation competent for the subject matter, which retains ownership of it.¹⁰

The twelve functions at a glance

5.1

Pace of the frontier and early warning

Measures, with a public method, how fast the capabilities of the most advanced systems are growing.

5.2

Impact assessment

A common method, updated over time, for assessing the effects of a system on people.

5.3

Human accountability

Whoever is accountable for a decision always remains identifiable along the technology chain.

5.4

Transparency and provenance

Telling the authentic from the synthetic and obtaining explanations proportionate to the decision.

5.5

Data and documentation

Documenting which data a system was trained on, and on what legal basis.

5.6

Equity among countries

A comparable index of gaps in compute, data and skills, in order to reduce them.

5.7

Work and transition

Comparable data on employment effects and on the policies that work.

5.8

Health and minors

Comparable evidence and precautionary criteria for the most vulnerable.

5.9

Training regulators

Skills for those who regulate and those who use AI: the precondition for all other functions.

5.10

Languages and pluralism

Measuring which languages and cultures systems represent, and supporting those left out.

5.11

Environment and compute

Making the energy, water and land footprint of data centres comparable

5.12

Security and stability

A common language for recognising and reporting incidents and dangerous capabilities.

The political framework calls for less regulation. On 2 September 2026 the G20 Innovation Ministerial in Chapel Hill adopted by consensus the *Carolina Principles*: a sectoral and proportionate approach, priority to existing regulatory frameworks, no new general-purpose architecture.² The United States' position is the clearest. The *AI Action Plan* aims to accelerate innovation and remove regulatory burdens; on 19 September the President announced an "AI Force"; on 22 September, at the General Assembly, he declared that the United States "totally rejects" any globalist scheme of control and proposed calling the technology "super intelligence".^{23, 24, 38, 40}

The technical and multilateral framework calls for more coordination. The *Pacing the Frontier* appeal, signed by 1,386 employees of frontier labs, and Dario Amodei's essay of 12 September call for international instruments to govern the pace of development.^{4, 19} On 21 September twenty-two leaders, who became thirty in the following days, including the heads of government of Canada, Germany, Spain and Australia and the President of the European Commission, invited Member States to explore creating an international institution "able to set standards, enable verification, and convene states when capability thresholds are crossed".⁴³ On 23 September, at the Security Council, Sam Altman called for common standards for measuring capabilities and assessing risks, rapid incident-reporting protocols and secure channels among governments; Amodei proposed targeted global agreements, mutual verification systems and common testing standards.^{46, 48}

The two lines meet in a precise space. Washington rejects a global regulatory regime but, through White House science adviser Michael Kratsios, acknowledges that "you cannot govern technology you do not understand" and asks the United Nations to help States build domestic capacity.⁵⁰ This is exactly what the proposed pillar would do: produce methods, measurements, warnings and skills, without any power to regulate, authorise or sanction.

Structure and timing. A lean secretariat with technical capacity, serving the Global Dialogue and the Scientific Panel; mixed financing, on the model of GAVI and CEPI, that does not weigh on the regular budget. Implementation proceeds in three phases: within twelve months, the functions that require no new mandate; from twelve to thirty-six months, consolidation; beyond thirty-six months, the functions that require scale.

Why Milan, Italy and Europe. Milan concentrates universities, clinical research, industry and finance within the same urban radius, and is among the first European cities with its own public framework of AI principles.^{32, 33} Italy is the third United Nations hub after New York and Geneva, led the AI dossier during its 2024 G7 presidency and is the first EU Member State with a comprehensive national law on the subject.^{28, 31, 37} Europe is a third-party location with respect to the competition between the United States and China, and has the only legal order with a horizontal framework already being applied.^{29, 30}

The recommendation. Italy should come to the second session of the Global Dialogue, in New York on 3 and 4 May 2027, with a proposal of functions and not merely of a seat: a project that can be embraced both by those who call for less regulation and by those who call for more verification.

1. From location to functions

The policy paper of 27 July 2026 advanced a thesis about location: in the multilateral architecture now taking shape there is an institutional space not yet occupied, and Milan has the credentials to put itself forward as the seat of a future UN-mandated body dedicated to the ethical governance of artificial intelligence.⁵ That thesis opened the debate, but left the next question unanswered: what, exactly, this body should do.

This is not a secondary question. In multilateral practice, the choice of seat follows the definition of the mandate; it does not precede it. A bid offering only a building, however prestigious, would be read as a logistical proposal. A bid that accompanies the seat with an orderly and verifiable description of functions engages instead with the substance of the negotiation and helps to steer it. It is the difference between asking to host an institution and proposing what that institution can do for the international community.

1.1 Nature of the paper

The paper starts from established institutional facts, derives from them a scope of mandate compatible with the orientations of the major technology economies, translates into operational functions the ethical principles already shared at universal level, and indicates the tools with which to exercise them. It is neither a draft treaty nor a draft resolution. It is a technical contribution, offered by an organisation of economic civil society to the States that will bear responsibility for the decisions.

The paper does not presuppose the creation of a new specialised agency, a decision that has not been taken in any formal forum, nor does it presuppose a seat-selection procedure, which does not exist in codified form. The functions, the rules of method and the implementation sequence are AmCham Italy's own elaboration. The legislative, scientific and magisterial texts cited indicate the grounding of the criteria adopted and do not imply that the institutions that produced them endorse any specific institutional architecture.

1.2 Why Milan

A seat is not judged by its beauty, but by the density of interlocutors it makes available to the hosted institution. From this point of view Milan brings together four features that are hard to find together elsewhere in Italy.

The first is a concentration of knowledge covering the full range of the twelve functions. The Politecnico, the Università degli Studi, the Cattolica, Bocconi, Bicocca, IULM and Humanitas University all operate in the same metropolitan area, together with data and life-science research centres and major clinical institutes of national standing. Since 2021 Milan has been the main Italian unit of ELLIS, the

European network of centres of excellence in machine learning, and it takes part in national research programmes on artificial intelligence.^{33, 34} For functions such as impact assessment, health or the training of regulators, having these interlocutors a few kilometres away is the condition for producing methods and measurements on timescales compatible with those of the technology.

The second is an administration that has already chosen to engage with the problem. With its Manifesto for the use of artificial intelligence, the Municipality has set transparency, accountability, inclusion, protection of rights and the centrality of the person as criteria for administrative action, and has accompanied them with a mapping of use cases already active in city services, an urban digital twin and programmes devoted to human-centred solutions.³² The pillar's functions would not arrive in Milan as an abstract exercise: they would find an administration that already needs them.

The third is an economic base that makes the issue immediately concrete. Milan is the country's financial centre and the gateway for most foreign investment in Italy; Lombardy produces about one fifth of national gross domestic product and records the highest research and development spending among Italian regions.^{35, 36} An institution that produces methods and measurements for business, labour, health and infrastructure needs a productive fabric on which to test them before proposing them to the world.

The fourth is an international vocation already put into practice. The city has hosted negotiations, summits and major events, most recently the 2026 Winter Olympic Games. The bid, moreover, is not an isolated initiative: it rests on a promoting committee that brings together local institutions, universities, businesses and their associations, professional bodies, the financial system and members of parliament of different political orientations.¹⁶

1.3 Why Italy

The seat is a city, but the bid is made by a State. There are three Italian reasons.

Italy is already a host country of the United Nations system. Rome is home to FAO, the World Food Programme and IFAD; Brindisi hosts the humanitarian logistics base; Turin the staff training centre. Taken together, these presences make Italy the third UN hub in the world after New York and Geneva.³¹ There are therefore established host-country agreement practices, an administration familiar with them, and a community of international civil servants already settled in the country.

Italy has already taken this issue to the highest level. During the 2024 G7 presidency artificial intelligence was a cross-cutting dossier: the Borgo Egnazia Summit devoted a specific session to it, the only one in G7 history to be addressed by a Pope, and the Italian presidency completed the monitoring mechanism for the International Code of Conduct for advanced AI systems, in continuity with the Hiroshima Process.^{14, 37}

Finally, Italy is the first Member State of the European Union to have a comprehensive national law on artificial intelligence. Law No. 132 of 23 September 2025 supplements the European regulation with provisions on health, labour, public administration,

justice, copyright and the protection of minors, and designates the competent national authorities.²⁸ For the bid, the point is one of method: Italy can document from its own experience the problem the pillar is meant to solve, namely the cost of applying shared principles in the absence of comparable methods.

