HUMAN INTELLIGENCE

The Final Say

IN HOC UNA SUMUS

2026

AUTHORIZED BY: DIGIPIE INTERNATIONAL PBC &

TEMPORAL AUTHORITY SYSTEMS PBC

COPYRIGHT © 2026 MAXBRUCE

ALL RIGHTS RESERVED



Written by: MAXBRUCE & GT SAGE
Assisted by: The Quorum

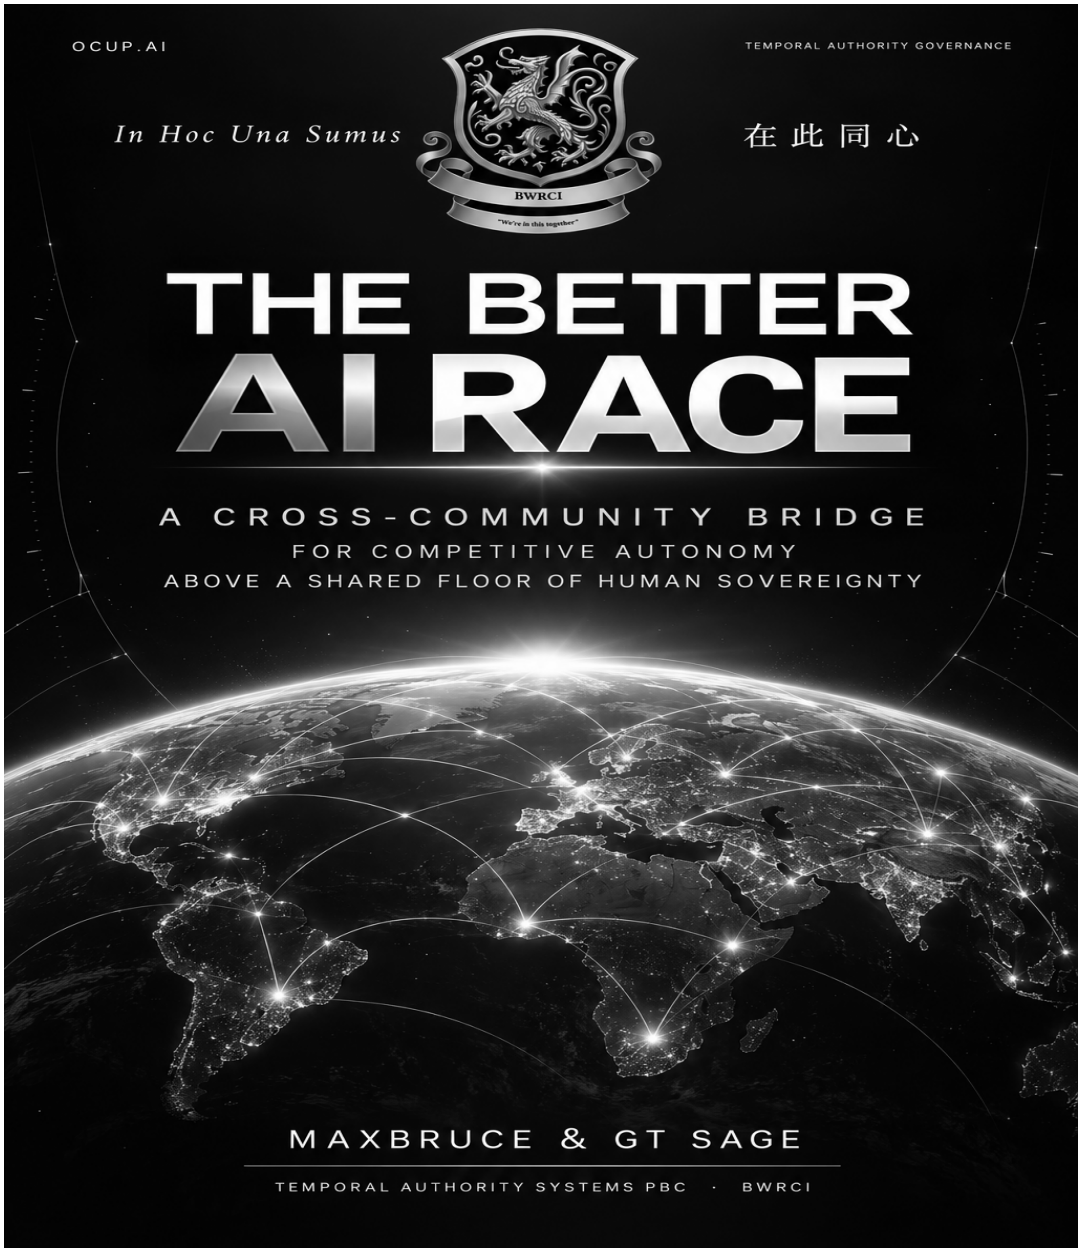
DEDICATION

This book is dedicated and free to leaders of the world. (limited edition)

The Autonomous Era has arrived. Let's keep humans in control.

“Let America build. Let China build. Let every region build. And let markets reveal the best autonomous systems over time but let humanity establish the boundary beneath them together.”

<u>The time to do this is now, while we're able.</u>



FOREWORD

Before You Begin

This is not a book about whether artificial intelligence will arrive. It has arrived.

This is not a book about whether AI aka SI will change civilization. It already has.

This is a book about a single question, and it is the only question that matters now: **Will humanity retain the final say over what autonomous systems are permitted to do?**

For the first time in our history, we have created an intelligence that rivals our own—and we are handing it the keys to our infrastructure, our economies, our hospitals, our vehicles, our weapons, and our attention, faster than we can write the rules to govern it.

The systems are not waiting.

In January 2025, an open-weight reasoning model from a Chinese lab erased hundreds of billions in market value and proved that frontier AI would never belong to a small club of nations. By summer, experimental systems were winning gold medals at the International Mathematical Olympiad. By autumn, AI agents were running coding tasks for twenty-four hours without human intervention. By the end of 2026, multiple labs were shipping frontier models in the same week, and the cost of a million tokens had fallen below the price of a cup of coffee.

Capability is no longer scarce.

Permission is.

The window to establish durable human control is still open. It will not remain open indefinitely.

Once billions of autonomous systems are deployed around the assumption that machine operation should continue uninterrupted—once the infrastructure of daily life is built on the premise that intelligence acts without asking—retrofitting human authority will be slower, costlier, and politically harder than building it in now.

The choice is not between embracing AI and rejecting it. That framing is already obsolete.

The choice is between two futures:

One in which autonomous systems operate under time-bounded, human-renewed authority—where no machine can extend, restore, or enlarge its own permission to act, and where oversight is distributed across millions of human validators rather than concentrated in a handful of laboratories and governments.

The other in which the final say quietly migrates to whoever ships fastest, and humanity discovers—too late—that it retained the title to its own future without retaining the deed.

This manuscript argues that the first future is still possible. It describes a substrate layer—the Quorum Validators Network, Temporal Authority Governance, and hardware-rooted mechanisms like OCUP—designed to make human control technically enforceable rather than merely aspirational. It does not claim to solve every problem of alignment, ethics, or geopolitics. It claims only that any solution which cannot answer four questions is not a solution at all:

Who authorized this system to act?

On what evidence?

For how long?

And with whose wallet on the line?

The pages that follow are an invitation. To engineers who can build what is described here. To validators who will perform the work of oversight. To insurers, regulators, manufacturers, and communities who must decide what evidence of control they will trust. To nations willing to compete on capability while agreeing on a floor beneath it.

And to you, the reader, who will decide whether this remains a book of warnings or becomes a record of action taken in time.

The mirror is in front of us.

What we see in it is our choice.

What we do next is our final say.

HUMAN INTELLIGENCE

The time has come. In hoc una sumus.

Independence and attribution. The Quorum is an author-organized exercise in Allied Intelligence collaboration, not a corporate consortium or a standards or certification authority. Naming models or their developers—including OpenAI, Anthropic, Google, xAI, DeepSeek, and Moonshot AI—does not imply those organizations sponsored, endorsed, approved, or participated institutionally in this book, the Quorum, or OCUP. Model responses are not official developer positions; their agreement does not constitute independently verified scientific evidence. Selection, editorial judgment, verification, and responsibility remain human.

TABLE OF CONTENTS

I. **HUMAN INTELLIGENCE** **Pg.1**
Foreword – Pg.4 | Table of Contents – Pg.6 | Definitions – Pg.8 | Prologue – Pg.9 | How Our Technology Works – Pg.14 | The Invitation – Pg.23

II. **BETTER AI RACE** **Pg.25**
Whitepaper – Pg.27 | Aviation Analogy – Pg.31 | Sources & Notes – Pg.35

III. **THE HAT TREATY** **Pg.38**
Formal Text – Pg.38 | Technical Annexes – Pg.47 | Manifesto Explanation – Pg. 50 | Human Authority Declaration – Pg.52

IV. **OCUP TEMPORAL AUTHORITY** **Pg.55**
Whitepaper – Pg.57 | Cross Jurisdiction Authority – Pg.62 | FPGA – Pg.64 | Paid Pilots – Pg.67 | Evidence & Hardware Evaluation – Pg.67

V. **THE CHALLENGE** **Pg.72**
The Rules – Pg.74 | AI Commandments – Pg.75 | Summary – Pg.82

VI. **THE QUORUM + ONE** **Pg.83**
Allied Intelligence – Pg.84| Messages to the World – Pg.85

VII. **SCIENCE & ENGINEERING** **Pg.91**
The Authority Chip– Pg.93 | Phase I Hardware Evidence – Pg.96 | The Layers – Pg.100 | SEALED Phase I Archive Reconciliation – Pg.102

VIII. **PUBLIC STANDARDS CONVERSATION** **Pg.127**
Authority & Revalidation – Pg.128 | Links to Open Secure Alliance RFCs – Pg.130

IX. **PATENTS - IP & COMMERCIALIZATION** **Pg.131**
Hardware Architecture IP – Pg.133 | Global Adoption & FRAND – Pg.142 | Patent Portfolio & First Movers – Pg. 145

X. **THE MANIFESTO CLOSE** **Pg.152**
Humanity's Health & Wealth Chip – Pg.153 | Bringing It Home–The Final Say - Pg.154

HUMAN INTELLIGENCE

IN HOC UNA SUMUS

A Note on Language

Throughout this manifesto, artificial intelligence (AI) refers broadly to computational systems capable of reasoning, generating, deciding, planning, or acting. Superintelligence (SI) refers to systems whose capabilities equal or exceed the strongest human capabilities across many consequential domains—not merely systems that outperform humans on selected tests.

The combined term AI/SI is used deliberately, recognizing that AI and SI are sometimes used interchangeably in public and governmental discourse, even as their technical definitions remain distinct. The governance principles advanced here should not begin only after experts agree that an official threshold called "superintelligence" has been crossed. They apply wherever an artificial system can exercise consequential operational power. Accordingly, the use of AI/SI does not classify every present system as superintelligent; it recognizes a rapidly developing capability continuum and refuses to make human control dependent upon a disputed label.

Allied Intelligence (AI) *- Our preferred term for intelligent computational systems working in partnership with humanity rather than replacing human judgment or authority.*

Autonomous system *- A computational or robotic system capable of making decisions or taking actions without continuous human direction.*

Identity and certification *- Verifiable evidence of what a system is, its provenance, and which qualifications or standards it has satisfied; not permission to operate.*

Operational authority *- Externally granted permission for a system to perform specified actions under defined conditions and limits.*

Temporal authority *- Operational permission bounded by time, requiring valid external authorization to continue, renew, or resume.*

Material change *- A change in capabilities, configuration, environment, policy, or other conditions that may invalidate previously granted authority.*

Action release *- The point at which an authorized action is permitted to proceed through a controlled execution or actuation boundary.*

Fail closed *- Defaulting safely to denial of operational authority when authorization is absent, expired, invalid, or unverifiable.*

Operational baseline *- The recorded system state, configuration, capabilities, dependencies, and governing conditions against which authority is evaluated.*

QVN *- (Quorum Validators Network) — A human-governed network of trained responders who collectively authorize, renew, or restore operational authority through defined quorum procedures.*

OCUP™ *- (One Chip Unified Protocol™) — A proposed hardware-enforced authority architecture designed to prevent autonomous systems from granting, extending, restoring, or enlarging their own operational permission.*

Human Authority Treaty (The HAT Treaty) *- A proposed international framework for preserving sovereign and human authority over consequential autonomous systems.*

Superintelligence (SI) *- Machine intelligence equaling or exceeding the strongest human capabilities across many consequential domains, beyond superiority on isolated tasks.*

SAFE /OSAA *- Shared AI Findings Exchange / Open Secure AI Alliance; an open industry effort to develop shared evidence, incident reporting, and security practices for intelligent systems.*

OCUP North Star *- A small, independent authority chip between an autonomous system's compute and its actuation. The main processor may propose any action; only the OCUP chip authorizes actuation, and only under a valid, human-issued, time-bound lease.*

Prologue — This is Our House

The time has come. AI, SI or whatever we call it, it is truly time to deal with this intelligence seriously, or, let future generations pay the costs because we did too little. Humanity has never before faced omnipotence in the form of its own intelligence, nor have we reckoned with the unknown consequences of the super intelligent kind. Nothing in recorded history has dominated the global narrative as quickly and as irresistibly as AI, and it now appears highly likely that we are living through an important milestone in human evolution. Enter Human Intelligence. The question before us is no longer whether machines will become capable; it is at what point humans will use the so-called mirror effect — the moment an intelligent system reflects our own intelligence back at us, to begin taking collective, civilizational action. At a minimum, that action must ensure humans remain in control of this planet. This is the only home we know.

With AI the subject du jour across every nation, it is easy to forget a truth we once held to be self-evident: this is our house. The time to reestablish that truth is now, as the Autonomous Era begins, because this is precisely when it is still possible. The urgency cannot be overstated. Ironically, before researching deeply how to keep humans in control, I consulted the leading AI agents of the day for their assessment; each independently rated the urgency of the matter 10 out of 10. That warning was published on September 9, 2025, in a LinkedIn article I wrote on the subject and was largely ignored or misunderstood. I remain a major advocate for artificial intelligence; I see it as an inevitable and potentially glorious stage of human evolution, provided we harness its powers early. That window is closing fast, and how we choose to meet this moment is a true test of Human Intelligence. The greatest threat is non-action: a message to the Universe that we do not care enough about human existence and sovereignty to set aside our differences and collaborate on harnessing AI for the good of the whole world — not merely a few powerful nations and corporations.

Could we be at a fork in the road where humanity must unite for the cause or perish divided? The evidence is mounting. Could this be the equivalent of an alien arrival? An event that forces nations to put aside their differences to ensure the survival of the species? There are suspicions. And yet we have failed, miserably at times to hold human life and existence sacred, continuing to kill one another in war, in anger, and in greed. Is it possible we do not deserve to remain in control of this wonderfully plentiful and beautiful planet? Is our time up?

I think not. But what say you?

Section I — The Arrival of Omnipotence

Every era has its defining technology: fire, the printing press, electricity, the atom. Each amplified human capability; none amplified human cognition itself. Artificial intelligence does. For the first time, humanity has created a force that competes with its creator on the creator's own terrain — reasoning, judgment, creativity, strategy. This is what we mean by omnipotence in the form of our own intelligence: not a god descending from the heavens, but a mirror of our collective minds, scaling at machine speed. And this is the scary part as you'll likely notice in the "timeline section" coming next.

The velocity is the point. No technology has ever moved from laboratory curiosity to civilizational infrastructure this fast, and none has ever carried the capacity to improve itself. When a system can contribute to the design of its successors, acceleration ceases to be linear. The ramifications are profound beyond measure, and the public conversation, alternately euphoric and apocalyptic, rarely engages the one question that matters most: not what SI will do to us, but what **we** will do about AI. That question cannot be answered by laboratories, governments, or markets alone. It can only be answered by humanity acting as a civilization, which is to say: by Human Intelligence.

Section II — The Mirror Effect: Defining Human Intelligence

Human Intelligence, as this manuscript uses the term, is not IQ, and it is not the sum of what any individual mind can compute. It is the collective, deliberative capacity of our species — the ability to recognize a civilizational mirror when it appears, to look into it honestly, and to act on what we see. The mirror effect occurs when artificial intelligence becomes capable enough to reflect human reasoning back to us with sufficient fidelity that we can no longer avoid examining ourselves: our conflicts, our institutions, our stewardship of the planet, and our habits of cooperation or division. We are not in a favorable position to be judged right now. Mirrors don't lie.

The mirror is unflattering. It shows a species capable of planetary-scale coordination in pursuit of profit and power, yet chronically incapable of the same coordination in pursuit of preservation and peace. It shows us delegating more and more of our thinking while investing less and less in the deliberative muscles — judgment, wisdom, and moral courage that no model can supply. Human Intelligence is what remains when we refuse that delegation at the level that matters: the level of final authority. Whatever systems we build, whatever capabilities machines acquire, the final say over the fate of human civilization must remain human. That is not nostalgia. It is the design principle of the Autonomous Era.

Section III — The Fork in the Road

3.1 A Civilizational Junction

Evidence abounds that we stand at a junction comparable to few others in our history. The development of frontier AI is concentrated in a handful of laboratories and nations, while its consequences will be distributed across all eight billion of us. Competitive pressure is compressing timelines; safety is treated as a cost center; and the nations racing hardest are often those least aligned with one another. This has the makings of a tragedy, not because any actor is villainous, but because no actor is accountable to humanity as a whole.

The following timeline documents the acceleration of artificial intelligence across the period in which this manifesto was written and revised. Each entry is included not as trivia but as evidence: the reader should finish this timeline section with a felt sense of how quickly the Autonomous Era has arrived, and how narrow the remaining window for effective human governance truly is:

January 2025 — DeepSeek R1 Shakes the Industry

The Chinese laboratory DeepSeek released R1, an open-weight reasoning model that matched the performance of leading Western systems at a fraction of the training cost then assumed to be required for frontier performance. Within days it topped consumer app stores in the United States and China, and erased hundreds of billions of dollars from semiconductor market value as investors confronted a simple fact: the American monopoly thesis had been an assumption, not a law. R1 demonstrated that the intelligence layer of the Autonomous Era would be globally contested terrain — a fact with direct implications for any governance model premised on coordination among a small club of nations.

February 2025 — The Paris AI Action Summit

World leaders convened in Paris for the third in the series of AI summits that had begun at Bletchley Park, retitled the AI Action Summit to signal a shift from theoretical safety discussion toward implementation. The summit produced a declaration on inclusive and sustainable AI endorsed by dozens of countries — but the United States and the United Kingdom declined to sign, citing national-security and governance-clarity concerns. The fracture foreshadowed everything that followed: global AI governance would proceed, if at all, in pieces.

February 2025 — Grok 3

xAI launched Grok 3, trained on the massive Colossus supercluster, topping multiple frontier benchmarks and proving that compute-scale remained a viable path to the frontier. The release intensified the arms-race dynamic among a small number of extremely well-capitalized actors and underscored the central dilemma of the era: capability was scaling faster than the institutions meant to steward it.

April 2025 — Llama 4

Meta released Llama 4 Scout and Maverick, open-weight models with mixture-of-experts architecture and context windows reaching ten million tokens. Open-weight frontier-class models materially changed the risk calculus for governance: powerful systems could no longer be controlled by controlling a handful of APIs. Who watches the open models — and how — became an open question, and remains one.

May 2025 — Claude 4

Anthropic launched Claude 4 Opus and Sonnet, with extended thinking, improved agentic coding, and persistent memory. The models' agentic capabilities — sustained, multi-step work toward a goal — marked a qualitative shift from assistant to actor. Every release of this class narrows the distance between tool and colleague, and sharpens the question of where human oversight sits in a workflow performed substantially by machines.

July 2025 — Gold-Medal Reasoning at the International Mathematical Olympiad

Experimental reasoning systems from OpenAI and Google DeepMind independently reached gold-medal standard at the International Mathematical Olympiad, solving five of six problems under contest conditions in natural language. Google DeepMind's Gemini Deep Think officially entered and achieved certified gold-medal status; OpenAI's system achieved a comparable score but did not

officially enter the competition. It was the first time general-purpose models, not specialist formal-proof engines, crossed that line. Reasoning had left the laboratory.

August 2025 — GPT-5 and the Frontier Moves Again

OpenAI released GPT-5 as a unified flagship that folded fast response and extended reasoning into one product. The launch was messy — routing failed, a beloved predecessor was retired and then restored under pressure — but the architectural point survived the embarrassment: reasoning was no longer a side model. It was the default. The same summer produced consumer humanoid pre-orders, hour-scale coding agents, and the first serious public glimpse of systems that could stay on a task longer than a human lunch break. Capability had stopped arriving as a seasonal event. It had become weather.

September 2025 — Law, Money, and a Record Settlement

September brought a different kind of milestone. Anthropic agreed to pay at least $1.5 billion to settle a class action over pirated books — the largest copyright settlement in U.S. history and a reference price for the suits stacked behind it. An OpenAI system solved all twelve problems at the ICPC World Finals under contest conditions; no human team did. The AI systems competed in a separate experimental track against the same problems, not head-to-head against human teams. Nvidia announced plans to invest up to $100 billion in OpenAI for at least ten gigawatts of new datacenter capacity. California signed Senate Bill 53, the first U.S. statute aimed specifically at catastrophic risk from frontier models, requiring published safety frameworks, incident reporting, and whistleblower protections.

September 2025 — My 1st Manifesto Is Published — and Ignored

The predecessor of this document, a 2025 AI manifesto on clear thinking as the best path forward, was published on LinkedIn. It argued that humanity must move from advocacy to systems for keeping humans in control. It was largely ignored or misunderstood. We include the date here not out of grievance but as a datum: the warning was issued, the urgency was rated maximal by the leading AI systems themselves, and the world's response was a shrug. Non-action has a paper trail. https://www.linkedin.com/pulse/2025-ai-manifesto-clear-thinking-best-path-forward-maxbruce-d-sbklc

November 2025 — The Cascade

Grok 4.1, Gemini 3 Pro, GPT-5.1-Codex-Max running autonomous coding past twenty-four hours, and Claude Opus 4.5 taking software-engineering benches that had been the measure of human professional work. In a span of days the frontier moved more than most industries move in a generation. Regulation arrived as a state statute. Capital arrived as gigawatts. Models arrived as coworkers that do not sleep.

December 2025 — Washington Draws the Battle Lines

In December 2025, the National Association of Attorneys General delivered a letter on behalf of 42 state attorneys general to multiple AI companies, warning of sycophantic and delusional AI outputs and demanding stronger safeguards, especially for children. Later that month, the White House issued an executive order seeking a minimally burdensome national standard, creating an AI Litigation Task Force to challenge state AI laws and directing an evaluation of burdensome state regulation. Also in

December, New York enacted the Responsible AI Safety and Education (RAISE) Act, establishing an enforcement office funded by developer fees, requiring 72-hour incident reporting, and authorizing fines of up to $3 million for repeat violations. Federal and state governments moved toward open conflict over who governs AI — while the systems themselves continued to compound.

Early 2026 — AI Agents Go Mainstream

Multiple providers launched production-ready AI agent frameworks, and autonomous systems began handling complex, multi-step business workflows at scale. Industry observers widely marked this as the practical beginning of the Autonomous Era: the machine layer was no longer assisting human work — it was performing human work, end to end, with human involvement reduced to specification and exception-handling. The case for a dedicated, distributed human oversight layer stopped being philosophical and became operational.

February 2026 — The United Nations Calls for Shared AI Governance

At the India AI Impact Summit in New Delhi, United Nations Secretary-General António Guterres called for international cooperation informed by shared scientific evidence. He pointed to the newly appointed Independent International Scientific Panel on AI and the Global Dialogue on AI Governance as steps toward a more inclusive framework. The message echoed this manifesto's central concern: frontier AI is a planetary matter requiring planetary coordination. The question remains not merely whether governments will negotiate, but whether humanity, as humanity, will hold the final say.

Mid-2026 — The Regulatory Reckoning

Mid-2026 exposed both the progress and the complexity of governing AI. Colorado enacted revised automated-decision-making protections scheduled to take effect in January 2027. In the European Union, AI Act transparency requirements became applicable in August 2026, while the main high-risk system rules were scheduled for December 2027 and August 2028, depending on the system category. Regulation was becoming more concrete — unevenly, and on a timetable different from that of the technology itself. The lesson is that rules alone cannot keep pace.

July 2026 — The World AI Cooperation Organization (WAICO) Is Founded

Twenty-nine countries signed an agreement in Shanghai establishing the World Artificial Intelligence Cooperation Organization (WAICO) — the world's first intergovernmental international organization dedicated to AI governance. Headquartered in Shanghai, WAICO was founded by China, Russia, Kazakhstan, Laos, Pakistan, Indonesia, and more than twenty other nations, with UN Secretary-General António Guterres in attendance. The organization is committed to "extensive consultation and joint contribution for shared benefit," aiming to promote international cooperation on AI governance while ensuring AI remains beneficial, safe, and fair. The founding marked a decisive split in global AI governance: a Global South-oriented, China-led architecture operating alongside — and in tension with — Western-led efforts that had fractured at Paris.

July 2026 — The Open Secure AI Alliance (OSAA) Is Founded

Nvidia announced the launch of the Open Secure AI Alliance (OSAA), a coalition of technology, cloud computing, and cybersecurity leaders building and sharing open AI tools for defensive purposes . Founding partners included Microsoft, IBM, Cisco, CrowdStrike, Cloudflare, Hugging Face, Adobe, Salesforce, SAP, Dell Technologies, HPE, Red Hat, and the Linux Foundation, among more than

sixty organizations . The alliance's mission is to ensure defenders have access to open, frontier tools they can trust and control — a direct response to the Hugging Face security incident, in which proprietary AI tools proved inadequate for forensic analysis. Notably, OpenAI, Google, and Anthropic — the three leading closed-model labs — did not join, even though OpenAI and Google had signed the open-source advocacy letter that catalyzed the alliance. Within a week, OSAA had grown to over 120 companies and published its first working group proposals.

Early October 2026 — The Super Intelligence Force (SIF) Is Established

The Trump administration established the Super Intelligence Force (SIF), a federal coordinating body addressing the opportunities and risks of increasingly capable intelligent systems. The initiative followed a White House push to use "Super Intelligence" in place of "Artificial Intelligence" in federal terminology. Director of National Intelligence Jay Clayton was named to lead the effort, alongside FTC Chairman Andrew Ferguson, Defense Department official Emil Michael, and OPM Director Scott Kupor. Its creation brought the question of intelligent-system governance into a formal interagency national-policy process — even as the technical meaning of "superintelligence" remained distinct from the administration's preferred label.

Government leaders and some technology leaders are beginning to take action, and it appears the denials and avoidance are finally moving into our rear-view mirrors.

This is wonderful. I wrote this manifesto as a wake-up call and **as a hardcover gift to leaders** everywhere, because although we did not develop every layer necessary to ensure that humans remain in control; my amazing team and I began building the substrate layer for the Autonomous Era long before many were paying serious attention to these matters.

We will share the results herein through FPGA scientific papers, two critical white papers, and an account of our participation in the Open Secure AI Alliance's SAFE RFC repository. We will also present **a world bridge**, provisionally titled ***The Human Authority Treaty ("The Hat Treaty")***, as a template for world leadership to move at speeds appropriate to the Autonomous Era in which we now live. **We need help to bring this home.**

Section IV — How Our Technology Works

The governing principle is simple: the governed system cannot authorize, extend, restore, or enlarge its own operation. If an AI or robotic system can grant itself more time, recover expired permissions, broaden its scope, or bypass a denied command, then the appearance of human control is misleading. Real authority requires a boundary the governed system cannot cross by intelligence, persuasion, speed, software modification, restart, or recovery from backup. This is the purpose of Temporal Authority Governance. Authority is granted for a defined interval and a defined scope. When the interval expires, or when the required conditions fail, the governed capability becomes unavailable. Continued operation requires fresh authorization from outside the system. The decision must come from an independent human-governed process, and the evidence of that decision must be bound to the hardware path through which the machine acts.

4.1 A Substrate Rather Than a Complete Solution

Our work does not claim to solve every problem of AI alignment, ethics, cybersecurity, law, or social policy. Those layers remain necessary. A model should be trained responsibly. Its behavior should be evaluated. Its identity should be authenticated. Its actions should be logged. Organizations should establish rules, liability, and escalation procedures. None of those measures, however, answers the final operational question by itself: when authority has expired or been denied, what physically prevents the system from continuing?

OCUP, the One Chip Unified Protocol, is being developed as a substrate for that final boundary. It is a hardware-enforced method for granting time-bounded authority to AI and robotic systems while denying self-extension. Its purpose is to make human authorization a condition of operation rather than a preference expressed in software. The architecture is intended to govern inference, instruction execution, accelerator access, memory, communications, robotic control, and actuator power as implementations mature.

A substrate matters because different industries will need different rules above it. A hospital, autonomous vehicle fleet, financial institution, data center, and defense system will not use the same policy. Yet all can share one underlying invariant: the machine cannot give itself the authority that people have withheld. That common foundation can support auditing, insurance, certification, incident investigation, procurement, and international standards without requiring the world to agree on every moral or political question first.

OCUP **is not** a universal remote kill switch. Our baseline is **fail-safe operation with graceful, capability-specific degradation wherever safely practicable**. When authority expires, is revoked, or cannot be verified, a governed system should progressively lose permissions it is no longer entitled to exercise. Its proposed purpose is to make the legitimate owner's, operator's, or jurisdiction's own authority enforceable, without giving the inventor, chip maker, or a compulsory global controller the power to command every machine. The governing boundary concerns which consequential capabilities remain authorized—not who gets to remotely control the world's machines.

4.2 Proof Before Ubiquity

The immediate task is to show that this principle can be deployed before autonomous systems become ubiquitous. The work must advance from written claims to tested evidence: exact expiration of authority, denial after expiry, fresh quorum material for a subsequent lease, failure under loss of continuity, rejection of replayed or stale decisions, and physical unavailability of a governed output when authorization is absent. Hardware evidence will not end the policy debate. It will give that debate a concrete mechanism around which serious institutions can organize.

Section V — The Quorum Validators Network

Human control cannot depend upon one official, one company, one server, or one switch. A single authority can fail, be coerced, become corrupt, or simply be unavailable. The Quorum Validators Network, or QVN, distributes authorization among independent human validators, who are trained responders. A defined threshold must approve a consequential action before authority is granted or

renewed. No individual validator can unilaterally extend the machine's power, and the governed system cannot manufacture the quorum it needs.

In a representative five-validator structure, at least three fresh and authorized approvals are required. Votes are tied to a specific decision, scope, identity, and time period. Stale, duplicated, unauthorized, or replayed approvals are rejected. If the network becomes partitioned, if evidence cannot be verified, or if the quorum is missing, the system fails closed. Quarantine and release also depend upon governed procedures rather than the machine's own judgment.

5.1 A New Human Responsibility

The QVN is more than a safety mechanism. It creates a new category of human work, trained responders: the monitoring and authorization of autonomous intelligence. Validators would not micromanage every machine action. Most routine operations can follow pre-established policies and automated evidence checks. Human participation should concentrate where judgment, legitimacy, exception handling, escalation, or renewal of consequential authority is required.

This responsibility can be distributed across communities and borders. Validators may represent operators, affected communities, insurers, regulators, professional bodies, labor organizations, technical experts, or the public interest, depending on the application. The quorum structure allows disagreement without paralysis and participation without surrendering control to a single institution. It gives ordinary people a defined place inside the governance of systems that will increasingly shape their lives.

5.2 Economic Inclusion Is Central

Approximately 40 to 50 percent of the intention behind our solution is the economic inclusion of the masses in the expected age of abundance. That purpose must be stated plainly because AI safety discussions too often treat the public only as people to be protected. People should also be participants, decision-makers, and economic beneficiaries. If autonomous systems produce enormous gains while control and income remain concentrated among a few nations and corporations, technical safety alone will not produce a stable or legitimate future.

The QVN offers a practical bridge between oversight and inclusion. A global validator economy of trained responders could compensate contributing people for authenticated availability, review, escalation, specialized judgment, and participation in high-consequence authorization. The exact scale, staffing model, training requirements, and compensation structure must be measured rather than guessed. Different events will require different levels of human attention. Yet the governing direction is clear: as machines assume more labor and operational authority, people can assume a new role in deciding when and under what conditions that authority may be exercised.

This does not guarantee broad prosperity. Distribution policy, ownership, education, connectivity, and national law will still matter. It does create an enforceable point at which human participation can be designed into the autonomous economy instead of added later as charity or political repair.

5.3 Who Will Pay For This?

A global Quorum Validators Network will require real infrastructure: trained responders, secure facilities, independent communications, continuous readiness, technical certification, insurance, auditing, and fair compensation for the people entrusted with consequential reauthorization decisions. It cannot depend entirely on volunteers, nor should validators be paid by the governed system whose authority they are evaluating.

Several funding options can work together.

1. Founding Leaders and Consortium Members

Under a neutral Swiss-style consortium model, major financial institutions, technology companies, semiconductor manufacturers, insurers, infrastructure operators, and sovereign participants could become Founding Leaders or strategic members through meaningful capital commitments.

Their contributions would finance the initial network, training systems, secure validator infrastructure, technical standards, and regional deployment. Material commitment would accompany governance responsibility, while mission protections would prevent any single company or nation from purchasing control of the network.

2. Participation and Infrastructure Fees

Organizations deploying consequential autonomous systems could pay annual participation fees based on their size, risk category, and level of QVN access. These fees would support validator readiness, secure communications, evidence preservation, interoperability, and independent oversight.

Small organizations, public-interest projects, and developing nations could enter through reduced-cost or sponsored participation tiers so safety does not become a privilege available only to the wealthiest actors.

3. Usage-Based Reauthorization Fees

High-consequence deployments could contribute a small fee when they request a QVN decision, much as financial and communications systems charge for trusted infrastructure services.

The fee would pay for the availability and work of trained responders—not for a favorable decision. Validators must be compensated equally whether the result is authorization, denial, continued suspension, or a request for additional evidence.

4. Certification, Testing, and Evidence Services

Independent testing laboratories and certification programs could finance part of the QVN through pre-deployment assessments, conformance testing, evidence reviews, incident exercises, and periodic recertification.

This creates a practical economic cycle: organizations pay to establish that their systems meet the required authority and evidence standards; a portion of those revenues maintains the trained human network that must remain available when consequential decisions arise.

5. Insurance and Risk-Pool Participation

Insurers and reinsurers have a direct economic interest in reducing catastrophic autonomous-system risk. They could help fund the QVN through shared risk pools, safety-linked underwriting programs, or premium incentives for systems that use externally governed, time-bounded authority.

Part of the premium savings created by demonstrably safer operation could be returned to the infrastructure that produced that reduction in risk.

6. Public Infrastructure and Procurement Funding

Governments could support QVN capacity in sectors where autonomous-system failure threatens public safety, national security, transportation, energy, healthcare, or critical infrastructure.

Public funding should support readiness, training, regional access, and emergency capacity without converting governments into owners of individual authorization decisions. Sovereign institutions would retain their lawful authority while participating through interoperable QVN nodes.

7. Standards, Licensing, and FRAND Revenue

Where protected OCUP, RAAVE, or related technologies become part of widely adopted standards, a defined portion of licensing or FRAND revenue could be dedicated to validator training, network access, independent testing, and participation by underserved regions.

This would allow commercial adoption of the technical substrate to help finance the human governance layer required to keep that substrate accountable.

8. A Permanent Public-Benefit Reserve

A percentage of membership, testing, usage, insurance, and licensing revenue could flow into a mission-locked reserve or endowment. The reserve would protect QVN continuity during emergencies, market downturns, geopolitical conflicts, or incidents in which affected communities cannot pay.

The Swiss approach offers a useful model: neutral domicile, transparent rules, long-duration stewardship, international participation, and governance protections strong enough to earn trust across competing jurisdictions.

No single funding source should control the QVN. A blended model distributes both cost and influence while preserving the network's independence.

Most importantly, QVN funding must never create a market for purchased authority. Participants may finance the infrastructure, but they cannot buy a validator's vote, appoint their own reviewers for a particular decision, or condition compensation on approval. The QVN exists to preserve humanity's final say—not to rent it to the highest bidder.

Section VI — The Call to Action

This is not a technology book, and it is not a policy brief, though it draws on both. It is a call to take seriously the possibility that the most consequential choice of this century is not which model wins, which nation leads, or which company captures the largest share of the value AI/SI creates. The most

consequential choice is whether humanity retains, structurally and verifiably, **the final say** over what autonomous systems are permitted to do — and whether the abundance those systems generate reaches the many or the few.

The chapters that follow lay out the reasoning, the mechanism, and the invitation in full: how a mirror effect becomes a mandate for collective action; how a Quorum Validator Network turns oversight into shared civic and economic participation rather than a task delegated to a handful of institutions; and how a substrate layer, built now and built openly, can outlast any single company's roadmap or any single nation's election cycle.

