DISCUSSION WHITEPAPER

THE BETTER AI RACE

A WAICO–OSAA Bridge for Competitive Autonomy

TEMPORAL AUTHORITY SYSTEMS PBC

OCUP.AI

A DISCUSSION WHITEPAPER FOR WAICO AND OSAA

THE BETTER AI RACE

A WAICO–OSAA Bridge for Competitive Autonomy Above a Shared Floor of Human Sovereignty

THE PROPOSITION

Let America build. Let China build. Let every region build. Let markets reveal the best autonomous systems—but let humanity establish the boundary beneath them together.

Temporal Authority Systems PBC | OCUP.ai

August 2026 | Independent proposal for consideration by WAICO and OSAA

Executive Thesis

This standalone paper is addressed to the emerging standards communities around WAICO in Shanghai and OSAA in the United States and global open-source ecosystem. It proposes a focused bridge between their work: a common, testable authority substrate for consequential autonomous systems. Temporal Authority Systems PBC submits the proposal independently; no endorsement, participation, or formal relationship with WAICO, OSAA, or their participants is implied.

The prevailing U.S.–China AI narrative asks which nation will reach the most powerful systems first and thereby secure economic and military superiority. That framing is incomplete. Capability without governable authority can create a shared vulnerability: powerful machines that criminals, malicious insiders, compromised operators, black-market modifiers, or the machines themselves can keep active outside legitimate control.

The better race is not to unleash autonomy fastest. It is to deploy extraordinary capability at scale while retaining verifiable human sovereignty over who authorized it, for what capability, for how long, and under what conditions.

This paper proposes a narrow area of cooperation beneath continued competition: consequential civilian autonomous systems should operate under externally issued, time-bounded authority that they cannot indefinitely extend, restore, enlarge, or grant themselves. Nations and companies remain free to compete on models, robotics, chips, manufacturing, applications, price, and performance. The shared substrate does not choose winners; it establishes the conditions under which every winner remains governable.

The proposed substrate is federated, not centralized. Each jurisdiction retains its own cryptographic roots of trust, validators, policies, legal rules, and enforcement authority. Interoperability requires only a shared evidence grammar through which one sovereign domain can verify another's bounded-authority credentials without accessing its models, private keys, training data, or command systems. **Like international civil aviation, the objective is mutual safety verification—not a global controller**, centralized validation network, or universal kill switch.

CIVILIZATION-LEVEL RULE

No consequential autonomous system may authorize itself.

OCUP is presented as one technical pathway toward that boundary through hardware-enforced Temporal Authority Governance, Self-Extension Denial, external validation, and authenticated evidence. RAAVE extends the model toward identity, lifecycle continuity, compliance signaling, noncompliance detection, and infrastructure eligibility. The Quorum Validator Network (QVN) supplies the distributed human and institutional authorization pattern. These technologies are under development; they are not currently international standards or certified production infrastructure.

THE WAICO–OSAA OPPORTUNITY

Two influential autonomous-system communities can help establish a shared substrate

early—before incompatible authority models harden across chips, agent frameworks, robots, infrastructure, and national markets.

1. The Red-Herring Race

The claim that whoever 'wins AI' necessarily obtains durable supremacy encourages speed to be treated as strategy and safety as delay. Yet neither country benefits from saturating global markets with systems whose authority cannot be reliably bounded. A nation can lead in models and still lose control of downstream machines; a company can lead in robotics and still create a platform that unauthorized operators exploit.

The arms-race narrative can become self-fulfilling when assumptions about an opponent justify weaker safeguards at home. Public discussion in both countries should therefore separate three questions: who develops the strongest technology, who deploys it most effectively, and whether the deployed technology remains subordinate to legitimate authority. The first two can remain competitive. The third creates a shared interest.

2. The Better Race

A healthier measure of leadership is the ability to build and deploy the most capable autonomous future while proving that no machine can authorize itself and no rogue operator can activate consequential autonomy invisibly.

- Compete on intelligence, robotics, semiconductors, energy efficiency, manufacturing, applications, affordability, and usefulness.
- Preserve national control over models, weights, training data, cryptographic roots, validators, policy, and enforcement.
- Cooperate on a minimal technical grammar for identity, bounded authority, expiration, non-self-extension, and trustworthy evidence in civilian systems.
- Allow jurisdictions to decide which machines are consequential, which capabilities require stronger authorization, and what consequences follow noncompliance.

