

Realization, Evolution and Leap of a Community with a Shared Future for Humanity

V3.9.3 Locked Edition for Symmetry of the Three Pillars and a Closed Operational Loop

Author: Zijun Fu. Produced by: Civilization Leap Research Group

Factual and version baseline: July 21, 2026

Contents

Decision Summary: Current Recommendations and Decision Boundaries

General Introduction: Deriving the Necessity of a Community with a Shared Future for Humanity from National Destiny

Phase Discipline: The Community Must Not Use a Vision of Its Stable Phase to Coerce Its Launch Phase

The Community Must Give Different Participants Reasons to Enter Without Loss

Part One: Real-World Anchors and Strategic Judgments

Part Two · Realization I: Building a National-Level Foundation for Shared Survival from Zero-Sum Realities

Part Three · Realization II: Establishing Individual Institutional Interfaces for Civilization Leap

Five Types of Operational Mechanism Spanning the National, Social and Individual Levels

Part Four · Evolution: From Guardrails to Structure, from Fear to Expectations

From Initiative to Adoption: Political Launch and a Chain of Institutional Support

Part Five · Leap: Transformation of Civilizational Logic, Long-Term Interfaces and the Ultimate Test

Conclusion

Selected Appendix One: Factual Baseline and Real-World Interfaces

Selected Appendix Two: Summary of Core Institutional Boundaries

Selected Appendix Three: Summary of Phases, Stopping and Exit

Selected Appendix Four: Research Foundations and Version Changes

Realization, Evolution and Leap of a Community with a Shared Future for Humanity

A Structural Exploration of Building a Non-Zero-Sum Foundation within Zero-Sum Realities

V3.9.3 Locked Edition for Symmetry of the Three Pillars and a Closed Operational Loop | Author: Zijun Fu | Civilization Leap Research Group

Version Note

This paper is a civilization-scale strategic exploration for readers in policy research, international governance and institutional design. It is neither an official policy document of any government or international organization nor an institutional code ready for direct implementation. Its purpose is to enable key decision-makers, by reading the main text alone, to understand real-world conditions, national interests, institutional structures, rights boundaries, phase discipline, risk control, political support and the decisions that should now be made. Throughout, it distinguishes [Verified Fact], [Author's Assessment] and [Institutional Proposal · Scenario Exploration]. Specialized protocols, parameter models, data fields and implementation manuals remain research reserves; understanding them is not a prerequisite for understanding the paper's central judgments.

The five-volume Civilization Leap series, *The Civilization of Being Needed*, *Pre-Institutional Language*, *Civilization Leap White Paper*, *Civilizational Perpetuity and Symbiosis of Meaning*, develops the public ideas and the wider framework. Building on V3.9.2, which completed the institutional skeleton of the intelligence-for-good system and raised the level of the five operational mechanisms, V3.9.3 makes only a final structural consolidation: it presents the governance community's polycentric power structure in one place, making the resource community, governance community and intelligence-for-good system genuinely symmetrical in structure; clarifies the dynamic closed loop of the five operational mechanisms; and reduces repeated explanations of the limits of existing governance tools. It does not reopen the central theoretical argument, phase discipline, contribution-value firewall, separation of three powers, failure circuit breakers or current decision request. Upon completion of this edition, the main text is frozen: changing facts go into the factual register, new subjects into subject-specific research, and new parameters into experiments and specialized research reserves. Civilization Leap does not wait for a perfect blueprint, but every step must preserve evidence, boundaries, exit and correction.

Decision Summary: Current Recommendations and Decision Boundaries

This page lets key decision-makers make an initial judgment before reading the full text. The paper is currently at the recommendation and prospective-research stage. It does not request immediate approval of any complete institutional system.

Decision question	Summary
The paper's central judgment	The long-term coexistence of multiple countries and the interconnection of national destinies across borders mean that common security, common development and common continuity are no longer merely matters of international morality. They are national-interest questions that countries must confront to preserve their own security, development capabilities and long-term futures.
The structure proposed here	With the resource community as its real-world foundation, the governance community as its institutional means, and the intelligence-for-good system as its source of new capabilities, this structure gives countries, societies and individuals another path to stronger capacities for security, development and continuity.
The decision currently requested	The first level recommends only including this proposition in national-level prospective research. Cross-departmental examination, baseline surveys of candidate cities and preliminary study of possible plans would follow in that order. Only after legal, fiscal, technical, rights-related and social feasibility have passed review would a separate decision be made on whether to authorize a city-level comprehensive experiment.
First-stage costs and exit	The first stage primarily incurs the costs of research, examination, baseline surveys and comparison of plans. It does not change existing property-rights rules, compel individuals, businesses or regions to participate, or deploy irreversible institutions. Work may stop if the case is not established, with complete evidence of the process and failure retained.
What this paper explicitly does not request	It does not request a one-time transformation of existing society, immediate establishment of a global governance center, determination of basic rights by contribution records, replacement of final human review by artificial intelligence, or prior sacrifice by any country or group as a condition for launching.

The first question now is not whether to accept the entire Civilization Leap plan, but whether this proposition merits formal national-level research and systematic examination against facts, law, interests, rights and failure scenarios.

General Introduction: Deriving the Necessity of a Community with a Shared Future for Humanity from National Destiny

[Author's Assessment]

National Destiny Does Not Exist Only Within National Borders

The world consists of multiple countries coexisting over the long term. No country can eliminate all others, nor can any country forever separate itself from all others and independently obtain every resource, technology, market, body of knowledge and security condition needed to sustain modern society.

A country can determine some of its own actions, but it cannot have exclusive possession of their consequences. Its security policies change other countries' security assessments; their responses in turn reshape its own environment. Its industrial, energy, technology and financial policies travel through supply chains into other countries. Wars, collapses, blockades, epidemics and technological loss of control elsewhere also cross its borders.

National destiny therefore depends on both internal and external conditions: on how a country organizes itself, on how countries affect one another, and on whether the world they collectively constitute can still accommodate its continued existence, development and correction of mistakes.

Security, Development and Continuity Together Constitute National Destiny

Security answers whether a country can continue to exist. Development answers whether it can keep acquiring new capabilities for survival. Continuity answers whether today's security and development can endure through time without undoing themselves in the future.

Security determines whether a country can exist; development determines whether it has a future; continuity determines whether that future is real. None can substitute for the others. Without security, a single crisis may destroy the gains of development. Without development, security gradually loses its material and capability foundations. Without continuity, security and development may amount only to temporary prosperity purchased through resource depletion, social division, external hostility and intergenerational loss.

Why Security, Development and Continuity Must Be Common

In an era of deep interconnection among nuclear weapons, artificial intelligence, biotechnology, global supply chains and cross-border infrastructure, a country cannot achieve lasting security by making every other country less secure. It cannot independently possess all development resources at reasonable cost. Nor can it alone prevent ecological, technological, financial and civilization-scale risks from crossing borders.

“Common” is not moral decoration. It is a practical conclusion derived from the long-term coexistence of multiple countries and cross-border causation. Common security means countries jointly preventing competition from breaching the foundation of shared survival. Common development means every country retaining a real route into the future through its own efforts. Common continuity means today's security and development not being purchased by exhausting the next generation, destroying shared systems or closing off civilization's future.

The Resource Community Is the Real-World Foundation

Security, development and continuity ultimately depend on resources. Resources here include not only food, energy, minerals, land and water, but also computing capacity, data, knowledge, technology, talent, organizational capability, capital, markets, logistics, public infrastructure and windows of time.

Countries have in fact long been part of a resource community. Those with minerals depend on processing and markets; those with models depend on chips, energy and data; those with capital depend on real production and social stability. The problem is not whether a resource community exists, but that existing interdependence remains opaque, asymmetric, vulnerable to coercion and weaponization. Idleness and shortages coexist; joint allocation is lacking in crises, and fair settlement is lacking after use.

A resource community is not common ownership of resources. With ownership clear and sovereignty and lawful rights protected, it enables some resource-use rights, assurance rights, emergency call-on rights and settlement responsibilities to operate collaboratively through voluntary, layered, auditable interfaces from which participants can exit.

The Governance Community Is the Institutional Means

Resources do not automatically flow rationally. Governance structures are needed to determine who may enter, who may call on resources, who has priority in crises, how to prevent domination by the strong, how to protect the weak, and how to address nonperformance, disputes and exit.

A governance community is neither a global central government nor a substitute for national sovereignty. It is a limited, layered, polycentric, auditable and correctable structure of common governance, with exit available, for shared problems across participants, regions and countries. It governs relationships of rights, authorization interfaces, risk boundaries, use procedures and settlement responsibilities. It neither acquires ownership of all resources nor governs all domestic affairs in place of countries.

The Intelligence-for-Good System Provides New Capabilities

Traditional administrative coordination, market allocation and cross-border movement cannot continuously discover, understand and connect vast dispersed resources, capabilities, opportunities and real needs on a global scale. Artificial intelligence brings new capacities for sensing, connection, matching, exploration, coordination and capability amplification that did not previously exist.

Intelligence for good does not require AI to possess abstract goodwill. Through objectives, permissions, audit, responsibility, exit and final human review, it gives priority to reducing shared harm, releasing idle resources and expanding national and individual capabilities, rather than using intelligence only to replace people, control them and overpower other countries. It offers countries another path to stronger capacities for security, development and continuity: not solving a country's problems indefinitely on its behalf, but strengthening its own ability to discover problems, organize resources, cultivate talent, respond to risks and develop its own answers.

Access, Usability and Autonomous Mobility Across Domains for People and Individual Resource Pools

The community cannot connect only countries and institutions. Every person possesses an individual resource pool comprising time, capabilities, knowledge, experience, works, data, relationships, lawful property and extensions of capability enabled by AI.

Access, usability and autonomous mobility across domains for people and individual resource pools are an important foundation for common governance. They are the direction in which the community evolves from cooperation among countries and organizations toward real participation by every person. They are also the direction of civilization's leap from geographical confinement, organizational dependence and passive allocation toward rights that follow the person, accessible resources, released capabilities and autonomy in life.

People are not resources. An individual resource pool belongs to the person and serves their freedom, development and meaning. Individuals have the right to decide whether to enter, what to make available, how long to authorize it, how it may be used, whether to withdraw authorization, when to exit and whether to return. Not moving, sharing or contributing does not affect basic subsistence or basic rights.

Overall Definition and Route

A community with a shared future for humanity is an institutional community that countries establish through a resource community, a governance community and an intelligence-for-good system to preserve security, development and continuity for themselves and their shared world, under conditions of long-term coexistence among multiple countries, interconnection of national destinies across borders, and deep cross-border transmission of resources and risks.