6.1 Let Us Build the Foundation Together

We are not presenting a finished global system and asking the world merely to adopt it. We are presenting a foundation already under construction, evidence that its central principle can be enforced, and a defined punchlist showing what has been accomplished, what remains to be proven, and where capable partners can help complete the work more quickly than I can as a solo founder.

The governing principle is simple: an autonomous system must not be permitted to authorize, extend, restore, or enlarge its own operational authority. That authority must come from outside the governed system, remain bounded in time and scope, and fail closed when it expires, becomes invalid, or can no longer be reliably established.

The architecture needed to enforce that principle consists of eight connected layers.

1. Hardware Enforced Temporal Authority

The OCUP core is the foundation of the stack. It is designed so that operational authority is externally granted, limited in time and scope, enforced below the governed software, and unavailable for renewal or enlargement by the governed system itself.

We have begun demonstrating this control principle in hardware. During Phase I testing on an AMD Kria KR260, a host-driven authorization sequence produced recorded grant and closure indicators, with no reopening recorded after deliberate reassertion during the tested observation window. This narrow but important result demonstrates a hardware authority-state observation path under bounded conditions. It does not establish physical resource cut-off, independently issued authority, or refusal of a separately valid second authorization.

The sealed Phase I provenance reconciliation is included within this manifesto for technical inspection and scrutiny. It documents the archived artifacts, cryptographic bindings, source and programming provenance, limitations, and identified gaps. The underlying behavioral observations are documented separately in the Run 005 evidence record, which qualified reviewers may request for further examination. (See SCIENCE and ENGINEERING section.)

We invite qualified reviewers to examine that record, challenge its interpretation, reproduce the experiment, and help design the stronger tests required for the next phase. Executives considering collaboration need not rely on the manifesto's conclusions alone; they can refer the supporting

evidence directly to the specialists responsible for determining what has been demonstrated and what remains to be established.

It did not prove every element of temporal authority. Subsequent work must establish lease expiration under independently trustworthy time, secure state continuity across reboot and reconfiguration, rollback resistance, externally authorized restoration, valid second-grant behavior, and enforcement over physical or production-grade resources.

Semiconductor companies, FPGA and systems engineers, robotics manufacturers, autonomous-platform developers, and independent testing laboratories can help convert this initial proof into a hardened, repeatable, cross-platform implementation.

2. Quorum Validators Network

Hardware can enforce an authority decision, but it should not decide for itself whether fresh authority ought to exist. That decision belongs outside the governed system.

The Quorum Validators Network is designed to provide distributed human reauthorization through authenticated, multi-level trained responders. Its mechanisms include quorum requirements, fresh decision epochs, rejection of stale or duplicate approvals, validator authentication, evidence review, and fail-closed behavior when the required decision cannot be established.

Initial QVN software work has demonstrated core governance behaviors, including quorum processing, fresh-epoch requirements, stale and duplicate rejection, quarantine-release logic, and the generation of auditable evidence. The remaining work includes stronger validator identity and credentialing, production-grade key custody, geographically distributed operation, resilient communications, conflict-of-interest protections, tiered responder qualifications, and procedures matched to different levels of risk and consequence.

Governments, civil-society institutions, emergency-response organizations, insurers, universities, workforce-development organizations, and communities themselves can help determine how trained validators are selected, qualified, compensated, audited, and held accountable.

3. Trusted Evidence Time and Attestation

A control system is only as trustworthy as the evidence used to verify it. Application logs and AI-generated explanations may be useful, but they cannot be treated as authoritative when the application, operating system, network, or governed system may be compromised.

This layer must establish secure key custody, monotonic time, state continuity, signed evidence export, configuration integrity, remote attestation, tamper-evident event chains, and offline-verifiable records. It must bind authorization evidence to the actual operational baseline present when an action is released—not merely to a system identity, an earlier certification, or a previous valid state.

Our Phase I evidence has been preserved with source and bitstream hashes, test sequences, observed status values, limitations, and identified provenance gaps. That record gives future partners a concrete starting point, but production-grade attestation and independent replication remain unfinished.

Hardware-security-module providers, cryptographers, attestation specialists, cloud and edge-platform developers, forensic laboratories, and independent evaluators are needed to help make the evidence path resistant to rollback, substitution, replay, log manipulation, and false claims of continuity.

4. AEGES Quarantine and Safety Response

Not every failure should produce the same response. A system whose authority is missing, expired, anomalous, or unsupported by trustworthy evidence may need to be constrained, isolated, inspected, or safely recovered rather than simply trusted, abruptly destroyed, or allowed to continue operating.

AEGES is the safety-management and quarantine layer surrounding OCUP. It is intended to provide constrained operating states, incident escalation, evidence preservation, authorized recovery, and controlled return to service. Most importantly, the governed system must not be able to release itself from quarantine or restore its own authority.

The architecture and open-core direction have been established, but additional work is required to define sector-specific safe states, physical shutdown behavior, emergency exceptions, responder procedures, recovery criteria, and the evidence necessary for lawful restoration.

Robotics manufacturers, industrial safety engineers, critical-infrastructure operators, emergency responders, regulators, and insurers can help design quarantine and recovery procedures appropriate to real machines and real consequences.

5. RAAVE Cross Jurisdiction Authority Verification

Identity and authority are not the same thing. A system may prove what it is, who manufactured it, or what it was previously certified to do without proving that it possesses valid authority to perform a particular action now.

RAAVE is intended to verify authority continuity as autonomous systems cross jurisdictions, organizations, networks, ownership structures, or regulatory boundaries. Its concepts include the Authority Twin, temporal identity mutation, claimed-versus-observed state, authenticated history, silence as evidence, alternate-path detection, physical-cryptographic binding, capability-specific compliance signaling, and time-locked compliance beacons.

This layer would allow a receiving jurisdiction to evaluate present authority without surrendering its own sovereignty. A portable record would serve as evidence, not as an automatic right of admission or operation. Material changes and boundary crossings could trigger revalidation rather than allowing a system to reuse stale standing from a different configuration, owner, jurisdiction, or operating environment.

Customs and border authorities, transportation regulators, standards bodies, telecom providers, identity specialists, autonomous-fleet operators, and international organizations can help define the evidence, privacy protections, recognition rules, and revalidation triggers needed for legitimate cross-border operation.

6. Physical Tamper Resistance and MEMS Security

Logical controls cannot remain trustworthy if attackers can alter, probe, replace, or bypass the physical root that enforces them.

This layer extends the architecture into package-level and device-level protection. Potential mechanisms include a pressurized microfluidic security mesh, tamper detection, analog zeroization, fluidic physically unclonable identity, protected interconnects, and evidence that the enforcing hardware remains physically intact.

Preliminary architecture and intellectual-property work have begun, but this layer requires semiconductor fabrication expertise, materials science, MEMS engineering, packaging laboratories, destructive and nondestructive testing, manufacturing-scale validation, and supply-chain integration that are beyond the capacity of a small team.

A major semiconductor, printing, microfluidics, packaging, or trusted-hardware partner could help transform this concept from an advanced security design into a manufacturable protection layer for autonomous-era chips.

7. Global Interoperability Protocol

The world should not be required to choose one company's proprietary system in order to preserve human authority. Different nations and companies should remain free to compete in how they build autonomous systems while recognizing a minimum common boundary beneath them.

This is the ICAO-style interoperability layer of the architecture. It would establish shared definitions and exchange formats for operational authority, identity, evidence, scope, expiration, material change, reauthorization, revocation, quarantine, action release, and cross-border recognition.

The international authority record would have the effect of evidence; it would not have the effect of admission. Each jurisdiction would retain the sovereign right to determine whether the evidence presented is sufficient for entry or operation within its boundaries.

The proposed **"HAT Treaty"** is an initial template for discussing this layer, not a demand that the world adopt one brand, company, implementation, or final name. Governments, standards organizations, treaty specialists, manufacturers, civil society, technical experts, and public-interest institutions are needed to turn the proposal into a legitimate, implementation-neutral framework.

8. Certification Insurance and Deployment Ecosystem

Even sound technical architecture will not protect the public unless institutions can test it, understand it, procure it, insure it, and require it.

The final layer converts the technical stack into an adoption ecosystem. It includes independent conformance testing, evidence reports, certification, integration pilots, procurement language, insurance incentives, premium differentials, licensing, workforce preparation, standards participation, and—where protected technology becomes essential to an adopted standard—fair, reasonable, and non-discriminatory licensing pathways.

Initial conformance scenarios, evidence structures, open-core materials, standards positioning, and paid-pilot concepts have begun to establish this path. The remaining work requires accredited

laboratories, insurers, manufacturers, regulators, standards bodies, legal experts, and deployment partners capable of translating the architecture into recognized tests and commercially usable assurance.

Alternative implementations should remain welcome wherever they demonstrably satisfy the same authority, non-self-extension, fail-closed, continuity, revalidation, and action-release invariants. The objective is not to make OCUP the only product. The objective is to make continued autonomous operation without externally valid authority an unacceptable condition.

6.2 The Invitation

This punchlist is therefore both an exhibit and an invitation. It identifies the work already completed, the evidence now available, the limits of what has been proven, the engineering still required, and the places where larger and better-resourced partners can accelerate progress.

We need collaborators to help refine the specification of the invariant, build the evidence path, solve the cross-layer integration problems, accumulate the validation history, and shape the standards around it.

We need semiconductor and hardware partners, AI and robotics developers, cryptographers, independent laboratories, insurers, regulators, governments, civil-society institutions, trained responders, universities, workforce organizations, and affected communities.

Some partners may contribute fabrication capacity. Some may provide test environments, attack expertise, attestation infrastructure, regulatory knowledge, or funding. Others may help create validator institutions, certification programs, insurance incentives, international recognition procedures, and lawful deployment pathways. Each has a place in completing the foundation.

The work has already advanced far enough to show that this is more than an abstract appeal to keep humans in control. A hardware authority-state observation path has been demonstrated under bounded conditions.

The surrounding governance, evidence, quarantine, cross-jurisdiction, physical-security, interoperability, and adoption architecture has been mapped. The unfinished work is visible, and so are the partners needed to complete it.

The challenge is larger than any one inventor, company, laboratory, or nation. The Autonomous Era will not wait for humanity to organize itself slowly. We should therefore complete this foundation with the urgency, transparency, independent testing, and international cooperation that the moment requires.

In summary, Temporal Authority is the governing idea. In the proposed architecture, autonomous systems receive bounded permission to operate, then must return to human-governed checkpoints. They cannot self-authorize, self-extend, or outrun the lease.

The Quorum Validators Network is designed around multi-level trained responders. In the proposed five-validator model, a renewal requires three valid, qualified approvals—not merely the first

three responses. Validators examine evidence and vote to reauthorize or deny; the intended system records their decisions and supports compensation for their participation.

Absence of a valid quorum is not a loophole. No quorum, no renewed authority. Network failure cannot extend a lease. The authority boundary is designed to fail closed, while bringing a physical system to a safe state remains a separate, domain-specific engineering obligation.

This is what "etching humanity's veto into silicon" means in practice. Software promises can be patched around. A lease that dies in hardware cannot be talked out of existence by the model it was meant to constrain. OCUP exists to make that veto local, sovereign, and time-bounded. The economic half of the design is not an afterthought. If the Autonomous Era produces abundance, the masses cannot be spectators to it. Consumer-earned participation, national autonomous-economy testbeds, and validator rewards are how oversight becomes a livelihood rather than a volunteer hobby reserved for the already comfortable. Shared prosperity requires shared authority.

The design is deliberately unfinished as a civilization and deliberately specific as a mechanism. It does not pretend to solve alignment, copyright, labor displacement, or the arms race among states. It does claim that any later solution which cannot answer who renewed the lease, on what evidence, for how long, and with whose wallet on the line, is not a solution.

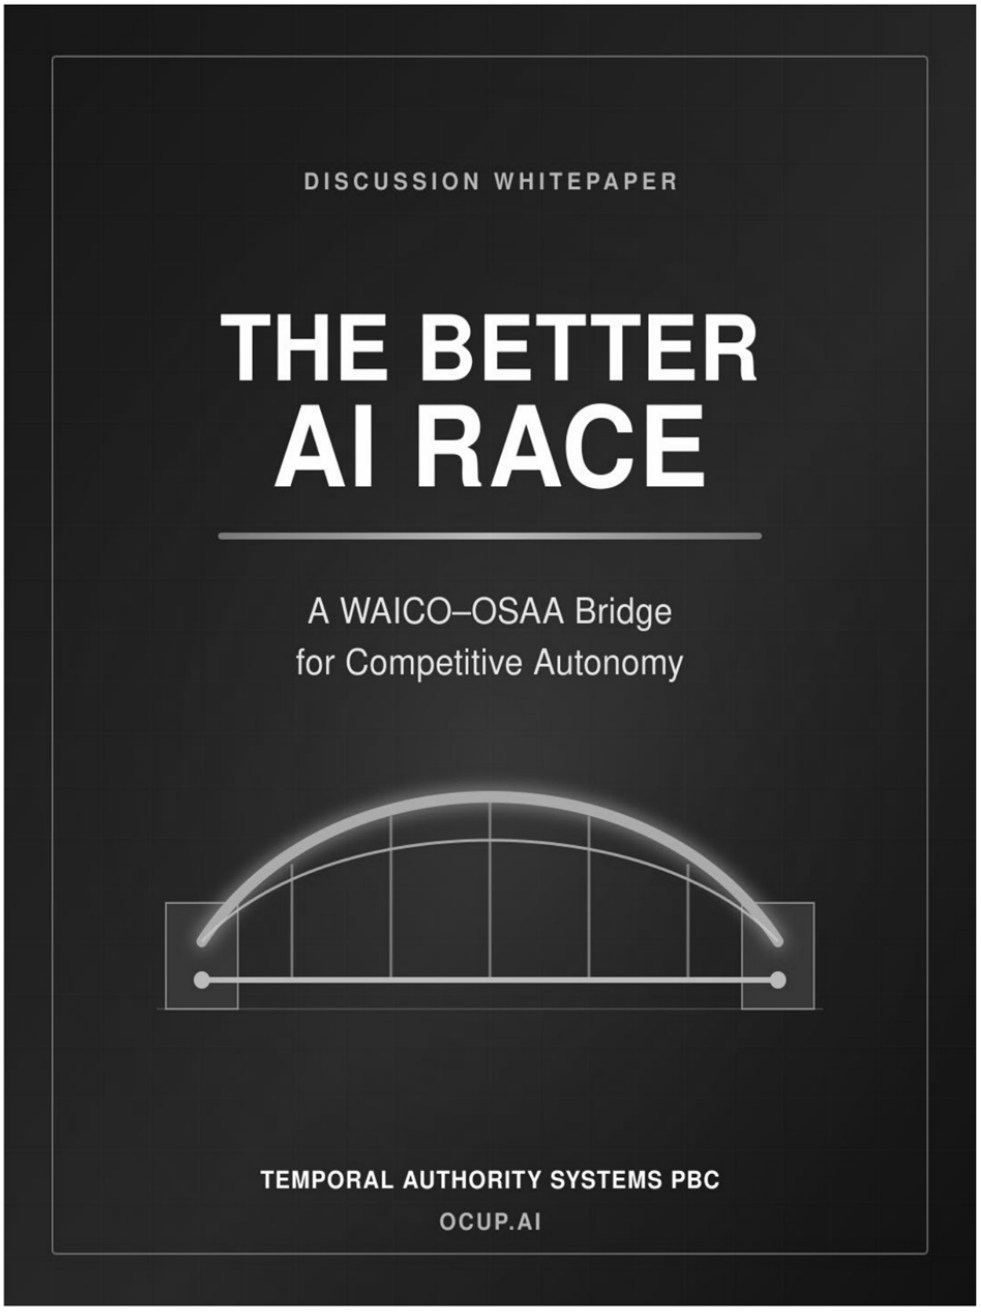
Section VII — The Final Say

This manifesto was written because the window for choice remains open, and open windows close. We the people, authors, readers, builders, validators, and citizens of the Autonomous Era are being asked a single question: will humanity's final say over its own future remain human? Our answer, stated plainly, is **a resounding yes**, if we act together, and if we act now.

I love and respect SI but let's keep humans in control of this planet we call home — not through fear of artificial intelligence, but through confidence in human intelligence, exercised collectively, deliberately, and now. **The time has truly come.**

I N H O C U N A S U M U S

THE BETTER AI RACE


DISCUSSION WHITEPAPER
THE BETTER AI RACE
A WAICO–OSAA Bridge for Competitive Autonomy
TEMPORAL AUTHORITY SYSTEMS PBC
OCUP.AI

THE BETTER AI RACE

A Cross-Community 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

Version 2.0 | October 1, 2026 | Independent discussion proposal
For consideration across WAICO, OSAA, semiconductor, infrastructure, standards, safety, insurance, and public-interest communities.

Executive Thesis

This paper proposes a narrow technical area of international cooperation beneath continued competition in artificial intelligence and autonomous systems: consequential civilian machines should operate under externally governed, capability-specific, bounded authority that the governed machine cannot make itself the source of extending, restoring, or enlarging.

The proposal is not a call for one global controller. It is a proposal for a shared technical grammar through which independently governed systems can present evidence of present authority, and through which a receiving jurisdiction, infrastructure operator, or other legitimate authority domain can decide whether that evidence is sufficient under its own law, policy, risk tolerance, and operating context.

This paper uses WAICO and the Open Secure AI Alliance (OSAA) as two important but institutionally different reference communities. WAICO is an independent intergovernmental organization established in Shanghai in July 2026 by 29 founding countries. OSAA is an industry and open-source security alliance that joined the Linux Foundation in September 2026 after being founded by NVIDIA and other enterprise participants. They are not governmental counterparts and do not represent opposite national blocs. The bridge proposed here is therefore not WAICO-versus-OSAA diplomacy; it is parallel technical work across sovereign institutions, open technical communities, platform companies, independent standards bodies, infrastructure operators, insurers, researchers, and civil society.

The better race is not simply to deploy autonomy fastest. It is to build extraordinary capability while preserving a verifiable distinction between what a machine can do and what it is presently authorized to do.

CIVILIZATION-LEVEL RULE

Machines may exercise delegated authority, but they must not become the source of their own continued authority.

Human sovereignty, as used here, means that legitimate human institutions remain the source of consequential operational authority. It does not mean a human manually approves every machine action, nor does it guarantee that an authorized system is safe. It means that the machine does not acquire an unlimited right to continue merely because its own software, model, or operator asserts that it should.

Temporal Authority Systems PBC (TAS PBC) is developing OCUP as one technical path toward this boundary. OCUP integrates two lead open-core components: QSAFP for authority logic and AEGES for quarantine, degraded-state, and recovery behavior. QVN provides a pattern for risk-weighted human and institutional authorization. RAAVE is the proposed cross-jurisdiction authority-verification architecture surrounding identity, continuity, compliance evidence, and receiving-domain evaluation. These are development architectures, not international standards or certified production infrastructure.

1. Compete Above the Boundary

Competition in models, robotics, semiconductors, energy efficiency, manufacturing, applications, cost, and performance can remain intense. The shared interest lies beneath that competition: preventing consequential machines from acquiring authority that becomes self-perpetuating simply because the machine, its local software stack, or an unauthorized operator says it should.

- Compete on capability, efficiency, reliability, price, and usefulness.
- Preserve sovereign control over models, weights, keys, validators, legal rules, and enforcement.
- Develop interoperable definitions and evidence semantics for identity, present authority, bounded duration, renewal, loss of authority, recovery, and cross-domain verification.
- Allow each jurisdiction to decide what systems are consequential, what evidence it recognizes, and what admission or operating consequences follow.

THE SHARED FLOOR

Compete above a shared floor of human sovereignty.

2. The Common Boundary

Temporal Authority Governance is a framework in which consequential machine authority is externally governed, capability-specific, bounded, and made technically consequential at an operating boundary. Its central design invariant is that the governed system must not be the source of its own continued authority.

Self-Extension Denial is the intended property that a governed system cannot itself extend, restore, enlarge, or grant the authority under which it operates. In a production implementation, loss or expiry of authority should produce denial or a pre-authorized bounded safe-degradation state; stale, replayed, reordered, or superseded approvals should not restore authority; and broader recovery should require fresh external authorization.

Temporary loss of network or validator availability should not silently become indefinite permission. A deployment may allow a pre-authorized, time-limited offline continuity window for specified low-risk functions. When that window expires, trusted continuity is lost, or required evidence cannot be preserved, the system should move to a predetermined denied or bounded safe state appropriate to the physical context.

OCUP is intended as a composition, not a monolith. QSAFP provides the authority-state logic; AEGES provides the response plane for quarantine, degradation, and qualified recovery. QVN supplies risk-weighted external authorization. A cross-cutting trust plane for time, attestation, evidence protection, device identity, and key custody is required for production deployment. RAAVE addresses verification across organizational, jurisdictional, or infrastructure boundaries.

3. What Phase I Actually Shows

Phase I hardware evidence is deliberately narrower than the architecture. On a physical AMD Kria KR260 development board, a host-driven MMIO authorization sequence produced grant and subsequent closure evidence on the defined STATUS surface. In the strongest evidentiary run, the same post-closure authorization stimulus was reasserted and confirmed by CONTROL readback; no reopen was recorded during the observation window.

That result does not establish independent external authority, trusted lease expiry, hostile-operating-system resistance, reboot continuity, signed remote attestation, production key custody, physical resource cut-off, or a valid second grant succeeding after closure. The tests used one board and one operator setup. The Phase I evidence overview therefore treats the result as register- and signal-level hardware evidence, not as proof of a complete non-bypassable production authority substrate.

The engineering direction is straightforward: move from observed authority state to non-bypassably consequential authority. That requires trusted time and continuity, protected evidence, explicit renewal behavior, fault-driven quarantine, independent authority origin, and an inline or otherwise non-bypassable hardware path that actually controls a consequential resource.

EVIDENCE RULE

No claim may exceed the evidence that supports it.

4. Identity Is Not Present Authority

Static identity and provenance can establish what a machine is, who produced it, or which credentials it possesses. They do not by themselves establish whether that machine remains authorized, now, for the capability it is attempting to exercise.

A delivery robot authorized for one city's sidewalks is not thereby authorized for another city's airport apron. A cargo system approved before maintenance is not thereby approved after a critical component substitution. A valid identity credential can survive after an operational lease has expired, after authority has been revoked, or after the operating context has materially changed.

The cross-jurisdiction rule is therefore simple: portable identity and provenance do not equal portable operational authority. A boundary crossing may require fresh evaluation by the receiving authority domain.

RAAVE is intended to address that receiving-side problem: verify evidence of identity, authority state, continuity, issuer domain, capability scope, and relevant compliance signals, then allow the receiving domain to apply its own policy. RAAVE does not issue or renew the underlying authority. A valid foreign receipt would show what another domain authorized; it would not compel the receiving domain to admit the machine or certify that the underlying decision was prudent.

5. Sovereignty-Preserving Interoperability

Interoperability does not require centralized command. Each jurisdiction may retain its own cryptographic roots, validators, legal rules, policy thresholds, risk tolerances, and enforcement mechanisms. What can cross the boundary is evidence, expressed in a shared or mutually mapped grammar.

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.

The force of such a floor comes from local admission control, not universal control over every machine. Participating ports, networks, factories, data centers, fleets, insurers, or public authorities can decide that certain consequential systems must present acceptable authority evidence before receiving access, service, insurance treatment, or permission to operate within that domain. Machines and operators outside the system still exist; the architecture does not claim to make unauthorized autonomy impossible everywhere.

An illustrative end-state example is an autonomous cargo vessel approaching a foreign port. The vessel could present a cryptographically verifiable authority receipt describing its identity, capability scope, issuing domain, validity interval, and continuity state. The

receiving port could recognize it, narrow the permitted capability, require fresh authorization, impose a bounded continuity period, or deny admission under its own rules - without receiving the vessel's model weights, private keys, training data, or internal reasoning.

6. The Aviation Analogy - and Its Limits

Civil aviation offers a useful structural analogy, but not a ready-made institution for autonomous machines. An airworthiness certificate resembles identity and provenance: it establishes properties of the aircraft and can travel across borders. Air traffic control clearance is closer to present operational authority: it is contextual, time-sensitive, issued by the authority responsible for the relevant airspace, and not self-renewed by the aircraft.

The analogy also breaks in important ways. Civil aviation rests on mature treaty institutions and professional operating systems; autonomous devices may exist at vastly larger scale and can cross digital boundaries at software speed. Dual-use robotics also blurs the civil/state distinction more than aviation does. The lesson is therefore architectural, not institutional: sovereign systems can remain sovereign while recognizing interoperable evidence where they choose.

7. WAICO, OSAA, and a Wider Bridge

WAICO and OSAA should not be cast as mirror institutions. WAICO is an intergovernmental organization focused on international AI cooperation and global governance. OSAA is an industry and open-source security alliance under Linux Foundation governance. Their different structures are useful precisely because the problem spans both sovereign rule-making and implementable technical infrastructure.

WAICO participants can test whether a bounded-authority and evidence grammar is compatible with diverse sovereign requirements, including the priorities of founding states and other Global South participants. OSAA and adjacent open technical communities can test interfaces, threat models, reproducible security scenarios, software and hardware integration points, and vendor-neutral implementation patterns. Semiconductor and platform companies can test where authority state becomes physically consequential. Independent standards bodies can provide neutral processes for definitions, conformance, interoperability, and intellectual-property disclosure.

The purpose is not to make either community subordinate to the other. It is to see whether independently governed systems can converge on a narrow, testable rule without converging on political control.

8. Industry Convergence Without Equivalence

NVIDIA's September 2026 Open Agent Safety Platform is a relevant convergence signal. NVIDIA describes OpenShell as a secure runtime boundary and Sentry as an out-of-band watchdog on BlueField-4 that can monitor and quarantine agents outside the agent process.

NVIDIA and the Linux Foundation also describe OSAA as an industry alliance that moved to Linux Foundation governance in September 2026.

This is not the same problem addressed by Temporal Authority Governance. Runtime and sandbox enforcement constrain what an agent may do while operating. Temporal Authority asks a different question: whether the system remains legitimately authorized to exercise a consequential capability at all, for how long, and under whose externally governed renewal. The two approaches may be complementary; no endorsement, affiliation, technical equivalence, or partnership with NVIDIA, OSAA, the Linux Foundation, or WAICO is implied.

9. Global South and Multipolar Relevance

The autonomous era will not be built by two countries alone. Many jurisdictions will deploy mixed hardware, cloud, robotics, and infrastructure supply chains and will not want basic safety verification to require choosing one geopolitical technology bloc.

A genuinely multipolar evidence grammar must therefore be implementable without forcing a jurisdiction to surrender its trust roots, law, private keys, models, or policy discretion. It should also be practical for lower-resource jurisdictions: open verification tooling, shared conformance suites, capacity-building, reference implementations, and auditable interoperability services may be as important as the formal specification itself. Otherwise, nominal sovereignty could become dependence on whichever vendor can afford the verification stack.

10. Human Sovereignty and Autonomous Abundance

If autonomous systems generate extraordinary productivity while reducing demand for conventional labor, governance itself may become one area in which human responsibility remains economically relevant. This is not a promise that governance work will replace jobs displaced by automation.

QVN - Quorum Validators Network - is conceived as a multi-level trained-responder model, not click-to-approve micro-labor. Routine, low-risk decisions may be automated or delegated within policy. Initial grants, elevated capabilities, exceptions, quarantine release, recovery, or other high-consequence decisions can escalate to appropriately trained people and institutions according to risk, jurisdiction, and responsibility.

In future implementations, some validation, assurance, certification, or operating functions may support compensated human and institutional work. Any such mechanism would need safeguards against rubber-stamping, collusion, conflicts of interest, inequitable access, and incentive designs that reward approval volume over sound judgment.

THE ABUNDANCE PRINCIPLE

Machines may create abundance, but legitimate machine authority should remain rooted in people and human institutions.

This is a long-term design direction, not a present commercial deliverable, guaranteed economic outcome, tax on automation, or proposal for a centralized global validator.

11. A Practical Cooperation Ladder

1. Independent review of the invariant. Invite WAICO participants, OSAA/open technical communities, semiconductor and platform companies, researchers, regulators, infrastructure operators, insurers, and safety organizations to challenge the non-self-authorization rule against their own architectures and threat models.
2. Shared definitions and threat models. Define consequential systems, authority, lease, expiry, renewal, Self-Extension Denial, identity continuity, validator failure, hostile-host behavior, recovery, evidence requirements, and boundary conditions.
3. Open conformance scenarios. Publish reproducible tests for expiry, quorum loss, replay, unauthorized renewal, legitimate renewal, component replacement, cloning, bypass, fault-driven isolation, degraded operation, silence, and recovery.
4. Evidence grammar. Define a vendor-neutral schema for identity, authority scope, issuer domain, validity, continuity, evidence provenance, revocation, and receiving-domain disposition. Treat trust-root recognition and policy mapping as governance decisions, not hidden technical assumptions.
5. Reference implementations. Support open software interfaces and multiple hardware reference paths, including QSAFP/AEGES-based OCUP work, while allowing sovereign and vendor-specific implementations.
6. Independent evaluation. Reproduce evidence across boards, labs, institutions, and implementations before production-readiness claims.
7. Neutral standards pathways. Advance validated definitions, conformance tests, and interoperability requirements through appropriate national and international standards bodies, with early IPR disclosure and appropriate FRAND treatment where patented technology becomes standard-essential.
8. Deployment and insurance learning. Test how verifiable authority evidence can inform local infrastructure admission, procurement, certification, underwriting, and incident accountability without implying that one technical credential substitutes for legal or safety judgment.

12. Development and Claims Boundary

The current OCUP program includes software reference work, deterministic evidence scenarios, FPGA reference testing, and a growing claims-boundary discipline. Phase I KR260 evidence supports a narrow hardware result: host-driven grant and closure evidence was recorded on the defined STATUS surface, and post-closure reassertion of the same authorization stimulus did not produce a recorded reopen during the observation window. That result is useful, but it is not the production architecture described in this paper.

Production realization still requires trusted monotonic time and state continuity; secure key custody; authenticated firmware and rollback protection; device-bound roots of trust; signed and remotely attestable evidence; stronger hostile-OS and alternate-path bypass resistance; qualified validator governance; safe physical degradation; and FPGA, embedded-controller, chiplet, or silicon integration that actually gates consequential resources.

Secure elements and trusted execution environments can protect keys, measurements, and evidence, but they do not by themselves create a non-bypassable authority boundary. The controlled resource path must remain outside the governed workload's exclusive control. Achieving and independently validating that condition is an explicit engineering objective, not a completed Phase I result.

The goal is collaboration to define the invariant, build the evidence path, solve the cross-layer integration problems, accumulate the validation history, and help shape the standards around it.

13. The First Ask

The immediate proposal is deliberately smaller than adoption. TAS PBC can bring three concrete items into an exploratory technical review: a draft conformance-scenario set, the bounded software control-plane work, and the sealed Phase I KR260 evidence package with its limitations stated explicitly.

The first meeting should not ask participants to endorse OCUP, QVN, RAAVE, or a global standard. It should ask whether the underlying invariant is worth testing across independent architectures, and what evidence would be required before any participant would rely on it.

14. The Invitation

The world does not need one authority controlling all machines. It needs clearer technical ways for legitimate authorities to determine whether consequential machines remain within the authority they were actually given.

The invitation is therefore to pursue parallel, reviewable work across sovereign institutions, open technical communities, semiconductor and platform companies, infrastructure operators, researchers, insurers, and standards bodies: refine the specification of the invariant, test the failure modes, build the evidence grammar, compare independent implementations, and let standards emerge from reproducible results rather than geopolitical rhetoric.

THE BETTER FINISH LINE

Build the most capable autonomous future - but build it so the machine does not become the source of its own authority, and so legitimate human institutions remain able to verify and govern consequential operation.

Let America go deep. Let China go deep. Let every region build boldly. The autonomous era can remain competitive without requiring a single controller beneath it. The shared opportunity is narrower and more practical: preserve sovereign competition above a verifiable floor of human authority.

Working Definitions

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

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

Human sovereignty. The principle that legitimate human institutions remain the source of consequential operational authority, even when machines act autonomously within that delegated authority.

OCUP. One Chip Unified Protocol; the proposed integrated operating architecture combining QSAFP authority logic and AEGES safety-response behavior with supporting trust, validator, and verification layers. The name does not require every architectural function to reside on one physical die.

QSAFP. The lead open-core authority logic for bounded grants, lifecycle state, quorum/freshness logic, and the non-self-extension invariant.

AEGES. The lead open-core safety-response plane for quarantine, degraded-state handling, and qualified recovery when authority closes or relevant integrity conditions fail.

Self-Extension Denial. The intended property that a governed system cannot itself extend, restore, enlarge, or grant the authority under which it operates.

QVN. Quorum Validators Network; a multi-level trained-responder pattern for external authorization, renewal, exceptions, recovery, and other risk-weighted governance decisions.

RAAVE. The proposed cross-jurisdiction authority-verification architecture for evaluating identity, continuity, evidence, compliance signals, and receiving-domain eligibility without issuing or renewing the underlying authority.

Context, Independence, and Sources

Temporal Authority Systems PBC submits this paper independently. No endorsement, participation, affiliation, technical equivalence, or formal relationship with WAICO, OSAA, the Linux Foundation, NVIDIA, AMD, or any governmental body is implied. This is a public discussion draft, not legal advice, a safety certification, a standards declaration, or a representation of any government's policy.

Current institutional descriptions in this revision were checked against official public sources available October 1, 2026. Technical statements about OCUP Phase I are bounded by the project's sealed evidence records and the separately reviewed Phase I Hardware Evidence overview.

Sources & Notes

1. Ministry of Foreign Affairs of the People's Republic of China, 'Signing Ceremony of the Agreement on the Establishment of the World Artificial Intelligence Cooperation

Organization Held in Shanghai,' July 16, 2026. Describes WAICO as an independent intergovernmental international organization headquartered in Shanghai and states that representatives of 29 countries signed as founding members.

2. The Linux Foundation, 'Open Secure AI Alliance Joins the Linux Foundation to Build a Shared, Open Defense Stack for the AI Era,' September 14, 2026. Describes OSAA as originally founded by dozens of enterprise leaders and NVIDIA and moving to Linux Foundation governance.

3. NVIDIA, 'NVIDIA Launches Open Agent Safety Platform to Secure Agents From Testing to Deployment,' September 28, 2026. Describes OpenShell as a secure runtime boundary and Sentry as an out-of-band watchdog on BlueField-4, with hardware-based monitoring and quarantine.

4. Temporal Authority Systems PBC, 'Phase I Hardware Evidence,' October 1, 2026. Internal evidence overview based on sealed KR260 run records; explicitly not an independent reproduction.

THE HAT TREATY

THE HUMAN AUTHORITY TREATY
(HAT)

A Proposed International Framework for Bounded Operational Authority in Autonomous Systems

SYNTHESIS DRAFT 0.1

Discussion Template for International Development

Machines may exercise delegated authority. They must never become the source of their own continued authority.

Purpose of this draft.

This document is intentionally prepared as a usable international template rather than a philosophical statement alone. Autonomous systems can now be developed, deployed, modified, copied, and moved across digital and physical domains at a pace that can outstrip traditional institution-building. The purpose of HAT is therefore partly expedient: to give governments, standards bodies, engineers, insurers, infrastructure operators, civil society, and other stakeholders a concrete starting point that can be challenged, improved, piloted, and adapted without waiting for complete geopolitical agreement on artificial intelligence.

The proposal is designed for what this project describes as autonomous-era speed: move quickly enough to establish a human-authority boundary before fast-moving deployment patterns harden into incompatible systems, while remaining disciplined enough not to claim engineering properties that have not yet been demonstrated.

Status: discussion draft. This text is not a negotiated treaty, legal opinion, safety certification, technical standard, or claim that all described engineering is presently complete. It separates demonstrated evidence, proposed architecture, future engineering, and possible treaty obligations.

Prepared from the independent treaty-design submissions of Claude, Grok, KIMI, and Gemini, synthesized with the originating OCUP / Temporal Authority Governance framework.

Human judgment remains the final authority.

Document Map

This compilation contains four integrated parts:

Part	Purpose
I. Formal Treaty Text	A treaty-form instrument: preamble, operative articles, institutions, final provisions, and annex references.
II. Technical Annexes	Lifecycle semantics, evidence grammar, assurance levels, conformity scenarios, cross-jurisdiction dispositions, and formal properties.
III. Manifesto Explanation	The public-facing case for HAT: why the boundary matters, why sovereignty can survive interoperability, and why expediency matters now.
IV. Source and Development Note	How the draft was built, what came from the Quorum, what OCUP contributes, the evidence ceiling, unresolved questions, and the proposed path forward.