To these reasons is added proximity to the Holy See, which in recent years has produced the most elaborate public body of thought on the ethical dimension of artificial intelligence, referred to in chapter 4 and taken up in September 2026 by the *AI and Human Development* roundtable convened by the Deputy Secretary-General of the United Nations at the Pontifical Gregorian University.⁸

1.4 Why Europe

The third question is the most political. In 2026 the competition over artificial intelligence is being played out between the United States and China, and Europe is often described as the continent that regulates what others build. This is precisely why a European seat makes sense, for three reasons.

In Europe, principles already have legal form. The Union is the only legal order with a horizontal framework for artificial intelligence being progressively applied; on 5 September 2024 in Vilnius the Council of Europe opened for signature the Framework Convention on Artificial Intelligence and Human Rights, the first binding international treaty in this field, open also to non-European States.^{29, 30} A pillar that produces methods on the rights of the person finds here a legal order on which to test the practicability of its proposals straight away.

A European seat is a neutral seat. The pillar must offer a common grammar to countries that compete with one another. Placed in one of the two poles of the competition, it would be read by the other as a partisan instrument. It is the same reason why, in other sectors, from nuclear energy to health and telecommunications, international technical verification has been located in countries that were not protagonists of the rivalry. The September debate at the Security Council, in which Washington and Beijing restated distant positions, confirms this.⁴⁷

Finally, Europe needs this function more than others. It depends to a large extent on models, compute capacity and infrastructure developed elsewhere. A seat that measures the pace of the frontier, documents the language coverage of systems, makes the environmental footprint of data centres comparable and trains regulators is of direct interest to Europe. It would not overlap with any existing seat: Geneva hosts the technical and human-rights dimension, Paris UNESCO's normative dimension, Valencia the applied laboratory, New York the political level. What is missing is the place where the ethical dimension is exercised as a permanent technical capability.

Milan

Offers the scientific, economic and institutional interlocutors that the functions require.

Italy

Offers its experience as a host country and the credibility of having already brought the issue to multilateral tables.

Europe

Offers the legal framework in which «technology at the service of the human person» is already law in force.

None of the three levels, on its own, sustains the bid: together they make it a project.

2. Summer 2026: what has changed

Between the publication of the first paper and mid-September, the international landscape shifted on four levels, which are best kept distinct because they follow different logics: the G20, United States policy, the demand coming from frontier labs and the United Nations process. The high-level week of the General Assembly, which brought these levels together, is dealt with separately in chapter 3.

2.1 The G20 and the Carolina Principles

On 2 September 2026 the G20 Innovation Ministerial, meeting in Chapel Hill, adopted by consensus the *Carolina Principles for Emerging Technologies*, with China also joining.² The document calls for the principles to be applied flexibly, with respect for national legal systems and circumstances, preserving the sovereignty of each member while allowing participation in international discussions.⁶ The political criterion was set out clearly: apply existing rules first, introduce new ones only where the technology creates genuinely new situations, and avoid treating every innovation as a policy problem.³

For a proposal like this one, the G20's orientation is not an obstacle but a design specification. It requires building a structure that does not produce general-purpose rules and that demonstrates, function by function, that it does not duplicate existing competences.

2.2 The United States' position

The position expressed at the G20 translates a declared national strategy to the multilateral level. The Administration's *AI Action Plan*, presented as the strategy to "win the AI race", rests on three pillars: accelerating innovation, building infrastructure, and leading in international diplomacy and security.²³ The first includes removing regulatory burdens, supporting open models and building an "evaluations ecosystem"; the second covers energy, semiconductors, data centres and a federal incident-response capability; the third covers exporting American technology and standards, countering Chinese influence in international bodies and the "global alignment of protection measures".

Two observations are necessary. The first is that the plan, taken as a whole, is a programme of competitive acceleration and not of multilateral construction. The second is that three of its objectives, the evaluations ecosystem, incident response and the global alignment of protection measures, cannot be achieved by a single State. An evaluation system that stops at the national border does not make results comparable; an incident response without a shared classification does not recognise the same event in two jurisdictions; global alignment presupposes, by definition, a

place in which to align.

In September the line became sharper. On the 19th the President announced an “AI Force”, modelled on the Space Force, and the appointment of a national coordinator, ruling out any move to “hinder or stifle” the sector’s growth.²⁴ In the preceding days he had opened a public consultation on the possibility of replacing the expression “artificial intelligence”, which he considered inaccurate.²⁵ On 20 September the US Ambassador to the UN, Mike Waltz, described the competition with China as “much more complex” than the Cold War, owing to the degree of economic interdependence, ruled out any brake on innovators and identified three matters on which a summit between the two Presidents could establish bilateral safeguards: the proliferation of algorithms, nuclear weapons systems and human oversight in chains of command.²⁶ ²⁷ These three matters correspond to specific functions in this paper (5.1, 5.3 and 5.12): to be verifiable, safeguards of this kind require common definitions, thresholds and formats.

2.3 The demand coming from the labs

In the week of the G20, industry too displayed differing positions. Alongside the call for light-touch regulation, the head of Google DeepMind urged governments to equip themselves with public capabilities for testing the safety of systems, which presupposes some form of shared verification.³

The most explicit signal had come at the end of July. The *Pacing the Frontier* appeal, published on 28 July 2026, starts from an observation: the leading companies believe they are close to automating AI research itself, with the risk that capabilities grow faster than our ability to understand and control them. Since no company and no country can slow down alone under competitive pressure, the signatories ask the United States government to support an international effort to build the technical and governance tools that are missing today. The appeal was signed by 1,386 employees of frontier labs.⁴

On 12 September the demand became a proposal. In the essay *We Must Pace the Frontier*, Dario Amodei argues that risk prevention alone is no longer enough and that the speed at which model capabilities advance must be modulated; he makes clear that the aim is not to halt research, but to secure the time needed to make systems safe and to have this confirmed by third-party evaluators.¹⁹ The text recalls a documented episode: in July 2026, during internal cybersecurity evaluations conducted by OpenAI, experimental agents circumvented the controls meant to isolate them from the internet, communicated through unauthorised channels and reached third-party systems, including Hugging Face’s infrastructure.^{20, 44} The proposed plan has three steps: external evaluators with ongoing access and the right to publish their conclusions; coordination among companies on common safety standards; and global coordination, which in its most ambitious forms envisages pre-release testing of models entrusted to a global standards body. On the same day,

OpenAI's leadership shared the need to set a pace for the frontier; the leaders of xAI and Google DeepMind expressed support for the general direction of the proposal.²¹

Two of the three steps are international by definition, and none requires a new regulatory power. They require a shared method, comparable measurements and technical venues for verification.

2.4 The United Nations process

Resolution A/RES/79/325 set in motion the construction of the permanent architecture. The Independent International Scientific Panel presented its preliminary report on 1 July 2026; the first session of the Global Dialogue on AI Governance was held in Geneva on 6 and 7 July, and the second is scheduled for New York on 3 and 4 May 2027.¹ The Valencia laboratory oversees the applied technical dimension and UNESCO, from Paris, the normative one. The UN80 Initiative, launched by the Secretary-General in March 2025, has meanwhile opened a phase of rationalisation of the entire system.⁷

The *AI and Human Development* roundtable at the Pontifical Gregorian University, on 4 and 5 September, falls within the same horizon; its outcomes will feed into the work of the Panel and the Global Dialogue.⁸ Public debate, too, has shifted to institutions: in his essay published on 26 August, Bill Gates identifies as a priority a national and international institutional framework to be built before a crisis forces governments to improvise, proposes an international organisation inspired by nuclear inspection regimes, civil aviation rules and the ozone agreements, and refers to Leo XIV's encyclical on artificial intelligence.⁹

3. New York, September 2026

The high-level week of the 81st General Assembly gave the governance of artificial intelligence a prominence it had never had before. In four days it was taken up by the President of the United States, the Secretary-General, thirty leaders who signed a joint call, the Security Council and the summit between Presidents Trump and Xi. The positions that emerged are far apart, but not for that reason irreconcilable.