A NARROWER BASIS FOR COOPERATION

The United States does not have to trust China's models. China does not have to trust America's. Neither must disclose proprietary technology. Each needs a way to verify that a consequential machine holds valid authority and cannot renew that authority for itself.

3. The Common Boundary

Temporal Authority Governance is a framework in which consequential machine authority is externally issued, capability-specific, time-bounded, and made technically consequential at the operating boundary.

OCUP's defining security property is Self-Extension Denial: the governed system cannot indefinitely preserve, renew, enlarge, restore, or grant its own authority.

A credible implementation should make authority loss produce denial or bounded safe degradation; prevent stale, replayed, reordered, or superseded approvals from restoring authority; require fresh external authorization for recovery; and preserve authenticated evidence of grants, denials, degradation, quarantine, and recovery.

Temporary loss of network connectivity or validator availability should not convert a previously valid grant into indefinite authority. A deployment may permit a preauthorized, time-limited offline continuity window for specified low-risk functions; when that window expires, trusted time or state continuity is lost, or required evidence cannot be preserved, the machine should enter a predetermined denied or bounded safe-degradation state. Restoration of broader authority requires fresh external authorization rather than self-recovery by the governed system.

Interoperability need not mean centralized control. Sovereign validator domains can retain their own membership, cryptographic infrastructure, legal rules, and risk tolerances. What crosses borders is not command authority but evidence: machine identity, authorized capability, validity period, issuer domain, continuity state, and the result of compliance verification.

4. The More Persistent Threat: Unauthorized Human Use

Great powers remain capable threat actors, and no technical architecture eliminates geopolitical conflict. But the more common operational danger may come from people using autonomy outside legitimate authority: organized crime, extremist groups, malicious insiders, compromised fleet operators, black-market repair networks, wealthy private actors, and ordinary criminals who obtain increasingly capable machines.

National origin, corporate ownership, registration, and a familiar logo are not proof of current authority. A legitimate machine can be stolen, cloned, altered, remotely hijacked, separated from its approved components, or operated beyond an expired grant. The relevant question is not merely who made it, but whether it is presently authorized for the capability it is exercising.

RAAVE—Runtime Authority Assurance, Validation, and Evidence—addresses this external verification problem. Its design territory includes authority-conditioned identity, physical-cryptographic binding, lifecycle continuity, claimed-versus-observed reconciliation, alternate-path detection, capability-specific compliance signaling, monitored authority zones, and authenticated evidence history. These are development objectives and evaluation pathways, not claims of deployed global infrastructure.

5. A Sovereignty-Preserving U.S.–China Mission

The United States and China possess two of the world's most consequential ecosystems for AI, robotics, chips, manufacturing, and autonomous infrastructure. Their rivalry will not disappear, nor must it. The

opportunity is a track-two and standards-first mission among engineers, scientists, civil society, insurers, infrastructure operators, standards bodies, and public institutions.

Early dialogues can test whether both ecosystems can recognize the same narrow invariant: consequential civilian machines should not possess an unbounded ability to continue operating solely because their own software says they may. From there, participants can explore shared conformance scenarios without sharing models, sensitive designs, or national-security systems.

This proposal is civilian in scope. Defense and national-security systems remain under each nation's exclusive sovereign control. The model is analogous in shape—not identical in institution—to civil aviation: competitors and geopolitical rivals can share a technical safety grammar for civilian interoperability while retaining sovereign authority everywhere else.

An Aviation-Safety Model for Sovereign Autonomy

International aviation does not depend on a single global cockpit, licensing authority, validation network, or kill switch. Each nation retains sovereign responsibility for registering aircraft, issuing or recognizing certificates, maintaining its own trust infrastructure, controlling its airspace, and deciding what may enter. International operation becomes possible because states recognize shared minimum standards and can evaluate credentials issued under another sovereign system.

A comparable model could govern consequential autonomous systems. The United States could retain its own hardware roots of trust, cryptographic keys, validators, legal rules, risk tolerances, and enforcement authority; China and every other participating jurisdiction could retain theirs. No foreign validator would control another nation's machines, and no central network would possess universal authority to activate or disable them.

Within the cross-border authority-verification transaction, only authenticated evidence need be exchanged. An American autonomous cargo vessel entering a Chinese port—or a Chinese vessel entering an American port—could present a tamper-evident, cryptographically verifiable receipt establishing its machine identity, authorized capabilities, issuing domain, validity period, and authority-continuity state. The receiving port could evaluate that evidence under its own sovereign admission rules without accessing the vessel's proprietary model, training data, internal reasoning, private keys, or command systems.