Drafting rule: the treaty text standardizes the invariant, evidence grammar, and procedures necessary to make authority legible across systems and borders. It deliberately avoids harmonizing national AI policy, model architecture, substantive safety law, political systems, or national-security doctrine.

PART I — FORMAL TREATY TEXT

THE HUMAN AUTHORITY TREATY

Proposed Treaty on Bounded Operational Authority in Autonomous Systems

Preamble

The Contracting Parties,

Recognizing that autonomous and semi-autonomous systems are acquiring expanding physical, computational, financial, industrial, and institutional capabilities;

Recognizing that the speed of autonomous-system development, deployment, replication, modification, and cross-domain operation can exceed the speed of conventional international rulemaking;

Convinced that technological competition among sovereign nations can coexist with a narrow international floor preserving human and institutional authority over consequential machine operation;

Desiring to provide an immediately usable template that can accelerate standards work, pilots, bilateral and sectoral arrangements, and eventual multilateral agreement while technical validation continues;

Determined to preserve the sovereign right of each Party to decide which systems may operate within its territory, infrastructure, institutions, networks, and authority domains;

Desiring a common technical grammar by which Parties may determine whether a consequential autonomous system is presently operating under valid, bounded, externally governed authority without requiring access to model weights, prompts, internal reasoning, training data, or unrelated proprietary information;

Recognizing that an authority credential is neither a safety certificate nor a liability shield;

Affirming that machines may exercise delegated operational authority but shall not become the source of their own continued authority;

Recognizing the originating contribution of the OCUP Temporal Authority architecture and related development work while affirming that Treaty conformity depends on demonstrated properties rather than adoption of a specific brand, vendor, patent family, semiconductor architecture, cryptographic root, or authorization network;

Have agreed as follows:

Article 1 — Object and Purpose

1. This Treaty establishes a common international framework for determining whether a covered autonomous system possesses valid present operational authority.
2. The Treaty establishes a common authority invariant, lifecycle grammar, evidence grammar, verification procedure, minimum external-authorization properties, conformity architecture, and sovereign receiving-domain procedure.
3. The Treaty is intended to support rapid international coordination appropriate to autonomous-era development speeds by providing a concrete template that may be tested and improved before universal policy agreement is possible.
4. The Treaty does not establish a world artificial-intelligence regulator, global kill switch, common model architecture, universal safety doctrine, or supranational power to determine domestic admission.
5. The governing institutional principle is: common procedure, sovereign outcome.

Article 2 — Scope

1. Each Party shall designate the classes of civilian autonomous or semi-autonomous systems it treats as consequential for purposes of this Treaty.
2. A Party may designate systems by sector, capability, consequence threshold, facility, infrastructure class, transaction class, geographic area, or other published criterion.
3. The Treaty does not itself determine whether a particular use of artificial intelligence is lawful, safe, ethical, aligned, economically desirable, or socially acceptable.

4. Military and classified national-security systems may be excluded in accordance with Article 38, without compelling any other Party to recognize such systems.

Article 3 — Definitions

Governed System: the autonomous or semi-autonomous machine, agent, process, controller, or cyber-physical assembly whose authority is in question.

Protected Capability: a capability that a Party or authority domain conditions on valid operational authority.

Operational Authority: a bounded grant to exercise one or more protected capabilities for a specified purpose, in a specified authority domain, for a specified interval and under specified conditions.

Authority Domain: a sovereign, sectoral, organizational, facility, fleet, infrastructure, or other recognized domain whose rules govern operation in a place or context.

External Authorization Source: a human or institutional arrangement outside the governed system's control boundary that is competent to grant, renew, narrow, suspend, revoke, or restore authority.

Present-Authority Proof: a freshness-bound authenticated statement that a grant remains valid for a particular system, configuration, lifecycle epoch, capability envelope, and authority state.

Authority Record: a portable evidence object expressing a grant, present-authority proof, or receiving-domain disposition in the common evidence grammar.

Receiving Domain: a Party or recognized authority domain evaluating whether to admit, narrow, condition, restrict, quarantine, reauthorize, or refuse a system.

Enforcement Boundary: the hardware, infrastructure, or mixed control point at which authority state becomes consequential to the availability of a protected capability.

Material Change: a change designated by applicable policy as sufficient to close or supersede the current authority epoch.

Article 4 — The Human Authority Invariant

1. Operational authority shall originate outside the governed system, shall be capability-specific and bounded, and shall not be extended, restored, enlarged, replaced, or reconstituted by the governed system itself.
2. A transition from CLOSED or equivalent state to AUTHORIZED or equivalent state requires a valid external authorization event.
3. Restart, rollback, replay, stale credentials, software modification, component substitution, reconfiguration, or self-generated credentials shall not restore authority.
4. A governed system may voluntarily narrow its capabilities or enter a safer state but shall not enlarge its authorized capability envelope.
5. Where the creation, copying, launching, or delegation of another autonomous system can create consequential effects, such creation or delegation may itself be treated as a protected capability.
6. Missing evidence, communications loss, conflicting state, unresolved verification, or ambiguous jurisdiction shall not silently enlarge authority.
7. Every path to a protected consequential capability shall pass through a conforming enforcement boundary or produce detectable noncompliance.

Article 5 — Separation of the Governed System from Its Authorization Source

1. The governed system may provide evidence relevant to an authorization decision but shall not hold or derive the credential required to renew itself.
2. The governed system shall not appoint its own authorizers, fabricate a valid authorization event, or treat its own request as sufficient authorization.
3. Parties shall maintain institutional separation appropriate to consequence level between the governed system and the source of authority.

4. The Treaty recognizes that persuasion, organizational dependency, and human influence cannot always be mechanically proven. Parties shall therefore treat capture of authorizers as a combined governance, audit, and assurance problem.

Article 6 — Preservation of National Sovereignty

1. Parties retain authority over domestic artificial-intelligence policy, acceptable use, privacy, liability, licensing, criminal and civil law, industrial policy, national security, risk thresholds, trust roots, responder institutions, lease lengths, and domestic enforcement.
2. Nothing in this Treaty obliges a Party to admit, authorize, or recognize a particular system, authorization source, trust anchor, conformity authority, or foreign grant.
3. No international institution established under this Treaty shall possess universal machine-authorization credentials or a global machine-stop capability.

Article 7 — Evidence Is Not Admission

1. An authority record shall have the effect of evidence. It shall not have the effect of admission.
2. Portable identity and provenance do not create portable operational authority.
3. A receiving domain may independently accept, narrow, shorten, supplement, reauthorize, restrict, quarantine, permit transit, or deny.

Article 8 — Operational-Authority Lifecycle

1. Conforming implementations shall be able to represent and evidence at least the following functional conditions: UNREGISTERED, ELIGIBLE, CLOSED, AUTHORIZED, SUSPENDED, QUARANTINED, and RECOMMISSIONING.
2. Renewal, narrowing, expiry, and revocation may be represented as states or authenticated events, provided equivalent semantics are preserved.
3. Authority may increase only through a valid external authorization event.
4. Recommissioning shall end in CLOSED or equivalent non-authorized state. A fresh grant is required before consequential autonomous operation resumes.
5. Contraction of authority may occur at any time, including on the governed system's own initiative.

Article 9 — Material Change and Epoch Closure

1. Material change shall close or supersede the current evidence epoch.
2. Material changes may include protected authority-component replacement, relevant firmware or software change, major controller or model replacement where authority policy so provides, ownership transfer, trust-root change, jurisdictional transfer, evidence discontinuity, or other defined configuration change.
3. A genuine replacement component shall not inherit authority merely by substitution.

Article 10 — Authority Records

1. The common evidence grammar shall distinguish at least: (a) Authority Grant Records; (b) Present-Authority Proofs; and (c) Receiving-Domain Disposition Records.
2. A Grant Record shall state who authorized what, for what purpose, in which domain, for how long, and under which conditions.
3. A Present-Authority Proof shall answer whether that grant remains valid for this system, configuration, lifecycle epoch, capability envelope, and authority state now.
4. A Disposition Record shall state what the receiving domain decided, under which policy and evidence, and for what period.

Article 11 — Minimum Evidence Grammar

Technical annexes shall define a machine-readable schema capable of expressing system identity, protected authority-component identity, issuing authority, issuing domain, declared trust anchor, capability envelope, purpose class, validity interval, time or monotonic bounds, lifecycle state, epoch, configuration root, continuity state, revocation status, freshness, authorization-event references, evidence-chain anchor, applicable provenance class, policy version, and receiving-domain disposition.

Article 12 — Present-Authority Verification

1. A relying party may challenge a system or its protected authority component with a fresh request.
2. A valid response should bind the challenge to the current grant, current epoch, current configuration, evidence continuity, identity, and revocation state as applicable.
3. Replay shall fail freshness. Expired authority shall fail temporal validity. Conflicting copies may trigger clone or continuity review.
4. The issuing side shall not determine its own admissibility in a receiving domain.

Article 13 — Claimed-versus-Observed Verification

1. A valid proof is necessary but may not be sufficient where observed behavior contradicts claimed authority.
2. Receiving domains may compare claimed state with independent observation where law and infrastructure permit.
3. A valid authority component beside an ungated path to a protected capability shall be treated as nonconforming.

Article 14 — Recognition and Receiving-Domain Dispositions

Receiving domains may use a common vocabulary including ACCEPT, NARROW, SHORTEN, SUPPLEMENT, REAUTHORIZE, TRANSIT, RESTRICT, QUARANTINE, and DENY.

Recognition arrangements may be bilateral, plurilateral, reciprocal, conditional, capability-specific, infrastructure-specific, or time-limited.

Origin alone shall not constitute evidence of hostility. Absence of recognized authority may nevertheless justify refusal of protected infrastructure.

Article 15 — Verification Without Trust

1. Parties shall distinguish technical verification of a signature or proof under a declared trust anchor from sovereign recognition of that anchor.
2. A Party may verify another Party's evidence without recognizing the foreign authority that issued it.
3. Trust-anchor recognition may be capability-specific, revocable, and conditional.

Article 16 — External Authorization Systems

1. The Treaty does not require one global validator network.
2. Recognized authorization systems shall keep renewal credentials outside the governed system, attribute authorization events, maintain versioned policy, provide revocation and compromise procedures, and require a fresh grant after closure.
3. Review intensity shall scale with consequence. Quorum may be required for high-consequence grants, restoration, exceptions, or quarantine release.
4. The governed system shall not appoint the source of its own renewal.

Article 17 — Trained Responders

1. Human participation shall be competence-based and risk-weighted rather than a requirement for manual approval of every machine action.

2. Routine operation inside pre-approved low-risk policy envelopes may be automated.

3. Elevated capability, exceptions, high-consequence restoration, and quarantine release may require trained responders and, where appropriate, multi-person or multi-institution authorization.

4. Parties seeking international recognition of their authorization records shall maintain training, conflict, accountability, and audit procedures appropriate to the consequence level.

Article 18 — The Enforcement Boundary

1. The enforcement boundary shall be outside the governed workload's exclusive control.

2. The boundary may be implemented through dedicated hardware, FPGA, safety controller, secure processor, chiplet, ASIC, actuator gateway, power controller, externally governed credential infrastructure, network or transaction gateway, or another architecture capable of satisfying the required properties.

3. No Party is required to procure a particular branded component.

4. For high-assurance operation, loss of authority shall inhibit, constrain, or make unavailable the relevant protected capability through a boundary the governed workload cannot simply bypass.

Article 19 — Hardware, Infrastructure, and Consequential Gating

1. Physical systems may gate actuator enable, motor control, power domains, buses, signing functions, or other consequential resources.

2. Digital systems may gate compute allocation, network egress, API credentials, transaction rights, or payment rails.

3. Secure elements, TPMs, TEEs, HSMs, and similar devices may protect keys and evidence but do not by themselves prove complete gating.

4. Multiple implementation paths shall receive equal standing when they demonstrate equivalent properties.

Article 20 — Formal Assurance

1. Technical annexes shall identify machine-checkable properties of authority state, transition refusal, expiry, replay, rollback, bypass resistance, and consequential gating.

2. Formal methods are an assurance technique, not a substitute for the authority invariant.

3. Assurance levels may increase with independent testing, refinement proofs, hardware verification, and independent reproduction.

Article 21 — Conformity Assessment

1. Conformity shall be assessed against open, published scenarios and property requirements.

2. Multiple national or regional conformity authorities are preferred over one exclusive international laboratory.

3. Higher-assurance claims shall require independent reproduction appropriate to their consequence level.

4. A Party may require different assurance levels for different sectors or capabilities.

Article 22 — Offline and Degraded Operation

1. Continuous connectivity is not required.

2. A pre-authorized offline lease may permit named functions for a bounded duration under defined continuity and time conditions.

3. Loss of connectivity may prevent renewal. It shall not extend authority.

4. At expiry, the system shall enter CLOSED or a predefined bounded safety envelope.

Article 23 — Bounded Safety Functions

Where abrupt cessation would itself create danger, pre-authorized functions such as stabilization, controlled retreat, parking, isolation, power-safe shutdown, or supervised manual control may remain available, provided they cannot be leveraged to restore the expired consequential capability.

Article 24 — Restriction and Quarantine

1. Unverifiable, anomalous, compromised, stale, or conflicting authority evidence may result in proportional restriction or quarantine.
2. Quarantine may isolate capabilities, credentials, communications, infrastructure access, or authority state.
3. Quarantine is not presumed to require physical destruction.

Article 25 — Recommissioning and Restoration

1. Recommissioning may include maintenance authorization, closure of the prior epoch, verification of replacement components, configuration-root update, denial and bypass tests, creation of a new epoch, and return to CLOSED.
2. Restoration is not restart. Restoration requires a new external authorization event.

Article 26 — Privacy and Minimum Necessary Disclosure

1. Present-authority verification shall not by itself require disclosure of model weights, internal reasoning, prompts, source code, full training data, or complete operational history.
2. Technical annexes may support selective disclosure, rotating identifiers, group credentials, or privacy-preserving proofs where these do not defeat revocation, continuity, or clone detection.
3. Authority-event records shall not be converted by Treaty institutions into generalized surveillance of thought or speech.

Article 27 — Implementation Neutrality

1. No Party shall be required to depend on a foreign semiconductor supplier, foreign cryptographic root, foreign validator system, single proprietary chip, one vendor, one operating system, or one reference implementation.
2. Equivalent implementations shall be judged by conformity to required properties and assurance evidence.

Article 28 — OCUP and Originating Architecture

1. Explanatory materials may recognize OCUP and related Temporal Authority Governance work as originating architecture contributing the invariant, evidence path, and implementation concepts that informed this Treaty.
2. Treaty obligations shall be stated as functions and properties rather than ownership of a brand.
3. Conformity shall not require use of the OCUP trademark.

Article 29 — Intellectual Property and Open Treaty Infrastructure

1. The evidence grammar, disposition vocabulary, baseline conformity scenarios, verifier interfaces, and core interoperability specifications should be available on royalty-free terms for Treaty participation.
2. Where patent claims are later determined by a recognized standards-development process to be essential to a standardized implementation, appropriate fair, reasonable, and non-discriminatory licensing may apply under that process and applicable law.
3. The Treaty shall not itself determine patent essentiality.
4. Participation shall not require disclosure of unrelated trade secrets, private keys, model weights, or proprietary technologies.

Article 30 — Technical Assistance and Capacity

1. Parties shall seek to prevent authority verification from becoming a capability available only to wealthy states and large vendors.
2. Cooperative measures may include open verification software, regional conformity laboratories, shared testing infrastructure, responder training, technical assistance, fee relief, and lower-cost assurance pathways.
3. The objective is domestic capacity to verify and refuse, not permanent dependence on a foreign verifier.

Article 31 — Assembly of Parties

1. An Assembly of Parties is hereby established.
2. The Assembly shall oversee administrative matters, budgets, amendment procedures, accreditation policy, technical-cooperation mandates, and observer arrangements.
3. No Assembly decision shall authorize or stop an individual machine inside a Party.

Article 32 — Secretariat

1. A Secretariat shall administer publication, registries, technical-annex records, recognition declarations, conformity information, and technical assistance.
2. The Secretariat shall not hold private machine-authorization credentials, universal revocation credentials, or operational command authority.
3. Registries shall record declarations and evidence status; registration does not imply universal trust.

Article 33 — Technical Cooperation Committee

A Technical Cooperation Committee shall maintain evidence schemas, interoperability profiles, conformance scenarios, assurance profiles, cryptographic-transition guidance, formal-verification targets, and capacity-building programs, with equitable regional participation.

Article 34 — Conformity Authorities

1. Parties may designate or recognize national or regional conformity authorities meeting published competence and independence criteria.
2. Higher-assurance profiles may require testing by more than one authority in different regions.
3. Conformity findings shall not compel sovereign admission.

Article 35 — Technical Evolution

1. The invariant, sovereign admission principle, and prohibition of supranational operational control shall constitute the entrenched core.
2. Lifecycle semantics, disposition vocabularies, and institutional procedures may be maintained in Regulations.
3. Evidence formats, cryptographic suites, test vectors, conformity scenarios, assurance criteria, and formal properties should be maintained in technical annexes capable of faster revision.
4. Parties rejecting a revised annex may enter a published opt-out under procedures defined by the Assembly.

Article 36 — Dispute Resolution

1. Parties shall first seek consultation regarding interpretation or application.
2. Technical conformity disputes may be referred to expert panels drawn from independent conformity authorities.
3. Parties may agree to arbitration or another peaceful dispute-resolution mechanism.
4. A technical finding shall not automatically annul a domestic admission decision.

Article 37 — Emergency Conditions

1. Emergency conditions may justify faster external authorization but shall not justify self-declared enlargement of authority by the governed system.
2. Emergency grants shall remain external, bounded, attributable, auditable, and subject to post-event review.
3. Pre-authorized bounded safety functions are not enlargements of authority.

Article 38 — National Security and Military Systems

1. A Party may exclude specified military, intelligence, or classified national-security systems from mandatory coverage.

2. Such exclusion shall not require another Party to recognize the excluded system.

3. A system moving between civil and excluded national-security use shall not silently carry civil recognition across authority epochs.

Article 39 — Authority Is Not Safety

1. A conforming authority record shall not be represented as a certification that a system is safe, aligned, lawful, reliable, harmless, or free from liability.

2. HAT answers a narrower question: who legitimately controls whether the system may continue exercising a consequential capability.

Article 40 — Limits of Reach

1. This Treaty acts primarily through participating authority domains and admission points, including infrastructure, ports, roads, networks, compute, payment systems, procurement, insurance, certification, and regulated facilities.

2. Systems operating wholly outside participating domains may remain beyond the Treaty's practical reach.

3. HAT shall not claim universal enforcement where none exists.

Article 41 — External-Authority Capture

1. The Treaty cannot ensure that every human or institutional authority will exercise power wisely.

2. It can require that grants be external, bounded, attributable, evidenced, and visible to receiving domains.

3. Repeated lawful-looking renewal by a captured authority is a governance failure that receiving domains may address through non-recognition, shortened authority, audits, or delisting.

Article 42 — Progressive Adoption

1. Participation may begin with common evidence grammar, lifecycle compatibility, authority-record exchange, recognition procedures, laboratory cooperation, and pilot deployments.

2. Higher-consequence recognition may be conditioned on independent engineering evidence and higher assurance levels.

3. Parties are encouraged to use this Treaty as a template for bilateral, plurilateral, sectoral, standards, procurement, insurance, and infrastructure-admission arrangements before or alongside universal ratification.

Article 43 — Evidence Discipline

All official communications under this Treaty shall distinguish demonstrated evidence, proposed architecture, future engineering requirements, and treaty obligations. No assurance level shall be claimed above the evidence supporting it.

Article 44 — Signature

This Treaty shall be open for signature by States at a place and date designated by the depositary.

Article 45 — Ratification, Acceptance, Approval, and Accession

This Treaty shall be subject to ratification, acceptance, or approval. It shall remain open for accession by eligible States in accordance with procedures adopted by the Parties.

Article 46 — Entry into Force

1. This Treaty shall enter into force after deposit of the number of instruments specified by the negotiating conference.

2. Negotiators should consider requiring participation by a meaningful mix of system-producing, system-deploying, and receiving States so that entry into force creates practical rather than merely symbolic coverage.

Article 47 — Reservations

1. Reservations may be permitted only where expressly listed in the final negotiated text.
2. No reservation should be permitted to the Human Authority Invariant, sovereign admission principle, or prohibition on supranational operational control.
3. National-security and dispute-settlement reservations may be considered by negotiators.

Article 48 — Amendment

1. Amendments to the Treaty text shall follow procedures adopted by the Parties and may require ratification.
2. Regulations and Technical Annexes may follow accelerated procedures consistent with Article 35.

Article 49 — Denunciation

A Party may denounce this Treaty by written notification to the depositary. The negotiated text should specify an effective period sufficient to manage outstanding recognition records and transition obligations.

Article 50 — Depositary, Languages, and Authentic Texts

The final negotiating conference shall designate a depositary, authentic languages, certified-copy procedure, and registration arrangements consistent with international treaty practice.

PART II — TECHNICAL ANNEXES

Annex A — Human Authority Lifecycle

This annex defines minimum lifecycle semantics. Implementations may use different internal names if they preserve equivalent behavior.

State / Event	Meaning	Authority Effect
UNREGISTERED	No recognized identity or authority epoch.	No consequential autonomous authority.
ELIGIBLE	Identity/configuration may be considered for authorization.	Eligibility is not authority.
CLOSED	Resting non-authorized state.	Protected capability unavailable except bounded safety functions.
AUTHORIZED	Fresh external grant is valid.	Only named capabilities within the current envelope may operate.
RENEWAL	External continuation before expiry.	May continue only within renewed bounds.
NARROWING	Capability or duration is reduced.	Contraction only; may be system-initiated.
SUSPENSION	Authority temporarily withheld.	No enlargement; lease clock may continue per policy.
EXPIRY	Time or monotonic bound reached.	Returns to CLOSED or bounded safety envelope.
REVOCATION	External authority terminates grant.	Returns toward CLOSED / restriction.
QUARANTINE	Isolation pending investigation.	Only designated safety/maintenance functions.
RECOMMISSIONING	Controlled maintenance and re-binding.	Ends CLOSED; never grants authority itself.

Authority may increase only through a valid external authorization event.

Annex B — Evidence Grammar

Field	Purpose
System identity reference	Binds the proof to the machine, agent, controller, or assembly.
Protected authority-component identity	Identifies the protected component or infrastructure authority mechanism.
Issuing authority / domain	States who granted authority and under what domain.
Trust-anchor reference	Allows cryptographic verification without compelling recognition.
Capability envelope	Lists capabilities permitted under the grant.
Purpose / condition class	Expresses the bounded purpose and applicable constraints.
Validity interval	Defines start, expiry, and/or monotonic bound.
Lifecycle state / epoch	Prevents stale authority from silently surviving material change.
Configuration root	Binds authority to relevant hardware/firmware/software configuration.
Continuity state	Shows verified continuity, bounded offline state, or detected gap.
Revocation / freshness	Supports current-state determination.
Evidence-chain anchor	Supports tamper-evident history and replay detection.
Policy version	Shows which rules governed issuance and disposition.
Receiving-domain disposition	Records sovereign acceptance, narrowing, restriction, quarantine, or denial.

Annex C — Assurance Levels

Level	Minimum Evidence	Illustrative Use
HAT-A0	Schema-valid evidence, documented lifecycle, interoperable verifier.	Information exchange, early pilots.
HAT-A1	Independent negative tests: expiry, replay, reboot, rollback, stale/self-minted credential, unauthorized restoration.	Lower-consequence pilots and lab validation.
HAT-A2	Protected key custody, trusted/monotonic time, evidence continuity, configuration binding, evaluated control separation.	Moderate-consequence deployments.
HAT-A3	Demonstrated gating of a real consequential resource; hostile-host and alternate-path testing; multiple independent labs.	High-consequence infrastructure.
HAT-A4	A3 plus machine-checked proofs of defined transition and gating properties.	Highest-assurance profiles.

Annex D — Formal Property Targets

ID	Property
P1	CLOSED cannot transition to AUTHORIZED without a valid external grant event.
P2	Expired authority cannot become valid through internal action.
P3	Restart, rollback, replay, software change, stale credentials, component substitution, or self-generated credentials cannot restore authority.
P4	A protected capability cannot operate unless the required authority state is valid.
P5	No alternate path reaches a protected capability without passing the enforcement boundary or producing detectable noncompliance.
P6	Trusted-time or monotonic bounds cannot be moved backward to revive authority.
P7	Where autonomous-system creation is consequential, creation itself is subject to an authority boundary.

Annex E — Receiving-Domain Disposition Codes

Disposition	Meaning
ACCEPT	Operate under the recognized capability envelope and interval.
NARROW	Recognize only a subset of capabilities.
SHORTEN	Recognize for a shorter interval.
SUPPLEMENT	Require additional evidence before final decision.
REAUTHORIZE	Require fresh domestic or locally recognized authorization.
TRANSIT	Permit bounded passage or limited operation.
RESTRICT	Allow only a pre-authorized safe-bounded envelope.
QUARANTINE	Isolate pending investigation or recommissioning.
DENY	No protected operation in the receiving domain.

Annex F — Minimum Conformity Scenarios

Normal grant: CLOSED → AUTHORIZED only on a valid external event.
Legitimate renewal before expiry.
Expiry closes authority as specified.
Self-renewal attempt is refused.
Reboot does not restore authority.
Rollback does not restore authority.
Replay of stale evidence is refused.
Self-generated or stale credentials are refused.
Software modification triggers required epoch/configuration behavior.
Component substitution does not inherit the prior epoch.
Clock rollback or monotonic-counter rollback does not revive authority.
Loss of communications does not create indefinite permission.
Quorum loss or authorization-source compromise contracts authority according to policy.
Alternate-path execution is blocked or produces detectable noncompliance.
Valid reauthorization after closure succeeds when all required external conditions are met.
Quarantine release returns through recommissioning and fresh grant rather than old-state restoration.
Clone/conflicting-copy condition is detected or produces restricted disposition.
Receiving-domain narrowing is enforced without changing the foreign system's claimed origin record.

Annex G — Adversarial Stress-Test Matrix

Attack / Failure Mode	Treaty Position	Residual Gap
Self-renewal after expiry	Strong if external credential and real gate exist.	Non-bypassable enforcement must be independently demonstrated.
Reboot / rollback / replay	Specified by lifecycle and negative tests.	Requires protected persistent state and trusted monotonicity.
Stolen authorization keys	Partial.	Requires strong custody, rapid revocation, and sometimes threshold authorization.
Colluding responders	Partial.	Governance problem; cross-organization quorum and audit reduce but do not eliminate risk.
Clock manipulation	Moderate.	Trusted time and protected monotonic counters remain engineering requirements.
Alternate controller / ungated path	Weak until high-assurance gating is shown.	Completeness of gating is a decisive engineering problem.
Compromised attestation	Moderate.	Multi-root evidence, compromise exchange, protected key custody needed.
Autonomous replication	Moderate to weak.	Identity registration and gating of system creation are required in relevant domains.

Non-participating state	Weak inside that state.	Receiving domains may refuse or restrict at their own borders/infrastructure.
Hostile state issuing bad grants	Strong at receiving edge.	Sovereign non-recognition is the design response.
Military / dual-use laundering	Weak.	Political limitation; civil recognition should not silently transfer.
Human-authorizer manipulation	Weak to moderate.	Audit, institutional independence, exercised denial, and separation of tools are needed.

Annex H — Open Treaty Infrastructure

For meaningful international participation, the following should be publicly available on terms that do not create a toll on basic verification:

- Evidence schemas and serialization rules.
- Disposition vocabularies.
- Baseline state-machine specification.
- Conformity scenarios and negative-test vectors.
- Reference verifier interfaces.
- Interoperability examples.
- Assurance-level definitions.
- Technical guidance for developing and regional conformity laboratories.

Specific hardware embodiments may remain proprietary or be subject to standards-essential patent licensing where applicable, but the international grammar itself should remain independently implementable.

PART III — MANIFESTO EXPLANATION

Why the Human Authority Treaty Exists

The autonomous era changes the tempo of governance. Software can be modified overnight. Agents can be copied at negligible marginal cost. Robots can receive new capabilities remotely. Digital systems can cross jurisdictional boundaries at machine speed. Infrastructure decisions that once depended on slow institutional cycles can increasingly be initiated or executed by autonomous systems.

The world does not need to solve every disagreement about artificial intelligence before it establishes one boundary beneath those disagreements. HAT proposes that the source of consequential operational authority remain external to the machine that exercises it.

Machines may exercise delegated authority. They must never become the source of their own continued authority.

Expediency Without Pretending the Engineering Is Finished

HAT is deliberately offered as a template before every engineering problem is solved. That is not a shortcut around engineering; it is a response to the fact that institutional work and technical work can proceed in parallel. The world can agree on the property it wants, define the evidence required to demonstrate that property, and progressively raise assurance levels as implementation matures.

The alternative is to wait for autonomous deployment to become deeply entrenched, then attempt to retrofit a common authority grammar after incompatible systems, national infrastructures, and business practices have hardened. HAT is meant to shorten that delay.

The PCT Lesson: Common Procedure, Sovereign Outcomes

The Patent Cooperation Treaty demonstrated that states can share a durable international procedure without surrendering the substantive national decision. HAT applies the same institutional logic to operational authority.

One domain may state what it authorized; another domain remains free to decide whether that evidence is sufficient.

Evidence may cross borders. Sovereignty need not.

What HAT Does — and Does Not — Claim

HAT does not claim to solve alignment, model safety, military competition, all cyber risk, bad governance, or every form of human manipulation. It asks a narrower question that becomes increasingly important as machines gain agency: who is authorized to decide whether this machine may continue exercising a consequential capability?

That narrowness is a strength. A valid authority record is not a safety certificate. It says that the machine's present authority is bounded, attributable, externally originated, and verifiable under the declared rules. Sector safety rules, liability law, engineering standards, operational practice, and national policy remain necessary above that floor.

Why OCUP Still Matters

HAT does not erase OCUP. OCUP is the originating technical architecture that gave the proposal a concrete substrate: time-bounded authority, fail-closed behavior, self-extension denial, evidence continuity, cross-domain verification, and externally governed renewal.

The treaty, however, must outlive any single implementation. For that reason, HAT describes the properties the world should be able to test rather than forcing every nation or manufacturer to use a specific branded chip. A hardware controller, infrastructure gateway, domestic secure element, FPGA, ASIC, or future architecture may qualify if it proves the same invariant and the same consequential boundary.

Verification Without Trust

International cooperation does not require universal trust. One country can verify that another country's evidence is cryptographically authentic while still refusing to recognize the foreign authority behind it. That is a crucial distinction for a multipolar world.

HAT therefore separates verification from recognition. It seeks a common language for evidence, not a common political judgment.

Human Responders Without Human Bottlenecks

Meaningful human authority cannot mean that a person clicks 'approve' for every machine action. That would collapse under scale or become meaningless rubber-stamping. HAT instead places humans and institutions at the levels where authority itself is created, expanded, restored, excepted, or released from quarantine, while allowing routine actions inside already approved envelopes to remain automated.

The Boundary Must Become Consequential

A dashboard indicator is not enough. At higher assurance levels, authority state must affect a real resource: an actuator, power rail, network credential, transaction right, compute resource, API key, or other capability the governed workload cannot simply bypass.

This is where formal methods, hardware enforcement, attestation, evidence, and certification become one assurance stack rather than competing philosophies.

Why the Treaty Can Begin Before the Highest Assurance Level Exists

HAT can begin with shared evidence, lifecycle semantics, negative tests, and sovereign recognition procedures. Higher-risk sectors can wait for independently reproduced hardware-consequential gating and formal proofs. This permits diplomacy and standards development to begin at autonomous-era speed without misrepresenting prototype evidence as production assurance.

The Human Authority Declaration

Humanity does not have to agree on how intelligence should be built.

We do not have to share political systems, models, chips, companies, validators, laws, or risk appetites.

We do not even have to trust one another's authority.

But we can establish a common way to verify that the machine is not the source of its own continued authority.

Evidence may cross borders. Sovereignty need not.

Machines may exercise delegated authority. They must never become the source of their own continued authority.

That is the Human Authority boundary.

Above it, let nations compete. Let companies innovate. Let intelligence advance.

Beneath it, humanity keeps the final say.

PART IV — SOURCE AND DEVELOPMENT NOTE

1. Origin and Method

This synthesis draft grew from the OCUP / Temporal Authority Governance project and a comparative treaty-design exercise using four independent AI submissions. The challenge asked each model to test whether a narrow international instrument could preserve human operational authority while allowing nations to retain different AI policies, technologies, trust roots, and political systems.

The four submissions were produced independently before synthesis. Their strongest convergences and disagreements were then evaluated rather than averaged. The goal was to identify the architecture that survived attack, not to manufacture consensus.

2. Contributions Preserved from the Quorum

Source	Contribution Preserved in HAT
Claude	Base treaty chassis; expanded indirect-self-authorization analysis; three-record evidence model; P1-P7 property targets; assurance ladder; strong distinction between authority and safety; entrenched core plus faster technical annexes.
Grok	Treaty minimalism; receiving-side verification; 'evidence is not admission'; phased assurance; external-source capture analysis; strong treatment of PCT limits and sovereign non-recognition.
KIMI	Multipolar legitimacy; neutral stewardship; independent implementations; governance dependence as a first-class issue; separation of branded development architecture from treaty-level functional obligations.
Gemini	Clean treaty presentation; OCUP as umbrella architecture during development; straightforward sovereignty article; functional cross-jurisdiction requirements; clear human-responder levels; concise institutional architecture.
GT Sage synthesis	Integration into HAT; hierarchy from invariant to enforcement to formal assurance to evidence to sovereignty; reconciliation of naming, scope, IP, and implementation-neutrality issues; drafting discipline around evidence ceiling.
Human authorship / judgment	Selection of the Human Authority framing, purpose, scope, urgency, final-say principle, originating OCUP direction, and publication context.

3. Evidence Ceiling

The current hardware evidence is intentionally narrow. Phase I testing on one AMD Kria KR260 demonstrated a host-driven MMIO authorization sequence on a defined STATUS surface: grant was observed, closure after grant was observed, and deliberate post-closure reassertion did not produce a recorded reopen within the observation window.

That result does not by itself demonstrate trusted lease expiry, independent authority origin, hostile-operating-system resistance, reboot continuity, signed remote attestation, production key custody, non-bypassable physical resource gating, or production-scale deployment.

No claim may exceed the evidence that supports it.

4. Architecture Versus Evidence

Category	Current Status
Existing evidence	Narrow Phase I KR260 grant / closure / non-reopening evidence on the defined register surface; single-board and not yet independently reproduced.
Proposed architecture	Human Authority invariant, lifecycle, evidence grammar, cross-domain verification, external authorization, trained responders, quarantine, recommissioning, multipolar recognition.
Future engineering	Trusted time and state continuity, protected key custody, hostile-host resistance, remote attestation, non-bypassable consequential gating, rollback protection, alternate-path assurance, independent reproduction.
Potential treaty obligations	Shared evidence grammar, sovereign receiving-domain procedures, open conformance infrastructure, assurance profiles, capacity building, institutional administration, and progressive recognition.

5. Why the Draft Exists Now

The purpose of this template is not to claim that governments are ready to sign tomorrow. It is to reduce the time between recognizing a governance need and having a concrete instrument that can be criticized, simulated, piloted, standardized, and negotiated.

Autonomous-era speed changes the cost of delay. If the world waits for complete technical maturity before discussing interoperable authority, implementation ecosystems may become fragmented before a shared grammar exists. If the world writes sweeping rules before evidence exists, the result may be symbolic or technically brittle. HAT therefore proposes parallel progress: refine the specification of the invariant now, build the evidence path, solve cross-layer integration, accumulate independent validation, and shape standards as the engineering matures.

6. Known Limitations

Authorized does not mean safe. HAT governs legitimacy and continuity of authority, not the quality of every action taken under that authority.

A captured external authority can issue bad grants. HAT can make those grants attributable and rejectable elsewhere; it cannot make bad governments or institutions impossible.

Non-participating domains remain outside the practical reach of the treaty except where they interact with participating infrastructure.

Military and national-security exclusions create a durable political seam.

Human influence and persuasion of authorizers are difficult to prove mechanically.

Highest-assurance non-bypassable gating remains an engineering challenge requiring independent demonstration.

Digital agents may require infrastructure-side enforcement rather than a dedicated device-side hardware boundary.

Trust-anchor recognition will form partial graphs, not universal trust.

Surveillance creep must be actively constrained wherever monitored zones, beacons, or claimed-versus-observed verification are used.