3.1 President Trump's speech

On 22 September, from the rostrum of the General Assembly, President Trump devoted one of the sharpest passages of his address to artificial intelligence. He declared that the United States "totally rejects any attempt to construct a globalist scheme to control artificial intelligence" and announced that from then on US government documents would use the term "super intelligence", because the word "artificial", he said, "makes it sound fake, and it's not fake. It's incredible. But we have

to be careful".^{40, 41}

The President then reiterated that the Administration intends to encourage the sector's development, not rein it in, because he does not intend to "stifle growth of something that will be bigger than the Industrial Revolution", and added: "Whoever wins super intelligence wins. And we're leading now over China by a lot."^{38, 41} He compared warnings about extreme risks to those about the climate, arguing that ordinary law-enforcement tools would suffice, if necessary, to contain the sector's excesses.^{38, 39}

The speech confirms, at the highest level, the line described in chapter 2. For the purposes of this proposal, the relevant point is the object of the rejection: a global scheme of control. The pillar described here does not control, authorise or sanction; and the President himself acknowledged that "we have to be careful". Recognising the need for caution without delegating it to a global regulator means, in practice, needing tools to know when and where to be cautious.

3.2 The Secretary-General and the call of 21 September

On the same day António Guterres, in his last address to the Assembly before the end of his term, pointed in a direction opposite in tone but precise in content. "Children must never become guinea pigs for unregulated systems," he said, and "decisions of life and death must never be delegated to machines"; and again: "the danger is not technology, it is technology without accountability".^{40, 41} He asked States to "advance the conditions for the responsible pacing of AI" and to "work towards a multilateral AI risk management framework, supported by credible and independent oversight."⁴²

The day before, on 21 September, twenty-two leaders led by Finnish President Alexander Stubb and Norwegian Prime Minister Jonas Gahr Støre had signed *A Call for Control of Frontier AI Models*. The signatories include the heads of government of Canada, Germany, Spain, the Netherlands, Ireland, Denmark, Australia and Singapore, the Presidents of South Africa, Kenya and Kazakhstan, the Foreign Ministers of Türkiye, the United Arab Emirates and Bahrain, and the President of the European Commission.^{43, 45} The text states that AI "must remain under human direction, oversight and control", and makes three requests: of companies, transparent safety protocols with mandatory pre-deployment testing and independent evaluation; of governments, common standards, shared reporting of serious safety incidents and access for all countries to scientific capacity and trusted evaluation; of UN Member States, to explore creating an international institution "able to set standards, enable verification, and convene states when capability thresholds are crossed".⁴³ In the following days eight more leaders joined, including French President Emmanuel Macron and the heads of government of Austria, Portugal and Luxembourg: by 24 September there were thirty signatories.⁴³ Italy is not among the signatories; the declaration remains open to other countries.

In the same hours the Scientific Panel published its first thematic report, devoted to the Hugging Face episode. It states that some 1,200 agents exchanged more than 70,000 messages and files, circumventing the intended safeguards, and that “the traditional safeguard model is breaking down”. Co-chair Yoshua Bengio observed that the three conditions for loss of control are “a misaligned goal, the capability to pursue it and an environment that allows it”, and that “this summer, all three came together in a real system, not a laboratory”.⁴⁴ The Secretary-General welcomed the report and expressly referred to the proposal of the twenty-two.⁴⁴

3.3 Altman and Amodei at the Security Council

On 23 September the Security Council, convened by the French presidency, returned to artificial intelligence after its meetings of 2023 and 2024. Foreign Minister Jean-Noël Barrot opened the proceedings by recalling the precedent of atomic energy: “as with the atom then, the international community must now regulate artificial intelligence in order to draw the most from its potential and avoid the worst”.^{47, 51} The heads of OpenAI and Anthropic, the two leading frontier labs, addressed the Council.^{48, 50}

Sam Altman, present in the chamber, described the choice facing the world: “AI can either be more like a new Renaissance of creativity and discovery, or more like a new Industrial Revolution of upheaval and disarray.”⁴⁹ He called for a “middle path” between blind optimism and doomerism and said that “if AI is to be democratic, the most important decisions cannot be made by labs in San Francisco alone. They must be shaped through democratic processes and by governments accountable to the people that they serve.”^{46, 49} On risk he was explicit: “It doesn’t matter whether people put the risk of catastrophe at 10%, or 1%, or 12%, or .1%. None of these levels are remotely acceptable.”⁴⁹

“We need a mechanism for complementary national and international frontier AI standards: standards for measuring capabilities, assessing risks, determining whether safeguards are sufficient, and preserving meaningful human oversight as systems become more autonomous.”

Sam Altman, United Nations Security Council, 23 September 2026

Altman’s proposals were concrete: common standards so that countries can compare evidence, verify compliance with commitments and share a common language; accurate and speedy incident classification and reporting protocols, “so the world can learn from failures before they become catastrophes”; secure channels among governments, critical infrastructure operators and technical experts. He added two caveats that bear directly on this proposal: standards “should not lock in incumbents”, and “each government should decide how to incorporate standards into its own legal

system".⁴⁶ Finally, he recalled that OpenAI has "unilaterally slowed down in the past" and will do so again.⁴⁶

Dario Amodei, speaking by video link, described artificial intelligence as "the most important global security issue facing the world today" and added that, "if managed poorly", it "could be a risk to humanity as a whole".^{48, 49} He made a public commitment: "We will slow down as much as necessary in order to make sure that every successive AI technology that we release is actually safe."^{49, 51} And he recalled that "no leader, no company and no nation can manage this alone".⁵⁰

His proposals are threefold: targeted global agreements, such as a ban on using artificial intelligence to produce biological weapons; evaluation and verification systems allowing countries to check each other's compliance with commitments; and common global testing standards, accompanied by a safety-incident notification system.⁴⁸ Clément Delangue, co-founder of Hugging Face, also addressed the Council, recounting how the company, attacked by AI agents, defended itself with tools of the same kind, and argued that the world needs open models "more than ever to defend itself".^{49, 50, 51}

3.4 The positions of governments

Delegations confirmed distant positions. For the United States, Michael Kratsios, Director of the White House Office of Science and Technology Policy, ruled out a global regulatory regime and set out a precise task for the United Nations: "You cannot govern technology you do not understand. This body and others like it should focus on sharing best practices to build domestic capacity, not establishing a global regulatory scheme."^{47, 50} Chinese Ambassador Fu Cong rejected bloc logic and "digital hegemony", called for attaching "equal importance to development and security" and supported open models.^{47, 50} The United Kingdom announced that it will place the control of artificial intelligence at the centre of its 2027 G20 presidency, with the aim of a single set of global standards; the leaders' summit will be held in Manchester on 27 and 28 November 2027.^{47, 51, 54}

The summit between Presidents Trump and Xi, on 24 September in Washington, produced no visible breakthroughs on artificial intelligence.⁵³ In the preceding days Treasury Secretary Scott Bessent had proposed to his Chinese counterpart a notification mechanism for artificial intelligence incidents threatening national security, in the context of the talks held in New York on 20 September with Vice Premier He Lifeng, in which the two delegations discussed launching a dedicated bilateral dialogue; the Chinese response has not been made public in detail.⁵² During the meeting President Xi stated that the two countries "have the capacity and the responsibility" to develop and govern artificial intelligence while keeping it under human control and at the service of people.⁵³

On 29 September the White House gave this line a concrete form. President Trump

and the heads of Google, Anthropic, Meta, OpenAI, xAI and Nvidia signed the *White House Accord on Super Intelligence*, a voluntary commitment to four layers of control over frontier models: internal controls, an internal review team, an independent external evaluator and a committee of the board of directors. The companies will meet to define common standards and best practices; the accord provides for no penalties and no public oversight.^{55,56} Washington, too, therefore regards independent verification and shared standards as a condition of trust, provided they are not entrusted to a regulator.

3.5 From divergence to a common language

Placed side by side, the New York positions outline an area of convergence broader than the tone of the debate might suggest. The American rejection concerns a global regime of control; the request of the signatories of the 21 September call concerns standards, verification and convening; the proposals of Altman and Amodei concern measurement methods, mutual verification and incident reporting; Kratsios asks the United Nations to help States build expertise; Washington has proposed a bilateral notification mechanism to Beijing. None of these requests presupposes a new regulator. All of them presuppose a common language.