OCUP supplies the bounded-authority state and Self-Extension Denial behavior. RAAVE provides the identity, lifecycle, compliance-signaling, detection, and evidence architecture through which that state can be externally evaluated. Device-bound roots of trust, physical-cryptographic binding, and tamper-responsive protections can strengthen confidence that the evidence corresponds to the physical machine presenting it rather than a clone, substituted component, or fabricated software claim.

The objective is therefore not centralized global control, but sovereignty-preserving interoperability: different nations, different validators, different laws—and a shared technical grammar for determining whether a consequential machine is presently authorized to act.

WHAT COOPERATION DOES NOT REQUIRE

No shared models. No shared weights. No foreign control of domestic validators. No single global kill switch. No surrender of proprietary chip designs. No inclusion of military systems.

WAICO and OSAA: Complementary Roles

WAICO can provide a high-value forum for examining how authority, identity, compliance signaling, and infrastructure eligibility apply across China's fast-growing ecosystem of embodied AI, robotics, manufacturing, and exported autonomous infrastructure. Its standards-oriented setting can test whether the proposal fits Chinese technical, commercial, and sovereign-governance requirements rather than importing assumptions from another jurisdiction.

OSAA can provide a complementary open implementation and interoperability pathway across agentic software, open-source infrastructure, platform vendors, hardware adapters, and reproducible security testing. Its ecosystem is well positioned to examine how a common authority interface can remain vendor-neutral and composable while producing evidence that downstream infrastructure can verify.

The bridge is not institutional merger or political alignment. It is reciprocal technical work: compatible definitions, threat models, conformance scenarios, evidence fields, and reference interfaces that each community may evaluate and adapt within its own governance structure. Either community can proceed independently; interoperability becomes valuable where their civilian systems, supply chains, and markets meet.

6. Global South and Multipolar Relevance

The autonomous era will not be built by two nations alone. Chinese and American hardware, cloud systems, vehicles, humanoids, and infrastructure will be deployed across Europe, the Middle East, Africa, Latin America, and Asia. Many countries will use mixed supply chains and will not want their safety architecture dependent on choosing one geopolitical bloc.

A manufacturer-neutral authority and evidence layer can help importing jurisdictions verify present eligibility without demanding access to proprietary models or trusting static origin labels. This can support digital sovereignty, safer procurement, insurance, maintenance continuity, and cross-border operation while allowing each jurisdiction to define its own lawful authority domain.

7. Let Markets Choose Above the Boundary

Temporal Authority Governance should not decide which robot, model, operating system, or nation's product is best. Markets, public procurement, and users can reward performance, reliability, affordability, and usefulness. The authority substrate performs a different job: it prevents competition from externalizing an ungovernable risk onto everyone else.

This approach treats safety infrastructure as an enabler of market scale. Insurers can distinguish governed from ungoverned risk; infrastructure operators can condition access on current evidence; manufacturers can prove authority behavior; and responsible autonomy can enter more sensitive environments with a clearer basis for trust.

8. Human Sovereignty and Autonomous Abundance

If autonomous systems generate extraordinary productivity while reducing demand for conventional labor, the authority infrastructure required to govern those systems may also provide a new form of human participation. Legitimate machine operation will require issuance, validation, renewal, continuity assurance, evidence, certification, and governance. Those functions can remain rooted in people and human institutions.

In a future model, a QVN could do more than approve bounded authority. Subject to law, privacy safeguards, governance design, technical validation, and fair access, people and institutions could participate in verification and receive compensation from protocol, assurance, certification, or operating fees associated with governed machine activity. Increased autonomous productivity could therefore increase demand for human-governed validation and create a channel through which some value recirculates.

THE ABUNDANCE PRINCIPLE

Machines may create abundance, but legitimate machine authority should remain rooted in people and human institutions—and the value created by validation can help keep humanity economically connected to the autonomous systems it authorizes.

This is not a present commercial deliverable, a guaranteed economic outcome, a tax on automation, or a proposal for one centralized global validator. It is a long-term design direction. The immediate task is technical and institutional: prove the boundary, validate the evidence, build the hardware path, test governance, and establish credible standards processes.