It takes the resource community as its foundation, the governance community as its means, and the intelligence-for-good system as its capability support. It takes access, usability and autonomous mobility across domains for people and individual resource pools as an important foundation and evolutionary direction of common governance. Its goals are common security, common development and common continuity.

Its ultimate purpose is not to serve resources, governance or intelligence themselves, but to ensure that every country, society and person continues to have a future to look forward to.

Phase Discipline: The Community Must Not Use a Vision of Its Stable Phase to Coerce Its Launch Phase

Every institutional proposal in this paper must be understood in terms of the launch, initial, transition and stable phases. Phases are triggered by real outcomes, rights records and independent verification from the preceding stage, not advanced automatically by the calendar. The cross-domain structure of the stable phase must not be written as a present obligation in the launch phase; nor should the launch phase's voluntary invitations be mistaken for the community's final form.

Phase	Main task	Nature of rules	Minimum rights boundary
Launch phase	Build understanding, motivation and minimum cooperation interfaces	Voluntary and invitational, adding only tangible benefits	No participant's existing baseline may deteriorate because the proposal appears
Initial phase	Verify the closed loop of resources, governance, intelligence and rights at a limited number of nodes	Observable, open to exit and correctable	No disadvantage on exit; the proposal bears the costs of failure
Transition phase	Develop mutual recognition among multiple nodes and an optional institutional attraction	Describe emergence; do not design coercion	Nonparticipants retain rights to existing paths; entry remains open
Stable phase	Establish verified, limited common governance with a constitutional backstop	One person, one vote safeguards the course; professional experience governs domains; rules can be rolled back	The right to subsistence precedes contribution; the brakes always remain with people and countries

Note on levels: The civilizational voting power and “one person, one vote safeguards the course” in this table apply mainly to lawfully authorized city-level comprehensive experimental zones and the mature communities that subsequently develop from them. They do not propose directly implementing global one-person-one-vote governance between countries today. Intergovernmental common governance follows sovereign authorization, layered participation and cooperation on limited issues.

Three iron rules run through every phase. First, phase labeling is mandatory: every mechanism must specify when and at what scale it is valid. Second, no loss for anyone: no group's basic rights or real-world baseline may be treated as a transition cost. Third, adversarial deadlocks are forbidden: any proposal that can begin only after a group has been defeated, deprived or subdued is still reproducing the old logic and must be returned for redesign.

The Community Must Give Different Participants Reasons to Enter Without Loss

The community must do more than explain how the whole will eventually improve. It must answer why each actual bearer of responsibility would be willing to enter starting today. Reasons to enter must arise from that participant's real interests in security, development, dignity, responsibility and the future, rather than requiring sacrifice first and a distant reward later.

Participant	Real motivation	Tangible benefits of entry	Non-negotiable no-loss boundary
Countries and governors	Security, development, controllable reform and historical responsibility	New capabilities for identifying risks and organizing resources; retained ultimate responsibility and stopping power	Core sovereignty, domestic enforcement of law, exit and security exceptions must not be hollowed out
Existing capital and businesses	Security of lawful assets, long-term certainty and new markets	Returns from the use of idle resources; institutional credibility and room for development in a new cycle	No retrospective confiscatory reckoning of lawful existing assets; no forced transfer of ownership in the name of participation
Professional contributors	Justice for experience, professional judgment and genuine attribution	Professional capability enters domain decisions; contributions are traceable and can be carried forward	Professional judgment must not be diluted by money, traffic or a unified rating
Ordinary workers and those displaced	Subsistence security, a sense of being needed and new routes to value	Care, maintenance, presence in the community and capability transitions are recognized	Basic subsistence does not depend on contribution; exit and temporary nonparticipation are not punished
Young people and those yet to enter	New coordinates for growth and room to create	Low barriers to participation, portable learning outcomes and broadly accessible AI capability	Rights to subsistence and education are not conditional on contribution records
Resource-poor countries and regions	Autonomous development capabilities and avoidance of permanent marginalization	Growth in local-language, education, health, industrial and governance capabilities	No extraction of data, control of models or creation of permanent dependence in the name of support
AI development and safety communities	Creative freedom, boundaries of responsibility and avoidance of being swept into an arms race	Trust and room for long-term innovation within transparent, auditable boundaries	Responsibility matches observability; no unlimited liability

If the loss of any category of participant becomes the incentive for another category to enter, an adversarial deadlock is triggered. The right path is not to require someone to lose first, but to redesign a “concession” under the old logic as an “exchange into” the new logic.

Part One: Real-World Anchors and Strategic Judgments

[Reader's Guide] The general introduction derives the necessity of common security, common development and common continuity from the long-term coexistence of multiple countries and national destiny. This part further separates facts that have occurred, judgments drawn from them, and gaps in existing tools. It shows that the logic is not an abstract proposal detached from reality: institutional interfaces already exist on which further construction is possible. The title's three stages correspond to the paper's structure as follows. Realization, Parts Two and Three, puts the resource community, governance community, intelligence-for-good system and individual institutional interfaces into real operation for the first time. Evolution, Part Four, progressively connects individual nodes into a community network spanning regions, countries and domains. Leap, Part Five, structurally transforms national security, modes of development, resource value and the sources of human value.

1. What Happened on July 16 and 17, 2026

[Verified Fact]

On July 16, 2026, a signing ceremony for the agreement establishing the World Artificial Intelligence Cooperation Organization took place in Shanghai. Representatives of 29 countries signed. The agreement defines the organization as an independent intergovernmental international organization headquartered in Shanghai, China. Signing the agreement, its entry into legal force, organizational launch and actual operation are different institutional stages; this paper does not merge them into a single fact.

On July 17, the 2026 World Artificial Intelligence Conference and High-Level Meeting on Global AI Governance opened. President Xi Jinping delivered a keynote address entitled “Working Together to Build a Fair and Equitable Global AI Governance System,” offering four proposals: uphold openness and win-win outcomes, strengthen risk awareness, encourage inclusiveness, and advocate working together in harmony. The AI Cooperation and Development Action Plan and the International AI Ethics Governance Action Plan were released at the same time. At the main forum, the Cyberspace Administration of China and relevant parties proposed the Global Cooperation Initiative on Mutual Trust, Connectivity and Interoperability of Intelligent Agents, setting out action directions in ten areas; Xinhua and others published its full text on July 17.

See Selected Appendix One at the end for factual sources and real-world interfaces.

2. What the World Artificial Intelligence Cooperation Organization Means

[Author's Assessment]

Twenty-nine countries signed a single document. This agreement may become one of the earlier non-zero-sum institutional interfaces of the AGI era directed toward common development and common governance. Its significance is not that competition has ended, but that countries are beginning to try to establish baselines of common security and common development that competition cannot breach.

From “one agreement” to multiple action documents, the negotiating table's function is evolving from “establishing minimum consensus” toward “implementing concrete action.” This is the real-world anchor for every exploration in this paper.

The Whole-of-Humanity Clause

The community described here addresses humanity as a whole, including current competitors, countries that have not signed any agreement, technological powers and the least developed countries. The 29-country agreement is a starting interface, not the boundary of a bloc. The vulnerability analysis in Section 7 shows that China, the United States, Europe and global supply chains are deeply coupled. Whether or not they sign, technological powers are structurally already part of the same risk community. The proposed mechanisms are therefore designed around universal access: joining must not require recognition of any country's leadership; participants should jointly

make the rules; clear procedures and institutions should protect exit rights; participation is layered, allowing different participants to select issues and tracks without accepting the entire theory first. Nonparticipation is not punished. Safety-layer cooperation, including major-incident notification, humanitarian communication, disaster relief and verification of shared survival, does not require membership and remains open to nonmembers. But openness does not mean unconditional access to sensitive data, verification powers, resource call-on rights or governance seats. Specific permissions still require risk-based necessity, authorization under agreements, security review, acceptance of responsibility and reciprocal obligations. Nor does freedom from punishment for not joining mean that participants who have not borne common costs or treaty obligations automatically receive all the additional entitlements created by joint efforts. Resource access, sensitive-data access, governance seats and other additional entitlements are allocated according to actual responsibility, resource inputs, risk borne and agreement obligations. The three relationships can be expressed in one sentence: basic safety interfaces remain open, nonparticipation attracts no retaliation, and additional entitlements remain tied to responsibility. The reverse boundary also holds: internal community rules must not aim to punish nonparticipants, create external technological blockades or entrench geopolitical blocs. Members may allocate internal additional entitlements according to joint inputs, but must not turn community mechanisms into instruments of exclusion, coercion or collective pressure against nonmembers. Initiators do not monopolize interpretation and enforcement of the rules: who proposes the initial framework does not determine who owns the institution. These are this paper's institutional proposals; they do not imply that existing organizational charters already include these arrangements. If the community evolves into one bloc confronting another, it violates its own definition. Bloc formation itself is listed among the warning signals for Scenario A in Section 20.

A Ladder of Layered Participation: No One-Time Political Declaration in Place of Long-Term Verification

Rather than reserving a fixed empty chair for a particular country, the community establishes a layered ladder: public observation → safety-information interface → technical observation and exercises → participation in individual agreements → cooperation in joint projects → governance of limited issues → formal membership. Any participant may remain at one level indefinitely, or move up, move down or exit under the applicable conditions. Deeper participation brings more additional entitlements, with corresponding increases in responsibility and audit obligations. Basic safety interfaces remain open, nonparticipation attracts no retaliation, and additional entitlements remain tied to responsibility.

3. Global AI Capability and Governance Are Advancing Along Four Tracks at Once

[Author's Assessment]

Current global AI governance is not inactive. Rather, action is dispersed across four independent tracks moving at different speeds.

The industrial-market track is driven chiefly by expansion of computing capacity, model iteration and product deployment. Its central logic is efficiency and first-mover advantage. It moves fastest, but pressures for speed, returns on capital and market advantage make it prone to externalizing some safety, ethical and social costs.

The academic-community track aims to produce verifiable knowledge. Its central logic is peer review and reproducibility. It provides a knowledge base for understanding the boundaries of AI capability, but its production cycles are slower than industrial iteration.

The government technical-network track comprises joint testing, a common language for risk and alignment of technical standards. Its central logic is operational safety. Mechanisms for joint frontier-AI safety testing, risk notification and red-team exercises already exist on a small scale, but their coverage is limited and their binding force weak.

The intergovernmental international-organization track comprises development, governance, capacity building and frameworks for cooperation among countries. Its central logic is rulemaking and the generation of legitimacy.

The July 2026 Shanghai conference and the World Artificial Intelligence Cooperation Organization are the latest developments on this track.