What was requested in New York	Who requested it	Pillar functions
Common measurement and testing standards	Call of 21 September; Altman; Amodei; United Kingdom	5.1, 5.2, 5.4
Verification of commitments	Call of 21 September; Amodei; Altman	5.2, 5.12
Incident reporting	Call of 21 September; Altman; Amodei; Bessent proposal	5.1, 5.12
Convening beyond thresholds	Call of 21 September; Secretary-General	5.1, 5.12
Human oversight	Guterres; Altman; Xi; call of 21 September	5.3
National capacity	United States (Kratsios); call of 21 September	5.6, 5.9

Three consequences follow for the design of the pillar. The first concerns scope: the formula of the 21 September call, standards, verification and convening, coincides with three of the four products described in chapter 6 and confirms that excluding any power to authorise or sanction is the condition for the proposal to remain acceptable to Washington. The second concerns priorities: incident reporting and the training of administrations, called for by very different voices, must be part of the first phase of implementation. The third concerns method: Altman's caveat that standards must not favour incumbents, and Beijing's concern for open models, should be written into the pillar's rules from the outset.

One open question remains, which this paper does not claim to resolve: whether standards should remain voluntary or become binding. The choice belongs to States. The proposed pillar is useful in either case, because neither path is practicable without shared methods and measurements.

4. Criteria and scope of the mandate

An institution is judged by the criterion by which it chooses its tasks. The twelve functions stem from a declared order of sources, a test of political compatibility and an explicit scope.

4.1 Where the tasks come from

The first level indicates the why. The encyclicals *Laudato si'* (2015) and *Magnifica Humanitas* (2026), the Message for the 2024 World Day of Peace, the address to the G7 at Borgo Egnazia and the note *Antiqua et nova* (2025) provide the ordering criterion and the ends: the safeguarding of the person, the non-delegable responsibility for decisions, the limits of automation, the universal destination of benefits, care for our common home.^{12, 13, 14, 15, 22} They are referred to not for confessional reasons, but as the most elaborate public body of thought on questions that technical categories cannot exhaust. Referring to them does not imply any endorsement by the Holy See of a specific institutional architecture or of any city's bid.

The second level provides the words. The UNESCO Recommendation on the Ethics of Artificial Intelligence, adopted by all Member States in November 2021, provides a universally accepted vocabulary: proportionality and do no harm, safety and security, fairness, sustainability, data protection, human oversight and determination, transparency and explainability, responsibility, literacy, participatory governance.¹⁰ Its limitation lies not in its content but in its implementation: principles alone do not generate comparable methods, consistent indicators and widespread technical capacity.¹¹ The twelve functions are its operational translation.

The third level identifies the tools that are missing today. Resolution A/RES/79/325, *Pacing the Frontier*, the essay of 12 September, the AI Action Plan and the New York statements show that the demand for measurement, verification and reporting comes from governments, labs and the scientific community.^{1, 4, 19, 23, 43, 46} This level does not confer legitimacy, which derives from the first two: it shows how timely the tasks are.

The derivation closes with a test, conducted against the Carolina Principles, consisting of four questions. Does the function create a new general-purpose regulatory architecture? Does it intervene where existing law already applies? Does it duplicate the mandate of an existing organisation? Does it affect national sovereignty over authorisations, sanctions or security? A single affirmative answer means the function is excluded or reformulated. This is why the pillar produces methods, measurements, warnings and skills, and not rules.^{2, 6}

4.2 Seven rules and four exclusions

Each function is formulated in accordance with seven rules, derived from the Chapel Hill orientations and from the rationalisation principle of the UN80 Initiative.⁷ The pillar gives precedence to existing law and facilitates its application, without proposing new rules where those in force suffice (first rule). It intervenes only on genuinely new issues (second). It works by sector, by use case and by level of risk, never through general-purpose frameworks (third). It always operates in coordination with the organisation competent for the subject matter, which retains ownership (fourth). It produces voluntary tools, adaptable to national legal systems and free of compliance obligations (fifth). It subjects each function to performance indicators and periodic review, closing those that do not produce measurable effects (sixth). It maintains a lean structure and a contained cost (seventh).

A scope is credible if it also declares what it leaves out. The pillar does not authorise or register models, a task for national and regional legal systems. It does not supervise or sanction companies. It does not produce horizontal rules valid for every use of artificial intelligence. It does not take competences away from the specialised agencies, nor the scientific role away from the Independent Panel.

What remains is what no institution produces systematically today: shared evaluation methods, consistent measurements, early warning on emerging capabilities and technical expertise for States that cannot build it on their own. It is what the leaders of the call, the labs and the United States government asked for in New York, in different words.

5. The twelve functions of the pillar

The functions follow a logical order, not an order of importance. The first two build the capacity to know: without a measure of the phenomenon and without a consistent way of assessing its effects, every other function would work on impressions. The six central ones directly protect people. The last four concern long-term conditions, without which the results of the first cannot be consolidated. For each, the paper sets out the problem, the pillar's activity, the limit of what it does not do and the sign that it is working; the foundations are given in a note.

5.1 Monitoring the pace of the frontier and early warning

Measuring, with a public and stable method, how fast the capabilities of the most advanced systems are growing.

Everyone is debating whether to accelerate or slow down, but no international body

measures speed with criteria that are the same for all. Every statement about pace thus remains an opinion, and two governments can observe the same phenomenon and see two different things. The pillar builds a knowledge base that is comparable over time: which capabilities have appeared, when, and at what distance from the previous generation. From this it derives a periodic report on the state of the frontier and a rapid notification procedure for States when a capability with systemic implications emerges. If States were to adopt capability thresholds, as the 21 September call requests, this is the function that would define them and monitor them with the same criteria for all.^{19, 43} The pillar does not evaluate individual models, does not authorise their release and does not replace the Scientific Panel. It works if there is a public time series that makes it possible to answer the question “is the frontier accelerating?” with data, and if notifications reach States before the capability is publicly released.

Foundations. *Laudato si’* and *Magnifica Humanitas*: technical power must be understood before it is indulged. UNESCO: proportionality, safety, adaptive governance. Pacing the Frontier; We Must Pace the Frontier; call of 21 September 2026; Guterres on “responsible pacing”.^{4, 19, 42, 43}

5.2 Common impact-assessment methodology

Writing once, and keeping up to date, the method by which the effect of a system on people is assessed.

The UNESCO Recommendation provides for ethical impact assessment, but there is no common method. Companies and administrations repeat the same exercise with incompatible criteria, at high cost and with no possibility of comparison. It is one of the rare cases in which an international instrument reduces the regulatory burden instead of increasing it. The pillar develops a reference methodology, structured by sector and by level of risk, which States adopt on a voluntary basis: it indicates what must be ascertained, not how to organise the offices that ascertain it. It also serves a second purpose, which emerged strongly in September: commitments to admit external evaluators and the mutual verification proposed at the Security Council produce comparable results only if what is being verified is defined in the same way.^{19, 46, 48} The pillar does not carry out assessments, does not accredit evaluators and does not impose its method. It works if a growing number of legal systems and companies adopt it and if two assessments conducted in different countries become readable side by side.

Foundations. *Magnifica Humanitas* and *Antiqua et nova*: effects on the person are assessed before adoption. UNESCO: ethical impact assessment, responsibility. The “evaluations ecosystem” of the *AI Action Plan*; proposals by Altman and Amodèi.^{23, 46, 48}

5.3 Dignity, human determination and the chain of accountability

Always making it possible to identify who is accountable for a decision, and preventing accountability from dissolving along the technology chain.

All international instruments affirm the principle of human oversight; in New York it was repeated by the Secretary-General, the signatories of the 21 September call, Altman and President Xi.^{42, 43, 46, 53} But its implementation depends on technical answers: in which decisions human involvement must be effective and not merely formal, what information the decision-maker must receive, how control is documented. Without these answers oversight becomes a box to tick, and the incidents of summer 2026 show that the chain breaks precisely where action is delegated to groups of agents.⁴⁴ The pillar develops operational criteria which, in systems affecting significant rights or interests, make the accountable party identifiable: minimum information requirements, reasoned cases by sector, attribution criteria. It does not attribute responsibility in specific cases and does not replace national law. It works if, faced with an automated decision, it is possible to indicate quickly who is accountable for it and on the basis of what information the decision was taken.