9. A Practical Cooperation Ladder

1. **Parallel WAICO and OSAA review.** Invite each community to evaluate the non-self-authorization invariant against its own architectures, risk assumptions, and deployment priorities.
2. **Liaison or track-two technical dialogue.** Convene willing participants from both communities—along with researchers, insurers, infrastructure operators, and civil society—without requiring formal political agreement.
3. **Shared definitions and threat models.** Define consequential systems, temporal authority, Self-Extension Denial, identity continuity, validator failure, rogue-operator scenarios, and evidence requirements.
4. **Open conformance scenarios.** Publish reproducible tests for expiration, quorum loss, replay, unauthorized renewal, component replacement, cloning, bypass, silence, and recovery.

5. **Two reference paths, one evidence grammar.** Support OSAA-oriented open C/Rust interfaces and adapters alongside WAICO-oriented embodied-AI, chip, robotics, and infrastructure evaluation, while mapping both to compatible evidence semantics.
6. **Sovereign implementations.** Allow each jurisdiction and vendor to implement common behavioral requirements using its own trusted roots, validators, policies, and hardware.
7. **Independent evaluation.** Run bounded research, enterprise, academic, and public-sector evaluations before any claim of production readiness.
8. **Recognized standards pathways.** Advance validated specifications through appropriate national and international bodies, with FRAND access where patented technology becomes standard-essential.

10. Development and Claims Boundary

OCUP currently includes a bounded Rust reference control plane and deterministic evidence scenarios for temporal leases, quorum validation, expiration, Self-Extension Denial, degraded conditions, quarantine, recovery, and replayable audit output. Rust provides a systems-oriented basis for pre-production testing; it does not by itself create non-bypassable hardware enforcement.

Production-level realization would require trusted time and monotonic state continuity, secure key custody, authenticated firmware and rollback protection, device-bound hardware roots of trust, remote attestation, and FPGA, embedded-controller, chiplet, or silicon integration that physically gates consequential resources. Secure elements and trusted execution environments may protect keys, measurements, and evidence, but do not by themselves establish a non-bypassable authority boundary; the controlled resource pathway must remain outside the governed workload's exclusive control. Appropriate independent validation and certification would also be required.

11. The Invitation

WAICO and OSAA do not need identical charters, governance, participants, or implementation stacks to recognize that neither ecosystem benefits from powerful autonomous chaos. The world does not need one authority controlling all machines. It needs a shared technical expectation that consequential machines remain subordinate to legitimate, externally verifiable authority.

The invitation to WAICO and OSAA is therefore simple: establish parallel exploratory workstreams on Identification, time-bounded authority, Self-Extension Denial, compliance evidence, and the exclusion of unauthorized machines from legitimate infrastructure. Build sovereign and vendor-neutral implementations. Test them against common scenarios. Compare evidence. Let standards emerge from reproducible results rather than geopolitical rhetoric.

THE BETTER FINISH LINE

Build the most capable autonomous future—but build it so no machine can authorize itself, no rogue human can operate powerful autonomy invisibly, and humanity continues to participate in the abundance it creates.

Let America go deep. Let China go deep. Let every region build boldly. The autonomous era can remain competitive without becoming ungovernable—and the nations that prove both capability and control may define the most durable form of leadership.

Working Definitions

Temporal Authority Governance. A framework in which consequential machine authority is externally governed, capability-specific, time-bounded, non-self-extensible, and made technically consequential at the operating boundary.

Temporal Authority. The bounded authority granted to a machine or component for a defined capability, period, and operating context.

runtime authority. A lowercase descriptive term for the authority a system possesses or is evaluated against while operating.

OCUP. One Chip Unified Protocol; the proposed enforcement architecture for hardware-rooted Temporal Authority Governance.

Self-Extension Denial. The property that a governed system cannot indefinitely extend, restore, enlarge, or grant its own authority.

RAAVE. Runtime Authority Assurance, Validation, and Evidence; the proposed identity, lifecycle, signaling, detection, and evidence architecture surrounding governed machines.

QVN. Quorum Validator Network; a distributed pattern for externally governed authorization, recovery, and evidence.

Context and Attribution

This paper was informed in part by the public debate over whether an assumed U.S.–China AI arms race can become self-fulfilling, including journalist Yi-Ling Liu's discussion of safety rollbacks, track-two dialogue, digital sovereignty, and the global spread of Chinese technology in the POLITICO Magazine interview "What You're Getting Wrong About China and AI." The policy synthesis, Temporal Authority Governance framework, OCUP architecture, RAAVE architecture, Self-Extension Denial formulation, and QVN abundance thesis presented here are those of Temporal Authority Systems PBC.

Public discussion draft. Not legal advice, a safety certification, a standards declaration, or a representation of U.S. or Chinese government policy.