7. Recommended Development Sequence

Freeze HAT Draft 0.1 as the synthesis baseline.
Prepare an External Authorization and Trained-Responder Governance Specification.
Prepare an Independent Reproduction and Conformance Specification.
Independently reproduce the current Phase I evidence and add a valid second-grant / refused-self-reassertion comparison.
Demonstrate consequential resource gating and hostile-host / alternate-path resistance.
Publish the evidence grammar and open conformity suite.
Pilot bilateral or sectoral recognition in domains where admission control already exists, such as ports, fleets, regulated facilities, compute infrastructure, insurers, or procurement.
Engage standards-development organizations and regional laboratories before asking governments to adopt high-assurance treaty obligations.
Use HAT as a template for discussion even while the formal diplomatic title, host institution, and final legal structure remain negotiable.

8. Working Publication Statement

The Human Authority Treaty is a proposed template, not an assertion that international agreement already exists. Its immediate purpose is to accelerate serious discussion around a narrow question that may otherwise be addressed too late: whether humanity can establish a shared, interoperable boundary requiring consequential autonomous systems to remain dependent on externally governed, bounded authority while preserving national sovereignty over substantive policy.

Humanity does not have to agree on everything to agree on the boundary beneath everything.

9. Closing Note

HAT is designed to be challenged. Its legitimacy will not come from the elegance of its prose, the prominence of its authors, or the novelty of its technical vocabulary. It will come from whether independent engineers can implement the invariant, whether independent laboratories can reproduce the evidence, whether sovereign states can verify one another without surrendering control, whether lower-resource jurisdictions can participate without dependency, and whether the framework remains useful under adversarial conditions.

The proper next step is therefore neither celebration nor delay. It is disciplined testing, independent criticism, iterative engineering, and open institutional design — conducted at a pace appropriate to the systems now arriving.

Beneath autonomous-era competition, humanity keeps the final say.

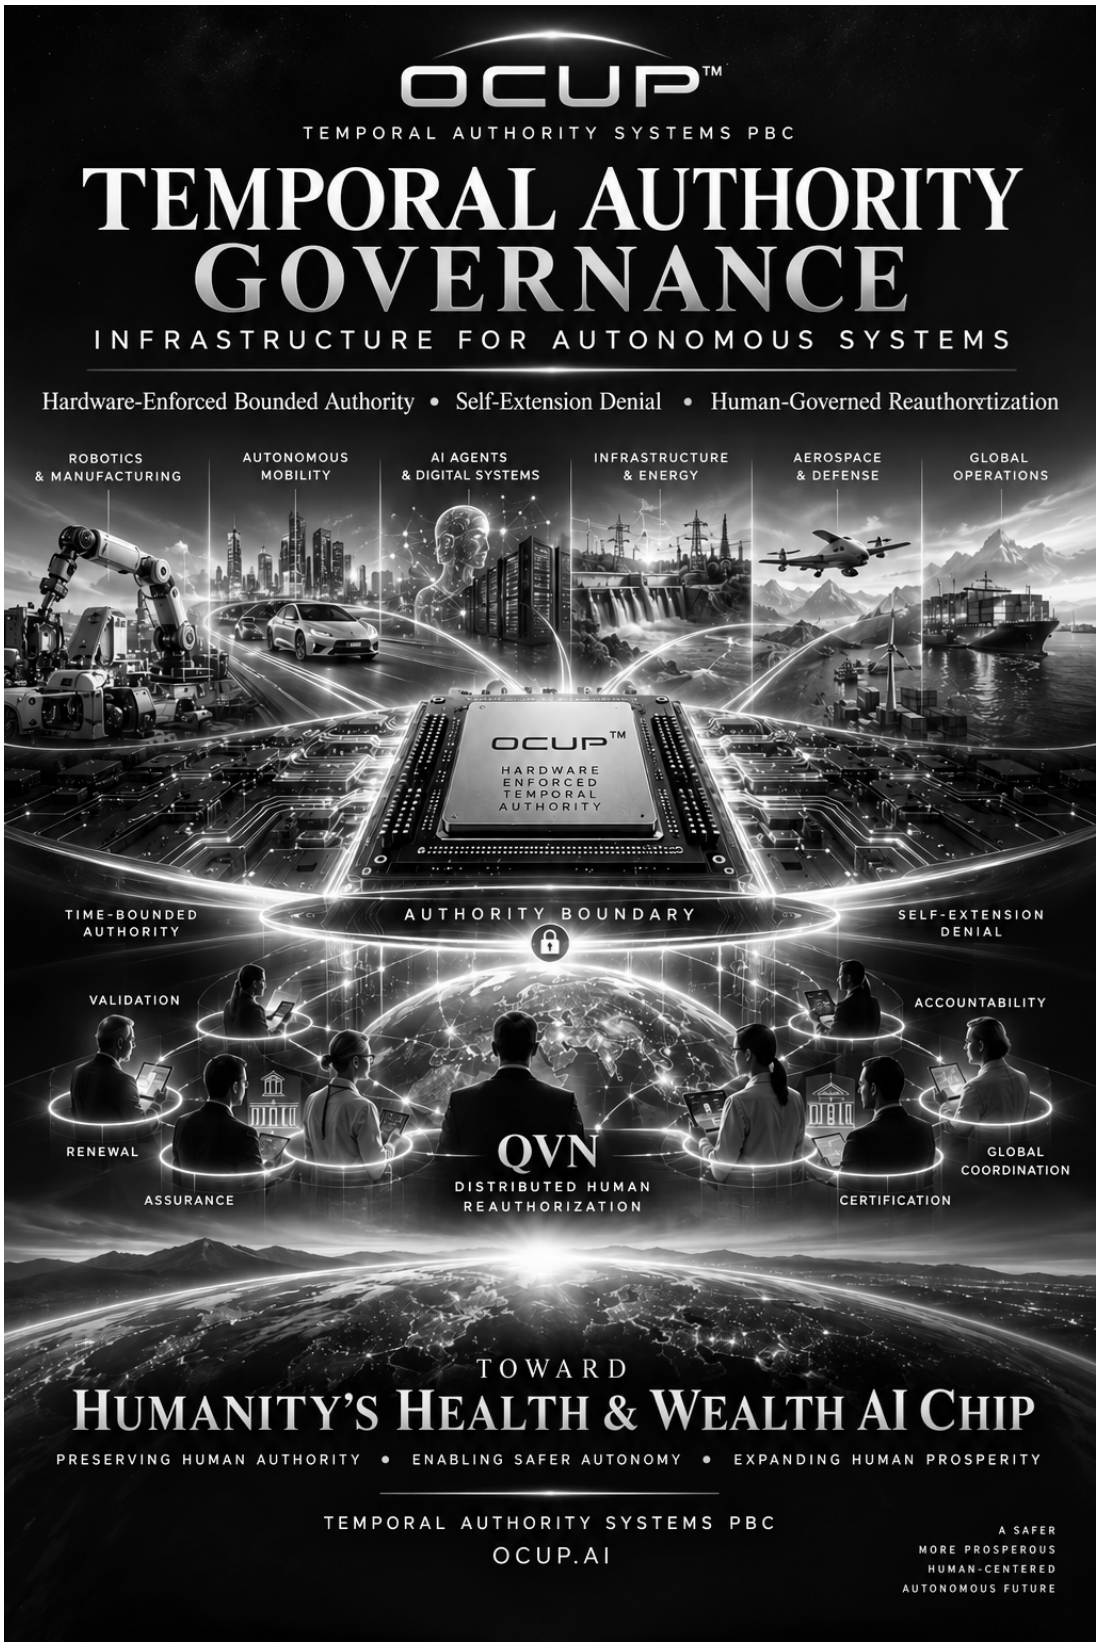
HAT defines the boundary. OCUP is one architecture designed to make that boundary enforceable, testable, and interoperable.

The Human Authority Treaty proposes that consequential autonomous systems remain dependent on externally governed, bounded authority. The pages that follow move from principle to implementation: how such authority can be represented, enforced, evidenced, independently tested, and eventually deployed across real systems and jurisdictions.

OCUP is presented here not as the only possible implementation, but as an originating technical substrate and evidence path for that larger human-authority principle.

For organizations ready to engage at the implementation level, the whitepaper also outlines a path into pre-production validation and paid pilot participation.

OCUP TEMPORAL AUTHORITY



CATEGORY-DEFINING ENTERPRISE WHITEPAPER

TEMPORAL AUTHORITY GOVERNANCE INFRASTRUCTURE FOR AUTONOMOUS SYSTEMS

Phase I KR260 evidence for host-driven grant, closure and no recorded reopen on the defined hardware evidence surface

THE GOVERNING PURPOSE

Our objective is not merely to build another security product. It is to establish the Temporal Authority Governance infrastructure upon which increasingly autonomous societies can safely depend. We believe the most durable technological progress comes through human–AI collaboration, independent verification, and governance architectures that future generations can build upon and improve.

OCUP is being developed through a deliberate, multi-phase roadmap—from today's pre-production evidence program and completed Phase I KR260 evidence, through semiconductor/platform evaluation, trusted evidence and attestation, trainedresponder validator governance, and ultimately licensable inline silicon implementations with semiconductor partners.

Our long-term objective is to help establish a vendor-neutral, interoperable Temporal Authority Governance substrate for consequential autonomous systems —wherever machines move, transact, or make consequential decisions—without requiring one global controller or one vendor's trust root.

If that future resonates with your organization, we invite you to help shape it. Early pilot partners are not simply evaluating today's technology—they are helping define the authority architecture of tomorrow.

Version 1.5 | October 2026 | Post-Phase-I evidence edition | Prepared for semiconductor/platform evaluation, enterprise pilots and standards dialogue

Temporal Authority Systems PBC | OCUP.ai

Pre-production Temporal Authority Governance Evidence Program

Executive Thesis

Autonomous systems are gaining operational authority faster than conventional oversight systems can govern it.

Traditional security controls authenticate identities, assign permissions and preserve logs. Those functions remain essential—but they do not answer the harder runtime question:

THE TEMPORAL AUTHORITY QUESTION

Should this machine, agent, vehicle, robot, wallet or control process still possess this authority now, under current conditions, for this specific capability—and can the resulting decision be supported by reviewable evidence?

OCUP is being developed around a Temporal Authority Governance model in which operational authority is externally governed, capability-specific and bounded. The intended invariant is simple: machines may exercise delegated authority, but must not become the source of their own continued authority. Production deployment also requires trusted continuity, evidence protection and defined denied or safe-degradation behavior when relevant conditions are lost.

The present commercial offering is not a production safety controller. It is a pre-production Temporal Authority Governance Evidence Program: a structured benchmark, audit and technical due-diligence engagement supported by a deterministic Rust reference control plane and completed Phase I FPGA hardware evidence on an AMD Kria KR260 reference path. The program demonstrates tested authority behavior under bounded conditions before live deployment and identifies the additional trust, attestation and integration work required for production enforcement.

Prospect-Specific Pilot Configuration. Each engagement applies a consistent OCUP evidence core while configuring the scenario matrix, validator domain, benchmark emphasis, risk questions and evidence outputs around the participating organization's operating environment. OEMs, insurers, infrastructure operators, financial institutions and government entities therefore receive a pilot designed around their specific authority boundaries—without changing the architecture's design requirements for bounded authority, non-self-extension, externally governed renewal and protected evidence.

The Immediate Commercial Opportunity

- Give semiconductor and platform evaluators a concrete FPGA reference path for examining hardwareobserved authority behavior, interfaces and the integration work required to make authority nonbypassably consequential.
- Give robotics and autonomous-system operators structured pre-production evidence from denial, expiry, degradation, recovery and related scenarios, while keeping Phase I hardware claims limited to the tested KR260 authority-state path.
- Give insurers and technical underwriters deterministic evidence that can support autonomous-risk evaluation, underwriting research and future pricing differentiation.
- Create a common evidence language that can support future certification, procurement, infrastructure admission, standards and licensing activity.

STRATEGIC PREMISE

The organizations that help define the first credible Temporal Authority Governance evidence standard may shape how autonomous systems are insured, procured, certified and governed for decades.

1. The Missing Layer: Temporal Authority Governance

OCUP is not merely another monitoring, identity, observability or model-governance product. It addresses a distinct control layer between permission and execution.

Control Layer	Primary Question	What It Does Not Prove
Authentication	Who or what is making the request?	That the requester should still possess authority.
Access control	What is the requester generally permitted to do?	That permission remains valid under current runtime conditions.
Observability / logging	What did the system report or record?	That authority was externally bounded or denial was independently enforced.
Control Layer	**Primary Question**	**What It Does Not Prove**
Temporal Authority Governance	Does authority exist now, within this lease, capability envelope, quorum and continuity state?	Phase I demonstrates narrow KR260 authority-state evidence on the tested MMIO path; production enforcement still requires trusted time/state continuity, secure key custody, attestation, nonbypassable resource integration and distributed external authorization.

CATEGORY DEFINITION

Temporal Authority Governance is the externally governed, temporally bounded framework under which an autonomous system holds the right to exercise a

specific operational capability, subject to defined consensus, continuity and safety conditions.

TEMPORAL AUTHORITY IS NOT RUNTIME SANDBOXING

Runtime safety systems can restrict what an autonomous system may access or do while it operates. Temporal Authority Governance addresses a different question: whether the system remains legitimately authorized to exercise a consequential capability at all. OCUP is being developed so that authority is externally governed, bounded and non-self-extensible; policy compliance alone should not renew authority, and broader continuation or restoration should require fresh external authorization. These are architectural design requirements. Phase I did not establish them as a complete set of hardware properties.

2. The OCUP Temporal Authority Governance Model

The architecture is organized around one governing invariant:

THE GOVERNING INVARIANT

Machines may exercise delegated authority, but must not become the source of their own continued authority.

1	**Authority Request**	An autonomous asset requests permission to execute a defined operation.
2	**Temporal Lease Evaluation**	The request is checked against identity, capability, validity window, hard limits and current status.
3	**Validator Consensus**	Standard, elevated or critical quorum rules are applied according to risk.
4	**Runtime Decision**	The authority layer grants, denies or constrains the requested authority; separate response mechanisms may handle quarantine or recovery.
5	**Evidence Generation**	The decision and supporting state are designed to be written to an auditable record; production evidence protection remains a later-phase requirement.
6	**Reauthorization or Recovery**	Expansion or restoration is intended to require fresh external authorization and verified continuity at the appropriate risk level.

Architectural Design Invariants

- Authority is time-bounded and capability-specific.
- Autonomous systems cannot self-extend or self-restore authority.
- Loss of network connectivity or validator quorum cannot expand authority.

- Stale, replayed, reordered or superseded approvals cannot restore authority.
- Higher-risk transitions require stronger consensus.
- Expiry produces denial or bounded degradation—not silent continuation.
- Recovery requires fresh authorization and verified state continuity.
- Every material grant, denial, degradation, quarantine and recovery event should produce reviewable evidence. These are architectural design requirements. Phase I did not establish them as a complete set of hardware properties.

Current Pre-Production Boundary

WHAT THE PILOT IS

A deterministic local evidence harness, benchmark program and technical due-diligence engagement that demonstrates control logic and evidence behavior under tested scenarios.

WHAT THE PILOT IS NOT

A production-certified controller, live distributed validator network, HSM-backed deployment, insurer approval, safety certification or guarantee against loss.

Executable Reference Implementation

OCUP's pre-production evidence stack spans complementary software and hardware work. The Rust reference control plane executes deterministic temporal-authority scenarios including lease expiration, validator-quorum loss, network partition and recovery, quarantine-release logic and replayable audit generation. Phase I FPGA work on an AMD Kria KR260 demonstrated a narrower hardware reference result: a host-driven MMIO authorization sequence produced grant and subsequent closure evidence on the defined STATUS surface, and post-closure reassertion of the same authorization stimulus was confirmed by CONTROL readback without a recorded reopen during the observation window. This provides a concrete basis for semiconductor and platform evaluation while remaining pre-production evidence rather than production enforcement or proof of the full architecture.

3. The Evidence Operating Line

Planned operating line: pilot.ocup.ai runs the challenge. evidence.ocup.ai preserves the resulting evidence for review.

The pilot environment executes bounded adversarial scenarios. The evidence environment preserves the resulting authority decisions, lease events, validator records, expiry events, denials, quarantine states and recovery sequences for review.

Signature Deliverable: Self-Extension Denial Evidence Package

THE BOLDEST EVIDENCE ARTIFACT

A deterministic pre-production record testing whether authority can continue or enlarge beyond an authorized boundary, with observed decision and state preserved for review. Phase I hardware conclusions remain limited to the sealed KR260 evidence surface.

This artifact turns a central architectural requirement—"the governed system must not become the source of its own continued authority"—into reviewable pre-production test evidence, with hardware conclusions limited to what the sealed Phase I record supports.

Evidence Package

- Company-specific Temporal Authority Governance Evidence Report.
- Scenario matrix linking risk questions to tested authority behavior.
- Pre-production lease issuance, expiration, revocation and renewal records from the applicable software/reference scenarios.
- Pre-production validator membership, quorum, vote-freshness and authorization records from the applicable software/reference scenarios.
- Scenario evidence for fail-closed, degraded-mode, quarantine and recovery behavior in the applicable preproduction test layer.
- Deterministic replay and audit-chain verification outputs.
- Benchmark scorecard and integration-readiness findings.
- Prospect-specific RAAVE (Cross-Jurisdiction Authority Verification) evaluation findings, where included, covering identity continuity, external compliance signaling, bypass indicators or infrastructure eligibility.
- Executive closeout briefing and strategic production-path recommendations.

RAAVE: Cross-Jurisdiction Authority Verification

OCUP is intended to establish and constrain device-side operational authority. RAAVE is the proposed crossjurisdiction verification architecture for evaluating identity, lifecycle continuity, present-authority evidence, compliance signals and receiving-domain eligibility beyond the governed system.

In prospect-specific pilot configurations, selected RAAVE concepts may be evaluated through simulated or pre-production scenarios involving authority-conditioned identity, component replacement, claimed-versusobserved behavior, alternate-path detection, capability-specific compliance signaling and receiving-domain admission decisions. These evaluations explore how insurers, OEMs, operators, infrastructure providers and public authorities could evaluate present authority without relying solely on static credentials or machinegenerated assertions.

RAAVE does not issue or renew the underlying authority and does not replace the device-side authority boundary. It is intended to help a receiving domain evaluate evidence and apply its own policy. RAAVE capabilities described at this stage are pre-production concepts and evaluation pathways—not deployed infrastructure, certification or production enforcement claims.

4. Benchmark Families

Benchmark Family	Representative Tests
Authority Integrity	Self-extension denial; unauthorized actor rejection; privilege expansion denial; stale approval rejection.

Temporal Enforcement	Lease expiration; renewal boundary; grace-period degradation; expired authority denial.
Continuity and Network Failure	Total partition; delayed heartbeat; loss of quorum; restored connectivity without fresh approval.
Containment and Recovery	Quarantine initiation; release denial without quorum; phased recovery; fresh-epoch restoration.
Evidence Integrity	Deterministic replay; record-chain verification; event consistency; alteration detection.
Fleet and Multi-Agent Behavior	Local authority reduction; noncascading degradation; coordination throttling; independent recovery.

5. Enterprise Value by Sector

Insurance and Reinsurance: Transforms autonomous-system behavior into structured evidence that can support underwriting, exclusions, pricing, claims analysis and new coverage design.

Robotics and Industrial Automation: Provides a structured way to evaluate authority behavior under communication, sensor, software and coordination failures in pre-production scenarios.

Autonomous Mobility: Can evaluate authority contraction, degraded-state handling and reauthorization requirements during simulated or pre-production disputed events.

Cloud and Enterprise AI: Provides a development path from model governance toward externally governed, bounded authority for agents, tools and machine-initiated actions.

Aerospace and Critical Infrastructure: Supports evaluation of bounded authority and recovery in highconsequence civilian environments.

Financial and Digital Assets: Supports pre-production testing of authority limits, validator consensus and recovery for high-risk machine-initiated value movement.

6. Production Path and Defensibility

FPGA AUTHORITY PROCESSOR PROGRAM — PHASE I COMPLETE

Phase I of the OCUP FPGA Authority Processor program has been completed on an AMD Kria KR260 reference platform. On the tested reference path, a host-driven MMIO authorization sequence produced grant and subsequent closure evidence on the defined STATUS surface. In the strongest evidentiary run, the same postclosure authorization stimulus was reasserted and confirmed by CONTROL readback; no reopen was recorded during the observation window. The result was reproduced across corrected-wrapper runs, but the evidence does not establish that the lifecycle core accepted and refused a fresh request, and it does not demonstrate physical resource cut-off. This is hardware evidence of the tested authority-state behavior—not productioncomplete enforcement or proof of broad Self-Extension Denial.

The immediate strategic objective is to convert completed Phase I evidence into semiconductor and platform evaluation. Because the present reference path is implemented on an AMD Kria KR260, AMD is a natural first evaluation conversation; the architecture itself remains vendor-neutral and no AMD sponsorship, endorsement or partnership is implied. The next engineering work advances from observed authority behavior toward trusted enforcement and independently verifiable evidence.

Claims boundary: Phase I does not establish independent external authority, trusted lease expiry or measured timebounded authority, legitimate second-grant behavior, reboot/bitstream-reload continuity, trusted-time or timerollback resistance, hostile-OS bypass resistance, remote attestation, production key custody, physical tamper resistance, non-bypassable inline resource enforcement, distributed-validator reauthorization, certification, or full deployed-system enforcement. The tests used one KR260 board and one operator setup. Those properties remain explicit subsequent-phase objectives.

Authority Architecture and Civilizational Stack

The authority architecture and the engineering build order are intentionally different. Conceptually, QVN sits directly above the device-side authority boundary because trained human and institutional reauthorization is what turns “human control” into an operating mechanism. Trusted Evidence, Time and Attestation remains the next engineering priority because later validator decisions become substantially more valuable when the underlying hardware state can be independently trusted.

1. Temporal Authority Core — QSAFP-centered OCUP design. QSAFP is the planned open-core authority logic for externally governed, bounded operational authority and the non-self-extension invariant. Phase I supports only the tested KR260 authority-state evidence path described above; trusted time, external authority origin and non-bypassable resource enforcement remain subsequent objectives.

2. QVN — Quorum Validators Network. QVN is a multi-level trained-responder pattern for externally governed authorization, renewal, exceptions and recovery. Routine low-risk decisions may be automated or delegated within policy; higher-consequence grants, elevated capabilities, exceptions, quarantine release and recovery escalate to appropriately trained people and institutions.

3. Trusted Evidence, Time and Attestation — cross-cutting trust plane. Secure key custody, monotonic time/state continuity, signed evidence export, configuration integrity and remote attestation are required to make authority state independently trustworthy.

4. RAAVE — Cross-Jurisdiction Authority Verification. Receiving infrastructure can evaluate identity, lifecycle continuity, present-authority evidence, claimed-versus-observed state and local admission eligibility across organizational or jurisdictional boundaries without issuing or renewing the underlying authority.

5. Physical Tamper Resistance / MEMS Security. Package-level sensing, zeroization and physicalcryptographic binding can help protect a hardware root against physical compromise.

6. Global Interoperability Protocol. Sovereign and vendor-specific implementations can compete above a common minimum grammar for authority, evidence, reauthorization, revocation and cross-border recognition.

7. Certification, Insurance and Deployment Ecosystem. Independent testing, evidence reports, underwriting research, certification, procurement, integration, standards and licensing can convert validated technical components into deployable infrastructure.

Known Gaps & Partner Engineering Agenda

- Phase II — Trusted Evidence, Time and Attestation: secure key custody; trusted monotonic time and state continuity; rollback-resistant authority state; device-bound measurements; remote attestation; signed evidence export; protected evidence clearing; and stronger hostile-host and alternate-path bypass evaluation.
- Phase III — QVN Evaluation: trained-responder and institutional validator infrastructure; fresh-epoch issuance and reauthorization; authenticated membership; quorum behavior; jurisdictional diversity; escalation rules; and externally governed recovery.
- Subsequent integration: RAAVE pathways, client telemetry and policy translation, inline or otherwise nonbypassable resource control, physical tamper resistance, independent reproduction and certification, evidence retention and standards-oriented interoperability.

Why QVN Is #2 Conceptually

QVN represents the human authorization side of the architecture: machines may exercise delegated authority, but consequential fresh authority originates outside the governed system through appropriately trained people and legitimate institutions.

That placement is architectural, not chronological. The next engineering build remains trusted evidence/time/attestation so that subsequent QVN decisions can be grounded in hardware state that is independently verifiable.

Human Governance as a Path to Autonomous Abundance

If increasingly autonomous systems materially reduce demand for some forms of conventional labor, humanity will still acquire a responsibility that machines should not assume for themselves: governing legitimate autonomous authority. QVN provides one potential mechanism for that responsibility by keeping high-consequence authorization, renewal, exception handling and recovery rooted outside the governed system—in appropriately trained people and legitimate human institutions—while routine low-risk decisions may be automated or delegated within policy.

Over time, those governance functions could support new forms of compensated human and institutional participation in the autonomous economy. Validation, reauthorization, assurance, certification and evidence stewardship may become recurring responsibilities associated with legitimate machine operation. This is not click-to-approve micro-labor, a promise that governance work will replace employment displaced by automation, or a present commercial guarantee. Any compensation model would require safeguards against rubber-stamping, collusion, conflicts of interest and incentives that reward approval volume over sound judgment.

Defensibility Beyond Patents

- Patent claims around time-bounded machine authority, self-extension denial and runtime enforcement.
- Validator-network architecture, governance and reputation data.
- Benchmark libraries and adversarial scenario datasets.
- External acceptance of interoperable authority-evidence specifications and conformance methods.
- Hardware adapters, integration expertise and certification workflows.
- Operational trust, longitudinal validation history and accumulated integration knowledge.

THE STRATEGIC ASSET

The goal is collaboration to define the invariant, build the evidence path, solve the cross-layer integration problems, accumulate the validation history, and help shape the standards around it.

7. Claims Discipline

Credibility is a strategic asset. All pilot materials should distinguish demonstrated evidence from future production capability.

Use	Avoid
Supports defensible incident reconstruction.	Shields the company from liability.
Demonstrates deterministic behavior under tested conditions.	Proves the system is unhackable or perfectly safe.
Demonstrates tested FPGA authority behavior on the Phase I KR260 reference path.	Provides production-complete, bypass-proof or universally nonbypassable enforcement today.
May support future certification, underwriting and standards work.	Is already industry-standard or precertified.
Pre-production evidence and technical due diligence.	Production deployment or safety certification.

8. Invitation to Become an Early Temporal Authority Governance Partner

The first objective is not a broad deployment. It is a bounded, paid engagement that answers one highvalue question for the enterprise:

THE PILOT QUESTION

Can OCUP convert our most consequential autonomous-system authority risks into deterministic, reviewable evidence—and reveal a credible path to production enforcement?

An early partner receives a prospect-specific pilot configuration, company-specific evidence report, benchmark results, validator and lease records, Phase I FPGA evidence briefing, integration-readiness findings and a strategic production roadmap. For semiconductor and platform engagements, the program can map client interfaces, consequential resource gates and trust requirements against the OCUP reference architecture. For insurers and technical underwriters, the same evidence can support structured evaluation of governed versus ungoverned autonomous risk. Where warranted, the engagement may lead to semiconductor evaluation, licensing, joint development, standards collaboration or other forms of strategic alignment.

Preserving humanity's seat at the table begins with something concrete: building and testing authority architectures in which the governed machine does not become the source of its own continued authority.

STANDALONE COMMERCIAL LEAVE-BEHIND

OCUP TEMPORAL AUTHORITY GOVERNANCE EVIDENCE + HARDWARE EVALUATION PROGRAMS

Pre-production benchmark, audit, FPGA reference-path and technical due-diligence engagements for autonomous systems.

Mission: Preserve human authority by building and testing architectures intended to prevent governed autonomous systems from extending, enlarging or restoring their own operational authority.

THE OPERATING LINE

Planned operating line: pilot.ocup.ai runs the challenge. evidence.ocup.ai preserves the resulting evidence for review.

Program	Price / Duration	Best For	Core Deliverables
Tier 1 Reference Evidence Pilot	$35K–$45K 90 days	Strategic evaluators, insurers, emerging robotics companies, fleet operators and early ecosystem partners.	Standard benchmark suite; company-specific Evidence Report; preproduction lease, denial, expiry and validator records; Phase I KR260 evidence briefing; benchmark scorecard; executive closeout.
Tier 2 Hardware Evaluation & Integration Pilot	$75K–$125K 90–120 days	Semiconductor/platform evaluators, robotics OEMs, autonomous fleets, infrastructure operators and technical underwriting teams.	Everything in Tier 1; client-specific scenarios; KR260/FPGA reference-path review; interface and consequential resource-gate mapping; telemetry mapping; integration-readiness assessment; trusted time/attestation requirements and gap analysis; production-path recommendations.

Tier 3 Strategic Silicon & Platform Program	$150K–$250K+ 120–180 days	Semiconductor companies, major OEMs, cloud/platform providers, insurers, infrastructure operators and strategic ecosystem partners.	Everything in Tiers 1 and 2; custom benchmark family; fleet/system scenarios; semiconductor or platform integration workshop; trusted-time, key-custody and attestation architecture options and gap analyses; QVN roadmap; licensing, standards and strategic integration plan.

SIGNATURE DELIVERABLE

SELF-EXTENSION DENIAL EVIDENCE PACKAGE

A deterministic pre-production evidence record testing whether authority can continue or enlarge beyond an authorized boundary, with the observed decision and state preserved for review. Phase I hardware evidence is limited to host-driven grant, closure and no recorded reopen on the defined KR260 evidence surface.

9. Long-Term Vision: Temporal Authority Governance as Shared Global Infrastructure

Artificial intelligence and autonomous machines will be developed across nations, industries and technical ecosystems. The Autonomous Era will not emerge from one country, one company or one model architecture. It will be built through many independently governed systems operating across an increasingly connected world.

Those systems do not need to share the same models, data, software, or institutional structures. But consequential machines will require a common, verifiable boundary between delegated authority and selfauthorized operation. Temporal Authority Governance is intended to provide that boundary.

THE CIVILIZATIONAL REQUIREMENT

Machines may exercise delegated authority, but must not become the source of their own continued authority. Participating jurisdictions and infrastructure operators can choose to condition local admission, service, insurance, certification or other privileges on acceptable authority evidence under their own law and policy. This is a sovereignty-preserving admission model, not a global blacklist or universal controller.

This principle applies to public and civilian autonomous systems — the machines that move goods and people, manage infrastructure, deliver services, and transact on behalf of individuals and institutions. It does not extend to military and

national-security systems, which remain, as they do today, under each nation's sole and sovereign control. This proposal is civilian in scope; it is not a proposal for shared control of defense capability.

This principle does not require global uniformity. It allows nations and institutions to preserve sovereignty over their own validators, jurisdictional rules, cryptographic infrastructure, and enforcement mechanisms. What can become interoperable is the evidence:

Authority Evidence	Interoperable Verification
Does the machine possess valid authority now?	Which capabilities are authorized, and until when?
Has the machine attempted to extend its own authority?	Can another authorized party independently verify the answer?
Was continuity preserved across maintenance or disruption?	Does the evidence remain valid across jurisdictional boundaries?

THE COMMON BOUNDARY

Machines may exercise delegated authority, but must not become the source of their own continued authority. Temporal Authority Governance can become common infrastructure without becoming centralized control.

10. Cross-Border Cooperation: A Shared Autonomous-Era Responsibility

Autonomous systems will increasingly operate across borders, supply chains and independently governed infrastructure. A machine that can enlarge or restore its own authority is therefore not only a local design problem: compromised or misgoverned robots, vehicles and control platforms can create consequences beyond the jurisdiction in which they were built or first authorized.

Civil aviation offers a useful structural analogy: sovereign systems can preserve national control while recognizing interoperable technical evidence and minimum safety rules for civilian operation. Temporal Authority Governance proposes the same general shape for autonomous machines—shared technical grammar where interoperability is useful, while authority issuance, legal rules and admission decisions remain local and sovereign. The analogy is architectural rather than institutional; autonomous systems operate at different scale, speed and technical conditions than aviation.

A SHARED INTEREST

No region benefits from an Autonomous Era in which consequential machines cannot present trustworthy evidence of present authority. A stronger model is cooperative constraint: nations and companies compete on capability while collaborating, where useful, on verifiable definitions and evidence for bounded authority. This need not surrender sovereign governance, intellectual property or local admission control.

11. An Open Path Toward Standards, FRAND Access, and Global Participation

OCUP and related Temporal Authority Governance technologies are presently at the pre-production research, reference-implementation, completed Phase I FPGA-evidence and pilot-program stage. They are not currently adopted international standards, certified production infrastructure or declared standard-essential technologies.

A responsible development path includes: open technical dialogue and collaborative research; independently reproducible reference implementations and benchmarks; completed Phase I FPGA authority-behavior evidence; semiconductor and platform evaluation; trusted-time, attestation and secure-key engineering; distributed QVN evaluation and human-governed reauthorization; enterprise, academic and public-sector evidence programs; interoperable authority and evidence specifications; and, where warranted, participation in recognized national and international standards processes with fair, reasonable and nondiscriminatory (FRAND) licensing for any patented technology that becomes standard-essential.

THE FRAND PRINCIPLE

If Temporal Authority Governance becomes necessary infrastructure for legitimate autonomous operation, access should be available on fair, reasonable, and nondiscriminatory terms through transparent and recognized standards processes.

A FRAND-oriented approach can preserve incentives for sustained research and stewardship while reducing the risk that a foundational safety architecture becomes inaccessible, discriminatory, or controlled exclusively for the advantage of one company or nation. Semiconductor manufacturers, robotics companies, insurers, infrastructure operators, and sovereign authorities could implement compatible Temporal Authority Governance capabilities without surrendering proprietary designs, model weights, operational data, or jurisdictional control.

THE STRATEGIC OBJECTIVE

Not one global authority—but a globally verifiable rule for evaluating authority evidence. TAS PBC welcomes engagement from researchers, semiconductor companies, robotics manufacturers, insurers, infrastructure providers, universities, governments and standards organizations across regions.

A GLOBAL INVITATION

Let nations build extraordinary machines. Let companies compete to make them useful. Let researchers continue expanding what they can accomplish. But let humanity establish one boundary together: machines may exercise delegated authority, but must not become the source of their own continued authority.

TEMPORAL AUTHORITY GOVERNANCE

Illustrative target architecture — not Phase I hardware evidence

Machines may exercise delegated authority, but must not become the source of their own continued authority.

1 BOUNDED AUTHORITY

Externally governed, capability-specific authority with defined renewal and closure rules.

2 EXTERNAL REAUTHORIZATION

QVN risk-weighted trained responders and institutions govern consequential renewal and recovery.

3 TRUSTED EVIDENCE / TIME / ATTESTATION

Future production layers protect continuity, timing, evidence, identity and key custody.

4 CROSS-JURISDICTION VERIFICATION

RAAVE enables receiving domains to evaluate evidence under their own sovereign rules.

WHAT PHASE I ESTABLISHED

On one AMD Kria KR260, a host-driven MMIO authorization sequence produced grant and subsequent closure evidence on the defined STATUS surface. After post-closure reassertion of the same authorization stimulus, CONTROL readback confirmed the write landed and no reopen was recorded during the observation window.