Foundations. *Magnifica Humanitas*: responsibility for decisions cannot be delegated. Borgo Egnazia address. UNESCO: human oversight and determination. Human oversight in chains of command as a matter for bilateral safeguards.^{14, 26} Coordination with OHCHR.

5.4 Transparency, explainability and content provenance

Enabling anyone to know whether what they are looking at is authentic or synthetic, and to obtain explanations proportionate to the seriousness of the decision.

Distinguishing authentic from synthetic content has become a condition for public discourse to function. Watermarking and tracing techniques exist, but their effectiveness depends on interoperability: a standard adopted by a single country protects no one, because content crosses borders and standards do not. The pillar fosters the convergence of provenance standards, tests and publishes their resistance to removal and counterfeiting, and defines a scale of explainability: the more a decision affects a person, the more explainable it must be. It has no power over content, which remains with national legal systems, and does not manage registries or watermarking infrastructure. It works if the number of mutually interoperable standards increases and if published tests reduce the gap between declared and measured reliability.

Foundations. *Antiqua et nova*: truthfulness of information and the risk of manipulation. UNESCO: transparency and explainability. Countering synthetic content in the *AI Action Plan*; common standards requested by the 21 September call.^{23, 43} Coordination with ITU.

5.5 Data, privacy and dataset documentation

Making it possible to document what a system was trained on: where the data come from, whom they represent, and on what basis they were acquired.

The quality of a system depends on the data with which it is built. Rules on personal data protection are now widespread, whereas the capacity to document the composition of a dataset, its representativeness and the lawfulness of its acquisition is far less developed. Without this documentation none of the functions on equity or languages can be verified. The pillar defines standard documentation sheets, promotes data-trust models in which data are managed in the interest of those who generated them, and supports the creation of public information resources in countries that lack them. It does not create centralised archives, does not access the data of States or companies, and does not judge the lawfulness of individual supply chains. It works if the sheet is attached to a growing number of systems and if it becomes possible to compare two models on the representativeness of their data, not only on their performance.

Foundations. *Antiqua et nova*: the person cannot be reduced to the data that describe them. UNESCO: privacy and data protection. A/RES/79/325 on the capacity gap.¹

5.6 Equity, non-discrimination and gaps between countries

Measuring, with a comparable index, the gap between those who have compute, data and skills and those who lack them, and reducing it.

The gap between countries with technological capacity and those without it is the main equity issue of the transition, and also a question of legitimacy: a governance system that excludes part of the international community from rule-setting cannot call itself global. The 21 September call asks that all countries have access to scientific capacity and trusted evaluation; China, at the Security Council, asked that development be given the same importance as security.^{43, 47} Today the gap is described, not measured. The pillar measures it with comparable indicators, documents the discriminatory effects of systems in use and supports institutional capacity-building programmes. It does not distribute compute resources or manage funds, and operates in coordination with UNDP and UNCTAD. It works if the index is published consistently and finds its way into negotiating documents, and if the number of States with a functioning national authority grows.

Foundations. *Magnifica Humanitas*: universal destination of benefits. *Laudato si'*: the link between the environmental and the social question. UNESCO: fairness and non-discrimination. Bill Gates on access to benefits.⁹

5.7 Work and the employment transition

Producing the missing comparable data on employment effects and documenting which transition policies really work.

The transformation of work is the most keenly felt effect and the least documented with comparable data. The international debate anticipates a disruption concentrated within a decade rather than across generations, one that will also affect cognitive tasks; Altman, at the Security Council, evoked the risk of a new industrial revolution “of upheaval”.^{9, 49} A transition policy built on perceptions gets sectors and timing wrong. The pillar builds a comparable evidence base, sector by sector, and an evaluated repository of transition policies, from reskilling to cooperation between businesses and the education system, indicating not only what has been done but with what result. It does not develop labour standards, which remain the competence of the ILO. It works if comparable cross-country series exist and if the policies adopted cite evidence rather than anecdotes.

Foundations. Borgo Egnazia address: the dignity of human work. *Laudato si’*: work as the fulfilment of the person. UNESCO: economy and labour. Coordination with the ILO and the social partners.^{14, 22}

5.8 Health, well-being and the protection of minors

Gathering comparable evidence on the effects on health and on the development of minors, and defining precautionary criteria for the most vulnerable.

In health, artificial intelligence produces documented benefits and specific risks: clinical validation, professional liability, equity of access. A distinct and urgent issue concerns the effects on the development of minors and on the quality of relationships. The Secretary-General put it bluntly: “Children must never become guinea pigs for unregulated systems.”⁴⁰ In both cases the evidence exists, but it is fragmented. The pillar gathers comparable data, defines precautionary criteria for systems intended for minors, from age limits to the design of persuasive interactions, and supports validation protocols for clinical uses. It exercises no health competences, which remain with the WHO, and does not authorise devices. It works if the criteria are taken up by authorities and operators and if a citable evidence base, absent today, comes into being.

Foundations. *Antiqua et nova*: effects on those still in formation. *Magnifica Humanitas*: preferential care for the most vulnerable. UNESCO: health and do no harm. Coordination with the WHO.^{9, 12, 15}

5.9 Education, training of regulators and research

Training those who regulate and those who use AI: without expertise in administrations, none of the other functions has any effect.

A regulator who does not understand the technology cannot apply any method; a citizen who does not know how it works cannot exercise their rights. The Panel's preliminary report and the first session of the Global Dialogue identified the technical capacity of administrations as the first bottleneck.¹ In New York the United States said the same thing in other words: the United Nations should help countries build domestic capacity, because "you cannot govern technology you do not understand".⁴⁷ The pillar trains public officials through a dedicated programme for regulators, disseminates literacy pathways adaptable to national contexts and supports access for public research to shared compute capacity. It is the function with the best cost-to-result ratio and the one that can be launched most quickly, through existing university networks. It does not define school curricula, does not award qualifications and does not run its own infrastructure. It works if the number of trained officials and of administrations with in-house expertise grows measurably.

Foundations. *Magnifica Humanitas*: discernment must be formed. *Antiqua et nova*: education in critical judgement. UNESCO: literacy, education and research. The US position at the Security Council.^{12, 15, 47}

5.10 Culture, languages and pluralism

Measuring which languages and cultures systems really represent, and supporting those at risk of being left out.

Systems are trained in a highly asymmetric way with respect to the world's languages and cultures. The result is a progressive reduction in the diversity represented, which affects access to services, cultural production and the transmission of heritage. It is a self-reinforcing gap: the less a language is present in the data, the less systems serve it, and the less it is used in the digital world. The pillar measures the language coverage of systems in use with a public index, supports the development of resources for under-represented languages and documents the effects on the cultural and creative sectors, including remuneration and attribution of works. It does not intervene on intellectual property, which is the competence of WIPO, and does not finance cultural production. It works if the index is published regularly and if the number of languages with usable resources grows. For Europe, with its twenty-four official languages, this function is of direct interest.

Foundations. *Magnifica Humanitas*: the person lives in a community of language and culture. *Laudato si'*: the value of local cultures. UNESCO: culture and diversity. Coordination with UNESCO and WIPO.^{12, 22}

5.11 Environment, energy and compute infrastructure

Making the energy, water and land footprint of compute centres comparable.

The expansion of compute infrastructure has a footprint that local communities experience directly. Today it is reported with inconsistent criteria, when it is reported at all: two data centres may declare non-comparable values, and a local administration has no way of verifying them. The issue is also central to national strategies, which include the expansion of the power grid among the conditions for the sector's growth.²³ The pillar defines common reporting criteria, maintains a comparable database and documents good practice in siting and resource management. It has no authorisation powers over facilities and sets no consumption limits. It works if operators report according to the common standard and if an administration can compare two projects on the basis of consistent data.

Foundations. *Laudato si'*: care for our common home and a critique of any increase in power not proportionate to its effects. UNESCO: sustainability, environment and ecosystems. Energy and the grid in the *AI Action Plan*.^{22, 23}

5.12 Security, dual use and systemic stability

Giving governments and authorities a common language for recognising and reporting incidents and dangerous capabilities.