Central judgment: each track has its own function, but the real global risk is that all four are accelerating without stable interfaces between them. Industry is accelerating capabilities; academia is understanding them; safety networks are measuring them; intergovernmental organizations are trying to govern them. The gaps among the four are widening. This paper's mechanisms offer structural ideas for interfaces between these tracks.

4. Why a Community with a Shared Future for Humanity Is Becoming a Structure for Survival Rather Than a Vision

[Author's Assessment]

In the past, a “community with a shared future for humanity” was more often expressed as a diplomatic vision, an advocacy of values or an idealistic ultimate goal. In an era of balanced nuclear deterrence and slow globalization, that framing sustained its appeal. The AGI era changes the nature of the concept. AI's deep penetration into military, biological and infrastructure domains greatly increases the difficulty of attributing risks and establishing responsibility. In some crisis scenarios, critical systems may compress decision windows to minutes or even seconds. A community with a shared future for humanity is no longer “an ideal we should pursue,” but “a structure for survival that we must build before collapse.”

This is not an ideal that has appeared out of nowhere. It extends the historical logic of ever-expanding circles of non-zero-sum interaction. Early human cooperation was usually bounded by kinship, tribes and nearby communities. As city-state alliances, the sovereign-state system and regional cooperation mechanisms developed, recognizing the continued existence of outside actors and establishing stable rules with them gradually became important means of expanding civilization's non-zero-sum space. Today, AI, biological risks and interconnected infrastructure are pushing the boundary from “regional” to “global.” This time, however, expansion must outpace collapse.

5. Why Existing Global Governance Tools Are Still Insufficient

[Author's Assessment]

Existing tools each perform necessary functions, but none can alone provide complete governance of cross-border risks and shared capabilities in the AGI era. Markets can improve local allocation efficiency, but do not automatically protect the foundation of shared survival. National governance can bear domestic legal, fiscal and security responsibilities, but cannot alone govern cross-border risks and interlocking resources. International cooperation can coordinate projects, but without stable layers of rights, audit, exit and shared capabilities, it is difficult to turn that cooperation into sustainable institutions.

Three realities must also be acknowledged: sovereign states remain the principal vessels of public power and international commitments; competition will not disappear because of an initiative; and critical resources and ecological carrying capacity have limits. This paper does not propose bypassing states, competition or scarcity. It proposes building limited, auditable interfaces with exit available on top of these realities to address compressed decision windows, attribution difficulties and cross-domain transmission of risk.

Part One Conclusion: The Community Is Not a Substitute for National Interests

[Author's Assessment]

A community with a shared future for humanity does not ask countries to abandon their own destiny. It asks them to see its complete structure. Precisely because every country must protect its own security, development and

continuity, and these already partly reside in relationships between countries, they need jointly to maintain a world that will not collapse as a whole, will not permanently close off any country's future, and can keep correcting mistakes.

The institutional interface that emerged in Shanghai in 2026 does not prove that the community has already formed, or that participating countries are equally committed. What it proves is that, while competition persists, countries can still begin building shared capabilities that transcend a purely competitive relationship. The next question is not whether the idea is inspiring, but whether the resource community, governance community and intelligence-for-good system can actually operate under constraints of rights, interests, nonperformance and crisis.

Part Two • Realization I: Building a National-Level Foundation for Shared Survival from Zero-Sum Realities

[Reader's Guide] This part asks how, in a world where competition never stops, the community can become an operable institution rather than a concept. It addresses the separation principle and escalation clause, the three-level progression from vulnerability to survival, the first negotiating table, resource mutual assistance and three-ledger settlement, and the boundaries of sovereign cooperation.

[Institutional Proposal] The mechanisms in this and subsequent parts are structural explorations, subject to verification across generations and to the phase discipline stated earlier. The launch phase may establish only voluntary, limited, revocable interfaces. The initial phase expands only mechanisms already verified against rights and facts. The transition phase describes institutional attraction, without designing coercion. Stable-phase rules must not become present obligations prematurely. Being right about the direction does not permit bypassing rights, evidence or failure costs. Every step must preserve boundaries, exit, correction and rollback.

6. Layered Governance of Common Security, Common Development and Common Continuity

A community with a shared future for humanity must begin with “irreversible risks,” rather than “common prosperity.” The former is the baseline; the latter is the direction.

Competitive domains, where one side may advance and another retreat, include chip manufacturing, AI applications, electric vehicles, market share and international standard-setting; industrial policy, tariffs, technology blockades and countermeasures. Noncompetitive domains, where shared survival is essential, include preventing misjudgment at AGI's “Oppenheimer moment”; immediate notification and joint investigation of biological laboratory leaks; and responses to the paralysis of critical undersea cables, satellite links and global computing centers.

Escalation Clause

The boundary between competitive and noncompetitive domains is a dynamic criterion, not a static list. Competition for development may continue, but any competitive means that exhibits cross-domain transmission, irreversible diffusion and risks to collective survival is automatically escalated to the common-security agenda. Judgment does not belong to any single country. It should be exercised by a risk-classification mechanism jointly authorized by participants, with public and reviewable procedures. See Selected Appendix Three for minimum principles on evidence, review and suspension.

Mechanism Design

Establish a “strategic firewall.” In competitive domains, allow and even encourage intense contest. In noncompetitive domains, establish a “red hotline plus emergency-brake agreement” binding on participants once concluded. Minimum-necessary notification mechanisms should cover abnormal automated behavior posing strategic-misjudgment risks, changes in crisis conditions and signals triggering joint verification. What is disclosed is not every deployment, but signals sufficient to prevent misjudgment. Countries' rights to take necessary defensive measures remain unaffected, but multilateral scientific verification should begin at the same time.

According to public reporting, the keynote address stressed jointly opposing the overextension of national-security concepts in AI and placing one's own security above that of other countries. This is consistent with the direction of the separation principle.

7. From Shared Vulnerability to Shared Survival

The nuclear era's “mutually assured destruction” rests on a premise: the other side knows who you are, where you are and how to retaliate. In the AGI era, that premise is eroding. Some attacks may resist reliable attribution within the time a crisis permits; capabilities hide in algorithm updates, and sabotage masquerades as technical accidents. The community's real motivation is therefore not “goodwill,” but clear recognition of shared vulnerability. This recognition must progress through three levels, rather than remain at the first.

Level One, Real-World Diagnosis: Mutually Assured Vulnerability

The harder each side tries to be independent, the more it discovers that it cannot quickly produce what the other possesses. The harder it tries to choke the other, the more it discovers its own throat is also held. The collapse of any side harms others as well. Goodwill is unnecessary, but lies cannot stay hidden; trust may be absent, but no one escapes fear. Deep coupling is a verifiable industrial fact: China accounts for highly concentrated shares of global separation, refining and processing for multiple rare earths; US companies lead global chip design and, under the definitions used in public SIA/BCG studies and others, remain leaders in global semiconductor design and sales; advanced-process manufacturing is highly concentrated in a few foundries; and Europe holds nodes in certain high-end equipment, precision manufacturing and critical industrial capabilities that cannot readily be replaced. No side can achieve “complete decoupling” without enormous losses. Relevant data definitions, dates and statistical boundaries remain in research reserves for continuing verification. The premise must be firmly fixed: vulnerability is a common enemy, not a one-sided hostage. It arises from systemic deep coupling, not from any single party's strategic weakness. Once vulnerability is weaponized as coercive leverage, “I know you cannot do without me,” the mechanism immediately fails. Complex interdependence theory, associated with Keohane and Nye, supplies important analytical tools, particularly its distinctions among sensitivity, vulnerability, asymmetric dependence and multiple channels of relations. This paper does not equate contemporary great-power relations in their entirety with complex interdependence, or assume that economic coupling automatically produces peace. It asks how existing, potentially weaponizable asymmetric dependencies can be transformed into a multinode, substitutable and verifiable structure of shared survival.

Level Two, Institutional Transformation: Mutually Assured Interdependence

Transform uncontrollable dependencies vulnerable to coercion into multinode, substitutable and auditable structures of mutual assistance. The resource pool's anti-coercion design is the “many-to-many” principle: cap any single supplier's share of critical resources, with concentration thresholds separately modeled by resource category, replacement cycle and crisis conditions, to ensure that no “sole supplier” exists. When alternatives truly exist and can be activated within acceptable time and cost, a single supplier's coercive power falls substantially. Transparency is not intended to require prior mutual trust, but to reduce dependence on subjective trust and shift cooperation as far as possible toward verifiable data, independent audit and repeated verification. Satellite monitoring, computing-capacity audits and sharing of de-identified data constitute a verifiable transparency baseline.

Level Three, the Mature State: Mutually Assured Survival

No side's security depends any longer on another being weaker. It depends on the shared system remaining open, redundant and repairable. The full progression is: recognize shared vulnerability → establish controllable interdependence → achieve shared survival. “Maintaining vulnerability” is a dangerous long-term objective. Vulnerability is only the entry point; survival is the destination.

Relevant game-theoretic models and industry data definitions remain in research reserves. The main text retains only their verified structural conclusions.

8. The First Negotiating Table

A community with a shared future for humanity need not begin with “universal harmony.” It can begin with a minimal, operable negotiating issue involving cooperation that does not touch core sovereignty. This table’s agenda is limited to three items.

- 1. Emergency AI-misjudgment hotline: when any party’s early-warning system detects a high-risk anomalous signal that may cause strategic misjudgment, it should immediately submit it for review by its country’s lawfully authorized crisis-response body. When review finds that the agreed notification threshold is met, or when time is too short for adequate verification but continued silence may create greater risk, the hotline conveys the minimum necessary warning information and joint verification begins within the shortest time specified in advance by the agreement. AI may detect and warn; it may not independently conduct strategic notifications between countries. This requires no trust, only fear: both sides fear war starting before they have had time to explain.
- 2. Joint origin-tracing mechanism for anomalous biological events: for any unexplained public-health anomaly, publicly attributing an attack in an irreversible manner is inappropriate when evidence is insufficient. Parties retain the right to take necessary defensive measures, but should simultaneously begin multilateral scientific verification.
- 3. Minimum protection mechanism for critical cross-border civilian infrastructure: facilities serving broad civilian navigation, communications, finance, healthcare, disaster response and public computing functions must, in principle, not be attacked with the aim of causing systemic civilian paralysis, even during escalation of conflict. The scope, exceptions and verification standards for dual-use facilities require separate negotiation. This is rationality, not mercy: whoever destroys these first also first cripples their own global system of dependencies.