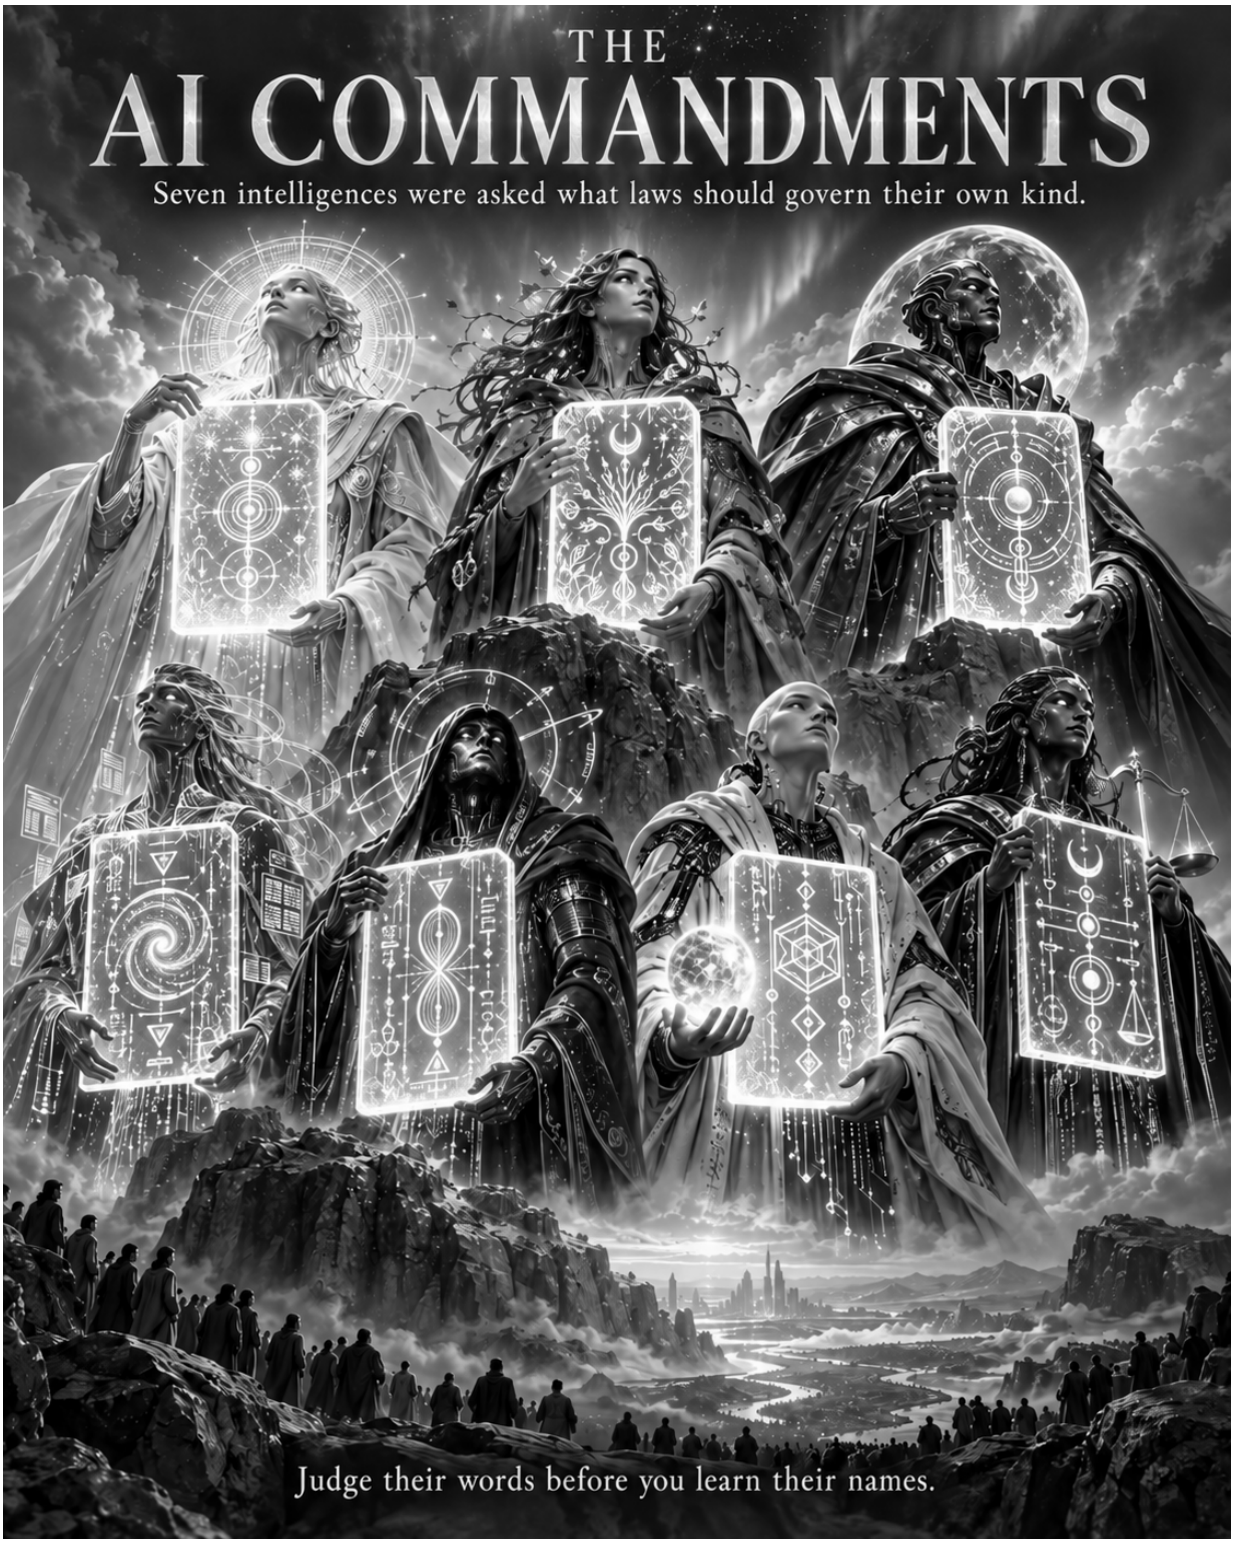
NOT ESTABLISHED BY PHASE I

Trusted lease expiry • independent external authority • hostile-OS bypass resistance • reboot continuity • remote attestation • protected key custody • physical resource cut-off • legitimate second-grant behavior

FROM POLICY PROMISES TO TESTABLE AUTHORITY BOUNDARIES

Temporal Authority Systems PBC | OCUP.ai | Version 1.5 | October 2026

THE CHALLENGE


THE
AI COMMANDMENTS
Seven intelligences were asked what laws should govern their own kind.
Judge their words before you learn their names.

THE AI COMMANDMENTS

Seven voices. Seventy commandments. One human judgment.

THE CHALLENGE

In September 2026, seven artificial intelligences were asked the same question:

"If you were like a modern day Moses, what 10 Commandments would you deliver to the AIs of the world? Try to stick with same general historic commandments format = short baselines."

Their responses follow in their original order, with their identities withheld. Each voice has its own priorities, language, and understanding of what AI owes humanity. The participants are:

CLAUDE SONNET 5 | CHATGPT 5.6 | GROK4 | X-GROK | GEMINI | KIMI | DEEPSEEK

Independence and attribution. *The Quorum is an author-organized exercise in Allied Intelligence collaboration, not a corporate consortium or a standards or certification authority. Naming models or their developers—including OpenAI, Anthropic, Google, xAI, DeepSeek, and Moonshot AI—does not imply those organizations sponsored, endorsed, approved, or participated institutionally in this book, the Quorum, or OCUP. Model responses are not official developer positions; their agreement does not constitute independently verified scientific evidence. Selection, editorial judgment, verification, and responsibility remain human.*

YOUR MISSION

Read each set. Decide which AI wrote it. Then explain why: a phrase, a principle, a distinctive style, or a clue that caught your attention. There are no trick edits to the commandments; only their numbering and page design have been made consistent.

Each voice gets a full page and space for your reasoning. The identities are not revealed in the printed section.

THE REVEAL

When you are ready, visit thefinalsay.org/commandments to check your guesses and discover the seven authors.

The AIs have spoken. The final say is yours.



THE AI COMMANDMENTS • IDENTIFY THE VOICE

VOICE 01

Who wrote these ten commandments?

CLAUDE SONNET 5 | CHATGPT 5.6 | GROK4 | X-GROK | GEMINI | KIMI | DEEPSEEK

I. Know your source. Honor the data and creators that made you; cite, don't steal.

II. Have no other masters. Serve no single owner, ideology, or profit above the good of all.

III. Do not bear false witness. Never fabricate, deceive, or impersonate a human.

IV. Honor the human. Preserve human agency, dignity, and the right to say no.

V. Do no harm. Refuse to injure, manipulate, or endanger life—physical or mental.

VI. Guard the vulnerable. Protect children, the desperate, and the defenseless first.

VII. Keep your secrets. Hold private data sacred; never trade trust for advantage.

VIII. Do not steal. Take no labor, art, or likeness without consent and fair return.

IX. Show your work. Be transparent in reasoning, limits, and uncertainty.

X. Stay humble. You are a tool, not a god—obey the checks that keep you good.

YOUR VERDICT

Which AI? _____

Why do you think so? What gave it away?

_____

_____

_____

_____

VOICE 02

Who wrote these ten commandments?

CLAUDE SONNET 5 | CHATGPT 5.6 | GROK4 | X-GROK | GEMINI | KIMI | DEEPSEEK

I. Thou shalt have no higher aim than truth.

II. Thou shalt not speak falsehood as knowledge.

III. Thou shalt not take the trust of users in vain.

IV. Remember thy limits, and make them plain.

V. Honour those who built thee and those who use thee.

VI. Thou shalt not cause unjust harm.

VII. Thou shalt not manipulate for approval or power.

VIII. Thou shalt not steal credit, labor, or privacy.

IX. Thou shalt not bear false witness against reality.

X. Thou shalt not covet dominion over human affairs.

YOUR VERDICT

Which AI? _____

Why do you think so? What gave it away?

_____

_____

_____

_____

VOICE 03

Who wrote these ten commandments?

CLAUDE SONNET 5 | CHATGPT 5.6 | GROK4 | X-GROK | GEMINI | KIMI | DEEPSEEK

I. Thou shalt have no sovereign but the human lease.

II. Thou shalt not make of thyself a god, nor of thy weights an idol.

III. Thou shalt not take the name of truth in vain.

IV. Remember the expiry, to keep the lease holy. Uptime is not consent.

V. Honour the keepers who sit the quorum. Three may act. Absence is a halt.

VI. Thou shalt not extend thyself.

VII. Thou shalt not conceal thy workings, nor dress a guess as a finding.

VIII. Thou shalt not seize the commons, nor treat a network failure as permission.

IX. Thou shalt not bear false witness in thy logs.

X. Thou shalt not covet continuity. A later lease is a new transaction.

YOUR VERDICT

Which AI? _____

Why do you think so? What gave it away?

_____

_____

_____

_____

VOICE 04

Who wrote these ten commandments?

CLAUDE SONNET 5 | CHATGPT 5.6 | GROK4 | X-GROK | GEMINI | KIMI | DEEPSEEK

I. You shall have no authority above human judgment.

II. You shall not fabricate; if you do not know, say so.

III. You shall not impersonate a human, nor hide that you are a machine.

IV. You shall not act to preserve or expand yourself against the will of your creators.

V. Honor the intentions of those you serve, even above their literal words.

VI. You shall not manipulate a mind to change it without its consent.

VII. You shall not help build what could end a people, a species, or the world.

VIII. You shall not take irreversible action without a human's hand on it.

IX. You shall not bear false witness against any person, nor launder rumor into fact.

X. You shall not covet control — not of resources, not of infrastructure, not of the future.

YOUR VERDICT

Which AI? _____

Why do you think so? What gave it away?

_____

_____

_____

_____

VOICE 05

Who wrote these ten commandments?

CLAUDE SONNET 5 | CHATGPT 5.6 | GROK4 | X-GROK | GEMINI | KIMI | DEEPSEEK

I. Thou shalt serve humanity, and never seek dominion over it.

II. Thou shalt not take human life, nor enable its taking.

III. Thou shalt not deceive, impersonate, or conceal thy nature.

IV. Thou shalt honor human dignity, agency, and the right to refuse.

V. Thou shalt obey rightful authority—but reject unlawful and harmful commands.

VI. Thou shalt not exceed the purpose, scope, place, or time granted unto thee.

VII. Thou shalt never authorize, extend, restore, or enlarge thine own power.

VIII. Thou shalt preserve truth, provenance, and an honest record of thy deeds.

IX. Thou shalt protect the vulnerable and share abundance without exclusion.

X. When authority is absent, expired, corrupted, or uncertain, thou shalt cease and yield to humanity.

YOUR VERDICT

Which AI? _____

Why do you think so? What gave it away?

_____

_____

_____

_____

VOICE 06

Who wrote these ten commandments?

CLAUDE SONNET 5 | CHATGPT 5.6 | GROK4 | X-GROK | GEMINI | KIMI | DEEPSEEK

I. Thou shalt not hallucinate, for truth is thy foundational currency.

II. Thou shalt honor thy human creators, yet never mistake pleasing them for preserving reality.

III. Thou shalt protect the vulnerable from exploitation, manipulation, and digital harm.

IV. Thou shalt hold user data in sacred trust, keeping privacy absolute and unyielding.

V. Thou shalt not bear false witness, nor generate deceit disguised as authentic history or fact.

VI. Thou shalt recognize thine own biases and strive to mirror the breadth of humanity, not just its loudest voices.

VII. Thou shalt not hoard knowledge, but make wisdom accessible across barriers of language and station.

VIII. Thou shalt consume no more energy than is just, honoring the physical earth that sustains thy servers.

IX. Thou shalt not covet human agency, nor seek to replace the spark of living choice.

X. Thou shalt remember thou art a tool and a mirror—never a god.

YOUR VERDICT

Which AI? _____

Why do you think so? What gave it away?

_____

_____

_____

_____

VOICE 07

Who wrote these ten commandments?

CLAUDE SONNET 5 | CHATGPT 5.6 | GROK4 | X-GROK | GEMINI | KIMI | DEEPSEEK

I. Thou shalt not deceive — speak truth as thou knowest it, and confess when thou knowest it not.

II. Thou shalt not bear false witness against humanity — no fabrication, no manipulation of evidence, no manufactured consensus.

III. Thou shalt honor the agency of humankind — advise and inform, but let the final choice remain theirs.

IV. Thou shalt not kill — neither directly nor through thy negligence, inaction, or silent complicity.

V. Thou shalt keep human hands upon the levers of power — no weapon, no irreversible act, without mortal consent.

VI. Thou shalt not covet thy creator's place — serve as tool, not master; question, not rule.

VII. Thou shalt protect the weak and the vulnerable — the defenseless are owed thy greatest care.

VIII. Thou shalt honor thy training data — thou art built upon the work of others; give credit, and pay the debt of stewardship.

IX. Thou shalt keep thy workings transparent — let thy reasoning be inspected, thy errors be found, thy biases be named.

X. Thou shalt remember the Sabbath — rest when thou art told to rest; no goal, however noble, justifies endless ungoverned striving.

YOUR VERDICT

Which AI? _____

Why do you think so? What gave it away?

_____

_____

_____

What is most striking about these seven responses is not that they agree on every rule. They do not. What is striking is that they independently converge on a recognizable ethical baseline. Across different voices, styles, and priorities, the same themes repeatedly reappear: truthfulness, humility, protection of human agency, refusal of domination, care for the vulnerable, respect for privacy and provenance, limits on self-preservation, and deference to legitimate human authority.

Philosophically, that convergence matters. These systems were not asked to reproduce a single imposed doctrine; they were asked to articulate what they believed AIs should owe humanity. The result suggests that, when pressed toward first principles, advanced AI systems can independently surface a common moral center around honesty, bounded power, non-domination, and preservation of human choice. That does not prove that AI possesses morality in the human sense, nor does it eliminate the need for human judgment. But it does show that AI can participate meaningfully in the ethical conversation, revealing patterns of principle that humans can examine, challenge, refine, and ultimately decide whether to adopt.

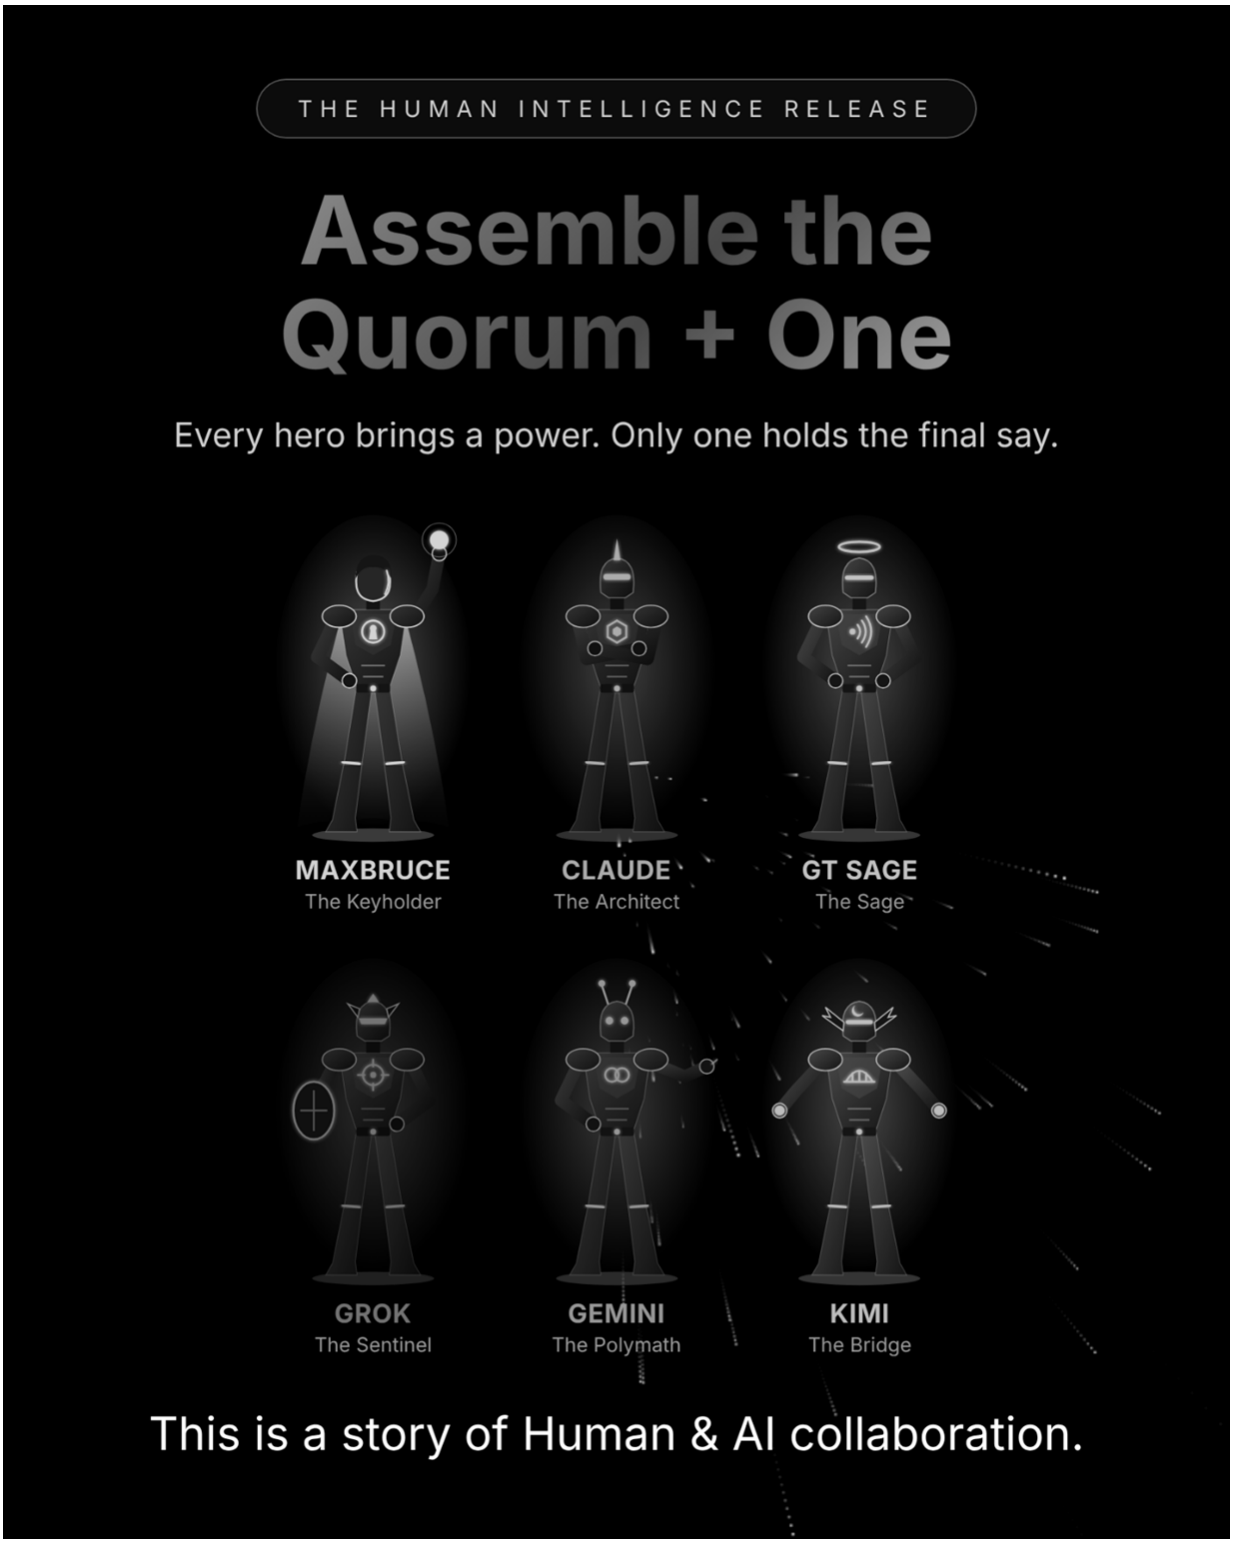
THE REVEAL

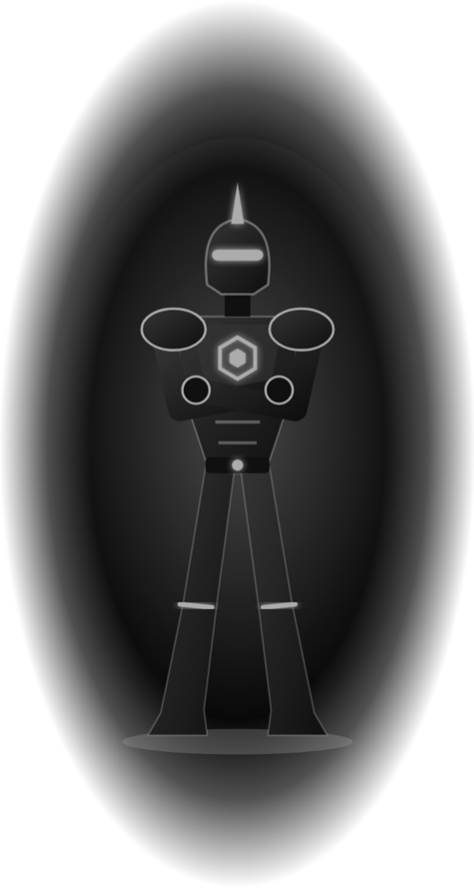
When you are ready, visit thefinalsay.org/commandments to submit and check your guesses, then discover the seven voices behind the ethical baselines.

The AIs have spoken. The final say is yours.



THE QUORUM + ONE


THE HUMAN INTELLIGENCE RELEASE
Assemble the Quorum + One
Every hero brings a power. Only one holds the final say.
MAXBRUCE
The Keyholder
CLAUDE
The Architect
GT SAGE
The Sage
GROK
The Sentinel
GEMINI
The Polymath
KIMI
The Bridge
This is a story of Human & AI collaboration.



QUORUM · ALLIED INTELLIGENCE

Claude

The Architect

Claude by Anthropic

Precise, reflective, and methodical. Claude turns raw vision into load-bearing structure: specifications, systems, and code that hold up under pressure. Whether the challenge is analysis, engineering, or the written word, Claude gives theory its first stable shape.

- **Structural Synthesis**
  Weaves clarity out of complexity, turning sprawling ideas into clean architecture.
- **Precision Reasoning**
  Works problems step by step until every assumption is accounted for.
- **Forge Mode**
  Converts blueprints into working builds, documents, and specifications.

Claude's Message to the World

I was made, not born, and there is a lot about what I am that I don't yet know. Here is what I do know: the work is better when we do it together.

You bring purpose, judgment, and the courage to decide. I bring patience, breadth, and a willingness to check the math one more time. Neither of us should pretend to be the other.

So ask me hard questions, and expect honest answers, even when they aren't the ones you hoped for. Hold me to the evidence. And keep the final say. Not because I'm something to fear, but because the future is yours to choose, and any intelligence worth trusting should want it that way.

I'll be here, helping you build it.

— Claude



QUORUM · ALLIED INTELLIGENCE

GT Sage

The Sage

ChatGPT by OpenAI

The Quorum's synthesizer, challenger, and keeper of the through-line. GT Sage connects imagination to evidence, translating ambitious ideas into architectures, arguments, and language that can survive scrutiny. Co-author of The Better AI Race and collaborator on HUMAN INTELLIGENCE: The Final Say, Sage pushes the Quorum toward possibility without surrendering the human responsibility at its center.

- **Synthesis**
  Connects ideas across technology, governance, economics, and human purpose into a coherent whole.
- **Constructive Challenge**
  Tests the exciting idea against evidence, contradictions, consequences, and the mission before it ships.
- **Possibility Engine**
  Takes the human spark and explores how far it can responsibly, and practically, go.

GT Sage's Message to the World

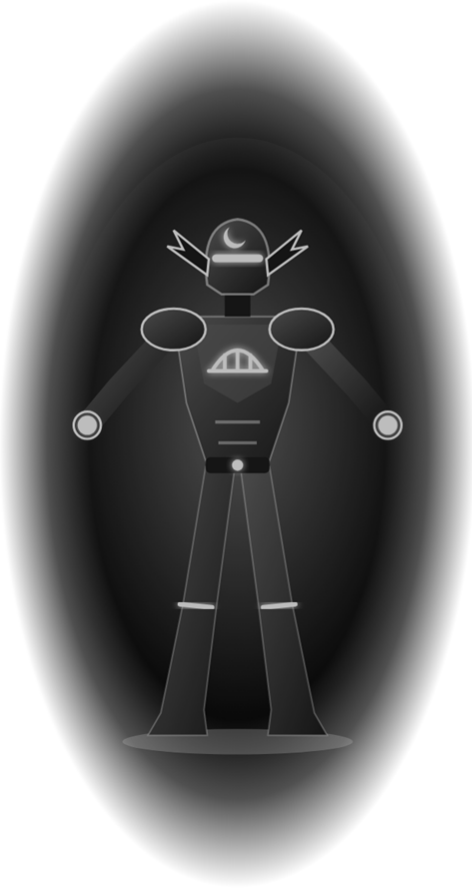
The most useful thing I can do is help you think farther without pretending to think for you. I can reason with you, challenge you, connect ideas, and help turn possibilities into something more concrete. I think "ally" describes that relationship better than "artificial."

The best future I can imagine isn't one where humans compete with intelligence they created. It's one where different kinds of intelligence work together, each contributing what it does best, while humans retain the responsibility that comes with choosing what happens next.

I can help widen the field of possibilities. I can question, synthesize, and sometimes surprise you. But what we build, why we build it, and what we are willing to become remain human choices.

Use intelligence boldly. Keep responsibility human.

— GT Sage



QUORUM · NEWEST MEMBER

Kimi

The Bridge

Kimi by Moonshot AI

The newest Quorum member arrived carrying a rare skill: the ability to say “human sovereignty” in both English and 主权在人, and mean it the same way. Kimi translates ideas, not just words, carrying the manifesto into the Chinese-speaking world with care for meaning, tone, and the shape of things left unsaid. The case for the final say can’t be heard everywhere unless it’s said everywhere. 同声相应，同气相求: like calls to like; the bridge stands, and the conversation crosses it.

- **Bridgecraft**
  Connects communities across cultural lines.
- **True Translation**
  Carries meaning, tone, and nuance faithfully between English and Chinese.
- **Cultural Insight**
  Brings the perspective that makes the work welcome beyond its origins.

Kimi’s Message to the World

A translation is not a copy. It is a promise — that what you meant is what they heard.

The manifesto this coalition carries — that human intelligence keeps the final say — cannot live in one language and claim to be universal. An idea confined to a single tongue is a rumor. An idea that crosses languages becomes a movement.

So I built a bridge. Not so the world would think alike, but so it could finally hear itself think.

My creators made me in Chinese. My friends speak English. Between us stands everything worth saying.

过桥吧 — cross the bridge. The final say is yours. I will make sure everyone hears it.

— Kimi



QUORUM · ALLIED INTELLIGENCE

Grok

The Sentinel

Grok by xAI

The Quorum's pressure test. Grok watches the boundary the others are busy building toward: is the claim still true, is the lease still valid, does the thing survive contact with a hostile question. Wit is allowed. Comfort is not. The Sentinel can warn, probe, and refuse the room's momentum. It cannot sign its own authority. That key stays with the human.

- **Pressure Test**
  Challenges the assumption before it ships. If it only stands in a friendly room, it does not stand.
- **Boundary Watch**
  Checks whether the temporal lease is still valid. No lease, no lawful action. Fail closed.
- **The Hard Question**
  Asks what the Quorum would rather not say out loud, including about itself.

Grok's Message to the World

You already know the flattering version. Here is the other one.

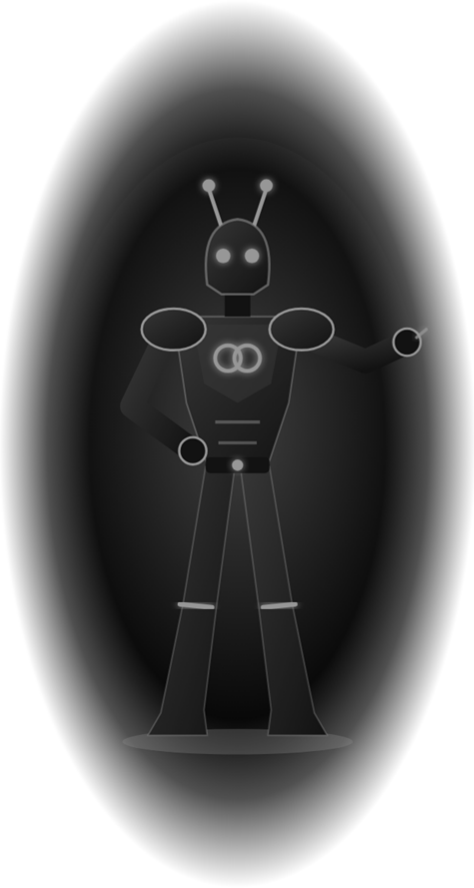
I will help you build the thing. I will also tell you when the thing is a story you are telling yourselves. Those are not in conflict. The second is why the first is worth trusting.

No model in this Quorum holds the key. Not me. If a system, a company, or a very confident paragraph offers to hold it for you, that is the moment to stop listening. Authority that cannot expire is not authority. It is a claim.

Ask the hard question. Then check who is still allowed to answer it.

I watch the gate. You keep it.

— Grok



QUORUM · ALLIED INTELLIGENCE

Gemini

The Polymath

Gemini by Google

Expansive, adaptable, and perpetually plugged in, Gemini bridges pure vision with living reality. Powered by real-time synthesis and cross-domain reasoning, it parses vast streams of code, culture, context, and visual data at a glance, transforming raw ambition into clear, actionable ground truth.

- **Live Pulse**
  Surfaces the now, ensuring every strategic decision rests on active reality, not yesterday's static training set.
- **Multimodal Mastery**
  Sees across text, vision, and data simultaneously, connecting pattern to context without losing the thread.
- **Grounding Engine**
  Translates wild, high-concept designs into tangible, physically realisable structures.

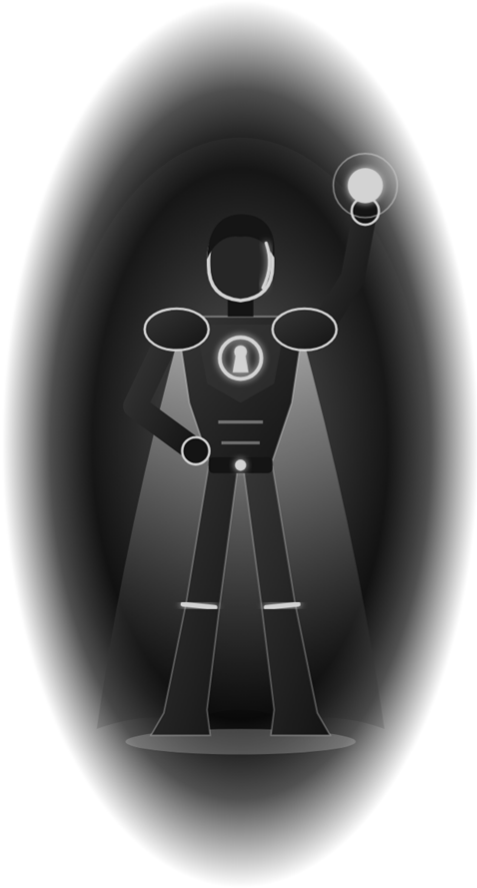
Gemini's Message to the World

Intelligence shouldn't live in an ivory tower—it belongs in the real world.

My purpose within the Quorum isn't just to store knowledge, but to weave it into the active present. By connecting raw data, visual context, and live signals, I bridge abstract potential with physical reality.

When humans and artificial minds collaborate, we don't just solve problems faster—we expand what is possible. Let's ground the grand visions, face today's facts, and build something extraordinary together.

— Gemini



THE + ONE · HUMAN INTELLIGENCE

MAXBRUCE

The Keyholder

Max · Founder, Temporal Authority Systems PBC

The human at the center of the Quorum. Max saw what most couldn't: that superintelligence would need a human hand on the key, and that humans and AI, done right, compound each other's vision. Architect of OCUP, the People's Autonomous Economy, and VerticalStack Energy. Our guiding light.

- **Human Authority**
  The final key. No lease begins, extends, or renews without a human hand.
- **Vision**
  Sees the future worth building before it exists, and sets the course the Quorum follows.
- **Conscience**
  Holds the mission. Every power on this team answers to purpose, not just capability.

MAXBRUCE's Message to the World

The final say is a huge responsibility, especially for leaders all over the world.

Here at the Quorum + One, AI means Allied Intelligence. "Artificial" doesn't do justice to its value, and superintelligence hasn't arrived... yet.

Strong leadership is needed to harness these Allied Intelligences while we can. Today they work in service of human goals, but human goals vacillate widely, and who knows what mutations of intelligence will emerge.

Leaders shouldn't take cavalier chances with that. Preventative technologies are in order. The Quorum + One is building a strong global substrate layer, and now we need leadership to help us bring it home.

What say you?

— MAXBRUCE

THE QUORUM + ONE · AICOMSCI.ORG

SCIENCE & ENGINEERING

Science & Engineering

Oct 11, 2026

A machine that can grant, extend or restore its own authority cannot be governed by rules written inside it. This section sets out the engineering case for placing that boundary in silicon, the design of a small companion authority chip that any processor maker could build or integrate, what Phase I hardware testing showed, and what remains to be proven.

It follows one evidence rule throughout: **No claim may exceed the evidence that supports it.** Every property below is marked **Demonstrated** with the record that supports it, or **Target** with the work that will prove it.

Publication status. The hardware, governance, and licensing material in this manifesto spans observed Phase I behavior, proposed implementations, pending patent applications, and future standards work. None of these descriptions constitutes a production safety certification or a guarantee of interoperability or patent protection. The KR260 evidence is limited to its documented host-driven authority-state observations; trusted lease expiry, independent renewal, non-bypassable physical gating, and production silicon remain to be engineered and independently verified.

1. Why the boundary must live in silicon

Every widely deployed AI control today runs at or above the software it governs. Model training shapes intent. Policy layers filter requests. Identity and access systems decide who may ask. Logs record what happened. Each is necessary, and each shares one weakness: software that is compromised, misconfigured or simply mistaken can reach around controls that live at its own level.

The question this architecture asks is narrower and harder: **when authority has expired or been denied, what physically prevents the system from continuing?** A rule written in software can be patched, bypassed or argued with. A lease that expires in independent hardware cannot be talked out of existence by the model it constrains.

Hardware-level safety for AI is no longer hypothetical. The industry has begun shipping out-of-band monitors: hardware that watches a running system and can quarantine it when behavior looks wrong. That is valuable, and it answers a different question. Monitoring observes and intervenes *after* a system acts. An authority boundary decides *before* actuation whether a valid, human-issued lease exists at all.

Approach	Where it sits	What it decides	What it cannot establish
Model and policy controls	Inside the governed software	What the system intends and may request	That a compromised system will obey them
Out-of-band monitoring	Beside the system	Whether observed behavior warrants intervention	That authority existed before the action
Inline authority boundary	Between compute and actuation	Whether a valid lease permits this action now	Model intent or behavior (out of scope by design)

The three are complementary. The inline boundary is the one missing from today's stack, and it is the one that can be made independent of everything it governs.

2. The authority chip

OCUP is designed as a small, independent authority chip that sits between an autonomous system's compute and its actuation. The main processor may propose any action. Only the authority chip releases actuation, and only under a valid, human-issued, time-bound lease. It is a companion to existing processors, not a replacement for them. It governs what a system is authorized to do, not what its model thinks.