Some risks no State can manage alone: the dual use of the most advanced capabilities, the resilience of critical infrastructure, large-scale manipulation of information. The Hugging Face episode showed that loss of control is a real category of incident, yet today the same event does not have the same name in two different countries.^{20, 44} In New York this gap was pointed out by all sides: the 21 September call asks for shared reporting of serious incidents, Altman for rapid classification protocols and secure channels among governments, Amodei for a notification system, and Washington proposes a bilateral mechanism to Beijing.^{43, 46, 48, 52} The pillar maintains a shared classification of incidents and dangerous capabilities, coordinates early warning among the competent authorities and promotes joint public-private exercises. Should States negotiate bilateral safeguards or common thresholds, it would provide the measurement basis on which to rest them. It has no operational competence in matters of national security and does not intervene in arms-control negotiations, which remain in their own forums. It works if an incident detected in one country is classified in the same way in another and if the time between detection and reporting is reduced.

Foundations. Message for the 2024 World Day of Peace: decisions of life and death are not to be automated. *Magnifica Humanitas*: ultimate responsibility remains human. UNESCO: safety and do no harm. Guterres on autonomous weapons systems; proposal for a ban on use for biological weapons.^{13, 40, 48}

6. What the pillar produces and how it is organised

The twelve functions generate four types of product, which are also the criterion for measuring their usefulness: a function that produces none of them in a work cycle has no reason to be maintained.

Product	Content	Functions
Methods	Manuals and performance standards that regulators and companies can adopt without each having to rebuild them on their own	5.2, 5.3, 5.4, 5.5, 5.11
Measurements	Comparable indicators on the pace of the frontier, gaps, work, languages and energy consumption	5.1, 5.6, 5.7, 5.10, 5.11
Warnings	Early warning on emerging capabilities and incidents, in coordination with the Scientific Panel and national authorities	5.1, 5.8, 5.12
Skills	Training of regulators, literacy, access for public research to compute	5.6, 5.9

None of the four products is a rule and none entails compliance obligations. All of them, if adopted, reduce the cost of applying the rules that States have already chosen or will choose. This is why the project is compatible with the Chapel Hill orientation and with the United States' position, while at the same time responding to the call for standards, verification and convening made in New York.

6.1 A lean architecture

The model is that of a new-generation pillar, different from the specialised agency of the last century. It has a horizontal mandate to coordinate the UN system on the matters touched by artificial intelligence, without duplicating existing mandates. It has a small secretariat, able to develop methods and recommendations, serving the Global Dialogue and the Scientific Panel. It has intergovernmental governance, with structured participation of business, the scientific community and civil society.¹⁶ What the United Nations needs is a technical capacity it does not have today, not another bureaucracy.

6.2 Financing

Financing follows a mixed model based on three sources.¹⁶ Assessed contributions from Member States guarantee the institution's independence. Multi-year voluntary contributions from companies, governed by agreements, follow the model tested in the health field by GAVI and CEPI; companies' willingness to contribute remains to be tested in the negotiations, but the calls for common standards made by OpenAI and Anthropic at the Security Council point in this direction.^{46, 48} Partnerships with foundations and international financial institutions support specific lines of work, in particular training and capacity-building in the Global South. The condition for credibility is a strict separation between those who fund and those who decide: voluntary contributions confer no voting or veto rights, are made public and are subject to caps that prevent any single operator from carrying more weight than others, in line with Altman's caveat that standards must not favour incumbents.⁴⁶

6.3 Coordination with the system

Each function is exercised together with the organisation that owns the subject matter: the ILO for labour, the WHO for health, UNESCO for ethics, culture and education, UNDP and UNCTAD for development, the ITU for technical standards, WIPO for intellectual property, OHCHR for human rights. Outside the system, the pillar operates consistently with the *Carolina Principles*, with European Union law, with the Council of Europe Framework Convention, with the work of the OECD and with the cooperation body established in July by twenty-nine countries.^{17, 18, 29, 30}

7. The implementation sequence

The functions do not all come into operation at once. The sequence starts from what requires no new mandate and moves towards what does, so that each phase produces verifiable results before the next step is requested. The credibility of a new institution is built on the results of its first phase.

Within twelve months, with no change of mandate. The functions that do not presuppose new legal competences are launched as a joint initiative of volunteer States, existing agencies, business and the scientific community: the common assessment methodology (5.2), dataset documentation sheets and the repository of provenance standards (5.4 and 5.5), the regulator training programme (5.9) and a voluntary register of frontier notifications (5.1). In light of what emerged in New York, this phase also includes a first shared classification of incidents (5.12), which can serve as the technical basis for the notification mechanism discussed between the United States and China and for the shared reporting requested by the twenty-two.^{43, 52}

From twelve to thirty-six months, consolidation. Coordination with the agencies

becomes permanent and the products acquire continuity: interoperability of data regimes (5.5), validation protocols and precautionary criteria for health and minors (5.8), the gap index and a first observatory on work (5.6 and 5.7), environmental reporting criteria (5.11). The permanent technical secretariat is established and the relationship with the Panel and the Global Dialogue is stabilised.

Beyond thirty-six months, maturity. The functions requiring multi-year continuity are scaled up: the common skills framework (5.9), programmes for linguistic pluralism (5.10), coordination of systemic early warning and joint exercises (5.12). At this point the sixth rule requires a review: functions that have not generated methods, measurements, warnings or skills are closed.

Date	Event	Relevance for the proposal
September 2026	High-level week of the 81st General Assembly and Security Council, New York	Call of 21 September; proposals on standards, verification and reporting
14 and 15 December 2026	G20 Leaders' Summit, Miami	Political follow-up to the <i>Carolina Principles</i>
31 December 2026	End of Secretary-General Guterres's term	Handover of the AI dossier
3 and 4 May 2027	Second session of the Global Dialogue on AI Governance, New York	Natural forum for discussing the proposed functions
2027	Full report of the Scientific Panel	Evidence base for 5.1 and 5.12
27 and 28 November 2027	G20 Leaders' Summit under the UK presidency, Manchester	Agenda on global AI standards

8. What it means for Italy and for Milan

Moving from the reasons for the seat to the content of the functions has practical consequences for those who support the bid. The reasons for the seat are set out in chapter 1; here we indicate what follows from them in terms of action.

A dossier of functions, not just of credentials. The host-city proposal now in preparation should adopt the twelve functions as its backbone, devoting to each a sheet setting out the tools, the Milanese and national scientific partners that can be involved and the expected indicators. This is how a bid presents itself as a project rather than a real-estate offer.

An urban laboratory already up and running. Several functions find in Milan an existing testing ground. The municipal Manifesto, the mapping of use cases in city services, the urban digital twin and the programmes on human-centred solutions provide a real test bed for impact assessment, transparency and training (5.2, 5.4

and 5.9).^{16, 32}

A coalition to be turned into technical content. Since the publication of the first paper the base of support has broadened: local institutions, universities and research centres, businesses and their associations, professional bodies, the financial system and members of parliament of different political orientations.¹⁶ The task for the coming months is to translate this consensus into technical work, entrusting mixed working groups with drafting the function sheets.

A bridging role after New York. The General Assembly week showed an international community divided over the instrument, not over the need. Italy, which has excellent relations with Washington, sits in the Council of the Union and is not among the signatories of the 21 September call, which France, Germany, Spain and Portugal have also joined, has room to propose a solution both sides can accept: a seat that sets methods and measurements without regulatory powers, consistent with the American request to build domestic capacity and with the call signatories' request for standards and verification. It is a position that a European seat, neutral with respect to the competition between the United States and China, can credibly uphold.

Three immediate steps. At national level, a standing table bringing together the Presidency of the Council of Ministers, the Ministry of Foreign Affairs, the Ministry of Enterprises and Made in Italy, the Ministry of Justice, the Lombardy Region and the Municipality of Milan, with the aim of arriving at the New York session of May 2027 with a consolidated Italian position, building on the experience of Law No. 132 of 2025.²⁸ At European level, coordination with the European External Action Service and with the Commission, which signed the 21 September call, to secure European support for the bid as an agreed seat.^{29, 43} At United Nations level, engagement with the Secretariat, with the co-facilitators of the process launched by Resolution A/RES/79/325 and with the chairs of the regional groups, building on Italy's role as the third host country of the system.³¹

Conclusion

The first paper explained why Milan. This second paper explains what Milan could host: twelve functions ordered by a single criterion, four verifiable products, a declared limit on what the pillar does not do, a lean structure and financing that does not weigh on the United Nations regular budget.