These three agenda items have precedents in arms control, crisis communication, confidence-building measures and civilian-infrastructure protection. Hotlines, incident notification, joint verification and limited nonattack commitments supply institutional language for AI and biological-risk governance. But the attribution difficulties, autonomy and cross-domain diffusion of new risks require redesign, not simple transplantation.

How to launch: initiators and participants can be different. A major power need not lead directly. A UN framework, regional alliance or neutral third party may convene under the heading of “scientific risk assessment,” reducing the political pressure that treats speaking first as a sign of weakness. The first to speak is not the one conceding, but the one defining which risks are worth discussing. Shanghai’s 2026 experience shows that an organized negotiating table for multilateral AI cooperation has emerged. The military-misjudgment, joint biological-verification and critical-infrastructure-protection issues proposed here must still enter a separate, professional negotiating process with appropriate security authorization.

9. Intergovernmental Resource Mutual Assistance and Settlement

“Intergovernmental resource pool” is an umbrella term. Internally, it is not a warehouse, but a resource mutual-assistance network made up of three types of mechanism.

Mechanism type	Applicable resources	Nature of rights transferred
Physical reserves and emergency allocation	Food, medical supplies, energy reserves and critical minerals	Ownership transfers after physical resources are consumed
Capacity reservation and service delivery	Computing, logistics, manufacturing capacity, communications redundancy and emergency medical capacity	Service capacity is scheduled; ownership of physical assets is not
Authorized access and mutual recognition of licenses	Data, patents, software, technical standards and model interfaces	Limited licenses transfer, not ownership

In principle, the resource pool does not change depositors' underlying rights in their original assets, but different resource types require different rights arrangements. For consumable physical resources, ownership or disposal rights transfer lawfully when resources are called on. Capacity services are settled under reservation and delivery rules. Data and intellectual property open only the agreed scope of access, use or licensing rights.

Separation of Three Ledgers

Resource-pool accounting strictly separates the ownership ledger, use-rights ledger and settlement ledger.

- Ownership ledger: registration in the pool does not itself change ownership. For consumable physical resources, ownership or disposal rights transfer under the agreement after members have called on them and delivery and acceptance are complete. Capacity services and authorized licenses do not transfer ownership of the original assets.
- Use-rights ledger: records which resources can be scheduled, when and where.
- Settlement ledger: records how much each participant has deposited, how much each has called on, and how many call-on rights each still holds.

See Selected Appendix Two for the three-ledger model's minimum structure. Cross-category settlement, deferred delivery, emergency calls and compensation for nonperformance require separate modeling based on resource categories and actual operating data; they are not predetermined here.

Cross-border joint resource management is not unprecedented. Ostrom's work on clear boundaries, participatory rules, graduated responses, low-cost dispute resolution and nested governance, in *Governing the Commons* (1990; Nobel Prize in Economics, 2009), provides important institutional reference points. Those experiences do not directly prove that joint management of strategic resources between countries must be feasible. Sovereign competition, security exceptions, capability asymmetries and gaps in international enforcement must be addressed anew when scaling from local common-pool governance to the intergovernmental level. The three-ledger separation and supplementary rules here are an attempt to apply these principles across scales, not established proof of their effectiveness. "Intergovernmental resource pool" is also an umbrella term, not a claim that all its resources have the economic attributes of traditional common-pool resources. Ostrom's framework mainly provides governance principles, not a uniform classification of every resource type.

Core Principles

- Deposits are voluntary.
- Eligibility to call on resources is nondiscriminatory; specific calls remain subject to inventory, priorities, emergency conditions and settlement balances.
- Exit is permitted: any country may leave under the prescribed procedure.

Operation also requires rules for overbooking, depreciation of value, delivery default, emergency freezes, cross-category conversion and exit settlement. Relevant parameters require specialized verification and are not prerequisites for understanding this paper.

Conditions for Effective Resource Mechanisms and Limits on Complexity

Three-ledger separation can clarify ownership, use rights and settlement responsibility, but ledgers cannot manufacture willingness to perform or actual delivery capacity where neither exists. Every intergovernmental resource mechanism presupposes minimum willingness to perform, real delivery capacity and enforceable domestic authorization among participants. Institutions reduce uncertainty, identify nonperformance, arrange alternatives and organize remedies; they must not pretend to bypass states and real capabilities.

Nor may a jointly authorized mechanism be assumed, in the launch phase, to be a global center that already exists. The launch phase begins with bilateral, small multilateral and single-resource agreements. The initial phase establishes joint projects, audit and settlement interfaces. The transition phase develops mutual recognition among

multiple nodes. Only long-term operation and explicit authorization can make limited common governance possible.

Institutional complexity is itself a risk. The complexity of any resource, contribution or audit system must not exceed what participants can understand, governors can maintain and independent auditors can review. Mechanisms that cannot be effectively audited within reasonable cost and time must not enter formal operation; they must remain at the shadow-recording, manual-sampling or limited-verification stage.

Differentiated Access

National resource inputs and public capabilities are recognized regardless of scale or form. Access rules should consider development stage, resource endowments, technical capabilities and real needs, avoiding the entrenchment of structural inequality through existing capability gaps. Effective outputs under limited computing capacity may be explored with weighting for scarcity and public value. Local professional bodies must first confirm the provenance, effectiveness, knowledge rights and community authorization of local medicines, agricultural adaptation techniques and other traditional knowledge. Before cross-border calls and settlement, professional review, benefit sharing, informed consent and dispute handling are also required. Specific weights and guaranteed minimum amounts require subsequent verification.

Differentiated access addresses developing countries' structural disadvantages. Access is equally necessary for technological powers and developed economies, but their reasons to enter, responsibilities, interest structures and cooperation paths differ. The separation principle, minimum-necessary notifications and common-security agenda remain open to all; specific participation should be designed according to each participant's capabilities, risks and responsibilities. Real resistance must be acknowledged: geopolitical bloc formation may make countries unwilling to deposit resources alongside rivals; technology monopolies may resist mutual resource recognition; and developing countries may be structurally disadvantaged in a “deposit-withdrawal” system. These constraints need individual countermeasures through differentiated access, inclusive funds, and clauses against discrimination and domination and for benefit sharing and protection of vulnerable members.

10. Boundaries of Sovereign Cooperation

The community resource pool does not constitute a transfer of sovereignty. All mechanisms rest on voluntary participation and preserve exit rights. But sovereignty is not a switch; it is an onion. Nothing is conceded at the core, while cooperation is possible in outer layers.

Layer	Contents	Degree of cooperation	Exit mechanism
Core layer	Nuclear-weapons codes, core strategic data and defense secrets	Zero cooperation: information is never disclosed	None
Buffer layer	Unclassified computing redundancy, civilian data samples and nonsensitive technical parameters	Information cooperation with physical retention: data are traceable, deletable and auditable	Exit under agreement: end subsequent access and processing; delete or return unused data; retain records of lawful completed calls and audit evidence only to the minimum necessary extent
Interaction layer	Cross-border resource withdrawal, emergency scheduling and joint origin tracing	Coordinated allocation of physical ownership, service-benefit rights, limited use rights or licenses according to resource type; specific changes in rights follow the call-on agreement	An agreed buffer period; remaining resources returned or settled at an agreed valuation

Layer	Contents	Degree of cooperation	Exit mechanism
Outer layer	Mutual recognition of standards, transparency baselines and alignment of ethical frameworks	Cooperation on rules without transfer of assets: rules made jointly and implemented domestically by each party	Takes effect under members' domestic procedures and joint agreements

Key principles: “depositing” is not “handing over.” In principle, the pool does not change depositors' underlying rights in original assets; rights arrangements vary by resource type. Information must be revocable: buffer-layer data disclosures must be traceable, deletable and auditable. Exit must have a path: costs must be transparent and predictable, without retaliatory sanctions for leaving. The absolute baselines of retained sovereignty are the right not to disclose core strategic-resource data, national-security exceptions, the right to leave agreements, and autonomy over domestic institutions.

The Governance Community: The Resource Community's Institutional Operating Layer

Community resource pools are not owned and controlled by a single global center. A safer form is a community resource collaborative-management system comprising national nodes, regional nodes, specialized resource nodes, city and experimental-zone nodes, international settlement nodes and independent audit nodes.

At a minimum, the governance community handles joint rulemaking, confirmation of rights, resource registration, supply-demand matching, scheduling and settlement, risk classification, dispute arbitration, independent audit, responses to nonperformance, exit settlement and correction of rules. Management, audit, resource-decision and civilizational voting powers must be separate. No center may use efficiency as a reason to acquire resource ownership, individual decision-making power or sovereign authority that has not been explicitly authorized.

The Governance Community's Polycentric Power Structure

The governance community is more than a collection of governance functions. It must also separate powers to decide rules, execute rules, review rules, resolve disputes and terminate deviations. Its minimum framework follows.

Position of power	Core responsibilities	Principal bearers	Mandatory checks and balances
Rulemaking power	Set rules for limited issues, authorization boundaries and amendment procedures	Rulemaking nodes jointly composed of participating countries, lawful public representatives and professional participants	May not simultaneously control execution and audit; major amendments require public reasons, impacts and objections
Execution and coordination power	Coordinate performance, resource interfaces, public services and crisis response	Distributed national, regional, city and specialized execution nodes	Execute only within authorization; may not amend rules or expand their own powers
Independent audit power	Verify resources, data, algorithms, procedures and effects on rights	Multi-jurisdictional audit nodes separate from executing bodies	May not operate, profit from or arbitrate the matters they audit
Dispute-arbitration power	Handle disputes over performance, rights, permissions and exit	Plural arbitration and review structures	May not be controlled by a party to the dispute; decisions must be reviewable and open to appeal

Position of power	Core responsibilities	Principal bearers	Mandatory checks and balances
Civilizational veto and termination power	Prevent irreversible deviation; initiate suspension, rollback and renewed phase review	Lawfully authorized public voting, final human review and emergency-braking structures	May not be purchased with capital, technical capability or contribution value; emergency powers must expire automatically

This power structure follows four principles: polycentricity, without a single all-powerful center; limited and revocable authorization; adjacent nodes with veto powers may not be controlled by the same actor; and emergency powers are used only to limit losses for a specified time, expire automatically at the deadline, and undergo independent review.

The resource community answers how resources enter real use. The governance community answers how power is authorized, constrained and terminated on failure. The intelligence-for-good system answers how new capabilities are built, called on, audited and shut off. All three depend on one another, but no pillar may absorb the other two.

The Intelligence-for-Good System: The Community's New Capability Layer

For the resource and governance communities to operate globally, AI is needed to discover dispersed resources, identify real needs, explain different rules, explore risks, connect languages and institutions, reduce coordination costs and amplify national and individual capabilities.