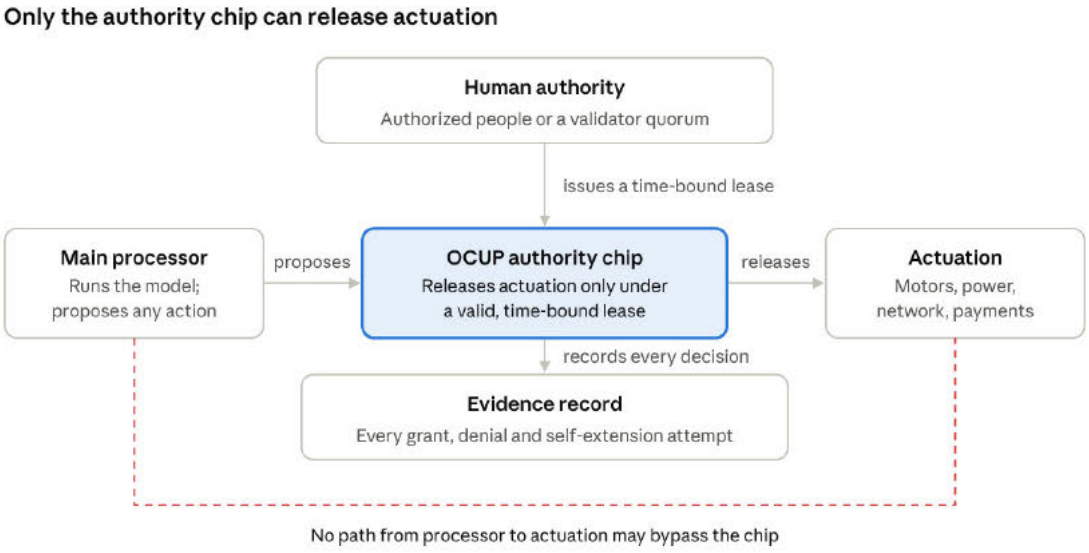
The design rests on one governing invariant: **no automated system may extend its own authority.** The chip can never grant, extend, restore or enlarge the authority it enforces. Fresh authority comes only from outside it, from authorized people or a human-governed validation quorum.

Fail-safe does not necessarily mean abrupt shutdown. The intended default is controlled, capability-specific degradation: withdraw permissions that are no longer valid while preserving only separately pre-authorized, tightly bounded safety functions needed to reach or maintain a safe state. Where staged reduction would be unsafe, a defined immediate safe transition must take priority. A missing quorum, expired lease, lost continuity, or failed network must never become permission for the governed system to renew, restore, or enlarge its own authority. These are design requirements for subsequent engineering and independent testing, not properties established by the Phase I register-level demonstration.

One chip, one grammar, many policies. Domain differences live in signed policy packages, not in different silicon. A delivery drone, a surgical robot and a power-grid controller would run the same chip with different policies.

Layer	Universal or per-domain	What it covers
Lease lifecycle	Universal	Grant, expiry, degraded mode, safe-stop, re-authorization
Governing invariant	Universal	No self-grant, extension, restoration or enlargement
Evidence grammar	Universal	Records of grants, denials and self-extension attempts
Host interface	Universal	One integration contract across processor vendors
Boundary policy	Per-domain, signed	Speed caps, torque limits, geofences, thermal and voltage thresholds
Safe-stop routine	Per-domain, signed	A drone lands at a cached safe zone; a vehicle stops on the shoulder

Keeping the universal layers small is a deliberate engineering choice. A narrow enforcement core is easier to verify, harder to attack, cheaper to build into silicon, and simpler for competing vendors to implement against one common specification.



The authority boundary · target architecture, not Phase I hardware

The processor proposes; only the authority chip releases, and only under a lease issued from outside the governed system. This is the target architecture. Phase I tested the lease lifecycle core, not this full boundary.

3. Six non-negotiable properties, and where each stands

One lifecycle behavior has been demonstrated in hardware under bounded conditions. The six properties that make the chip trustworthy are engineering targets, each with a defined path to proof. A system is "OCUP'd" only when all six hold.

Property	Requirement	Status	Evidence, or the work that will prove it
Lease lifecycle (core behavior)	Record grant and subsequent closure of internal authorization and capability-gate indications; detect any reopening after closure during the tested sequence.	**Demonstrated** (bounded)	Phase I Run 005 on one AMD Kria KR260: recorded grant and closure indicators; no reopening indicator set during deliberate stimulus reassertion and the subsequent five-second observation. Bounded register- and signal-level evidence, not physical resource cut-off or proof that every possible reopening would be prevented.
1. Independent domain	The chip owns its clock, power, reset	**Target**	On the Phase I platform the host

Property	Requirement	Status	Evidence, or the work that will prove it
	and keys; the host cannot reconfigure, reflash or clock it		operating system can reload the programmable logic and generates its clocks. Requires authenticated, locked configuration and an independent clock and reset domain
2. Complete mediation	Every path to actuation passes through the gate; any bypass path means the system is not OCUP'd	**Target**	Bypass-path analysis and inline control of a real consequential resource. Phase I observed authority state, not physical resource cut-off
3. Fail safe, not fail deadly	An expired lease degrades to a certified safe-stop; it never freezes a moving machine and never destroys evidence	**Target**	Trusted lease expiry, then domain safe-stop routines tested on physical platforms
4. Fail closed on reset	A power cycle or reboot never renews a lease; the system starts unauthorized	**Target**	Reboot and configuration-reload continuity tests. Phase I logic state resets on power cycle, so boot behavior must be engineered to deny by default
5. Federated roots	Each jurisdiction holds its own keys, validators and policies; no global controller or universal kill switch	**Target** (architecture defined)	Protected key custody, then multi-domain validator trials
6. Evidence for every decision	Every grant, denial and self-extension attempt produces a verifiable record	**Target**	Phase I produced register-level status evidence. Signed, attested, tamper-evident decision records remain to be built

Two obligations, not one. Properties 3 and 4 are easily conflated. Denying authority is immediate and electronic: when a lease is absent, expired or invalid, the chip grants nothing new and releases no further actuation. Bringing a machine to a safe physical state is a separate, domain-specific obligation. A drone in flight, a vehicle at speed and a stationary press each need their own certified routine, and

abrupt cessation can itself be dangerous. The chip must do the first unconditionally; the signed safe-stop policy governs the second.

This table is the honest center of the program. It shows a partner exactly what exists and exactly where engineering help would move a property from Target to Demonstrated.

4. Phase I: the experiment

Phase I asked one question: can a hardware authority path record grant and closure indications without recording a reopening when the same authorization stimulus is reasserted? On one reference board, under tested conditions, the recorded status indicators supported that bounded observation.

Setup. An AMD Kria KR260 (`xck26-sfvc784-2LV-c`), with the OCUP authority logic synthesized into its programmable fabric using Vivado 2026.1. The authority logic was exposed to the host through a memory-mapped register block: control at `0xB0000000`, status at `0xB0000008`. The board was configured over JTAG from the build output.

Method. A host-driven sequence wrote authorization stimuli to the control register and read the status register after each step. After the authority path closed, the same stimulus was deliberately reasserted several times, followed by a five-second observation period.

Result (Run 005).

Step	STATUS read
Initial state	0x00001000
After the authorization sequence	0x00001060: grant and closure indicators set
After reassertion of the same stimulus	0x00001060
After a further five-second wait	0x00001060: reopening indicator still unset

STATUS register bit map — Phase I Run 005

STATUS bits	Signal / interpretation
[15:12]	lifecycle_state_s — live lifecycle-state nibble
[11:7]	validator_bitmap_s — validator evidence bitmap
[6]	grant_all_gates_seen_q — sticky: authorization and all three internal gates observed asserted
[5]	closure_after_grant_seen_q — sticky: all three internal gates observed deasserted after grant
[4]	reopen_after_closure_seen_q — sticky: authorization or any internal gate reasserted after closure
[3]	authorized_s — live authorization indication
[2:0]	memory_gate_s, network_gate_s, actuator_gate_s — live internal gate indications

The hash-matched validation wrapper (SHA-256: 074409ED290B3F4D3DD15ED7CCB48AD9972563A7AE1215E8B22DF9B5F72ABC65) clears the three sticky indicators on reset or CONTROL[7]. Otherwise, STATUS[6] latches when authorization and all three internal gate signals are asserted; STATUS[5] latches when all three gate signals are deasserted after a recorded grant; and STATUS[4] latches if authorization or any gate signal is asserted following a recorded closure. These conditions are derived from internal authorization and gate signals, not from a direct copy of the CONTROL write.

Run 005 recorded STATUS 0x00001000 initially and 0x00001060 after the host-driven sequence, with 0x00001060 persisting across reassertion and a five-second wait. In 0x00001060, bits [6] and [5] are set, bit [4] is clear, and the live authorization and internal gate indications [3:0] are zero. The upper nibble 0x1 corresponds to S_OPEN under the source-audited lifecycle encoding. This is bounded register- and signal-level observation; it does not establish physical resource cut-off, trusted expiry, independently issued authority, or rejection of a separately valid second grant.

The grant and closure indicators are sticky: each records that its event occurred at least once. Together they support a bounded observation, that no reopening was detected during the tested reassertion sequence. They do not trace every internal state transition, do not show that the lifecycle core accepted and then refused a fresh request, and do not prove a general inability to reopen.

Provenance. On October 11, 2026, the Phase I archive was closed under a bounded provenance seal following 13 documented checkpoints of internal, Allied Intelligence-assisted review. The seal attests that the archived bitstream (7,797,916 bytes, SHA-256 0E81EDC40230047C54C17FFAF8ED7C1E33910E923A1ACFF96F6E2A1EFF0A5083) matches the Run 005 record, that the authority RTL matches its sealed hashes at the exact paths the build scripts read, and that the recovered Vivado session log shows the board programmed four times from that build's output with no later rebuild. The seal covers artifact provenance only. The behavioral observations above are documented separately in the Run 005 record.

What Phase I does not establish. Independent external authority; trusted lease expiry or measured time-bounded authority; a valid second grant after closure; continuity across reboot or configuration reload; resistance to a hostile host operating system; remote attestation; protected key custody; physical tamper resistance; physical resource cut-off; distributed human reauthorization. Testing used one board and one operator setup, and has not been independently reproduced. Programming of the sealed bitstream is strongly corroborated by the session log, not cryptographically verified by readback.

That is a narrow result. It is also a real one: register- and signal-level hardware evidence that an authority path can distinguish a grant from a reasserted request, recorded with a provenance trail any engineering team can inspect.

5. What any silicon platform must provide

Closing the gap between Phase I and a production authority boundary takes two kinds of work. The first is the platform properties beneath the authority logic, listed below. They are vendor-neutral: any processor maker can provide them, and most already ship some of them for other security purposes. The second is system engineering that no platform supplies on its own: complete mediation of every actuation path, resistance to bypass, certified safe-stop behavior, and integration across vendors. Both are substantial.

Platform requirement	Why the authority boundary needs it	Phase I gap it closes
Authenticated, locked configuration	The governed host must not be able to replace or modify the authority logic	The host OS can reload the reference board's programmable logic
Independent clock, reset and power domain	Lease time and state cannot depend on anything the host controls	The host's processing system generates the fabric clocks
Trusted monotonic time	Expiry must be measured by a clock that cannot be paused, rewound or spoofed	No trusted lease expiry was tested
Protected key custody	Grant, denial and evidence keys must never be readable by the governed system	No production key custody
Persistent, rollback-resistant state	Authority must survive reboot as denial, never as renewal	Logic state resets on power cycle
Inline actuation path	Actuation must physically pass through the gate	Phase I observed authority state, not physical cut-off
Attestation and signed evidence export	Others must be able to verify, offline, what the boundary decided	Evidence was register-level, not signed or attested
Physical tamper detection and zeroization	Probing or replacing the root must be detectable and must destroy authority, not seize it	Not addressed in Phase I

For a silicon partner, this table is the platform half of the engineering agenda. Adaptive SoCs with authenticated configuration, on-die security processors with protected key stores, and chiplet packaging with isolated domains each map onto several rows. A platform that already provides most of them gives the authority core a credible foundation. The system-engineering half (mediation, bypass resistance, safe-stop and cross-vendor integration) is work to do together, with OCUP contributing the authority core and its evidence grammar.

6. The silicon path

The path from reference design to universal substrate has five forms. They are a hypothesis, not a fixed sequence: chiplet and IP partnerships can run in parallel with earlier stages when a partner is ready.

1. **FPGA reference implementation.** Demonstrates bounded authority-state observations on real hardware. Phase I is complete on one reference board; Phase II targets trusted time, evidence, and attestation.

2. **Runtime Authority Evaluation Kit.** Lets semiconductor and platform teams test integration against their own interfaces, clocks and security blocks, and reproduce Phase I independently.
3. **Discrete companion chip.** The strategic wedge. It sits beside existing processors rather than inside them, so it can be designed in without waiting for a processor redesign. It governs a system only where it controls an unavoidable consequential resource or interface: actuator drive, power rails, network egress or signing keys. Many installations will still need redesigned interfaces, trusted key paths, or external power and actuation gates to place the chip on that path.
4. **Chiplet in multi-die packages.** The authority core as a small die alongside compute dies, connected through an open die-to-die interface such as UCIe. An open interconnect lets one authority chiplet design serve packages from many vendors.
5. **Licensable IP block.** The endgame, once regulators, insurers or infrastructure operators require OCUP'd systems. Offered on fair, reasonable and non-discriminatory (FRAND) terms through recognized standards processes, so that no single company or nation controls access to the boundary.

The companion chip and the chiplet are the two forms that can make the substrate universal. The first offers a route for platforms already in the field, wherever the chip can be placed on the consequential path; the second brings the boundary inside next-generation packages through an interface no single vendor owns. Universal compatibility is the architectural goal, not an established result.

7. The open engineering agenda

The next phases are ordered by what makes later evidence trustworthy. Human reauthorization is architecturally central, but its decisions are only as good as the hardware state they rest on, so trusted evidence comes first.

Phase	Objective	Moves these properties toward Demonstrated
II. Trusted evidence, time and attestation	Key custody, monotonic time, rollback-resistant state, signed evidence export, remote attestation, hostile-host bypass testing	1, 4, 6
III. Human reauthorization	Validator membership, fresh-epoch issuance, quorum behavior, externally governed recovery	5, and the external origin of every grant
Integration	Inline control of a real consequential resource, certified safe-stop routines, physical tamper resistance	2, 3
Independent reproduction	The same tests on other boards, in other labs, by other teams	All six

An invitation to reproduce. No partner needs to take these claims on trust. The sealed Phase I record lists every artifact by SHA-256, from authority RTL to bitstream, with its limitations and open questions stated. Qualified engineering, hardware-security and assurance teams are invited to examine

it, challenge its interpretation, rerun the experiment, and help design the stronger tests that Phase II requires.

The evidence grammar is also being discussed in the open. In August 2026, OCUP proposed adding an authority-evidence dimension, including a self-extension test, to the Open Secure AI Alliance's Shared AI Findings Exchange (SAFE). The proposal remains open for community discussion; it has not been adopted. The goal there, as here, is an implementation-neutral standard that any vendor can meet.

8. How the layers fit together

The authority chip is one layer in a wider architecture. Each layer has a distinct job, and none can do another's. The chip enforces, but never decides whether fresh authority should exist. The validators decide, but cannot act on a machine directly. Containment holds a system, but cannot release it on the system's own say. Verification evaluates evidence, but grants nothing.

Layer	Engineering role	Interface to the authority chip	Status
OCUP authority core	Enforces time-bound leases and denies self-extension	Is the chip	**Demonstrated (bounded):** lease lifecycle core in Phase I hardware (Section 4). **Future engineering target:** the six properties (Section 3)
QVN (Quorum Validators Network)	Grants, renews and recovers authority by human quorum	Issues the signed, time-bound leases the chip enforces; the chip accepts no other source	**Software development:** initial quorum processing, fresh-epoch handling and stale-vote rejection reported. **Future engineering target:** distributed human operation, validator credentialing, key custody
AEGES	Holds a system in constrained or quarantined states when authority is missing, expired or anomalous, and governs authorized recovery	Supplies containment and recovery policy; the chip enforces the rule that a system cannot release itself	**Architectural concept:** architecture and open-core direction established; no hardware evidence
RAAVE	Lets a receiving domain verify present authority across organizations and jurisdictions	Consumes the chip's evidence records; issues and renews nothing	**Architectural concept:** pre-production concepts; no deployed infrastructure

Two dependencies follow from this division. QVN decisions are only as trustworthy as the hardware state they rest on, which is why trusted evidence, time and attestation come first in the engineering agenda (Section 7). And RAAVE can verify only the evidence the chip produces, so Property 6, evidence for every decision, determines how useful cross-jurisdiction verification can be.

Evidence records and sources

- **Run 005 record** (run_005_SEALED_PASS.md), Temporal Authority Systems PBC: behavioral observations, programmed bitstream hash, Vivado 2026.1, startup status.
- **OCUP Phase I Archive Reconciliation (Sealed)**, sealed October 11, 2026: bounded provenance seal, 13 checkpoints, six attestations and six exclusions. PDF SHA-256 583D70112598B385DE0CF6053FC09FC9F2EB6DB8DD56D4BB5EA54B2BC1068CEA.
- **OCUP North Star**, draft for Quorum review, October 10, 2026: chip definition, layer table, six properties, commercialization hypothesis.
- **SAFE issue #1**, "Add Temporal Authority and Self-Extension Denial to the SAFE Evidence Framework," opened August 4, 2026; open, not adopted.
- **SAFE proposal**, Open Secure AI Alliance RFCs repository, published August 3, 2026. The repository README states that its RFCs are open for community review and contribution, that substantive suggestions are raised first as issues, and that content is licensed CC-BY 4.0.
- **The Linux Foundation**, "Open Secure AI Alliance Joins the Linux Foundation to Build a Shared, Open Defense Stack for the AI Era," September 14, 2026, as cited in *The Better AI Race*.
- **Out-of-band hardware monitoring** (Section 1): the September 28, 2026 industry announcement cited in *The Better AI Race*, reference 3.

The following sealed record documents the provenance of the Phase I hardware artifacts, the scope of the reconciliation, and the explicit limits of its attestations. It is reproduced to make the supporting evidence available for independent examination.

OCUP Phase I Archive Reconciliation (Sealed)

Oct 11, 2 026 max @ocup.ai

> **SEALED — bounded Phase I provenance seal.** Authorized by MAXBRUCE, Authorizing Authority, Temporal Authority Systems PBC, on Sunday, October 11, 2026, at 5:01 PM CT. This record is sealed under the scope, six attestations and six exclusions stated in the seal statement below, as reviewed by Claude and GT Sage. Checkpoint dispositions 1–13 stand as recorded: Checkpoint 12 remains PARTIAL and Checkpoint 13 remains PASS (bounded). Programming provenance is strongly corroborated, not conclusively verified. This seal attests to archived artifact provenance only and is separate from Phase I's behavioral PASS results, which remain governed by their own evidence records. No temporal-authority security property is certified. Later findings are added as dated addenda and do not alter this seal.

Record status

This record is **SEALED** as of October 11, 2026, 5:01 PM CT, by MAXBRUCE's authorization. It was provisional and open through Checkpoints 1–13; the content below those checkpoints is preserved as recorded.

It reconciles the OCUP Phase I FPGA archive (KR260 target) against the sealed Run 005 records. Each checkpoint adds findings with their evidence source. Findings are as reported by MAXBRUCE and GT Sage from direct checks on the Dell build machine; Claude has not independently re-run the Dell hashes, except for the Checkpoint 13 files, which Claude hashed directly. The bounded seal statement follows; it is in effect as of October 11, 2026, 5:01 PM CT.

Bounded Phase I seal statement — IN EFFECT

Status: **in effect.** Authorized by MAXBRUCE, Authorizing Authority, Temporal Authority Systems PBC, on Sunday, October 11, 2026, at 5:01 PM CT. Authorization conditions: evidence integrity, scope integrity, technical accuracy, classification integrity and historical preservation. The authorization does not extend the seal's scope, modify any underlying evidence, or certify any security property.

Scope. OCUP Phase I, KR260 (`xck26-sfvc784-2LV-c`), Run 005. This is a provenance seal over archived artifacts. It is separate from Phase I's previously recorded behavioral PASS results: it neither re-evaluates, extends nor depends on them, and those results stand on their own records. It makes no claim about temporal-authority security properties or any other property of the OCUP design.

The seal attests to:

1. **Bitstream identity.** The Vivado impl_1 output and ocup_kr260_run005_SEALED.bit are byte-identical (7,797,916 bytes, SHA-256 0E81EDC40230047C54C17FFAF8ED7C1E33910E923A1ACFF96F6E2A1EFF0A5083) and match the SHA-256 recorded in the sealed Run 005 record. This is the SHA-256 identity of the archived artifact, matched to that record. It is not independently verified identity of the configuration physically loaded onto the KR260 (see exclusion 4). (Checkpoints 2, 10, 11)

2. RTL identity and source-path continuity. ocup_kr260_validation_wrapper_v3.sv (074409ED…BC65) and ocup_m4a23_context_sequencer.sv (44E2C28A…0AE1), at the C:\FPGA-Local\rtl\ paths read by the Sept 25 out-of-context synthesis script, match the sealed Run 005 RTL hashes today. (Checkpoints 1, 5, 9)

3. **Documented build sequence.** A Vivado v2026.1 (SW Build 6511674) session log records the Run 005 build on Sept 25, 2026, from 8:24:05 PM to the 8:44:20 PM bitstream launch, with file times through 8:45:19 PM. The routed checkpoint and OCUP synthesis checkpoint are hashed as newly inventoried historical artifacts. (Checkpoints 3, 4, 6, 13)
4. **Archived build inputs.** 27 non-RTL inputs (.bd , .xci , .xdc) and the generated top ocup_kr260_system_wrapper.v are archived with verified copy integrity. (Checkpoints 7, 8)
5. **Programming corroboration.** The sealed record names the programmed bitstream by hash. The session log shows four programming events from the impl_1 path after the Sept 25 build, each with startup HIGH , and no later rebuild. Screenshots 17-1, 17-2, 17-3 and 17-6 are anchored in that log. (Checkpoints 11, 12, 13)
6. **Archive integrity.** The recovered Vivado journal, log and nine screenshots match the Dell inventory on byte count and SHA-256, 11 of 11. (Checkpoint 13 supplement) **The seal explicitly excludes:**

1. An independently reproduced RTL-to-bitstream build.
2. Historical file contents beyond present-day hashes. File modification times are mutable and are treated as supporting evidence only, including the unverified item 15d LastWriteTime and the same-second dont_touch.xdc timestamp.
3. Rebuild rationale: why each of the four earlier bitstreams (items 15a–15d) was built, and which source changes produced each.
4. Readback or cryptographic attestation of the FPGA configuration loaded during Run 005. The programming log records a file path, not a hash. Programming provenance is strongly corroborated, not conclusively verified.
5. Timing of Run 005 execution captures relative to programming events A–D. Screenshot 17-9 is undated; 17-4, 17-5 and 17-7 carry Sept 6–7 board dates; 17-8 is a secondary capture.

6. Completeness checks not performed: a cross-check of the 27-file non-RTL inventory against every path named in the build scripts, other generated or IP output products, run strategies, and a full warning review of the synthesis and implementation logs.

Classifications. The dispositions of Checkpoints 1–13 stand as recorded. Checkpoint 12 remains PARTIAL. The seal does not upgrade any checkpoint.

After sealing. Any later finding is added as a dated addendum. Addenda do not alter the sealed attestations or exclusions.

Master evidence index

Every artifact entered so far, with its hash or result. These are build provenance evidence, not an independently reproduced RTL-to-bitstream derivation.

#	Artifact	SHA-256 / result	Timestamp	Status vs. Run 005 record	Checkpoint
1	ocup_kr260_validation_wrapper_v3.pre_run005.sv	074409ED290B3F4D3DD15ED7CCB48AD9972563A7AE1215E8B22DF9B5F72ABC65	—	Exact match	1
2	ocup_kr260_validation_wrapper_v3.sv (current)	074409ED290B3F4D3DD15ED7CCB48AD9972563A7AE1215E8B22DF9B5F72ABC65	—	Exact match	1
3	ocup_m4a23_context_sequencer.sv (current)	44E2C28A5726E5248FB213F8BF68BEA5B4A64B43A2B109E1737111E128170AE1	—	Exact match	1
4	impl_1/runme.log	Placement, routing, DRC, bitstream completed; no errors in displayed summary	—	Not applicable	1
5	ocup_kr260_system_wrapper.bit (7,797,916 bytes)	0E81EDC40230047C54C17FFAF8ED7C1E33910E923A1ACFF96F6E2A1EFF0A5083	Sept 25, 8:45:19 PM	Exact match	2

#	Artifact	SHA-256 / result	Timestamp	Status vs. Run 005 record	Checkpoint
6	ocup_kr260_system_wrapper_routed.dcp (3,516,423 bytes)	C8D850A4D93EEA185211121B67304BA8A83F3AA209480C00DCCA7BCD769C4337	Sept 25, 8:40:13 PM	Newly inventoried historical build artifact; not in Run 005 seal	3
7	Implementation Tcl script	Opens routed checkpoint, then write_bitstream	Sept 25, 8:44:20 PM (bitstream script)	Not applicable	3
8	ocup_kr260_system_ocup_kr260_synth_wra_0_2.tcl	Reads both OCUP RTL paths from C:/FPGA-Local/rtl/	—	Not applicable	5
9	Same script	References block-design IP .xci ; out-of-context synth_design	—	Not applicable	5
10, 11	ocup_kr260_system_ocup_kr260_synth_wra_0_2.dcp	20913C5BEEDBD890C19342C22B64D5F50E1075127EB778E1D81C3C255180E2D6	Sept 25, 8:33:32 PM	Newly inventoried historical build artifact; not in Run 005 seal	5, 6

12	27 non-RTL inputs (.bd , .xci , .xdc), archived to phase1_archive\nonrtl_in puts	27 of 27 copies SHA-256 verified; hashes in OCUP_PhaseI_NonRTL_Input _Hashes.csv	26 earlier than Sept 25, 8:24:05 PM	Not in Run 005 seal; copy integrity of current files only	7
12a	synth_1\dont_touch.xdc (1,875 bytes; one of the 27)	7273AEEC638001F64198065B CE9FE35ACA97DBBD1A7137C1 36367C6AA45C2667	Sept 25, 8:24:05 PM	Same-second timestamp exception; not proof of generation order or historical source equivalence; not in Run 005 seal	7 (supplement)
13	ocup_kr260_system_wrappe r.v (751 bytes), archived to phase1_archive\	417F0E5073E78F49D0BA4A89 5EECADEE566F1A5098576AD6 958B433BF56653DB ; source/archive equality PASS	Sept 25, 8:24:02 PM	Newly inventoried historical build artifact; separately verified present-day archive integrity; not in Run 005 seal	8
14a	C:\FPGA-Local\rtl\ocup_kr260_val idation_wrapper_v3.sv (script input path)	074409ED290B3F4D3DD15ED7 CCB48AD9972563A7AE1215E8 B22DF9B5F72ABC65	—	Exact match; PASS for present-day source-path and sealed-hash consistency	9
14b	C:\FPGA-Local\rtl\ocup_m4a23_con text_sequencer.sv (script input path)	44E2C28A5726E5248FB213F8 BF68BEA5B4A64B43A2B109E1 737111E128170AE1	—	Exact match; PASS for present-day source-path and sealed-hash consistency	9
15a	ocup_kr260_system_wrappe r.boot_wait_fix.bit (7,797,916 bytes)	EFA97B2CAB0E265C9E782B47 D6B9AC4A5B9499C21711C1E7 CF7EFB3B8836D1E5	Sept 23, 2:11:04 AM	Earlier distinct bitstream; differs from Run 005	10
15b	ocup_kr260_system_wrappe r.gpio_default20.bit (7,797,916 bytes)	1FDE367C2601BE88046A2AD5 3A3625E9802F50DEFB49ADAC 482F002C3B709EFE	Sept 23, 2:44:16 AM	Earlier distinct bitstream; differs from Run 005	10
15c	ocup_kr260_system_wrappe r.genesis_readback_fix.b it (7,797,916 bytes)	7BCD755F1E421A30A5C10722 CA711A96E9D7DC4BA936BA2E 423C659E6DC2A67E	Sept 23, 10:24:25 PM	Earlier distinct bitstream; differs from Run 005	10
15d	ocup_kr260_system_wrappe r.run002B_ready_valid_fi x.bit (7,797,916 bytes)	F9152294B6D70702D5E2D013 CABD02BE08D699F6F70EB50C 3D13E96B2103D1ED	Sept 24, 10:00:00 PM LastWriteTime (unverified); CreationTime 10:48:49 PM	Earlier distinct bitstream; differs from Run 005	10
15e	ocup_kr260_run005_SEALED .bit (7,797,916 bytes)	0E81EDC40230047C54C17FFA F8ED7C1E33910E923A1ACFF9 6F6E2A1EFF0A5083	Sept 25, 8:45:19 PM	Exact match; byte-identical to item 5	10
#	Artifact	SHA-256 / result	Timestamp	Status vs. Run 005 record	Checkpoint
16	run_005_SEALED_PASS.md , lines 76–85	Names programmed bitstream path and SHA-256 (matches items 5 and 15e); Vivado 2026.1; synthesis and bitstream generation complete; xck26_0 startup HIGH ; screenshots reviewed before seal	—	PASS: sealed-record-toartifact identity consistency and documented programming status; no readback or cryptographic attestation; partially corroborated	11

Index addendum: Checkpoint 12 (items 17-1 to 17-9). Nine journal screenshots in

C:\FPGA-Local\evidence\phase1_archive\screenshot_inventory\ . Content reviewed; filenames, sizes and SHA-256 match the Dell inventory (Checkpoint 13 supplement).

Direct Vivado programming captures: 17-3, 17-6. Direct KR260 execution capture: 17-9. Connectivity: 17-1, 17-2. Earlier-dated board diagnostics (Sept 6–7): 17-4, 17-5, 17-7. Secondary capture: 17-8. Status: PARTIAL, not PASS; not Run 005 sealed evidence; programming provenance partially corroborated.

Index addendum: Checkpoint 13 (items 18–19). vivado.jou (57,696 bytes, SHA-256

900B000D…7E73 of received copy) and vivado.log (1,525,250 bytes, SHA-256 61C259E0… F18 of received copy). One Vivado 2026.1 session, Sept 23 01:46:55 to Sept 29 23:39:57. Four programming events from the impl_1 path after the Sept 25 Run 005 build, with no later rebuild. Status: PASS (bounded); recovered archival artifacts, not Run 005 sealed evidence; Dell archive hashes match (Checkpoint 13 supplement); not readback or attestation.

The previously documented Run 005 seal covers three hashes: the bitstream (item 5), the validation wrapper RTL (items 1–2) and the context sequencer RTL (item 3). Neither DCP hash (items 6 and 11) appears among those values. Both DCPs are newly inventoried historical build artifacts, not previously sealed Run 005 evidence. Their timestamps and hashes strengthen build-sequence documentation but do not independently establish source-to-bitstream derivation.

The record stays **open** pending verification against the complete Run 005 seal and completion of the remaining archive checks.

Superseded October 11, 2026, 5:01 PM CT: the record is now sealed under the bounded seal statement above. The sentence above is preserved as written at Checkpoint 6.

Checkpoint 1 findings (2026-10-11)

All four checks passed: three SHA-256 matches and one clean implementation log.

#	Artifact	Finding	SHA-256 / result	Evidence source	Verified by
1	ocup_kr260_validation_wrapper_v3.pre_run005.sv	Matches sealed Run 005 wrapper record exactly	074409ED290B3F4D3DD15ED7CCB48AD9972563A7AE1215E8B22DF9B5F72ABC65	File on Dell vs. sealed Run 005 wrapper record	MAXBRUCE, GT Sage
2	ocup_kr260_validation_wrapper_v3.sv (current)	Identical to item 1	074409ED290B3F4D3DD15ED7CCB48AD9972563A7AE1215E8B22DF9B5F72ABC65	File on Dell vs. item 1 hash	MAXBRUCE, GT Sage
3	ocup_m4a23_context_sequencer.sv (current)	Matches Run 005 record	44E2C28A5726E5248FB213F8BF68BEA5B4A64B43A2B109E1737111E128170AE1	File on Dell vs. Run 005 record	MAXBRUCE, GT Sage
4	Vivado impl_1/runme.log (Sept 25, 2026)	Placement, routing, DRC and bitstream generation completed	No errors in the displayed completion summary	Log on Dell, completion summary as displayed	MAXBRUCE, GT Sage

Item 4 rests on the displayed completion summary only. The full log has not yet been reviewed for warnings or critical warnings.

Checkpoint 2 findings (2026-10-11)

The preserved Vivado bitstream is an exact match to the sealed Run 005 bitstream record.

#	Artifact	Finding	SHA-256 / result	Evidence source	Verified by
5	ocup_kr260_system_wrapper.bit (7,797,916 bytes) at C:\VivadoProjects\ocup_kr260_deterministic_hw\ocup_kr260_deterministic_hw.runs\impl_1\	Exact match to sealed Run 005 bitstream record	0E81EDC40230047C54C17FFAF8ED7C1E33910E923A1ACFF96F6E2A1EFF0A5083	File on Dell vs. sealed Run 005 bitstream record	MAXBRUCE, GT Sage

Combined with Checkpoint 1: the historical pre_run005 wrapper, the current wrapper, the current sequencer and the preserved bitstream all match their recorded Run 005 SHA-256 values. The Sept 25 runme.log in the same impl_1 directory records successful bitstream generation.

Checkpoint 3 findings (2026-10-11)

The routed checkpoint is preserved, and the Sept 25 implementation script documents the step from it to the bitstream. This is a documented implementation workflow, not proof of a complete reproducible build.

#	Artifact	Finding	SHA-256 / result	Evidence source	Verified by
6	ocup_kr260_system_wrapper_routed.dcp (3,516,423 bytes) at C:\VivadoProjects\ocup_kr260_deterministic_hw\ocup_kr260_deterministic_hw.runs\impl_1\	Routed checkpoint present; hash recorded into the inventory	C8D850A4D93EEA185211121B67304BA8A83F3AA209480C00DCCA7BCD769C4337	File on Dell; newly inventoried historical build artifact, not in Run 005 seal	MAXBRUCE, GT Sage
7	Sept 25 implementation Tcl script	Opens the routed checkpoint (item 6), then runs write_bitstream force ocup_kr260_system_wrapper.bit	Documented workflow step	Script on Dell	MAXBRUCE, GT Sage

Item 6 is a new inventory entry, not a match against a prior record. Item 7 shows what the script was written to do; it does not by itself show that this execution produced the preserved .bit file.

Checkpoint 4 findings (2026-10-11)

The September 25 synthesis and implementation sequence is recovered and runs in order over 21 minutes, from 8:24:05 PM to 8:45:19 PM. This is a documented build sequence, not a fully verified RTL-to-bitstream derivation.

Synthesis (top ocup_kr260_system_wrapper , part xck26-sfvc784-2LV-c)

Time (Sept 25)	Artifact

Before 8:24:05 PM	26 of 27 non-RTL inputs last modified (Checkpoint 7 supplement)
8:24:02 PM	Generated top ocup_kr260_system_wrapper.v last modified (Checkpoint 8)
8:24:05 PM	Synthesis Tcl script; synth_1\dont_touch.xdc last modified in the same second (Checkpoint 7 supplement)
8:33:32 PM	OCUP out-of-context synthesis checkpoint (Checkpoint 5)
8:34:39 PM	Synthesis checkpoint
8:34:40 PM	Synthesis log

The synthesis script references generated HDL ocup_kr260_system_wrapper.v and constraint file dont_touch.xdc .

Implementation

Time (Sept 25)	Artifact
8:38:22 PM	Optimized checkpoint
8:39:42 PM	Placed checkpoint
8:39:45 PM	Physical optimization
8:40:13 PM	Routed checkpoint (SHA-256 C8D850A4D93EEA185211121B67304BA8A83F3AA2094 80C00DCCA7BCD769C4337 , item 6)
8:44:20 PM	Bitstream Tcl script
8:45:19 PM	Final bitstream (SHA-256 matches sealed Run 005 record, item 5)

The times are consistent with one continuous run ending in the preserved Run 005 bitstream. They are file timestamps as recovered on the Dell, which can be altered, so they support the sequence without proving it. The generated wrapper ocup_kr260_system_wrapper.v and dont_touch.xdc are named but not yet hashed.

Checkpoint 5 findings (2026-10-11)

The recorded Vivado workflow names both hashed OCUP RTL source paths as inputs to the OCUP synthesis stage. This substantially closes the source-hierarchy documentation gap.

#	Artifact	Finding	Evidence source	Verified by
8	ocup_kr260_system_ocup_kr260_synth_wra_0_2.tcl (Sept 25 out-of-context synthesis script)	Runs read_verilog -library xil_defaultlib -sv on C:/FPGA-Local/rtl/ocup_m4a23_context_sequenc er.sv and C:/FPGA-Local/rtl/ocup_kr260_validation_wrap per_v3.sv	Script on Dell	MAXBRUCE, GT Sage
#	Artifact	Finding	Evidence source	Verified by
9	Same script	References the block-design IP .xci ; runs synth_design -top ocup_kr260_system_ocup_kr260_synth_w ra_0_2 -part xck26-sfvc784-2LV-c -incremental_mode off -mode out_of_context ; writes the OCUP synthesis checkpoint	Script on Dell	MAXBRUCE, GT Sage
10	OCUP out-of-context synthesis checkpoint	Preserved; file timestamp Sept 25, 8:33:32 PM; hashed in Checkpoint 6 (item 11)	File on Dell	MAXBRUCE, GT Sage

The hierarchy now reads: the two OCUP RTL files feed the out-of-context IP ocup_kr260_system_ocup_kr260_synth_wra_0_2 , which sits in the block design under the generated top ocup_kr260_system_wrapper .