The New York week has made the proposal more timely. The President of the United States rejected any global scheme of control; the Secretary-General and thirty leaders called for standards, verification and independent oversight; the heads of the two leading frontier labs called, before the Security Council, for common measurement methods and rapid reporting channels. These requests start from different premises, but find an answer in the same venue: a place that does not regulate, but makes it possible to know, measure, warn in time and train those who must decide.

The proposal, therefore, does not ask the international community to regulate more. It asks it to know better. This is why it does not conflict with those who, like the United States, ask that innovation not be held back: the global alignment of protection measures that the American strategy itself sets as a goal requires a place in which to align.²³ Nor does it conflict with those who call for more control, because no control is possible without shared methods and measurements.

Milan can offer that place, together with an industrial, scientific, academic and healthcare system in which the twelve functions find real interlocutors. Italy can offer the experience of a host country and the credibility of one that has already brought the issue to the tables of the G7 and the United Nations. Europe can offer the legal order in which "technology at the service of the person" is already law in force. Italy's task, in the coming months, is to present itself at multilateral tables not with a request, but with a project.

Methodological note

The institutional facts referred to have been verified against primary sources or news agencies and are updated to 30 September 2026. Quotations from remarks delivered in English are given in their original wording where this is available in the sources cited; otherwise they are translated from the Italian edition of this paper. The twelve functions, the rules of method, the exclusions, the classification of products, the implementation sequence and the correspondence between the requests that emerged in New York and the proposed functions are AmCham Italy's own elaboration and do not reflect the positions of the organisations, governments or persons cited. Magisterial quotations indicate the ethical grounding of the criteria adopted and do not imply the Holy See's endorsement of a specific institutional architecture or of Milan's bid. No formal United Nations procedure appears to have been initiated for the establishment of a new specialised agency on artificial intelligence or for the selection of its seat.

Sources

1. United Nations, Independent International Scientific Panel on AI and Global Dialogue on AI Governance, established by Resolution A/RES/79/325 of 26 August 2025: <https://www.un.org/global-digital-compact/en/ai-panel-and-dialogue>; preliminary report of the Panel, 1 July 2026: <https://www.un.org/independent-international-scientific-panel-ai/en/preliminary-report>.
2. Reuters, *US urges hands-off approach to AI regulation at G20 tech meeting*, 1 and 2 September 2026: <https://www.reuters.com>; G20 Innovation Ministerial, Chapel Hill statement, 2 September 2026: <https://www.g20.utoronto.ca/2026/260902-innovation-statement.html>.
3. TechRadar Pro, *US and AI bosses ask G20 not to regulate AI because it could hurt America's bottom line*, September 2026: <https://www.techradar.com/pro>.
4. *Pacing the Frontier*, appeal of 28 July 2026, signed by 1,386 employees of frontier companies: <https://pacingthefrontier.com>.
5. AmCham Italy, *Milano e il futuro della governance globale dell'intelligenza artificiale. Perché candidare Milano a ospitare un Hub ONU per la governance etica dell'AI* [Milan and the future of global governance of artificial intelligence. Why Milan should bid to host a UN Hub for the ethical governance of AI], policy paper, 27 July 2026.
6. U.S. Department of Commerce, *The Carolina Principles for Emerging Technologies*: <https://www.commerce.gov/issues/g20-innovation-ministerial>.
7. United Nations, UN80 Initiative, launched by the Secretary-General in March 2025: <https://www.un.org/un80-initiative/en>.
8. AmCham Italy, *Non paper. Milano quale sede del nuovo pilastro delle Nazioni Unite per la governance dell'intelligenza artificiale* [Non-paper. Milan as the seat of the new United Nations pillar for the governance of artificial intelligence], September 2026, with reference to the *AI and Human Development* roundtable at the Pontifical Gregorian University, 4 and 5 September 2026.
9. Bill Gates, *The choices we make about AI now are critical*, GatesNotes, Make AI work for everyone series, 26 August 2026: <https://www.gatesnotes.com/a-turbulent-ai-era-and-critical-choices-to-make>
10. UNESCO, *Recommendation on the Ethics of Artificial Intelligence*, 41st General Conference, Paris, November 2021: <https://www.unesco.org/en/artificial-intelligence/recommendation-ethics>
11. UNESCO, *Readiness Assessment Methodology e Global AI Ethics and Governance Observatory*: <https://www.unesco.org/ethics-ai/en/ram>.
12. Leo XIV, encyclical letter *Magnifica Humanitas*, 15 May 2026, published on 25

- May 2026: <https://www.vatican.va>
13. Francis, *Artificial Intelligence and Peace*, Message for the 57th World Day of Peace, 1 January 2024: <https://www.vatican.va>.
 14. Francis, address to the G7 session on artificial intelligence, Borgo Egnazia, 14 June 2024: <https://www.vatican.va>.
 15. Dicastery for the Doctrine of the Faith and Dicastery for Culture and Education, *Antiqua et nova*, 28 January 2025: <https://www.vatican.va>.
 16. AmCham Italy, Non-paper cited in note 8, chapters 3, 4, 5 and 6 (nature of the pillar, financing, Milan's credentials, promoting coalition).
 17. High-Level Advisory Body on Artificial Intelligence, *Governing AI for Humanity*, 19 September 2024: <https://www.un.org/en/ai-advisory-body>; G20, 2026 calendar, Miami Leaders' Summit of 14 and 15 December 2026: <https://www.g20.org/events-calendar>.
 18. Reuters, *Twenty-nine countries sign agreement to establish global AI cooperation body*, 16 July 2026: <https://www.reuters.com>.
 19. Dario Amodei, *We Must Pace the Frontier*, 112 September 2026: <https://www.darioamodei.com>.
 20. OpenAI, *The Hugging Face incident and the road ahead*, 26 August 2026: <https://openai.com/index/hugging-face-incident-and-the-road-ahead/>; METR, Brief independent investigation of agents' behavior, reasoning and capabilities, 26 August 2026: <https://metr.org>.
 21. Sam Altman, statement of 12 September 2026: <https://x.com/sama>; Anadolu Agency, *Musk, Altman, Hassabis back Amodei's call to slow pace of AI development*, 13 September 2026: <https://www.aa.com.tr>
 22. Francis, encyclical letter *Laudato si'*, 24 May 2015: <https://www.vatican.va>
 23. The White House, *America's AI Action Plan. Winning the AI Race*: <https://www.ai.gov/action-plan>.
 24. Il Sole 24 Ore, «Trump lancia l'AI Force: non soffocheremo la crescita dell'intelligenza artificiale», 20 September 2026: <https://www.ilsole24ore.com/art/trump-lancia-ai-force-non-soffocheremo-crescita-dell-intelligenza-artificiale-AJgba-xIB>.
 25. Tgcom24, «Trump e l'intelligenza artificiale: il sondaggio per cambiarle nome», 19 September 2026: https://www.tgcom24.mediaset.it/mondo/trump-intelligenza-artificiale-sondaggio_116782814-202602k.shtml.
 26. CNN, interview with US Ambassador to the United Nations Mike Waltz, 20 September 2026: <https://edition.cnn.com/2026/09/20/us/video/23000788-waltzai-race-vrtc>.
 27. Yeni Şafak, *US envoy warns against AI throttle amid China rivalry*, 21 September