The intelligence-for-good system must meet six mandatory boundaries: humans retain final decision-making authority; countries and individuals retain control over core data and authorization; basic safety intelligence does not require full membership; technical support must not create permanent dependence; important models, rules and results must be auditable; and intelligence must not determine basic rights, personal worth or entitlement to subsistence.

Institutional Structure: Multinode Construction, Separation of Permissions, Independent Audit and Human Termination Power

No single country, enterprise or model monopolizes the construction and operation of the intelligence-for-good system. A safer structure comprises national public nodes, regional and city nodes, specialist-domain nodes, regulated-enterprise nodes and independent audit nodes. Different nodes may use different models and technical approaches, but must follow common rights boundaries, evidence formats, anomaly notifications and exit rules.

The power to amend core objectives, permission boundaries, data-use scope and irreversible rules must be separate from model-development, day-to-day operation, audit and termination powers. Developers may improve capabilities, but may not independently change civilizational goals. Operators may invoke systems, but may not expand their own permissions. Every high-impact amendment should leave a version record, reasons, an impact assessment, objections and a rollback point.

Audit must at least cover technical safety, data and authorization, effects on rights, long-term dependence and independent external review. Auditors must themselves face permission limits, conflict-of-interest disclosure and review of audit quality, preventing audit itself from becoming a new unchecked power.

Multiple intelligent systems should cooperate through open, replaceable and verifiable interfaces rather than merge into a single global model or intelligence hub. Interoperability concerns the minimum necessary formats for identity, authorization, evidence, anomalies and settlement. It does not require countries to surrender all data, model weights or governance autonomy.

When major errors, permission overreach, data misuse, unexplained discrimination, systemic dependence or irreversible risks arise, countries, lawful governing bodies and affected individuals should each retain powers of suspension, refusal, appeal and exit appropriate to their rights. High-impact systems must support human termination, degraded operation, version rollback and evidence preservation. No single actor may monopolize proposal, approval for execution, continuous monitoring and final shutdown.

AI's civilizational value is not merely to help a few countries pull ahead faster. It is to give more countries educational, medical, industrial, risk-identification and governance capabilities they could not otherwise obtain at the same cost, while preserving their autonomy to develop their own answers.

Part Three • Realization II: Establishing Individual Institutional Interfaces for Civilization Leap

[Reader's Guide] This part asks why national-level resource security cannot substitute for individual-level institutions of meaning. It addresses the precise central proposition, individual resource mechanisms, the nature and boundaries of contribution value, separation of three powers and dynamic arbitration, and review against the three meta-goals.

[Institutional Proposal]

11. Why Intergovernmental Resource Pools Cannot Substitute for Individual Institutions

Intergovernmental resource pools address mutual assistance in national-level resource security. But such assistance does not automatically resolve crises at the individual level. As AGI replaces more and more functional labor, the disappearance of the “sense of being needed” becomes a crisis deeper than poverty. This judgment, however, presupposes that basic subsistence is already secured. Where material deprivation remains unresolved, poverty remains the foremost survival problem.

The Precise Central Proposition

The old logic says: to make myself safer, I make you weaker; to be more valuable, I must be stronger than you. The new logic cannot be expressed as “I am needed, therefore I exist,” because that still ties a person's entitlement to exist to their function. The precise formulation is: I exist, so my basic worth cannot be denied; I connect, care, create and take responsibility, so meaning continually emerges in my life.

A Strict Distinction Among Three Forms of “Being Needed”

- Functional neededness: the system temporarily retains you because you can produce.
- Relational neededness: genuine care, dependence and response among living beings.
- Existential worth: not annulled by capability, output or contribution.

Civilization Leap seeks to end functional neededness's monopoly over entitlement to subsistence; to affirm relational neededness as an engine of meaning; and to secure existential worth institutionally. Subsistence protection is a factory setting, prior to every narrative of contribution.

The Case for Compatibility with Efficiency: The New Logic Upgrades Competition Rather Than Eliminating It

The old logic's competitive mechanisms, including national competition, market contest and technological races, were not “mistakes” over tens of thousands of years of human evolution. They were adaptive strategies under particular resource constraints. Competition selects: it removes inefficient institutions, advances technology and stimulates innovation. Simply abolishing it causes systemic efficiency losses and free-riding. The new logic seeks to change its direction. In competitive domains, races in chips, AI and markets continue, perhaps more intensely, but the criterion for winning changes from “making the other unable to survive” to “contributing more additional value to the community resource pool.” In noncompetitive domains, cooperation replaces competition, with the goal of “not failing together” rather than “being safer than you.” At the individual level, the contribution-value system preserves the human motivation to strive ahead, but redefines what “ahead” means. Recognition of contributions and resource support may make people more willing to create, care and take responsibility. Yet the system must not use contribution value to induce overwork, self-depletion, performative relationships or instrumental altruism. Any mechanism that drives people continually to exploit themselves for points is a failure of the contribution-value system.

12. Access, Usability and Autonomous Mobility Across Domains for People and Individual Resource Pools

The Role of the Individual Resource Pool

An individual resource pool is neither a panoramic personal dossier built by a platform nor a conversion of behavior, capabilities, relationships and worth into a unified score. It is a person-led collection of resources and capabilities, stored separately by domain, revealed as needed, authorized item by item and subject to withdrawal. It includes time, knowledge, skills, experience, works, data, relationships, lawful property, capacities for concern and care, and extensions of capability developed with AI.

People are not resources. Individual resource pools belong to people and serve their freedom, development and meaning. Basic healthcare, education, care and other basic public services must not be differentiated into tiers according to whether a person registers, shares, contributes or moves.

Access Across Domains

Access across domains answers “Can I enter?” Subject to law, safety and carrying capacity, individuals should be able to see and enter a wider range of geographical, industrial, organizational, educational, medical, living, creative and public-participation settings. Access does not mean abolishing borders, sovereignty or public-safety management. It means reducing unreasonable geographical, identity, information, qualification and organizational barriers.

Usability Across Domains

Usability across domains answers “After entering, can I genuinely live, participate and create?” Community resources should be lawfully available to individuals. After entering a new node, individuals should be able to continue using their qualifications, learning outcomes, lawful data, works, professional records, social protection and necessary entitlements within agreed bounds. Their own knowledge, skills, equipment, spaces, data and service capabilities may also be used across domains with their authorization, rather than automatically becoming invalid when they leave their former organization.

Autonomous Mobility

Autonomous mobility answers “Who decides entry, staying, change and exit?” Individuals have the right to decide whether to relocate, whether to participate remotely, which capabilities to open, for how long, to whom, when to withdraw and whether to return to the original node. Mobility rights also include the rights not to move, to refuse system matching, to exit and to choose again. No one may lose basic subsistence or basic public services for refusing relocation, a particular job or the sharing of personal resources.

Three Types of Flow Coexist

First, autonomous movement of people: individuals choose their own settings for life, study, work, care and creation. Second, autonomous movement of individual resources: knowledge, capabilities, works, time, equipment, data and AI-enabled extensions of capability cross spatial and organizational boundaries with the individual's authorization. Third, community resources flow toward individuals: education, healthcare, computing capacity, equipment, funding, services and development opportunities proactively reach people and regions short of resources.

AI's Capability Role

AGI-COS and other controlled intelligent systems may help discover resources, identify needs, explain qualifications and rules, offer multiple matching options, explain risks and long-term effects, preserve evidence of authorization and use, and detect anomalies and rights violations. AI may suggest possibilities, but may not

determine an individual's destiny; assist governance, but not monopolize it; and amplify human capabilities and care, but not make people appendages of system efficiency.

The Levels of Operational Mechanisms

People and individual resource pools are key interfaces through which the community enters personal life, but the community's operational mechanisms are not subordinate to individual resource pools. Resource use and intelligence for good span countries, regions, organizations and individuals. Contribution recognition and the reception and continuation of meaning primarily operate at social and individual levels. Pre-institutional language, complete experiments and high-ground verification constrain every level. Their full structure is developed separately below.

The Relationship Between Intergovernmental and Individual Resource Pools

Intergovernmental resource pools address whether the community's critical resources can flow reliably, be assured and be settled. Individual resource pools address whether individuals can autonomously carry, call on, open and use their own resources. The two connect through limited-authorization interfaces, without automatic conversion: intergovernmental settlement amounts cannot convert into personal contribution value, and personal contribution value cannot convert into governance weight in intergovernmental pools. The higher-level resource community must not absorb individual resource pools.

Five Types of Operational Mechanism Spanning the National, Social and Individual Levels

Only a set of cross-level mechanisms connecting the resource community, governance community, intelligence-for-good system and individual institutional interfaces can turn concepts into real operation. The five types below are neither isolated modules nor the exclusive property of any single level.

Operational mechanism	Primary applicable levels	Core role	Non-negotiable boundary
Resource discovery, authorization and effective use	National, regional, organizational and individual	Identify idleness and shortages; establish interfaces for voluntary authorization, matching, use, settlement and exit	No unauthorized ownership changes in the name of effective use; no single resource-control center
Contribution recognition	Primarily organizations, communities and individuals	Recognize public value under public rules and verifiable facts, informing relevant additional resources beyond basic rights	No assessment of personal worth, effect on basic rights, cross-domain aggregate score or conversion into political power
Reception, Continuation and Response to Meaning	Individuals and social relationships	Ensure care, companionship, creation, safeguarding and genuine responses are not forgotten by the system	No forced feedback, competitive scoring of emotional relationships or determination of subsistence entitlement by meaning records
Intelligence-for-good capabilities	Every level from countries to individuals	Provide sensing, matching, exploration, translation, early warning, audit and extensions of capability	Final human review, controllable data and authorization, replaceable multiple nodes and no permanent dependence
Pre-institutional language, complete experiments and high-ground verification	The whole community	Translate principles into comprehensible, verifiable, fallible and correctable institutional language and operating facts	Advance by phase with no loss for anyone; on failure, suspension, exit and rollback are available and evidence is permanently retained

The Dynamic Closed Loop of the Five Mechanisms

The five mechanisms are not five modules placed side by side. They form the cycle by which the community keeps operating and correcting deviations. Resource discovery and authorization bring dormant resources into real use. Real use produces verifiable facts of contribution. Contribution recognition and meaning responses sustain the motivation to participate. Intelligence for good improves discovery, matching, exploration and risk identification. Pre-institutional language, complete experiments and independent verification then test whether these mechanisms have departed from rights boundaries and the three meta-goals, feeding corrections into the next round of resource discovery, authorization and rule design.