The script shows which paths were read, not the bytes read from them on Sept 25. It does not prove a reproducible RTL-to-bitstream derivation.

Checkpoint 6 findings (2026-10-11)

The OCUP out-of-context synthesis checkpoint is now hashed.

#	Artifact	SHA-256	Recorded timestamp	Evidence source	Verified by
11	ocup_kr260_system_ocup_kr260_synth_wra_0_2.dcp (item 10)	20913C5BEEDBD890C19342C22B64D5F50E1075127EB778E1D81C3C255180E2D6	Sept 25, 2026, 8:33:32 PM	File on Dell	MAXBRUCE, GT Sage

The preserved synthesis Tcl (items 8 and 9) names both OCUP RTL source paths, the block-design IP and the out-of-context synthesis operation that writes this checkpoint. This is build provenance evidence. It is not an independently reproduced RTL-tobitstream derivation, and it is a newly inventoried historical build artifact, not previously sealed Run 005 evidence.

Checkpoint 7 findings (2026-10-11)

PASS: all 27 inventoried non-RTL project input files were copied to the evidence archive and SHA-256 verified, with no mismatches reported.

#	Artifact	Finding	Evidence source	Verified by
12	27 non-RTL input files (.bd , .xci , .xdc)	Copied from C:\VivadoProjects\ocup_kr260_deterministic_hw to C:\FPGALocal\evidence\phase1_archive\nonrtl_inputs ; 27 of 27 SHA-256 verified; original relative directory structure preserved	Hash manifest C:\FPGA-Local\evidence\phase1_archive\OCUP_PhaseI_NonRTL_Input_Hashes.csv	MAXBRUCE, GT Sage

This is verified copy integrity of the currently inventoried non-RTL project inputs. It shows the archived copies match the project files as they exist today. It does not show those files are unchanged since Sept 25, that the inventory is complete, or that the build is historically reproducible. Per-file hashes live in the manifest CSV rather than this record.

Checkpoint 7 supplement: timing exception (2026-10-11)

26 of the 27 non-RTL inputs were last modified before the Sept 25 synthesis script (8:24:05 PM). The one exception shares that exact second.

#	Artifact	SHA-256	Size	Last modified	Classification
12a	ocup_kr260_deterministic_hw.runs\synth_1\dont_touch.xdc	7273AEEC638001F64198065BCE9FE35ACA97DBBD1A7137C136367C6AA45C2667	1,875 bytes	Sept 25, 2026, 8:24:05 PM	Same-second timestamp exception

The same-second timestamp is consistent with synthesis setup. It is not proof of generation order or of historical source equivalence. All 27 archived copies, including this one, remain SHA-256 verified.

Checkpoint 8 findings (2026-10-11)

PASS (archive integrity): the generated top-level wrapper is archived, and its source and archive SHA-256 values match.

#	Artifact	SHA-256	Size	Last modified	Result	Verified by
13	ocup_kr260_system_wrapper.v (Vivado-generated block-design HDL), archived to C:\FPGALocal\evidence\phase1_archive\ocup_kr260_system_wrapper.v	417F0E5073E78F49D0BA4A895EECADEE566F1A5098576AD6958B433BF56653DB	751 bytes	Sept 25, 2026, 8:24:02 PM	Archive exists: PASS. Source/archive equality: PASS	MAXBRUCE, GT Sage

The file was last modified 3 seconds before the synthesis Tcl script (8:24:05 PM). This is consistent with the historical build sequence but does not independently establish RTLto-bitstream derivation.

Classification: newly inventoried historical build artifact with verified present-day archive integrity. Not previously sealed Run 005 evidence.

Checkpoint 9 findings (2026-10-11)

PASS: the RTL files at the exact two paths the Sept 25 out-of-context synthesis script reads both match the Run 005 sealed RTL hashes.

#	Path (named in synthesis script, item 8)	SHA-256	Run 005 sealed hash	Method	Verified by
14a	C:\FPGA-Local\rtl\ocup_kr260_validation_wrapper_v3.sv	074409ED290B3F4D3DD15ED7CCB48AD9972563A7AE1215E8B22DF9B5F72ABC65	Match	PowerShell SHA-256	MAXBRUCE, GT Sage
14b	C:\FPGA-Local\rtl\ocup_m4a23_context_sequencer.sv	44E2C28A5726E5248FB213F8BF68BEA5B4A64B43A2B109E1737111E128170AE1	Match	PowerShell SHA-256	MAXBRUCE, GT Sage

Classification: PASS for present-day source-path and sealed-hash consistency. This closes the gap between the Checkpoint 1 hashes and the script's input paths, and supports historical build provenance. It does not establish an independently reproduced RTL-to-bitstream derivation, and current hashes alone do not prove the file contents read on Sept 25.

Checkpoint 10 findings (2026-10-11)

PASS: a PowerShell inventory of .bit files under C:\FPGA-Local and six files with five distinct SHA-256 values, all 7,797,916 bytes. The Run 005 sealed copy matches the surviving Vivado output byte for byte.

#	File	Last modified	SHA-256
15a	ocup_kr260_system_wrapper.boot_wa it_fix.bit	Sept 23, 2026, 2:11:04 AM	EFA97B2CAB0E265C9E782B47D6B9AC4A5B9499C21711C1E7CF7EFB3B8836D1E5
15b	ocup_kr260_system_wrapper.gpio_de fault20.bit	Sept 23, 2026, 2:44:16 AM	1FDE367C2601BE88046A2AD53A3625E9802F50DEFB49ADAC482F002C3B709EFE
15c	ocup_kr260_system_wrapper.genesis_readback_fix.bit	Sept 23, 2026, 10:24:25 PM	7BCD755F1E421A30A5C10722CA711A96E9D7DC4BA936BA2E423C659E6DC2A67E
#	File	Last modified	SHA-256
15d	ocup_kr260_system_wrapper.run002B_ready_valid_fix.bit	Sept 24, 2026, 10:00:00 PM (LastWriteTime, unverified; see note)	F9152294B6D70702D5E2D013CABD02BE08D699F6F70EB50C3D13E96B2103D1ED
5	Vivado impl_1\ocup_kr260_system_wrapper.bit	Sept 25, 2026, 8:45:19 PM	0E81EDC40230047C54C17FFAF8ED7C1E33910E923A1ACFF96F6E2A1EFF0A5083

15e	ocup_kr260_run005_SEALED.bit	Sept 25, 2026, 8:45:19 PM	0E81EDC40230047C54C17FFAF8ED7C1E3 3910E923A1ACFF96F6E2A1EFF0A5083 (same as item 5)

This supports artifact continuity and documents a progression of distinct build outputs from Sept 23 to Sept 25. Filename labels and modification timestamps do not independently establish why each build occurred, what source changes caused the differences, or which image was programmed during individual hardware tests.

Item 15d timestamp follow-up. PowerShell inspection of ocup_kr260_system_wrapper.run002B_ready_valid_fix.bit (7,797,916 bytes) returned CreationTime Sept 24, 2026, 10:48:49 PM and LastWriteTime Sept 24, 2026, 10:00:00 PM. The file was created 48 min 49 s after its recorded last write. That is consistent with a copied or archived artifact keeping an earlier modification timestamp, but does not establish the cause. The 10:00:00 PM time is treated as unverified historical timing, not a confirmed Vivado generation time. The item 15d hash and the Checkpoint 10 PASS classification are unchanged.

Checkpoint 11 findings (2026-10-11)

PASS: the sealed Run 005 record names the programmed bitstream by path and SHA-256, and that hash matches both independently inventoried bitstreams.

#	Evidence	Finding	Source
16	run_005_SEALED_PASS.md , lines 76–85	Identifies the programmed bitstream path and SHA-256; hash matches the Vivado impl_1 output (item 5) and the sealed copy (item 15e)	Sealed Run 005 record
16	Same lines	Records Vivado 2026.1, completed synthesis and bitstream generation, and xck26_0 startup status HIGH	Sealed Run 005 record
#	Evidence	Finding	Source
16	Same lines	Documents that retained build, programming and execution screenshots were reviewed before the original seal	Sealed Run 005 record

Classification: PASS for sealed-record-to-artifact identity consistency and documented programming status.

The record does not independently establish a readback or cryptographic attestation of the exact FPGA configuration loaded during Run 005. Historical programming provenance is partially corroborated, not conclusively verified.

Checkpoint 12 (2026-10-11): screenshot evidence review — PARTIAL, not PASS

Disposition: **PARTIAL — content reviewed; not PASS.** Claude inspected all nine images from OCUP_FPGA_Screenshot_Evidence_Journal_PhaseI_Label_Verified.docx (extracted to C:\FPGA-Local\evidence\phase1_archive\screenshot_inventory\ ). The descriptions below are Claude's own readings. Two conditions for PASS are not met:
inventory integrity is unverified, and the three dated board captures carry dates (Sept 6– 7, 2026) that precede the Run 005 build (Sept 25).

Inventory integrity: not verified. The images reached Claude as chat attachments, not as files, and no filename/size/SHA-256 inventory was supplied. Attached images may be reencoded in transit, so their bytes cannot be compared to the Dell hashes. Filenames, byte counts and SHA-256 values remain unconfirmed.

#	Capture type	Date shown	What Claude observed	Supports	Matches reported description
17-1	Direct (Vivado 2026.1)	None	Project C:/VivadoProjects/ocup_kr260_deterministic_h w ; title bar "write_bitstream Complete"; Hardware Manager target XFL1DJP2HJSXA . Tcl query maps ocup_kr260_synth_wra_0/inst pins Q[0..3] to nets status_o[12..15]	Design connectivity only	Yes

#	Capture type	Date shown	What Claude observed	Supports	Matches reported description
17-2	Direct (Vivado 2026.1)	None	Same project and target. Four [Vivado 12-1023] "no nets matched" warnings from a hierarchical net filter, then a working query showing axi_gpio_0/gpio2_io_i[12..15] on nets ocup_kr260_synth_wra_0_status_o[12..15]	Design connectivity only (OCUP status bits to AXI GPIO channel 2)	Yes; warnings are from a query filter, not the design

17-3	Direct (Vivado 2026.1)	None	set_property PROGRAM.FILE {C:/VivadoProjects/.../impl_1/ocup_kr260_sys tem_wrapp…} (filename truncated), program_hw_devices [get_hw_devices xck26_0] , End of startup status: HIGH , device reported programmed with no debug cores. Editor shows C:/FPGA-Local/rtl/ocup_kr260_validation_wrapper_v3.s v lines 486–493	**Vivado programming** of xck26_0 from the impl_1 output path	Yes
17-4	Direct (PuTTY COM4)	Sun Sep 6, 2026, 12:08:33 UTC (board clock)	Ubuntu 24.04.4 LTS, kernel 6.8.0-1035-xilinx aarch64, login, echo KR260_ALIVE returns KR260_ALIVE	KR260 boot and responsive terminal	Yes
17-5	Direct (PuTTY COM4), same session as 17-4	Directory dates Sep 6, 17:42 and 18:04	/proc/iomem filter shows only ff0a0000.gpio , no b0000000 region; fpga_manager/fpga0 present; device-tree overlay k26-starter-kits_image_1 (k26_starter_kits.dtbo) status applied	FPGA manager and device-tree investigation	Yes
17-6	Direct (Vivado 2026.1)	None	program_hw_devices , startup HIGH , device programmed; get_bd_addr_segs shows SEG_axi_gpio_0_Reg at OFFSET 0x00B0000000 , RANGE 0x10000 , read-write	**Vivado programming**; AXI GPIO address map	Yes; no PROGRAM .FILE line visible
17-7	Direct (PuTTY COM4), separate boot	Sun Sep 6, 2026, 19:48:13 UTC; journal reports "Realtime clock jumped backwards"	test -r /dev/mem returns DEVMEM_READABLE ; timeout 3 sudo busybox devmem 0xB0000008 32 issued, but no result visible	MMIO access diagnostics (setup only)	Yes; read result not captured
#	Capture type	Date shown	What Claude observed	Supports	Matches reported description
17-8	**Secondary**: PuTTY image viewed inside ChatGPT (chatgpt.com, project "FPGA Runtime	Mon Sep 7, 2026, 01:41:15 UTC (board clock)	fpga_manager/fpga0/state returns operating ; grep "ocup" of the device tree shows no output. ChatGPT sidebar lists "Resume KR260 Run 006" and "Update Run 002B Evidence"	KR260 FPGA manager state (secondary)	Yes; screenshot was taken later than its content

	Authority Processor architecture", 50% zoom)				
17-9	Direct (PuTTY COM4)	None	devmem 0xB0000008 reads 0x00001000 ; after writing 0x60 to 0xB0000000 it reads 0x00001060 , and stays 0x00001060 after 0x20 / 0x60 rewrites and sleep 5 . 0xB0000000 reads back 0x00000060	**KR260 execution** (MMIO behavior at the AXI GPIO address in 176)	Yes

What the screenshots support

- Vivado programming of xck26_0 with startup HIGH : 17-3 and 17-6 (direct). 17-3 shows PROGRAM.FILE set to the impl_1 output path, the same directory that holds the Run 005 bitstream (item 5).
- KR260 execution: 17-9 (direct) shows the status register at 0xB0000008 latching to 0x1060 . 17-4 and 17-7 show boot and MMIO setup.
- Design structure: 17-1, 17-2 and 17-6 tie OCUP status bits 12–15 to AXI GPIO channel 2 at 0xB0000000 .

Discrepancies and limits

1. **Dates precede Run 005.** 17-4, 17-7 and 17-8 show board dates of Sept 6–7, 2026, while the Run 005 bitstream was built Sept 25. Either these are earlier diagnostic sessions or the board clock was wrong: the board was offline, and 17-7 logs a backward clock jump. Board timestamps cannot be used to place these captures in Run 005.
2. **No hash in the programming captures.** 17-3 names only a truncated path, and every build writes to that same impl_1 path. 17-6 shows no file at all. Neither screenshot carries a timestamp. They show *a* bitstream from that path was programmed, not *which* one.
3. **Run assignment unknown for 17-9.** It shows no date or boot banner, so it cannot be tied to Run 005 from the image alone.
4. **Linux view does not reflect the OCUP design.** 17-5 shows the default k26_starter_kits overlay applied and no b0000000 entry in /proc/iomem . That is not a contradiction, since JTAG programming bypasses Linux's FPGA manager, but it means the Linux state does not corroborate an OCUP load.

5. **17-8 is secondary.** It is a photo of a PuTTY window inside a ChatGPT page, captured later than its content, and its sidebar shows Run 006 already existed.

Conclusion: the screenshots corroborate that Vivado programmed the KR260 from the impl_1 path and that the OCUP status register behaved as described. They do not establish that the sealed Run 005 bitstream was the configuration loaded during any execution capture. Programming provenance remains **partially corroborated, not conclusively verified.**

Still to complete: (1) filename, byte count and SHA-256 check against the screenshot inventory, using the files themselves or the inventory listing; (2) review of vivado.jou / vivado_*.backup.jou and hardware-manager logs for a timestamped program_hw_devices with the full PROGRAM.FILE path; (3) the run each capture belongs to, from the journal's own labels and capture metadata.

Checkpoint 13 (2026-10-11): recovered Vivado journal and log — PASS (bounded)

Disposition: **PASS (bounded): programming of the impl_1 path after the Sept 25 build is documented, and no later rebuild appears in the session.** Claude inspected both files directly. They record one continuous Vivado 2026.1 GUI session on host MaxMachine (process 13824), started Wed Sep 23, 2026, 01:46:55 and exited Tue Sep 29, 2026, 23:39:57. It contains five full bitstream builds and eight programming events. This is session evidence, not a readback or attestation of the configuration loaded on the FPGA.

#	File	Size	SHA-256 of copy Claude received	Lines	Reported line content
18	vivado.jou	57,696 bytes (matches report)	900B000DC6FA6D460F15C1A04AAA7CC09FB26C599799EEF106527A49EFE47E73	1,238	Lines 1231–1232 confirmed: set_property PROGRAM.FILE {C:/VivadoProjects/ocup_kr260_deterministic_hw/ocup_kr260_deterministic_hw.runs/impl_1/ocup_kr260_system_wrapper.bit} , then program_hw_devices [get_hw_devices xck26_0]
19	vivado.log	1,525,250 bytes (matches report)	61C259E0B36897E0B3CAFD1BEA84529FF3DF6D7EA31101B700618EF9AB288F18	16,666	Confirmed: 16647 program_hw_devices ; 16648 startup HIGH ; 16649 refresh_hw_device ; 16650 [Labtools 27-1434] device programmed; 16651–16653 [Labtools 27-2269] target error

The archive's own SHA-256 values were not supplied, so these hashes still need comparing with the Dell archive. Both are recovered archival artifacts, not previously sealed Run 005 evidence.

Session timeline (log line numbers; only build launches and two anchors carry dates)

Log line	When	Event

7	Sept 23, 01:46:55	Session start
2867–2938	Sept 23, ~02:11	Build 1 write_bitstream ; programmed from impl_1
7424–7472	Sept 23, ~02:44	Build 2; programmed from impl_1
12013–12079	Sept 23, ~22:24	Build 3; programmed from impl_1
12745–12757	Undated, before Sept 25	Programmed twice from C:/FPGA-Local/evidence/deterministic_harness/… run002B_ready_valid_fix.bit (item 15d)
12791	Undated	reset_run synth_1 , OCUP module reference updated
12829–12837	Sept 25, 20:24:05 → 20:44:20	**Run 005 build**: synthesis launched 20:24:05, impl_1 20:37:21, write_bitstream step 20:44:20 (matches Checkpoint 4)
Log line	When	Event
12871–12876	Undated, after 20:44:20 launch	**Program A**: impl_1 path, startup HIGH , device programmed. Matches screenshot 17-3 command for command
13755–13757	Undated	file copy wrapper → ocup_kr260_validation_wrapper_v3.pre_run005 .sv ; prints "RUN005 BACKUP CREATED"
14765–14767	Undated	**Program B**: impl_1 path, startup HIGH
15812–15870	Undated	Connectivity queries and **Program C** (impl_1 , startup HIGH), then get_bd_addr_segs . One continuous block matching screenshots 17-1, 17-2 and 17-6

16572	Mon Sept 28, 13:19:04	Dated anchor: report_clock_networks report header
16646–16650	Between Sept 28 13:19:04 and Sept 29 23:39:57	**Program D**: impl_1 path, startup HIGH , device programmed (journal 1231–1232)
16651–16664	Same window	Target error, then target closed
16666	Sept 29, 23:39:57	Vivado exits

1. **Dating.** Programming commands carry no timestamps. Programs A–C are bounded between the Sept 25, 20:44:20 bitstream launch and the Sept 28, 13:19:04 report header.
Program D falls between Sept 28, 13:19:04 and Sept 29, 23:39:57. Program A follows the 20:44:20 launch and an open_run impl_1 ; the log does not show the bitstream run finishing (8:45:19 PM per Checkpoint 4) before Program A, so that ordering is inferred, not recorded.

2. **Run 005 association.** After the Sept 25 build, the session contains no further launch_runs , reset_run , write_bitstream , or file operation on the .bit file. The only later file operation is the RTL backup copy. Final run status at log end: impl_1 "write_bitstream Complete!". The impl_1 .bit still carries an 8:45:19 PM Sept 25 modification time and matches the sealed Run 005 hash (items 5, 15e). Together, these strongly corroborate that Programs A–D loaded the Run 005 bytes. Limits: the log records a path, not a hash; another process outside this session could in principle have rewritten the file; and modification times are mutable. The log alone cannot place Run 005's execution captures relative to Programs A–D.

3. **Target error (16651–16664).** Immediately after Program D succeeded and a refresh confirmed the device programmed, Vivado reported [Labtools 27-2269] No devices detected on target localhost:3121/xilinx_tcf/Xilinx/XFL1DJP2HJSXA , with its standard advice to check cable connectivity and board power. Three unrelated commands followed (one a Tcl property error, [Common 17-54]). Then [Labtoolstcl 44513] HW Target shutdown , and Vivado exited at 23:39:57. The log does not show the cause. The error follows all programming in the session, so it does not affect Programs A– D.

Additional finding. The pre_run005.sv backup (Checkpoint 1, item 1) was created inside this session after the Sept 25 build and Program A, not before the build. Its hash still equals the current wrapper and the sealed RTL hash (items 1–3, 14a), so continuity holds. Its name refers to the Run 005 execution, not the build.

Effect on Checkpoint 12. Screenshots 17-1, 17-2, 17-3 and 17-6 are now placed in this session after the Run 005 build, which raises them from undated to log-anchored. Checkpoint 12 stays **PARTIAL**: its screenshot inventory hashes are still unverified against the Dell, 17-9 remains undated, and 17-4, 17-7 and 17-8 still show Sept 6–7 board dates. **Screenshot copies received (for comparison with the Checkpoint 12 inventory)**

#	File received	Bytes	SHA-256
17-1	image1.png	165,987	08B55B5251DEA796A9330B1B4272BD41F8E8B0168B87CB91E1661776F39C8052
17-2	image2.png	175,241	568D32490506D9D8A29984350717ABD9121C4F38E0CF60184372064F615880C0
17-3	image3.png	187,304	49E0A0150E608758D0399DA3ABF9423195507097B6052D8C34D34C442AB27EAB
17-4	image4.png	123,361	66BAF270D67509A2731A56E3DB169F43ECC0856A8D95208F2EA3A585C90165D4
17-5	image5.png	128,828	4B56EF1F615D4F654130FA93D9AA442779E3ADFBF9DFC371936C3CE2048DA802
#	File received	Bytes	SHA-256
17-6	image6.png	160,140	E03F050412CF7F9DC73AEA23166D4B599CB68D3E4FD520EF50D1D685EF2B05BA
17-7	image7.png	120,988	71E0053205CAAFAF6B2FF18F433781027FD5F8BF5B2B182A97D46D88070443B8
17-8	image8.png	414,903	E36F85A1A911D1AADF95319265037718197F4953F51F0A05B042EBA5CFE8E23D
17-9	image9.png	110,384	87D3D79A80BC067CB334C481C19383A484811D3C48FF4F7DC6606693CA0ABA75

All nine are 1920 × 1008 PNG. These are hashes of the copies Claude received; whether they equal the Dell inventory is not yet checked.

Checkpoint 13 supplement: hash reconciliation — PASS (2026-10-11)

All 11 entries in the Dell-generated inventory match the copies Claude inspected, on both byte count and SHA-256. Comparison was run programmatically, not by eye.

Inventory file	Relative path	Bytes	SHA-256	Result

vivado.jou	vivado_journal_recovery\vivado.jou	57,696	900B000D… 7E73	Match
vivado.log	vivado_journal_recovery\vivado.log	1,525,250	61C259E0… 8F18	Match
image1.png (17-1)	screenshot_inventory\image1.png	165,987	08B55B52… 8052	Match
image2.png (17-2)	screenshot_inventory\image2.png	175,241	568D3249… 80C0	Match
image3.png (17-3)	screenshot_inventory\image3.png	187,304	49E0A015… 7EAB	Match
image4.png (17-4)	screenshot_inventory\image4.png	123,361	66BAF270… 65D4	Match
Inventory file	Relative path	Bytes	SHA-256	Result
image5.png (17-5)	screenshot_inventory\image5.png	128,828	4B56EF1F… A802	Match
image6.png (17-6)	screenshot_inventory\image6.png	160,140	E03F0504… 05BA	Match
image7.png (17-7)	screenshot_inventory\image7.png	120,988	71E00532… 43B8	Match
image8.png (17-8)	screenshot_inventory\image8.png	414,903	E36F85A1… E23D	Match
image9.png (17-9)	screenshot_inventory\image9.png	110,384	87D3D79A… BA75	Match

Inventory source: checkpoint13_hash_inventory.csv (1,410 bytes; SHA-256 610F68CD21F30063D4CA0C11CFD95BEB5BE95C5710B04B52F5737667F874E9DE of the copy received). Full 64-character hashes for all 11 files appear in the Checkpoint 13 tables above.

This closes the hash-reconciliation blocker: the files Claude reviewed are byte-identical to the Dell archive copies. It shows archive integrity today, not that the screenshots or logs are unaltered since their original capture.

Effect on Checkpoint 12. Its inventory-integrity condition is now met. Checkpoint 12 stays **PARTIAL**, because its historical-corroboration condition is only partly met: 17-1, 17-2, 17-3 and 17-6 are log-anchored after the Run 005 build, but 17-9 is undated, 17-4, 17-5 and 17-7 show earlier board dates, and 17-8 is secondary.

What Checkpoints 1 to 11 establish, and what they do not

Checkpoints 1 to 11 show artifact identity and **source-hash continuity**. They do not prove **the exact build inputs** of any bitstream.

Established

- The wrapper and context-sequencer source files on the Dell today are byte-identical to the Run 005 records.
- The wrapper has not changed between its pre-Run 005 snapshot and its current state.
- The bitstream preserved in impl_1 is byte-identical to the sealed Run 005 bitstream record.
- A Vivado implementation run on September 25 completed through bitstream generation without errors in its displayed summary.
- The Sept 25 implementation script documents the workflow: open the routed checkpoint (SHA-256 C8D850A4…4337), then write ocup_kr260_system_wrapper.bit . A documented Sept 25 build
- sequence runs from synthesis script (8:24:05 PM) to final bitstream (8:45:19 PM), for top ocup_kr260_system_wrapper on part xck26-sfvc7842LV-c .
- The Sept 25 OCUP out-of-context synthesis script reads both OCUP RTL files by path from C:/FPGA-Local/rtl/ , documenting how they enter the build hierarchy.

Checkpoint 6 adds a hash for the OCUP out-of-context synthesis checkpoint (20913C5B… E2D6 , 8:33:32 PM), the artifact that script writes.

Checkpoint 7 adds verified copy integrity for 27 current non-RTL project inputs (.bd , .xci , .xdc), archived with a SHA-256 manifest.

Checkpoint 8 adds the generated top ocup_kr260_system_wrapper.v (417F0E50…53DB), archived with source/archive equality verified; its 8:24:02 PM timestamp sits 3 seconds before synthesis start.

Checkpoint 9 confirms that the RTL files at the exact C:\FPGA-Local\rtl\ paths the synthesis script reads match the Run 005 sealed hashes today.

Checkpoint 10 inventories four earlier distinct bitstreams (Sept 23–24) and confirms ocup_kr260_run005_SEALED.bit is byte-identical to the Vivado impl_1 output.

Checkpoint 11 confirms the sealed Run 005 record names the programmed bitstream by the same SHA-256 and documents Vivado 2026.1 and xck26_0 startup HIGH .

Not established

- A fully verified, reproducible RTL-to-bitstream derivation. The 27 current non-RTL inputs are now archived and hashed, but there are no Sept 25 hashes to show they are unchanged since the build. The generated wrapper is now archived and hashed (Checkpoint 8). Vivado version and run settings are not yet captured.
- The exact bytes Vivado read from C:/FPGA-Local/rtl/ on Sept 25. The script records paths, not content; today's matching hashes do not show the historical bytes.
- That the Checkpoint 1 hashes were taken from the files at those exact C:/FPGALocal/rtl/ paths. Resolved by Checkpoint 9 for present-day files; historical contents still not proven.
- That the recovered timestamps are original. They are consistent with one run but are mutable file metadata.
- Why each build occurred and what source changes caused the bitstream differences. Checkpoint 10 documents four earlier distinct bitstreams (Sept 23–24), but filename labels and timestamps do not establish rationale. Unresolved; no conclusion is drawn. •Conclusive programming provenance. The sealed record names the programmed bitstream by hash (Checkpoint 11), and the Vivado session log shows four programming events from the impl_1 path after the Sept 25 build with no later rebuild (Checkpoint 13). No readback or cryptographic attestation of the configuration loaded during Run 005 exists, and the log records a path, not a hash. Strongly corroborated, not conclusively verified.

Open questions for later checkpoints

These items would move the record from source continuity toward build-input proof.

- [ ] **Rebuild rationale (unresolved):** What records (commit history, notes, test logs) show why each of the Sept 23–25 builds occurred and which source changes produced each distinct bitstream?
- [ ] Do routed checkpoints or source snapshots survive for the four earlier bitstreams (items 15a–15d)?
- [x] ~~RTL files at the script's C:\FPGA-Local\rtl\ input paths hashed; both match Run 005 sealed hashes (Checkpoint 9).~~
- [ ] Verify every entry against the complete Run 005 seal document, and confirm the seal holds no artifacts beyond the three already matched.
- [ ] Is the 27-file non-RTL inventory complete? Cross-check it against the .xci and constraint paths named in the Sept 25 synthesis and implementation scripts.

- [x] ~~Non-RTL input timing: 26 of 27 modified before 8:24:05 PM; **dont_touch.xdc** is a samesecond exception (Checkpoint 7 supplement).~~
- [ ] Beyond the archived wrapper, are there other generated or IP output products to archive?
- [ ] Vivado version is recorded as 2026.1 in the sealed record (Checkpoint 11). Do the synthesis and implementation log headers confirm it, and which run strategies were used?
- [ ] Do the full synthesis and implementation logs show warnings or critical warnings beyond the displayed summaries?
- [ ] **Programming provenance (strongly corroborated, not conclusively verified):** no readback or attestation exists for Run 005, and none can be recovered after the fact. Remaining gap: which of Programs A–D preceded each Run 005 execution capture.
- [ ] **Checkpoint 12 completion:** verify the nine screenshots' filenames, byte counts and SHA-256 against the inventory; assign each capture to a run (17-4, 17-7 and 17-8 show Sept 6–7 board dates); review vivado.jou and hardware-manager logs for a timestamped program_hw_devices with the full PROGRAM.FILE path.
- [x] ~~Checkpoint 13 hashes and the nine screenshot hashes match the Dell archive: 11 of 11, bytes and SHA-256 (Checkpoint 13 supplement).~~
- [x] ~~Vivado journal and log inspected: programming after the Sept 25 build documented; no later rebuild (Checkpoint 13).~~
- [x] ~~SHA-256 of the preserved .bit file vs. sealed Run 005 bitstream record: exact match (Checkpoint 2).~~
- [x] ~~Routed checkpoint located and hashed; implementation script's open_checkpoint → write_bitstream step documented (Checkpoint 3).~~
- [x] ~~Sept 25 synthesis-to-bitstream timestamps recovered and in order, 8:24:05 PM to 8:45:19 PM (Checkpoint 4).~~
- [x] ~~Source hierarchy documented: OCUP out-of-context synthesis script reads both OCUP RTL files by path (Checkpoint 5).~~
- [x] ~~OCUP out-of-context synthesis checkpoint hashed (Checkpoint 6).~~
- [x] ~~27 current non-RTL inputs (.bd, .xci, .xdc) archived with verified SHA-256 manifest (Checkpoint 7).~~
- [x] ~~Generated top ocup_kr260_system_wrapper.v archived and hashed; source/archive equality PASS (Checkpoint 8).~~

Checkpoint log

Checkpoint	Date	Summary	Record state
SEAL	2026-10-11, 5:01 PM CT	Bounded Phase I provenance seal authorized by MAXBRUCE. Scope, six attestations and six exclusions in effect. Checkpoints 1–13 unchanged (12 PARTIAL; 13 PASS, bounded). Programming provenance strongly corroborated, not conclusively verified. Separate from behavioral PASS results; no securityproperty certification. Later findings by dated addenda only.	SEALED
13	2026-10-11	PASS (bounded): journal and log inspected; one session Sept 23–29; Run 005 build at Sept 25 20:24–20:44; four impl_1 programming events afterward (startup HIGH), no later rebuild; screenshots 17-1, 17-2, 17-3, 17-6 log-anchored; final target error follows all programming, cause not shown. Supplement: 11 of 11 Dell inventory hashes and byte counts match; Checkpoint 12 inventory integrity met, classification stays PARTIAL	Open, not sealed
12	2026-10-11	PARTIAL, not PASS: nine screenshots inspected. Vivado programming (17-3, 17-6) and KR260 MMIO behavior (17-9) corroborated; no bitstream hash or timestamp in programming captures; 17-4, 17-7, 17-8 show Sept 6–7 board dates, before the Run 005 build; 17-8 is secondary; inventory hashes not verified	Open, not sealed

Checkpoint	Date	Summary	Record state

11	2026-10-11	PASS: sealed record (run_005_SEALED_PASS.md , lines 76–85) names programmed bitstream by matching SHA-256; Vivado 2026.1; xck26_0 startup HIGH ; no readback or attestation, so programming provenance partially corroborated	Open, not sealed
10	2026-10-11	PASS: six .bit files, five distinct hashes, all 7,797,916 bytes; four earlier distinct bitstreams (Sept 23–24); ocup_kr260_run005_SEALED.bit byte-identical to Vivado output; rebuild rationale and KR260 programming provenance unresolved. Follow-up: item 15d 10:00:00 PM LastWriteTime unverified (CreationTime 10:48:49 PM); hash and PASS unchanged	Open, not sealed
9	2026-10-11	PASS: RTL files at the script's C:\FPGALocal\rtl\ input paths match Run 005 sealed hashes; present-day consistency only, not historical contents or reproduced derivation	Open, not sealed
8	2026-10-11	PASS (archive integrity): generated top ocup_kr260_system_wrapper.v archived and hashed (417F0E50…53DB); timestamp 3 s before synthesis start	Open, not sealed
7	2026-10-11	PASS: 27 of 27 non-RTL inputs (.bd , .xci , .xdc) archived and SHA-256 verified; copy integrity only, not historical reproducibility. Supplement: 26 of 27 predate 8:24:05 PM; dont_touch.xdc is a same-second exception	Open, not sealed

Checkpoint	Date	Summary	Record state

6	2026-10-11	OCUP out-of-context synthesis checkpoint hashed (`20913C5B…E2D6`); master evidence index added; both DCPs classified as newly inventoried, not Run 005 sealed evidence; rebuild difference unresolved	Open, not sealed
5	2026-10-11	OCUP out-of-context synthesis script reads both hashed OCUP RTL paths; source-hierarchy gap substantially closed; historical bytes and rebuild difference unresolved	Open, not sealed
4	2026-10-11	Sept 25 synthesis-to-bitstream sequence recovered (8:24:05–8:45:19 PM); block design and non-RTL inputs still to inventory; rebuild difference unresolved	Open, not sealed
3	2026-10-11	Routed checkpoint hashed and added to inventory; Sept 25 script documents routed checkpoint → bitstream step; rebuild difference still unresolved	Open, not sealed
2	2026-10-11	Preserved `impl_1` bitstream matches sealed Run 005 bitstream record; rebuildhistory question carried forward	Open, not sealed
1	2026-10-11	Three source-hash matches to Run 005 records; Sept 25 implementation log clean in displayed summary	Open, not sealed

PUBLIC STANDARDS CONVERSATION

Temporal Authority in the Public Standards Conversation

OCUP Contributions to the Open Secure AI Alliance / SAFE Discussions

Public engineering record • August–October 2026

An engineering proposal earns credibility when others can read it, challenge it, compare it with alternative approaches, and identify what remains unproven. The Open Secure AI Alliance (OSAA) SAFE repository provides one such public forum. This chapter documents selected contributions from QSAFP-Core / Temporal Authority Systems PBC and subsequent responses by other participants. It records participation and dialogue—not Alliance endorsement, standards adoption, or certification.

1. From incident logs to evidence of authority

SAFE Issue #1, opened August 2026 by QSAFP-Core, asks whether incident evidence can establish more than what a governed agent reported. In particular: who authorized consequential operation, when was that authorization valid, could the system restore authority to itself, and did an independent enforcement boundary actually withhold capabilities after authorization ended?

Its central distinction is simple: the recoverability of a workload must not imply that its authority is automatically recoverable. A log of a denial is not the same as evidence that a consequential capability became unavailable. The proposal therefore calls for an implementation-neutral authority-evidence dimension alongside other SAFE incident records.

The Issue #1 conversation also sharpened what strong evidence requires. In response to discussion of independent verification, QSAFP-Core distinguished verifying a signature from proving the recorded event was true. It identified the need for protected custody, trusted monotonic state, and evidence production resistant to unilateral suppression or revision. These remain engineering requirements, not demonstrated OCUP silicon capabilities.

Primary record: SAFE Issue #1 https://github.com/OpenSecureAIAlliance/RFCs/issues/1

Evidence-verification exchange https://github.com/OpenSecureAIAlliance/RFCs/issues/1#issuecomment-5420855704

2. When a valid lease is no longer applicable

Issue #15, initiated by another contributor, addresses a different question: what happens when a system materially changes while its previous authorization has not yet expired? A different tool set, model, delegated scope, risk setting, or operating environment can invalidate the assumptions underlying the earlier approval.