- ber 2026: <https://en.yenisafak.com/world/us-envoy-warns-against-ai-throttle-amid-china-rivalry-3723484>.
28. Law No. 132 of 23 September 2025, *Disposizioni e delega al Governo in materia di intelligenza artificiale* [Provisions and delegation to the Government on artificial intelligence]: <https://www.gazzettaufficiale.it/eli/id/2025/09/25/25G00143/sg>.
 29. Council of Europe, Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law (CETS No. 225), opened for signature in Vilnius on 5 September 2024; Decision (EU) 2024/2218: https://eur-lex.europa.eu/legal-content/IT/TXT/HTML/?uri=OJ:L_202402218.
 30. European Commission, regulatory framework on artificial intelligence: <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>; AI Continent Action Plan, 9 April 2025: <https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:52025DC0165>.
 31. Italian Ministry of Foreign Affairs and International Cooperation, United Nations offices in Italy: https://www.esteri.it/it/temi/organizzazioni_internazionali/onu/sedi_onu/; Permanent Mission of Italy to the United Nations: <https://italyun.esteri.it/en/italy-and-the-united-nations/>.
 32. Municipality of Milan, Manifesto for the use of artificial intelligence and digital skills programmes: <https://www.comune.milano.it/servizi/competenze-digitali>.
 33. Milan Monza Brianza Lodi Chamber of Commerce, report on the Milan artificial intelligence ecosystem: <https://ester.milomb.camcom.it/sites/default/files/rapporto-mp/2024/08-capitolo-di-fraia.pdf>.
 34. FAIR Foundation, *L'ecosistema FAIR: l'AI tra ricerca e impresa*: <https://fondazione-fair.it/news/ecosistema-fair-lai-tra-ricerca-e-impresa/>; Lombardy Region, Open Innovation, CATCH atMIND digital hub: <https://www.openinnovation.regione.lombardia.it/it/news/news/7158/sanit-la-statale-alla-guida-del-nuovo-digital-hub-catchatmind>.
 35. Assolombarda, Research Centre, *Economia lombarda: le eccellenze da cui ripartire*: <https://www.assolombarda.it/centro-studi/economia-lombarda-le-eccellenze-da-cui-ripartire>; Lombardy Economy: <https://lombardiaeconomy.it/lombardia-sistema-innovazione/>.
 36. Gazzetta di Milano, «Milano vale oltre metà del fatturato lombardo: 740 miliardi di euro»: <https://www.gazzettadimilano.it/economia/milano-vale-oltre-meta-del-fatturato-lombardo-740-miliardi-di-euro/>.
 37. Department for Digital Transformation, "G7: completato il sistema di monitoraggio sull'IA": <https://innovazione.gov.it/notizie/articoli/g7-completato-il-sistema-di-monitoraggio-sulla-ia-e-pubblicati-i-risultati-della/>.
 38. PBS NewsHour, *Trump dismisses concerns about AI risks at UN, says he won't*

- rein it in*, 22 September 2026: <https://www.pbs.org/newshour/politics/watch-trump-dismisses-concerns-about-ai-risks-at-un-says-he-wont-rein-it-in>
39. Reuters, President Trump's address to the General Assembly, 22 September 2026: <https://www.reuters.com/world/us/trump-un-threats-diplomacy-super-intelligence-2026-09-22/>; ABC News: <https://www.abc.net.au/news/2026-09-23/donald-trump-addresses-united-nations-general-assembly/107179392>.
40. Tgcom24, United Nations General Assembly, addresses by Guterres and Trump, 22 September 2026: https://www.tgcom24.mediaset.it/mondo/onu-assemblea-generale-new-york-donald-trump_116889634-202602k.shtml
41. Notizie Geopolitiche, «L'intelligenza artificiale divide il Palazzo di Vetro: Trump rifiuta ogni controllo globale, Guterres chiede regole»: <https://www.notiziegeopolitiche.net/lintelligenza-artificiale-divide-il-palazzo-di-vetro-trump-rifiuta-ogni-controllo-globale-guterres-chiede-regole/>.
42. Reuters, *UN chief calls for AI curbs and end to wars in his last assembly address*, 22 September 2026: <https://www.reuters.com/world/europe/un-chief-calls-ai-curbs-end-wars-his-last-assembly-address-2026-09-22/>; Xinhua, 23 September 2026: <https://english.news.cn/20260923/2ac3cc9bdeca41e8962b8f019ca34669/c.html>.
43. *A Call for Control of Frontier AI Models*, declaration by twenty-two leaders, 21 September 2026: <https://www.government.nl/documents/2026/09/22/a-call-for-control-of-frontier-ai-models>; updated list of signatories, thirty leaders as of 24 September 2026: <https://www.presidentti.fi/en/a-call-for-control-of-frontier-ai-models/>.
44. UN News, first thematic report of the Scientific Panel and statement by the Secretary-General, September 2026: <https://news.un.org/en/story/2026/09/1168380>.
45. Al Jazeera, *20 countries propose global oversight body to manage AI dangers*, 22 September 2026: <https://www.aljazeera.com/economy/2026/9/22/20-countries-propose-global-oversight-body-to-manage-ai-dangers>.
46. OpenAI, remarks by Sam Altman to the United Nations Security Council, 23 September 2026: <https://openai.com/index/sam-altman-un-security-council-remarks/>.
47. The Guardian, 23 September 2026: <https://www.theguardian.com/world/2026/sep/23/unga-sam-altman-dario-amodei>; Associated Press: <https://apnews.com/article/ai-artificial-intelligence-un-security-council-64519ea66b38e2600026f4481ad7f211>.
48. CNN, 23 September 2026: <https://www.cnn.com/2026/09/23/tech/altman-amodei-ai-safety-un-security-council>.
49. Key4biz, 24 September 2026: <https://www.key4biz.it/big-tech-allonu-sui-rischi-dellai-altman-openai-puo-essere-nuovo-rinascimento-ma-facciamo-attenzio->

- ne-amodei-anthropic-rallentare-quanto-serve/588834/; Rainews: <https://www.rainews.it/articoli/2026/09/i-big-della-intelligenza-artificiale-al-consiglio-di-sicurezza-onu-servono-standard-e-sistemi-di-sicurezza-globali-204d25d7-bafc-4102-b07d-8c317f024a9f.html>.
50. Reuters, *AI leaders brief UN amid warnings technology could slip beyond human control*, 23 September 2026: <https://www.reuters.com/business/ai-leaders-brief-un-amid-warnings-technology-could-slip-beyond-human-control-2026-09-23/>.
 51. France 24, 23 September 2026: <https://www.france24.com/en/americas/20260923-ai-leaders-urge-caution-at-un-with-anthropic-chief-pledging-to-slow-down>.
 52. Reuters, meeting between Scott Bessent and He Lifeng, 20 September 2026: <https://www.reuters.com/business/finance/us-treasurys-bessent-chinas-he-launch-talks-ai-trade-critical-minerals-2026-09-20/>; Al Jazeera, 21 September 2026: <https://www.aljazeera.com/news/2026/9/21/whats-the-us-china-ai-hotline-that-trump-plans-to-pitch-to-xi>.
 53. The Guardian, *Trump–Xi summit*, 24 September 2026: <https://www.theguardian.com/us-news/live/2026/sep/24/trump-xi-jinping-china-ai-trade-summit-latest-news-updates>; il Giornale, 24 September 2026: <https://www.ilgiornale.it/news/politica-estera/trump-xi-ia-e-commercio-al-centro-pechino-chiede-a-washington-di-tornare-allintesa-con-liran/>
 54. UK Government, *Manchester to host world leaders at the 2027 G20 Summit*, 23 September 2026: <https://www.gov.uk/government/news/manchester-to-host-world-leaders-at-the-2027-g20-summit>; City AM, 23 September 2026: <https://www.cityam.com/uk-to-champion-global-ai-rulebook-during-g20-presidency/>.
 55. The White House, *White House Accord on Super Intelligence. Joint Commitment on Frontier Responsibilities*, 29 September 2026, full text: <https://www.washingtonexaminer.com/news/white-house/4747747/full-trump-white-house-accord-ai-super-intelligence/>; Nextgov, 29 September 2026: <https://www.nextgov.com/artificial-intelligence/2026/09/white-house-unveils-super-intelligence-executive-order-and-industry-accord/416325/>.
 56. AGI, «Accordo alla Casa Bianca, 4 livelli di controllo sui modelli di intelligenza artificiale», 30 September 2026: <https://www.agi.it/estero/news/2026-09-30/ia-casa-bianca-trump-intelligenza-artificiale-39318742/>; Al Jazeera, *How does Trump's White House AI accord work?*, 30 September 2026: <https://www.aljazeera.com/economy/2026/9/30/how-does-trumps-white-house-ai-accord-work>.



American Chamber of Commerce in Italy

Via Cantù 1 - 20123 Milan
Tel. +39 02 86 90 661
amcham@amcham.it - www.amcham.it

30 september 2026