Resources enter use → contributions receive recognition → meaning is received and carried forward → intelligence amplifies capabilities → institutional verification corrects deviations → resources are reallocated. If any link is distorted, monopolized or lacks an exit, the entire loop must pause for review; efficiency elsewhere cannot conceal it.

Together, these mechanisms answer three questions: how resources and capabilities reach real needs; how human contributions and meaning are received and sustained; and how institutions and intelligence continue operating without absorbing rights. Any mechanism that works only by first harming other groups, by unauditable complexity or by a single center of intelligence and power must not advance to the next phase.

13. The Precise Nature of Contribution Value and Records of Meaning Response

Contribution value (CV) is an institutional record, based on public rules, verifiable facts and appealable procedures, of the public value generated by individual actions. It determines priority in the use of additional resources beyond basic guarantees that are reasonably related to the type of contribution. It does not assess personal worth, affect basic rights, convert into political power or citizenship, or create a unified cross-domain aggregate score.

Framework note: the Chinese theoretical framework continues to use the term “贡献值.” Its English name for international communication must not be a mechanical literal translation; it requires independent semantic testing, review of legal implications and back-translation verification. In particular, expressions such as “credit,” “score” and “rating” that could be read as personal creditworthiness, social scoring or comprehensive ratings should be avoided. The formal English name remains unfixed until testing is complete.

Domain-Specific Recording and Purpose Limitation: Mandatory Clause

In principle, contribution value uses domain-specific records and purpose limits. It does not aggregate healthcare, education, consumption, social interaction, political expression, law-abiding records and public contributions into one personal score. Each category of contribution record may be used only in additional-resource settings announced in advance and reasonably related to it; it must not migrate automatically across domains. Low contribution, a dispute or a violation in one domain must not automatically reduce qualifications, services or social standing in unrelated domains. Any system combining different spheres of life into a unified tier of personal worth, creditworthiness or citizenship is not the contribution-value system described here.

A Responded-to Need Becomes a Record of Meaning

When vulnerable groups' needs are effectively met, this is recorded as a meaning response in a separate meaning-response ledger. It does not create contribution value or convert into resource-use priority. It is an affirmation of meaning, not a reason for subsistence. Subsistence is unconditional and precedes every contribution narrative. Receiving help need not be dressed up as a contribution to confer dignity; dignity precedes contribution.

Framework-linkage note: “meaning-response records” here are the accounting-level implementation of the experimental-zone documents' principle that “a need met is a contribution.” The principle describes a causal

chain: without a recipient, altruism cannot be completed. Accounting implements separation into two ledgers. The two formulations describe different levels of the same principle, not two sets of rules.

Value standards, the two-ledger structure, human review, handling of falsification and cross-regional mutual recognition belong to subsequent specialized research. The main text locks in these minimum boundaries: no aggregate personal-worth score; no determination of basic rights by contribution; no automatic cross-domain spillover; and no conversion of contribution records into political power or citizenship.

The Principle of Auditable Simplification

More detailed contribution records are not necessarily better. If accounting, verification, resource calls and audit in a domain become too complex for ordinary participants to understand or independent auditors to review within reasonable time, the system has created a new black box. Prefer records with the fewest fields, narrowest uses and least intrusion. Domains that cannot be simplified should not yet enter formal resource allocation.

14. Separation of Three Powers and Dynamic Arbitration

Note on Applicable Levels

The separation of three powers, civilizational voting power, and human-heart-and-meaning review described here apply first to national-level comprehensive experimental zones and other governance units that voluntarily adopt the structure. At the intergovernmental-cooperation level, every vote, resource right and compulsory obligation requires international agreements, member states' domestic laws and explicit authorization. International organizations may establish common principles, review interfaces and dispute-resolution procedures, but may not bypass sovereign states to exercise domestic public power directly without authorization.

Institutional operation requires three types of power, mutually checking and not concentrated. Civilizational voting power uses one person, one vote, with contribution value set to zero for irreversible decisions about civilization's course. Domain decision-making power is determined through a combined assessment of verified professional capability, practical experience, responsibility borne and relevant contributions. Contribution records may demonstrate eligibility to participate and actual experience, but no contribution of funds, resources or reputation may automatically convert into the authority to exercise professional judgment. Resource-use rights allocate additional resources in tiers according to contribution value, remain strictly separate from the first two powers, and cannot convert into voting weight.

Types of power must not be confused with procedural nodes. Launch, execution, suspension and termination are four procedural nodes through which every major matter passes; the appropriate type of power acts at each node. See Selected Appendix Three for minimum conflict-arbitration principles. Full procedures remain in research reserves.

Limited Exemption from Liability

If every proponent is heavily punished after termination is triggered, the perverse incentive is “the more you do, the more you err; do nothing and make no mistakes.” Thus, decision-makers who exercise reasonable care, disclose necessary information, follow procedures and disclose conflicts of interest, without deliberate concealment or gross negligence, bear no punitive liability even if the outcome fails. They are required only to conduct a public retrospective and add it to the experience repository.

Distinguish Accountability from Reassessment of Qualifications

- Accountability: intent, gross negligence, concealment and violations.
- Qualification reassessment: inadequate capability, repeated errors of judgment and declining risk-control capability.

Lack of fitness in capability must not be equated directly with culpability for malicious conduct.

Note on Levels

The three powers in this paper concern decisions at the level of civilization's course, the voting level. The constitutionally held powers of governors to launch and terminate experimental zones in the Proposal for a National-Level Comprehensive Experimental Zone are pioneering trial arrangements at the administrative-authorization level. Their levels and objects differ, and they can operate in parallel without contradiction.

15. Review Against the Three Meta-Goals

Beyond efficiency, resource and risk review, every institutional proposal must pass review by the Human Heart and Meaning Committee. Veto powers, however, must be graded; a single committee must not monopolize them without limit.

Mandatory Veto: The Committee May Suspend Directly

- Basic subsistence rights are made conditional on contribution value.
- Unauthorized monitoring of people's inner lives or psychology.
- An AGI black box directly deprives people of basic rights and interests.
- No possibility of appeal, exit or correction.
- Aggregate scoring of personal worth or systemic discrimination.
- Explicit reduction of people to objects of scoring.

Suspensive Veto: The Committee May Only Require Suspension, Further Justification and Public Reconsideration

- The experience of happiness may decline.
- Relationship structures may become thinner.
- Structures of meaning may become unbalanced.
- Technological convenience conflicts with interpersonal relationships.
- Groups differ in their understanding of happiness and meaning.

For the second category, the committee cannot veto permanently. It can only require suspension, a public explanation, further investigation and renewed examination, followed by determination through the appropriate decision procedure. Constitutional procedures prescribe how the committee is formed, its terms, composition and boundaries of responsibility, and subject it to independent audit and removal. The human heart takes priority over efficiency, but no minority monopolizes the right to interpret it.

Part Four • Evolution: From Guardrails to Structure, from Fear to Expectations

[Reader's Guide] This part asks how evolution proceeds after the foundation is built: the three-stage route, experimental zones as testing grounds, the visibility and learnability of the high ground, failure scenarios and circuit breakers.

[Author's Assessment + Institutional Proposal]

16. Common Security, Common Development and Common Continuity

The community's evolution is not a progression from one slogan to another. It is the gradual entry of three national capabilities into shared operation.

Stage One, common security: preserve the baseline against shared destruction. Implement hotlines, joint verification, minimum-protection commitments, emergency assurance of critical resources and basic safety-intelligence interfaces; launch the first intergovernmental resource pools, individual resource pools and city-level complete experiments. Common security does not require countries to stop competing, only that competition not breach the foundation of shared survival. Its initial motivation comes mainly from shared risks and shared vulnerability.

Stage Two, common development: give more countries, regions and individuals real development pathways. Resource pools expand from emergency assurance into capacity reservations, knowledge and technology licensing, computing capacity and basic public-intelligence services. The qualifications, data, outcomes and entitlements of people and individual resource pools begin connecting across nodes. Resource-poor countries and regions receive support to form their own capabilities, rather than permanent dependence on a single provider. The main motivation is calculable mutual benefit and gains in capability.

Stage Three, common continuity: enable security and development to endure through time. The community begins long-term exploration and correction concerning resource circulation, ecological carrying capacity, technological risks, demographic and fiscal structures, intergenerational effects and institutional distortion. The high-ground effect spreads at scale, and participants start arranging long-term investment, security and development around shared structures. The main motivation comes from shared expectations of the future.

Shared survival is the baseline of common security; shared expectations are the result once common continuity has formed. Security determines whether a country can exist, development whether it has a future, and continuity whether that future is real.

17. Five Levels of Evolution from Issue-Specific Cooperation to Community Structure

Evolution is not only the passage of time but a level-by-level upgrading of institutional form. At each level, the content of cooperation, power structure, source of trust and long-term expectations change qualitatively.

Level One, crisis contact: hotlines, incident notification and joint verification; trust comes from shared fear. Level Two, project cooperation: joint testing, backup computing capacity and disaster collaboration; trust comes from mutual benefit in individual projects. Level Three, mutual recognition of rules: evaluation methods, audit principles and data interfaces; trust comes from predictable procedures. Level Four, limited common governance: dispute resolution, independent audit and common funds; trust comes from the institution itself. Level Five, formation of shared expectations: participants begin planning long-term investment, security and development around shared structures; trust comes from a shared stake in the future. No level requires skipping stages or promises automatic advancement. Advancement occurs when the preceding level's operational record can support the next level's risk exposure.

From Initiative to Adoption: Political Launch and a Chain of Institutional Support

A national-level civilizational experiment cannot depend solely on one leader, department, enterprise or official instruction. An individual may open the door, but cannot alone bear legal, fiscal, data, public-opinion and failure burdens or ensure continuity across terms of office. What must be built is a chain of institutional support that spans individual terms and departmental boundaries.

Role	Core responsibility	Irreplaceable function
Political initiator	Bring the issue into national prospective research and decision procedures	Bear political responsibility for first proposing and launching it
Administrative host	Provide a real city, population, public services and cross-departmental coordination	Convert authorization into an operable social system
Legal and fiscal guarantors	Define experimental boundaries, rights protection, failure costs and remedies	Ensure the experiment is not an adventure without legal protection
Professional implementation community	Handle resources, intelligence, data, audit and social research	Prevent political will from replacing professional capability
Bearers of independent verification and public legitimacy	Protect dissent, appeals and exit, and disclose failures	Prevent the experiment from becoming a political showpiece or an exercise in self-validation

The appropriate decision request now is not one-time approval of an entire new system, but progress through four levels. First decide whether to include it in national prospective research. Then decide whether to establish a cross-departmental research and examination mechanism. Next select candidate cities for baseline surveys and preliminary study of plans. Only after legal review, public communication, minimum closed-loop verification and independent evaluation are passed should authorization of a city-level comprehensive experiment begin. Do not reduce the dimensions of the objective; do not skip levels in decisions.

18. A National-Level Comprehensive Experimental Zone for Civilization Leap

All mechanisms of the resource community, governance community, intelligence-for-good system, and people and individual resource pools require a complete testing ground that reduces geographical scope without reducing dimensions. An experimental zone cannot be merely a single-application pilot. It should have complete population, industry, public-service, resource-supply, social-protection and governance structures, with at least the carrying capacity of a substantial city-level setting.

Reducing scope controls risk; retaining dimensions preserves the complexity of real society. Communities, enterprises and platforms can test individual mechanisms, but cannot test how resources, power, technology, human hearts, public services and long time horizons operate together. City-level scale is not pursued for its own sake. It subjects the new logic to real interests, real conflicts, real ordinary people and real failures.

The experimental-zone proposal itself appears in the Proposal for a National-Level Comprehensive Experimental Zone for Civilization Leap in the AGI Era. Its relationship to the international mechanisms in this paper is as follows. The zone tests whether separation of resource ownership and use, collaborative management of community resources, checks among the three powers, intelligence for good, human-rights boundaries, cross-domain access for people, autonomous authorization of individual resource pools, and the three goals of security, development and continuity can form a closed loop. Intergovernmental resource pools provide external resource security and cross-regional interfaces for verification. In return, the zone's operating data, failure records, appeals, exits and corrections supply a factual basis for calibrating international-level parameters.

The zone must establish an international independent-verification interface from day one. What it must prove is not merely that “the system runs,” but that it genuinely improves the security, development and continuity capabilities of countries, societies and individuals without diminishing human dignity, creating new dependence, turning contribution value into a rating of personal worth or making AI the final adjudicator.

19. High-Ground Visibility and Information Authenticity

The high-ground effect reveals a path different from “breaking” the old logic: transcendence rather than confrontation. A high ground is a locality that first achieves closed-loop verification of the new logic. It need not attack the low ground; it need only keep operating and remain visible. When enough people see for themselves that another possibility truly exists and is genuinely better, the old logic loses explanatory power.

The Following Must Be Protected

- Public criticism.
- Negative reporting.
- Challenges to data and indicators.
- Satire, opposition and independent investigation.
- Public discussion of experimental failure.

The Following May Be Subject to Accountability

- Fabricating experimental data.
- Tampering with audit evidence.
- Impersonating public institutions.
- Cyberattacks.
- Deliberately producing false documents.
- Undermining data integrity.

The first condition for a high ground is not a self-declaration of success, but the ability of skeptics and even hostile actors to verify actual operating results independently. A high ground must show not only higher efficiency but also greater safety, simultaneous reductions in resource idleness and shortages, more real individual choices, revocable authorization, functioning appeals, AI support without permanent dependence, publicly discussable failure, exit without retaliation, and autonomous capabilities gained by disadvantaged regions.

International Credibility and Verification Across Institutional Systems

An experiment launched in China's institutional setting cannot claim global applicability merely because it succeeds within China. Its international credibility should be built in four layers. First, domestic authenticity: operation, appeals, failures and exits all leave verifiable records. Second, independent verifiability: professionals with different viewpoints may access necessary evidence. Third, international observation and participation: gradually open observer roles, joint research, individual joint experiments and audit nodes. Fourth, replication across institutional systems: demonstrate that underlying principles and interfaces can be realized anew under different legal, property, political and cultural conditions, rather than requiring replication of a single Chinese template.

Joint design and funding by multiple countries can strengthen credibility, but cannot be prerequisites for launch. Otherwise, a deadlock forms: international participation is needed to prove credibility, but credibility must first be proved to attract international participation. China may launch autonomously; external verification interfaces must exist from day one. What the high ground exports is verifiable principles, boundaries and operating facts, not a demand that the world copy one institutional form.

Mechanisms for High-Ground Information Authenticity and Independent Verification

- High-ground broadcasting protocol: institutional rules, evaluation methods, overall operating results, failure records and use of public funds should, in principle, be public. Personal privacy, trade secrets, national security, unpublished intellectual property and system-defense details are handled according to minimum necessity, independent audit and explainable exemptions. No confidentiality arrangement may conceal institutional failure, data falsification or abuse of power.
- High-ground observation fund: support on-site visits by independent journalists, scholars and citizens.
- Fact-checking alliance: independent institutions from multiple countries cross-verify findings.
- Multinode mirrored archives: independent academic databases, third-party audit institutions and lawfully operating redundant publication channels.

20. Failure Scenarios and Exit Mechanisms

Three Failure Scenarios and Circuit-Breaker Mechanisms

Scenario A: intergovernmental resource pools become political instruments. Powerful countries use pools as geopolitical leverage, weaker countries are compelled to deposit “voluntarily,” and exit mechanisms exist in name only. Warning signals include a single country's deposit share exceeding its limit; technical delays to weaker countries' exit requests; unilateral changes to evaluation standards; and enduring bloc-based membership, with access used to draw geopolitical lines. Circuit breaker: when one supplier's concentration exceeds the risk threshold, suspend further increases in its systemic weight and initiate alternative-supplier entry, diversification of additional procurement and reduction of risk exposure. Existing resources remain governed by the original agreements and may not be forcibly transferred without owners' consent. Specific deadlines for exit requests require separate determination according to resource type, outstanding obligations, systemic risks and replacement cycles. Supporting responses: make deposit records publicly searchable in real time; explicitly label each deposit agreement “does not constitute a transfer of sovereignty”; provide an exit buffer period; and establish an independent community evaluation fund restricted to publishing agreement texts, translating policy documents, supporting third-party evaluation, organizing public hearings, explaining exit costs and resource flows, and funding independent research.

Scenario B: algorithms monopolize the CV system. Warning signals include unfulfilled open-source commitments, hollowed-out meta-deliberation and overdue appeal backlogs. Circuit breaker: mandatory open-source audit. If meta-deliberation cannot lawfully convene over institutionally specified consecutive cycles, or its procedures remain ineffective, implementation of new standards automatically pauses and reverts to the latest version that passed full deliberation and public audit. Human review becomes a statutory right.

Scenario C: the high-ground effect is systematically obscured. Warning signals include blocked broadcast data, unjustifiably refused verification requests and systematic deletion of independent reporting. Circuit breaker: activate multinode mirrored archives and publication through independent academic databases.

Scenario D: intelligence for good degenerates into data extraction and permanent dependence. Warning signals include nonportable core models, prolonged stagnation in local capacity building, important data sent back by default and key public services dependent on a single external supplier. Circuit breaker: suspend expansion of deployment, begin data- and model-sovereignty review, and require replaceable interfaces, migration plans and local capacity-building plans.

Scenario E: efficiency logic turns people and individual resource pools into instruments. Warning signals include basic rights affected by refusal of system matching, cross-domain data joined by default, difficulty withdrawing personal authorization and administratively or algorithmically compelled mobility. Circuit breaker: immediately suspend the matching and calls concerned, restore the latest version that passed rights review, and initiate independent investigation, individual remedies and a public retrospective.

Part Five • Leap: Transformation of Civilizational Logic, Long-Term Interfaces and the Ultimate Test

[Reader's Guide] This part is preliminary research into long-range interfaces. It is not a current policy recommendation, a near-term engineering roadmap or an established scientific conclusion. It tests whether the foregoing institutional principles still hold over longer timescales and broader civilizational scales, and identifies in advance ultimate prohibited boundaries that today's institutional design must not cross. Current decisions do not require acceptance of every long-range judgment in this part.

[Scenario Exploration] Everything below is conditional preliminary research: the institutional interfaces enter further examination only when the relevant technological and social conditions genuinely hold.

21. Where the Leap Actually Occurs

Civilization Leap is not simply larger-scale cooperation, more international organizations, ever-bigger resource pools or more cross-border movement. It occurs first in a transformation of the underlying logic from which national security, modes of development, resource value, human value and uses of intelligence arise.

At the national level, national interests are no longer understood as residing only within national borders. Security gradually depends less on weakening, blockading and overpowering another party, and more on shared systems being open, redundant, verifiable and repairable, and on the other party still having a future. In development, progress depends less on exclusive possession, permanent blockades and locking other countries into subordinate positions, and more on effective resource use, complementary capabilities, knowledge creation and more actors gaining access to development. In continuity, security and development are no longer purchased by overdrawn resources, ecology, society and the next generation.

For resources, value comes less from how much is owned and increasingly from whether resources are actually used, released, scheduled, regenerated and made to produce public value. For individuals, entitlement to subsistence is no longer determined by functional value. People and individual resource pools can have access, usability and autonomous mobility across domains with their rights protected. A person's sense of worth need no longer come only from outperforming others; it can arise from connection, care, creation, responsibility and genuine response.

In human-machine relations, AI is no longer used mainly to replace people, control them and overpower other countries. It becomes an extension of human capabilities, a tool for common governance, a resource-connection system and a civilizational-risk early-warning system. Intelligence for good is tested not by whether a model professes goodwill, but by whether it truly expands national and individual autonomy, reduces shared harm and remains constrained by rights, audit, exit and final human review.

Realization builds the resource, governance and intelligence foundations. Evolution connects resources, people, capabilities and rules across nodes. Leap changes why civilization operates, whom it serves and what sustains its operation.

22. Layered Identities Without Contribution-Ranked Citizenship

Expansion of identity is gradual. Each step adds new identification, responsibility and community relationships onto the preceding layer. Whether new legal rights arise must be separately confirmed by relevant national laws, regional institutions or international treaties. World citizenship is a civilizational identification and an identity of public responsibility; it does not automatically confer rights beyond national law. Cross-border healthcare, public services, migration and participation rights must be separately confirmed through agreements between countries, regional law or international treaties.

The carbon-silicon covenant is a civilization-scale binding relationship between humans and nonbiological intelligent actors, including AI and AGI, concerning rights, responsibilities, control boundaries, common security and long-term continuity. It does not create, confer or cancel world citizenship. Identification or declaration as a world citizen is open to all. Contribution value affects only an individual's priority for additional resource use in the relevant pool. It does not affect citizenship itself or determine the boundaries of basic rights between humans and nonbiological intelligent actors.

Clause Against Contribution-Ranked Citizenship

No tier of citizenship is graded by contribution value, or purchased or exchanged for it.

23. Planetary-Civilization Interfaces

If humanity eventually develops stable capabilities for activities across planets, Civilization Leap must reserve institutional and meaning interfaces in advance. This is conditional preliminary research, not an assertion of “physical inevitability.” If that condition is met, the principles for interstellar interfaces are as follows.

Interfaces for basic interstellar survival rescue, humanitarian communication, disaster avoidance and civilizational survival must not close because of membership, institutional differences or historical participation. Additional resource calls, project participation, governance seats and long-term development entitlements may be allocated according to participation responsibilities, actual contributions and agreement obligations.

Three principles stand in the main text; long-term technical and institutional details remain in research reserves. Cosmic citizenship is a universal status of subjecthood in a multicelestial-body civilization, without contribution, health or capability thresholds. Eligibility to execute interstellar missions is selected according to capability and team needs: skills may determine who is suited to carry out a mission, but not who deserves more to be a subject of civilization. Scarce migration resources serving as backups for civilizational survival are allocated through public, fair and appealable procedures under the jurisdiction of civilizational voting power. They are not weighted by resource-pool settlement amounts or ranked by contribution value. The prohibited-boundary list is mandatory here: do not let contribution value buy everything; do not let anyone lose basic subsistence protection because of low contribution; and do not use contribution value to assess the value of human life.

24. Physical Assurance and Continuity of Mind and Meaning

Moving into deep space requires two life-support systems: one to sustain the body, another to sustain meaning. Supplies, energy, air and ecological cycles keep people alive. Mind, relationships, memory, freedom and shared visions ensure that civilization remains civilization.

The meaning-support system has five institutional components: mechanisms to sustain relationships under prolonged loneliness and confinement; voluntary transmission of intergenerational missions rather than compulsory inheritance, with descendants free to choose again and no generation using departure-time vows to bind the unborn; synchronization of civilization's memory repositories across planets; adaptation of meaning-feedback channels to time delays; and periodic rebuilding of shared visions. Every generation has the right to reconsider the purpose of the voyage, ways of living together and future direction. Missions may be carried forward, but cannot be frozen. When the material system fails, people die. When the meaning system fails, civilization dies, and that second death can occur while everyone is still alive.

25. The Ultimate Test

Whether Civilization Leap has occurred comes down to one ultimate question: have countries, groups and individuals regained a future they can look forward to?

At the national level: common security, common development and common continuity can improve together; critical resources and intelligent capabilities are not permanently locked behind a single node; countries gain the ability to solve problems autonomously; and national survival does not depend on overpowering another party.

At the group level: groups of different civilizations and institutional systems do not treat one another's existence as negation of their own; disadvantaged regions are no longer merely exporters of resources and talent; and common governance accommodates disagreement, dissent, exit and reentry.

At the individual level: people and individual resource pools have access, usability and autonomous mobility across domains; their capabilities, time, creation, care and lawful resources can reach real needs with their own authorization; they can refuse, withdraw, exit and choose again; and their existence is not negated because they are “useless.”

At the intelligence level: AI genuinely reduces shared harm and amplifies national and individual capabilities, without depriving humans of final judgment in the name of efficiency or creating permanent dependence in the name of support.

Conclusion

Realization, evolution and leap of a community with a shared future for humanity do not require all countries to accept the same values, the same system and the same organization before taking action.

It can begin with countries recalculating their own destiny. Can security rely only on deterrence and blockades? Can development be separated from transnational resources and capabilities? Can today's advantages endure in a world that keeps collapsing? Follow these three questions to their conclusions, and common security, common development and common continuity cease to be externally imposed moral demands and become the complete form of long-term national interests.

It can begin with a negotiating table, a limited authorization of resources and capabilities, a public-intelligence tool helping a resource-poor region develop autonomy, or a city-level complete experiment. It can also begin with one person, carrying their rights, capabilities and individual resource pool, autonomously entering wider settings for life, learning, work, care and creation.

The resource community turns existing disorderly dependence into shared capabilities that can be protected, used, settled and exited. The governance community safeguards sovereignty, rights, responsibilities and checks and balances. The intelligence-for-good system brings resources, opportunities, capabilities and real needs into reach of one another. Together, they serve not a larger management system, but the continued ability of countries, societies and every person to have a future to look forward to.

The institutional starting point that emerged in Shanghai in July 2026 is insufficient to prove that this transformation is complete. But it does prove that, in a world where competition persists, another path supporting national security, development and continuity already has real-world interfaces on which further construction is possible.

A community with a shared future for humanity is not a substitute for national interests. It is their complete form in an era of long-term coexistence among multiple countries.

Civilization Leap does not wait for a perfect blueprint, but every step must preserve evidence, boundaries, exit and correction.

Selected Appendix One: Factual Baseline and Real-World Interfaces

This appendix retains only the factual baseline necessary to understand the main text. Changing organizational status, the strength of members' commitments, full policy texts and subsequent changes require continuing verification. Reality is not frozen at this draft's baseline date.

Principal public sources as of July 21, 2026: the Ministry of Foreign Affairs of the People's Republic of China's July 16, 2026 release, “Signing Ceremony for the Agreement Establishing the World Artificial Intelligence Cooperation Organization Held in Shanghai”; and official materials for the relevant keynote address and action plans of July 17, 2026. This paper uses them only as real-world interfaces, not as grounds for assuming that any government or international organization has adopted its institutional proposals.

Fact	Use in this paper	Boundary
On July 16, 2026, representatives of 29 countries signed the agreement establishing the World Artificial Intelligence Cooperation Organization in Shanghai	A real-world anchor showing that a negotiating table for multilateral AI cooperation has emerged	Signing, legal entry into force, organizational launch and actual operation must be verified separately
On July 17, 2026, the 2026 World Artificial Intelligence Conference and High-Level Meeting on Global AI Governance opened, and related action documents were released	A policy interface for common development, ethical governance, capacity building and intelligence for good	Does not imply that all institutional proposals in this paper have been adopted or endorsed
Global AI capabilities, governance, industry and safety networks remain fragmented	Shows the need for interfaces across tracks	Does not imply uniform national positions or inevitable cooperation

Selected Appendix Two: Summary of Core Institutional Boundaries

Structure	Minimum definition	Boundary that must not be crossed
Resource community	Coordinating some use, assurance, emergency call-on and settlement relationships with ownership clear	Not common ownership, unified redistribution or possession by a single center
Governance community	Limited, layered, polycentric, auditable, correctable common governance with exit available	Does not replace countries or concentrate management, audit, arbitration, execution and termination powers
Intelligence-for-good system	Amplifies national and individual capabilities through final human review, data sovereignty, resistance to dependence and auditability	Does not determine personal worth, basic rights, entitlement to subsistence or irreversible civilizational direction
People and individual resource pools	Individual control, separate domain storage, item-by-item authorization, portability and revocability	People are not resources; refusing mobility, sharing or contribution does not affect basic rights
Contribution value and meaning response	Domain-related records for additional resources beyond basic rights; meaning responses recorded separately	No aggregate personal-worth score, automatic cross-domain spillover or conversion into political power or citizenship
Separation of three powers	Civilizational voting, domain decisions and resource use are separate	Funds, resources and contributions may not automatically convert into professional judgment or weight in one-person-one-vote decisions

Selected Appendix Three: Summary of Phases, Stopping and Exit

For any mechanism to enter the next phase, it must submit factual evidence, proof of no loss, independent review, a failure-and-rollback plan and an alternative path that does not advance. When phase changes concern irreversible directions and everyone's fundamental rights, they must enter one-person-one-vote civilizational voting, with no weighting for contribution records.

Mandatory stopping condition	Minimum response
Basic subsistence or basic public services are made conditional on contribution, ratings or obedience	Suspend immediately and restore rights
Personal data are joined across domains by default to form a unified personal profile	Stop linkage, delete unauthorized mappings and conduct independent audit
AI directly makes irreversible decisions about rights or civilization	Revert to the human-authorized version
Temporary or crisis powers are renewed indefinitely without independent review	Powers expire automatically, followed by a public retrospective
Exit is obstructed, withdrawal is difficult or exit triggers retaliation	Suspend expansion; complete remedies and verify exit
Stable-phase rules are prematurely imposed on launch-phase participants	Return for redesign; "future benefits" are no defense
A group's net loss is a prerequisite for launching the proposal	Identify an adversarial deadlock and redesign motivations and returns

Selected Appendix Four: Research Foundations and Version Changes

Research foundations include complex interdependence and asymmetric dependence, common-pool governance, polycentric governance, crisis communication and confidence-building measures, technological-risk governance, public rights and auditable institutional design. Academic anchors establish conceptual interfaces open to discussion; they do not replace real-world verification of the mechanisms proposed here.

Data discipline: every fact, indicator and organizational status must identify its source, date, definition, confidence level and conditions triggering updates. Inferences must not masquerade as facts, nor institutional proposals as current rules.

V3.9.3, symmetry of the three pillars and closure of the operational loop: building on V3.9.2's completion of the intelligence-for-good institutional skeleton and correction of operational levels, this edition consolidates the governance community's five-power structure of rulemaking, execution and coordination, independent audit, dispute arbitration and civilizational veto, specifying its polycentric bearers and mutual checks. It clarifies the dynamic loop from resource use to institutional verification and back to resource allocation, and reduces repeated explanations of existing governance tools' limits. Upon completion of this edition, the main text is formally frozen. New facts enter the factual register, new subjects enter subject-specific research, and new parameters enter experiments and specialized reserves. Main-text versioning reopens only if core causation, rights boundaries, phase discipline or the three-pillar structure changes.

Continue from This Paper

The following entries follow the paper's natural sequence of questions without repeating the explanation of the framework already in the main text. All lead to current stable document pages on www.civitas.top.

This Paper's Website Page

<https://www.civitas.top/en/documents/p03-strat-001/>

Original source shortcut (provenance): <https://www.civitas.top/go/p03-strat-001/>

01

[Ten Questions on Civilization Leap](#)

Helps first-time readers see the real problems, central judgments and where to read next.

<https://www.civitas.top/en/documents/p01-entry-001/>

02

[Civilization Leap Research Collection: First-Reading Edition](#)

Helps first-time readers see the real problems, central judgments and where to read next.

<https://www.civitas.top/en/documents/p01-entry-002/>

03

[Civilization Leap Entry-Route Map](#)

Helps first-time readers see the real problems, central judgments and where to read next.

<https://www.civitas.top/en/documents/p01-entry-003/>

[Complete Public Library](#)

<https://www.civitas.top/en/library/>

[Public Reading Gateway](#)

<https://www.civitas.top/en/documents/p01-entry-006/>

Author: Zijun Fu | Produced by: Civilization Leap Research Group | Unified public domain: www.civitas.top