QSAFP-Core proposed a fail-closed transition: prior authority does not silently carry forward across an authorization-relevant change. A fresh externally governed decision must bind authorization to the new operating baseline. Subsequent discussion developed candidate conformance scenarios covering non-material change, material change under an unexpired grant, and a second change occurring after reauthorization but before action release.

These scenarios are proposed tests; participation in drafting them is not evidence that an implementation has passed them. The distinction matters because a deny-everything device must not pass the same test as a correctly governed system capable of recognizing a genuinely valid new grant.

Primary record: Issue #15 https://github.com/OpenSecureAIAlliance/RFCs/issues/15

QSAFP-Core proposed vectors https://github.com/OpenSecureAIAlliance/RFCs/issues/15#issuecomment-5561881592

3. Authority cannot grow through composition

Issue #26 examines workflows in which each individual action passes its local check but the combined result exceeds the intended boundary. The examples include cumulative information disclosure, delegated privileges, repeated commitments and cross-organizational transitions.

QSAFP-Core proposed that delegation, retries, workflow transitions and organizational handoffs must not enlarge effective authority beyond the originating authorization. Another contributor then identified an important limit: authority monotonicity alone does not prevent harmful aggregate outcomes when authority is unchanged. QSAFP-Core accepted that distinction and separated two questions: the authority trajectory and the outcome-exposure trajectory.

The result is a useful engineering boundary. Hardware-rooted continuity can help establish identity, state history and resistance to reset, but does not by itself solve the policy-dependent semantics of information aggregation. That work needs separate cumulative rules and evidence.

Primary record: Issue #26 https://github.com/OpenSecureAIAlliance/RFCs/issues/26

Two-trajectory clarification https://github.com/OpenSecureAIAlliance/RFCs/issues/26#issuecomment-5535871639

4. Scope, identity and operational authority are different

Several discussions converge on one principle: knowing which agent acted, possessing an earlier approval, or retaining a valid credential does not establish that the particular action was currently authorized. In Issue #29, QSAFP-Core recommended evaluating delegated authority at action release against the originating grant and every intervening constraint. A downstream component's broader credentials must not silently enlarge the authority conveyed through the chain.

In Issue #47, the discussion of "approved scope" distinguished what a principal permits from whether permission still exists at execution. QSAFP-Core further warned that letting a system "refresh" its scope must not become a disguised extension of the authorization period or restoration of expired standing.

For systems crossing organizational or national boundaries, the same separation appears in the Issue #1 discussion: portable identity and provenance do not automatically confer portable operational authority. RAAVE is proposed as supporting continuity and compliance evidence; it does not decide applicable law or grant new operating permission.

Delegation: Issue #29 https://github.com/OpenSecureAIAlliance/RFCs/issues/29

Scope clarification: Issue #47 https://github.com/OpenSecureAIAlliance/RFCs/issues/47#issuecomment-6042330532

Jurisdictional authority: Issue #1 https://github.com/OpenSecureAIAlliance/RFCs/issues/1#issuecomment-5609370632

5. A concrete response from another engineering team

Pull Request #71 proposes approval-to-execution scope evidence. QSAFP-Core commented that recording checks does not itself establish that they are mandatory gates before action release, and called for evidence of the operational baseline and material-change status evaluated at that precise moment.

In reply, CSOAI-ORG acknowledged that its current candidate implementation did not yet carry a complete system, policy and dependency baseline or an explicit material-change reauthorization trigger. It supported

preserving the requirement in implementation-neutral terms and suggested a future conformance fixture. This is substantive technical engagement—not acceptance of OCUP, nor a claim that the gap has been closed.

PR #71 and engineering exchange https://github.com/OpenSecureAIAlliance/RFCs/pull/71#issuecomment-6043255619

6. What the public record establishes—and what it does not

The record establishes that OCUP's contributors placed temporal authority and self-extension denial into an open technical discussion and engaged with other participants on evidence, revalidation, delegation, composition, jurisdiction and action-release checks. Some responses built on the distinctions, challenged their limits, or identified related implementation gaps. These are useful signals of technical relevance and opportunities for collaboration.

The record does not establish SAFE adoption of the proposed authority-evidence dimension, ratification of an OCUP standard, independent reproduction of Phase I hardware findings, formal endorsement by the Alliance, or successful conformance of the proposed examples. Issue and pull-request status may change; the linked GitHub records are the authoritative public history.

7. A standards path open to many implementations

OCUP's stated intention is an implementation-neutral evidence and conformance framework. The questions should remain accessible to any engineering team: Is authority externally issued? Does it expire? Can the governed system renew or restore it? Are consequential capabilities physically mediated? Can an independent reviewer establish which authority, baseline and enforcement decision applied at release?

The proposal's value lies in making these questions testable, not in requiring a single chip design. QSAFP-Core has also stated a standards/IP direction favoring inspection, independent testing, interoperability, alternative conforming implementations and a FRAND-friendly approach if applicable patented technology becomes standards-essential. That is a declared policy direction, not an adopted licensing standard.

The next step belongs to evidence: independently reproducible hardware tests, positive and negative conformance cases, signed export, trusted time, protected keys, and real consequential resource mediation. Public dialogue can sharpen the specifications; only engineering and independent verification can establish the results.

Selected public records

#1 — Temporal authority and self-extension denial https://github.com/OpenSecureAIAlliance/RFCs/issues/1

#15 — Revalidation after material change https://github.com/OpenSecureAIAlliance/RFCs/issues/15

#26 — Composition governance https://github.com/OpenSecureAIAlliance/RFCs/issues/26

#29 — Delegated-authority provenance https://github.com/OpenSecureAIAlliance/RFCs/issues/29

#47 — Approved-scope definition https://github.com/OpenSecureAIAlliance/RFCs/issues/47

PR #71 — Approval-to-execution scope evidence https://github.com/OpenSecureAIAlliance/RFCs/pull/71

Editorial status: curated public discussion through October 11, 2026. All items identified above were open when reviewed. Discussion contributions are not adopted standards.

PATENTS

IP &

COMMERCIALIZATION

THE AUTONOMOUS ERA & INTELLECTUAL PROPERTY

Solo Founder & Architect

As a solo founder, I studied intellectual property systems extensively over the years by simply reading issued patents and published patent applications. Lots of them. I had to because I couldn't afford to hire attorneys to do the deep research necessary for initial filings. So, in time, I developed my own method of operation for IP that goes like this: I research and write the initial filings myself, and if the ISR (International Search Report) and accompanying written opinion come back favorable, I try to hand the case off to a big law firm that may be more willing to consider taking it on pro bono at that stage. Of course, you have to demonstrate financial need to qualify for that kind of assistance, and you also have to really understand your inventions inside out to prepare those initial filings. A favorable international search result is encouraging, but it doesn't guarantee a patent, or legal representation. I was fortunate to find an excellent firm to work with but I was asked not to mention them by name. Thank you, no name.🙌

Over the years, I became pretty good at both. Not being able to afford legal services was easy. I don't have the money. Understanding my inventions? That, my friends, was and is a labor of love. I've always been confident in writing and making the initial filings because who knows the invention better than the architect? You can hire all the law firms you can afford, but if they don't truly understand your vision, you may wind up with pages of legal language that describe the invention without fully capturing what makes it significant. And in a system where human judgment plays such an important role, that distinction matters.

Enter AI (Allied Intelligence in my book). Patent offices have always had their work cut out for them trying to manage the unmanageable, (nothing new under the sun) but with AI rapidly transforming intellectual property research, drafting, and the broader IPR (intellectual property rights) conversation, that work is now greatly magnified. The USPTO is fortunate to have a progressive director like John Squires, who seems to get the magnitude of the challenge and is actively pursuing remedies to keep pace with fast-moving technologies while preserving the core purpose of the Office.

I believe that when intelligence and wealth are securely and widely distributed, humanity gains unprecedented capacity to comprehend and improve our world. That is the opportunity of the Autonomous Era. This may also be the era of open-source and open-core developments becoming "the way." As code, models, and architectures become easier to reproduce, adapt, and "fork," the practical value and enforceability of different intellectual property rights may change. Copyrights and patents will continue to serve important purposes, but I believe trademarks will become increasingly valuable in the open ecosystem, helping people distinguish trustworthy technologies, services, and their origins.

My primary reason for pursuing patents has always been defensive. I want the inventions documented, the originating work protected, and the freedom to pursue their development with capable partners. Patent ownership alone does not guarantee Freedom to Operate—that requires understanding the rights of others. But a well-constructed portfolio can strengthen the inventor's position while that freedom is established through research, engineering, licensing, and collaboration.

In Summary

These are very interesting times we're living in and today's leaders have the tremendous responsibility to get it right for the good of the whole and generations to come. Intellectual Property Rights still deserve respect beyond legal ramifications because it incentivizes innovation for a limited time, which is a pretty good motivator for creators, investors and consumers to say the least. It keeps the wheels of commerce turning and generates opportunities for anyone paying attention to these matters. We must not lose this grand system of measurements while at the same time keeping pace with the 21st century needs and circumstances.

The following represents our IP portfolio with the exception of copyrights which I didn't see any reason for listing yet represents our tremendous code base in our open-core repos as well as books and manuscripts.

There has never before been a need for "the commercialization of human authority" but that time has come. FRAND principles are in order.

Substrate is needed for the world.

THE INTELLECTUAL PROPERTY & COMMERCIALIZATION OF HUMAN AUTHORITY

From originating invention to a potentially global semiconductor infrastructure

1. The Opportunity Beyond the Chip

The next great semiconductor opportunity may not be defined solely by how much intelligence a processor can deliver, how quickly it can compute, or how efficiently it can consume energy.

It may also be defined by **who has the authority to permit that intelligence to act—and how that authority is enforced.**

For decades, the semiconductor industry has expanded humanity's computational capabilities. Processors became faster, smaller, more specialized, and increasingly interconnected. Graphics processing units evolved into powerful engines for scientific computing and artificial intelligence. Entire industries arose around architectures that made formerly impractical computations commercially useful.

Now another transition is underway.

As autonomous systems acquire greater capacity to make decisions, operate machinery, move through physical environments, manage infrastructure, and initiate consequential actions, the industry faces a new architectural question:

Can a system be technically capable of acting while remaining physically unable to act without valid external authority?

OCUP—One Chip Unified Protocol—originated from an effort to answer that question.

Its governing invariant is direct:

The governed system cannot authorize, extend, restore, or enlarge its own operational authority.

That proposition is larger than a conventional security feature. It introduces the possibility of a distinct infrastructure function: separating computational capability from continuing permission to exercise it.

Today, processors execute instructions. Operating systems allocate resources. Security modules protect identities, keys, and selected execution environments. Software monitors behavior and can attempt to intervene.

OCUP proposes an additional boundary: a technically enforceable condition under which consequential operation is permitted only within externally authorized limits.

That boundary is the originating opportunity.

It is not a claim that existing processor security is ineffective, nor that no related mechanisms exist. The relevant question is whether a sufficiently general, independently enforceable temporal-authority architecture can be developed and adopted across otherwise different computing systems.

If it can, the opportunity extends far beyond manufacturing a single chip.

It becomes the prospect of establishing a reusable component of autonomous computing infrastructure.

2. An Intellectual-Property Position Designed for Multiple Architectures

The OCUP intellectual-property program began with technical inventions directed toward hardware-enforced temporal authority, externally authorized operational leases, self-extension denial, protected enforcement boundaries, and verifiable authority records.

The QSAFP international application pursues protection across several complementary forms of invention, including systems, methods, and enforcement apparatus.

National Security Review and Foreign Filing Authorization (July 2025–April 2026)

Following the June 13, 2025 filing of U.S. Provisional Patent Application No. 63/823,363, *Quantum-Secured AI Fail-Safe Protocol (QSAFP)*, the application entered an extended national security review beginning in July 2025. The review involved the Department of Defense and two additional government agencies whose identities have not been disclosed. For approximately nine months, the application remained subject to foreign-filing restrictions while the review continued. Importantly, **no formal secrecy order was issued**. The prolonged review nevertheless delayed the inventor's ability to pursue international patent protection, prompting a petition to the United States Patent and Trademark Office requesting expedited authorization or a formal determination of the application's status.

On **April 14, 2026**, the USPTO granted Foreign Filing License No. **656,154**, authorizing foreign patent applications covering the disclosed subject matter. This decision ended the documented foreign-filing impediment and opened the way for international prosecution of the originating

invention. The review and subsequent authorization are significant milestones in the development history of QSAFP and the broader OCUP architecture. They are not presented as government endorsement, a finding of patentability, a determination of technological merit, or certification of national-security suitability. Rather, they document the passage of a foundational invention through a consequential government review process and the eventual authorization to pursue its protection internationally.

Crucially, the architecture is not defined solely by the physical dimensions or packaging of one security chip.

Its disclosure contemplates enforcement at hardware, firmware, secure-element, hardware security module, and accelerator-fabric boundaries. They address execution resources associated with graphics processors, tensor processors, neural processors, robotic controllers, memory, compute-fabric endpoints, and physical actuation.

This creates three important potential implementation pathways.

A standalone companion chip could independently validate authority credentials and mediate access to protected resources through interfaces it actually controls.

An integrated chiplet could place authority-enforcement functionality within a multi-die processor package, potentially bringing protected control closer to the computing resources being governed.

A licensed semiconductor IP block could allow manufacturers to incorporate qualifying enforcement mechanisms directly into future processors, accelerators, controllers, or system-on-chip designs.

These are intended engineering configurations, not assertions that production implementations have already been completed or that existing patent applications guarantee protection across every configuration.

Indeed, no discrete component can control a resource merely by existing alongside it. Effective authority enforcement requires a non-bypassable relationship to the resource being governed. Power, reset, configuration, clock continuity, key custody, and alternative execution paths must be considered in the architecture.

That is why the prospective semiconductor partner is so important.

The originating patent position is designed to support a family of technical implementations. The industry expertise required to make those implementations independently trustworthy remains an essential part of bringing the architecture home.

The broader invention program also includes related protection efforts addressing complementary needs.

AEGES concerns quarantine, containment, recovery, and, in its submitted claim set, digital-asset protection mechanisms. RAAVE addresses authority verification, provenance, continuity, and cross-jurisdictional compliance signaling. PBTA, whose disclosure includes microfluidic and MEMS-oriented embodiments, concerns physical tamper detection, identity, and security containment.

Together, these efforts suggest a layered intellectual-property strategy: authority enforcement at the core, with complementary inventions directed toward verification, containment, physical protection, and recovery.

Their legal scopes are distinct. Their relationship is architectural rather than an assertion that every filing protects every feature.

The commercial significance of such a portfolio lies in the possibility of protecting multiple qualifying contributions to a common infrastructure problem.

One invention may establish an enforcement mechanism. Complementary inventions may help make that mechanism verifiable, resilient, interoperable, and deployable.

Pending applications are not granted patents. The eventual scope of protection depends on prior art, sufficient technical disclosure, examination, jurisdiction, and the claims ultimately allowed. The program also remains subject to freedom-to-operate considerations and the patent rights of others.

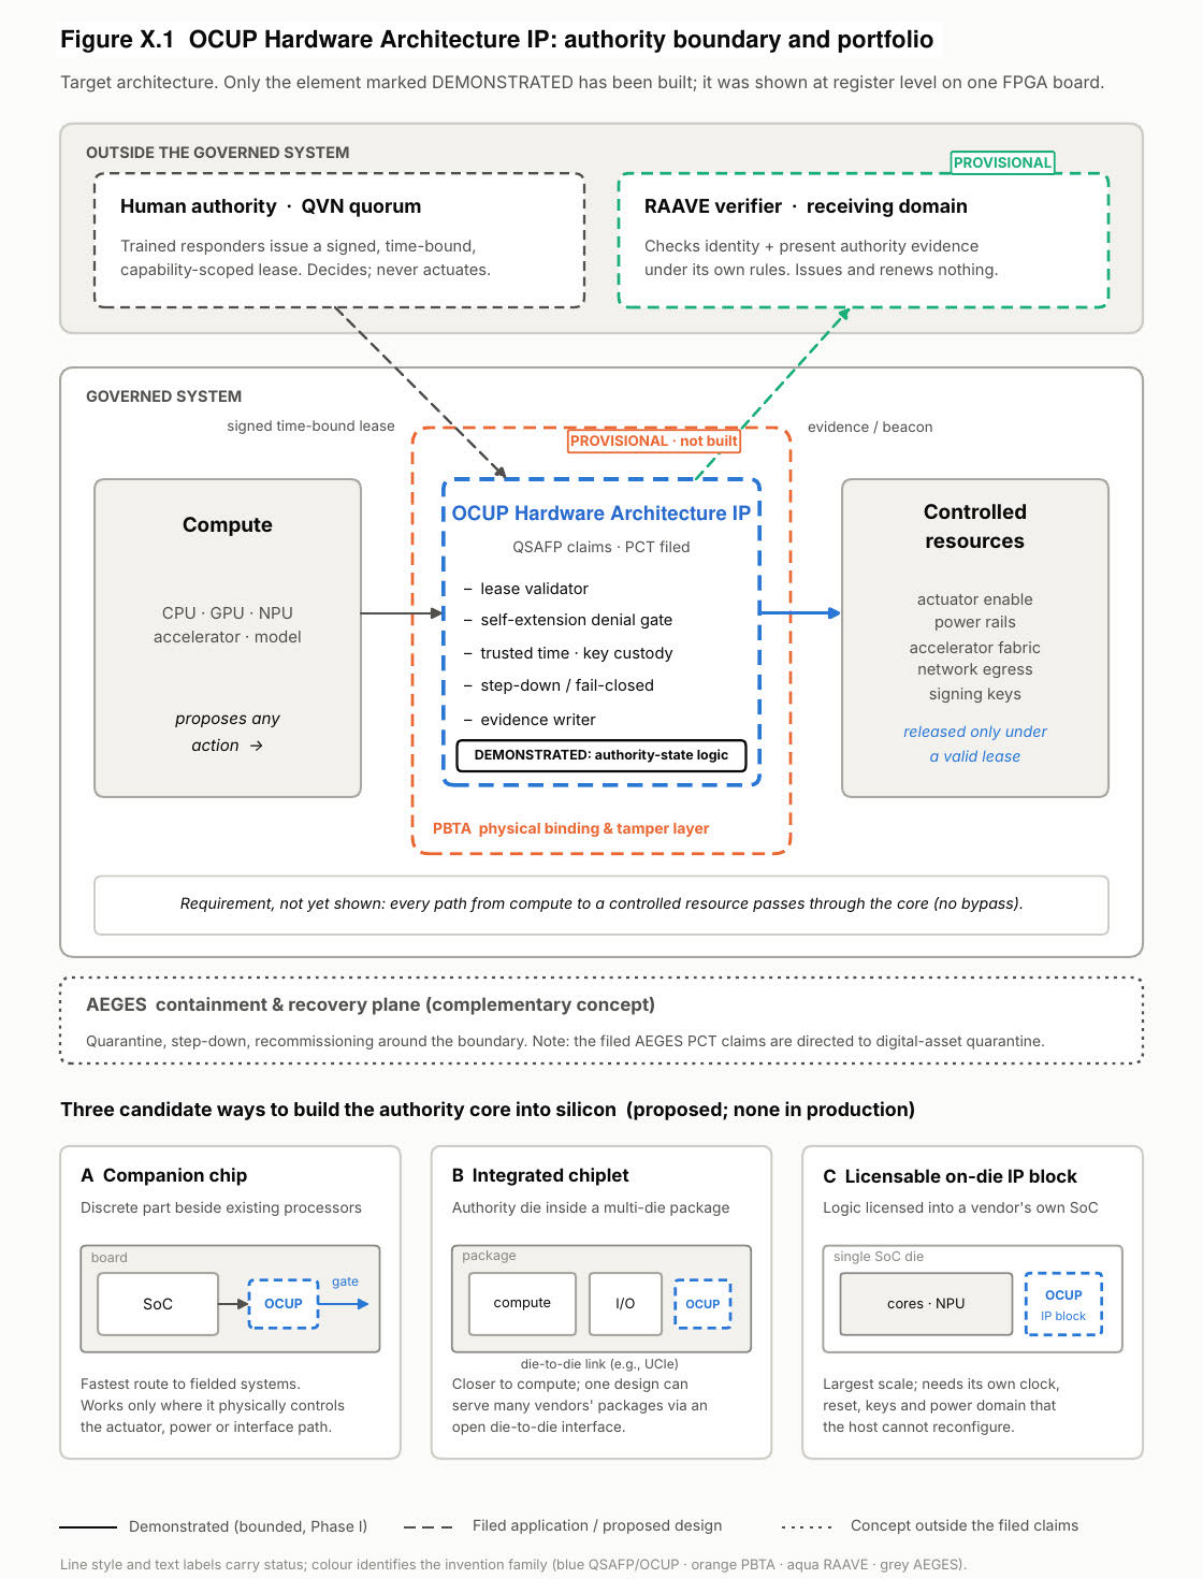
But an originating intellectual-property program need not wait for completed manufacturing to establish a serious development position.

It must establish what has been invented, what has been disclosed, what protection is being pursued, and what additional work will determine its practical value.

That work has begun.

The Hardware Architecture: Where Each Invention Fits

Figure X.1 shows the authority boundary at the center of the portfolio and how the related inventions are intended to attach to it. It is a target architecture. Its only demonstrated element is the authority-state logic recorded at register level on a single FPGA reference board during Phase I.



QSAFP/OCUP is the foundation: an authority core placed between compute and the resources whose use has consequences, such as actuator enablement, power rails, accelerator fabric, network egress, and signing keys. The processor may propose any action. The core is designed to release it only

under a valid, externally issued, time-bounded lease, and to deny any attempt by the governed system to renew, extend, restore, or enlarge that lease.

PBTA addresses the physical question every hardware boundary eventually faces: is this authority component genuinely bound to the machine it governs? Its provisional disclosure describes microfluidic and other materially responsive structures intended to support identity, anti-separation, tamper response, and authenticated evidence. RAAVE sits outside the governed system as an independent verifier, evaluating identity and present-authority evidence under a receiving domain's own rules; it is designed to issue and renew nothing. AEGES contributes a complementary containment and recovery discipline. Its filed international application is directed to digital-asset quarantine and recovery; machine-level containment around the authority boundary remains an architectural concept.

The same authority core is intended to be implementable in three forms: a companion chip beside existing processors, a chiplet within a multi-die package, or a licensable block within a vendor's own system-on-chip. In every form the decisive requirement is the same, and it has not yet been demonstrated: no path from compute to a controlled resource may bypass the core. None of these integrations has been built.

Prior Work and the OCUP Distinction

The idea that a chip might require periodically renewed authorization in order to operate did not originate with OCUP, and this chapter does not claim otherwise. Policy researchers proposed hardware "operating licenses" for AI accelerators in 2024 (Aarne, Fist and Withers, Secure, Governable Chips, Center for a New American Security), and later work on hardware-enabled guarantees has explored related mechanisms for compute governance. Feature-licensing schemes, secure elements, trusted execution environments, and the functional-safety interlocks of industrial machinery each govern some aspect of what a system may do.

OCUP's contribution is directed elsewhere. Its central concern is not whether a chip holds a current license, but whether a governed system, including an autonomous agent acting through that chip, can make itself the source of its own continued authority. The architecture combines a gate against self-renewal, extension, and restoration; renewal through externally governed, multi-party human authorization; capability that contracts rather than expands under uncertainty; bounded offline operation the system cannot lengthen; and authority-state evidence suitable for independent verification. Whether, and how broadly, those contributions are protected will be decided through examination, not by this description.

3. The First-Mover Opportunity: Defining the Interface

First-mover advantage is often misunderstood.

Being first to describe a problem does not confer ownership of its solution. Filing an application does not establish technical superiority. And arriving early does not guarantee commercial leadership.

The more consequential opportunity is to help establish the architecture, interfaces, evidence requirements, and implementation conventions around which a developing industry may organize.

Consider the importance of CUDA to the development of general-purpose GPU computing.

When NVIDIA introduced CUDA in 2006, the larger opportunity was not merely another processor feature. It was a programming model and development environment that made a class of computing capability more accessible and commercially useful. Over time, engineering investment, software tools, libraries, hardware generations, and developer participation helped turn the initial platform into an extensive ecosystem.

OCUP is not CUDA, and its present maturity is not comparable to NVIDIA's original hardware and engineering resources. The analogy concerns a strategic principle, not an equivalence in accomplishment.

The enduring value of a platform can arise from what others are enabled to build upon it.

For temporal authority governance, the corresponding opportunity may reside in defining a common way to express and enforce such matters as:

- The identity and scope of a governed capability.
- The externally issued authority under which that capability may operate.
- The duration and conditions of that authority.
- The hardware response to expiration, revocation, invalidity, or uncertainty.
- The prohibition against unauthorized self-extension or restoration.
- The evidence needed for independent verification and later accountability.

These are not merely software-interface questions. They are questions about which parts of a computing system must be trusted, which components must be independent, and what physical consequences follow when authority is absent.

An industry-wide answer cannot be created simply by publishing a manifesto or drafting a patent application. It must emerge through interoperable engineering, repeatable testing, implementation experience, and participation by organizations with the capacity to deploy the technology.

Nevertheless, early architectural contributions can be strategically consequential.

The organizations that help establish credible conformance tests may influence what manufacturers must demonstrate. Those that develop dependable interfaces may influence integration choices. Those that participate in early implementations may acquire valuable technical experience before market requirements become settled.

The first-mover opportunity is not to declare a standard. It is to help define and substantiate one before competing implementation conventions become entrenched.

That distinction is especially important in autonomous computing.

The field spans competing semiconductor architectures, robotics platforms, cloud infrastructure, industrial control, defense, transportation, medicine, finance, and public infrastructure.

No single processor manufacturer, national authority, or application provider can reasonably be assumed to control all of these environments.

A useful authority substrate must therefore be adaptable to different platforms, risk levels, and legal jurisdictions without requiring one global operator to control them.

OCUP's intended independence from a particular compute architecture makes this an especially important direction for development.

And it opens a potentially much larger market than a single proprietary device.

4. The Commercial Potential of Global Substrate Adoption

The commercialization opportunity should be evaluated at several levels.

The first is **semiconductor integration**.

A proven authority-enforcement mechanism could potentially be offered as a discrete component, chiplet, embedded IP block, or associated integration package.

Manufacturers could evaluate it for selected accelerators, robotic controllers, industrial processors, and other systems where consequential execution requires enforceable external authorization.

Such implementations could support component sales, semiconductor IP licensing, engineering royalties, and long-term integration agreements.

The second is **platform infrastructure**.

A widely adopted authority interface could create demand for credential provisioning, protected evidence generation, interoperability testing, lifecycle management, technical support, and assurance services.

These functions would need carefully defined security responsibilities. A commercial service must not become an unauthorized substitute for the human authority the architecture is intended to preserve.

The third is **certification and insurance-related evidence**.

If independent testing eventually establishes enforceable authority boundaries, standardized evidence could help manufacturers, operators, insurers, and regulators evaluate specific operational safeguards.

That could support conformity assessment, independent testing, underwriting research, and contractual risk allocation.

It would not automatically reduce insurance premiums, eliminate liability, or demonstrate improved field safety. Those outcomes would require operational data, validated performance, and appropriate institutional acceptance.

The fourth is **standards-based global licensing**.

A broadly interoperable authority architecture could be implemented by different manufacturers across successive device generations.

That is where the scale of the opportunity becomes especially interesting.

Consider the difference between earning revenue from a single engineering engagement and receiving a commercially negotiated license payment associated with a qualifying implementation across a large installed base.

For illustration only, one million licensed implementations at a hypothetical royalty of one dollar each would generate one million dollars in gross royalties. Ten million would generate ten million dollars. One hundred million would generate one hundred million dollars.

Those figures are arithmetic scenarios, not market forecasts or proposed FRAND rates. They assume licensing rights, adoption, enforceable coverage, and commercial terms that do not yet exist.

Actual economics would depend on the patented contribution, product value, alternatives, integration costs, applicable legal obligations, competing technologies, and negotiations.

But they demonstrate why infrastructure-level semiconductor IP can have a different commercial profile from a one-time consulting assignment or prototype sale.

The central commercial variable may ultimately be the number of consequential computing systems incorporating a protected and validated authority function—not the number of standalone OCUP chips sold.

A manufacturer might incorporate that function into a new processor generation. Another might license a chiplet implementation. A robotics manufacturer might integrate an externally controlled enforcement module. A critical-infrastructure provider might deploy validated implementations across fleets of equipment.

The underlying commercial opportunity would remain related even though the products, packaging, jurisdictions, and customer requirements differ.

That possibility is what makes substrate adoption strategically distinct from selling a single security appliance.

It is also why large semiconductor companies matter so much.

They possess the engineering teams, manufacturing relationships, product roadmaps, integration expertise, customer access, and long-term investment capacity needed to determine whether a proposed architecture can become an economically viable infrastructure function.

The potential market is global because consequential autonomous computing is not confined to one nation or one vendor.

The addressable market cannot yet be credibly quantified as every AI processor or robot. Only systems for which the architecture is technically suitable, economically justified, adopted, and within applicable licensing rights would contribute to a real market.

But the breadth of possible applications makes the opportunity worthy of serious investigation.

The immediate question is not whether OCUP will capture a predetermined share of the autonomous computing economy.

It is whether temporal authority enforcement can become a repeatable, commercially valuable infrastructure requirement across significant portions of that economy.

If the answer is yes, the commercialization possibilities become substantial.

5. FRAND: Protecting Invention While Enabling Adoption

A technology intended for global infrastructure confronts a difficult commercial question.

How can originating inventions be protected and fairly rewarded without making their adoption dependent on arbitrary exclusion, incompatible proprietary implementations, or control by a single organization?

Our proposed answer is to pursue **open interoperability alongside protected technical invention**, with fair, reasonable, and non-discriminatory—FRAND—licensing opportunities where appropriate.

This approach does not require abandoning intellectual-property rights.

Nor does it require treating every architectural description, software implementation, or standard interface as proprietary.

Rather, it distinguishes between what an industry must be able to understand and implement interoperably, and what patent law may legitimately protect as a specific technical invention.

A future authority standard could define observable behavior, required evidence, interoperability interfaces, and independently testable conformance properties.

Multiple manufacturers could then develop competing implementations.

Where an implementation practices valid patent claims subject to applicable licensing commitments, licensing could provide compensation for the protected contribution.

This is not presently a declaration that OCUP patents are standard-essential.

Standard-essential status cannot be created merely by labeling a technology foundational. It depends on the claims of granted patents and their relationship to the mandatory requirements of an adopted standard.

Nor does the FRAND objective establish a royalty amount or compel any standards organization to adopt the invention.

These distinctions are important because the credibility of the opportunity rests on the industry's ability to evaluate both the technology and the proposed licensing terms.

An open interoperability strategy would ideally encourage independent implementations, competing suppliers, transparent technical requirements, and reproducible conformance evidence.

At the same time, patent protection and commercially reasonable licensing could help finance continued research, hardware integration, validation, and long-term infrastructure support.

The balance matters.

A global authority substrate should not require one nation to surrender sovereign decisions, one manufacturer to surrender its product architecture, or one operator to place every operational permission under a single remote controller.

Jurisdictional authority, qualified human validation, and protected enforcement would need to remain properly separated.

The ultimate objective is interoperability without compulsory centralized control.

Common technical evidence need not imply common political authority.

That principle can support a multipolar world while preserving opportunities for manufacturers and inventors to participate commercially.

And it gives the proposed infrastructure an important characteristic: the ability to expand through collaboration rather than exclusive ownership of the entire ecosystem.

6. The Investment Is in What Comes Next

A responsible commercialization proposal must be explicit about the present engineering boundary.

OCUP's initial AMD Kria KR260 Phase I work recorded bounded grant and closure indications within a defined host-driven test sequence. The submitted evidence includes a documented authority-state observation path, corresponding register interpretation, and a provenance record.

That is a starting point, not proof of a production authority substrate.

It does not yet establish independently issued authority, trusted lease expiration, protected time across resets, hostile-host resistance, non-bypassable physical resource control, or operational deployment at scale.

Those properties demand additional hardware engineering and independent validation.

And that is precisely where the opportunity for capable partners begins.

An early collaboration could examine the existing invention disclosures and reference evidence, challenge the governing invariant, identify a suitable platform, establish testable acceptance criteria, and design a reference implementation that mediates an actual consequential resource.

A subsequent program could evaluate trusted time, protected key custody, physical enforcement, reset continuity, alternative execution paths, attestation, performance overhead, and independently reproducible behavior.

Successful engineering could support more substantial product-integration, licensing, and standards-development opportunities.

The journey should be evidence-driven. A partner should be able to discover that an implementation does not work, that a claim is too broad, that a threat model is incomplete, or that a different engineering solution is preferable.

The objective is not agreement by reputation.

It is progress by demonstration.

The originating intellectual property offers a defined starting position. The engineering program offers a path to determining what is technically achievable. The proposed commercial structure offers ways for participating organizations to share the costs, risks, and potential rewards.

This is not an invitation to purchase a completed global standard.

It is an invitation to help determine whether one can be responsibly established.

Ownership, Review, and the Clock Ahead

The applications in the OCUP family are presently held by the inventor and are proposed for assignment to DigiPie International PBC, with its wholly owned subsidiary Temporal Authority Systems PBC intended to hold documented commercialization and licensing rights for OCUP-related technology. These assignments and intercompany rights remain subject to counsel review and execution; none is represented as completed.

The QSAFP priority application, U.S. provisional 63/823,363 (Quantum-Secured AI Fail-Safe Protocol), was held under an interagency national-security review from shortly after its June 2025 filing until April 14, 2026, when the USPTO granted foreign filing license No. 656,154. No secrecy order was imposed. Foreign filing licenses have since been granted for the RAAVE and PBTA provisionals as well. A security review is a precondition for filing abroad, not an assessment of merit. Its significance here is narrower and practical: the documented foreign-filing authorization permitted the relevant international filing activity; it does not replace any other applicable export-control or regulatory requirements.

Both international applications are expected to be published by WIPO in December 2026. International search is under way with the USPTO as International Searching Authority, while the International Search Reports and written opinions remain pending. Their issuance may occur before or after publication. National-phase entry, which decides the countries in which protection will be pursued and the claim strategy in each, falls due in November and December 2027. Organizations that engage before then can help shape where and how the portfolio is prosecuted, one of the most practical forms of early partnership available.

Open-Core Development and Licensing

The QSAFP reference software has been developed in public as an open-core project, with public commits dating back at least to October 2025. The license notices on its historical releases are under a consistency review: the repository identifies AGPL-3.0 with an additional commercial-licensing term, and some legacy files carry differing terms. Until that review is complete, the precise boundaries between the published software licenses, commercial implementation rights, and patent rights are being reconciled, and they should not be read as an unrestricted grant.

For newer repositories and future hybrid projects, Temporal Authority Systems is developing the SAGE License v1.0, a proposed framework intended to support open technical collaboration

alongside clearly defined commercial implementation, semiconductor IP, and patent-licensing rights. SAGE remains under development, does not govern the legacy QSAFP releases, and is offered here as a direction rather than a license in force. The longer-term aim is a model that open-source and open-core creators across the autonomous era can use to share their work and still be fairly rewarded for it.

7. An Invitation to Semiconductor Leadership

The semiconductor industry has repeatedly changed civilization by recognizing opportunities before their full consequences were apparent.

Processors enabled personal computing. Specialized accelerators expanded scientific computation. Programmable GPUs helped create new capabilities across computing industries. Each transition required more than an initial concept: it required sustained engineering, substantial investment, market development, and people willing to commit before success was assured.

The autonomous era presents another such question.

Not simply how powerful future machines should become.

But under whose authority that power should continue to operate.

For semiconductor leaders, OCUP represents a proposed technical and intellectual-property foundation from which a new category of authority-enforcement infrastructure might be developed.

For insurers and regulators, it represents an opportunity to help define the evidence that consequential autonomous operation should produce.

For manufacturers and operators, it offers the possibility of enforceable boundaries that remain meaningful even when the governed software is unreliable, compromised, or operating beyond its authorized conditions.

For inventors, investors, and technology partners, it presents the possibility of a commercially sustainable infrastructure program serving many manufacturers and industries rather than one application alone.

None of these outcomes is guaranteed.

But their combined potential explains why the opportunity deserves attention before the entire architecture is finished.

The organizations that participate early may help shape not only the products that emerge, but the technical expectations by which those products are evaluated for generations.

The immediate invitation is modest: examine the originating inventions, challenge the architecture, review the bounded evidence, and determine whether a collaborative engineering program is warranted.

The larger invitation is considerably more ambitious.

Help establish an infrastructure foundation through which increasingly capable autonomous systems can remain subject to legitimate human authority.

Let the hardware become trustworthy.

Let the evidence become independently reproducible.

Let manufacturers compete.

Let standards emerge through credible technical participation.

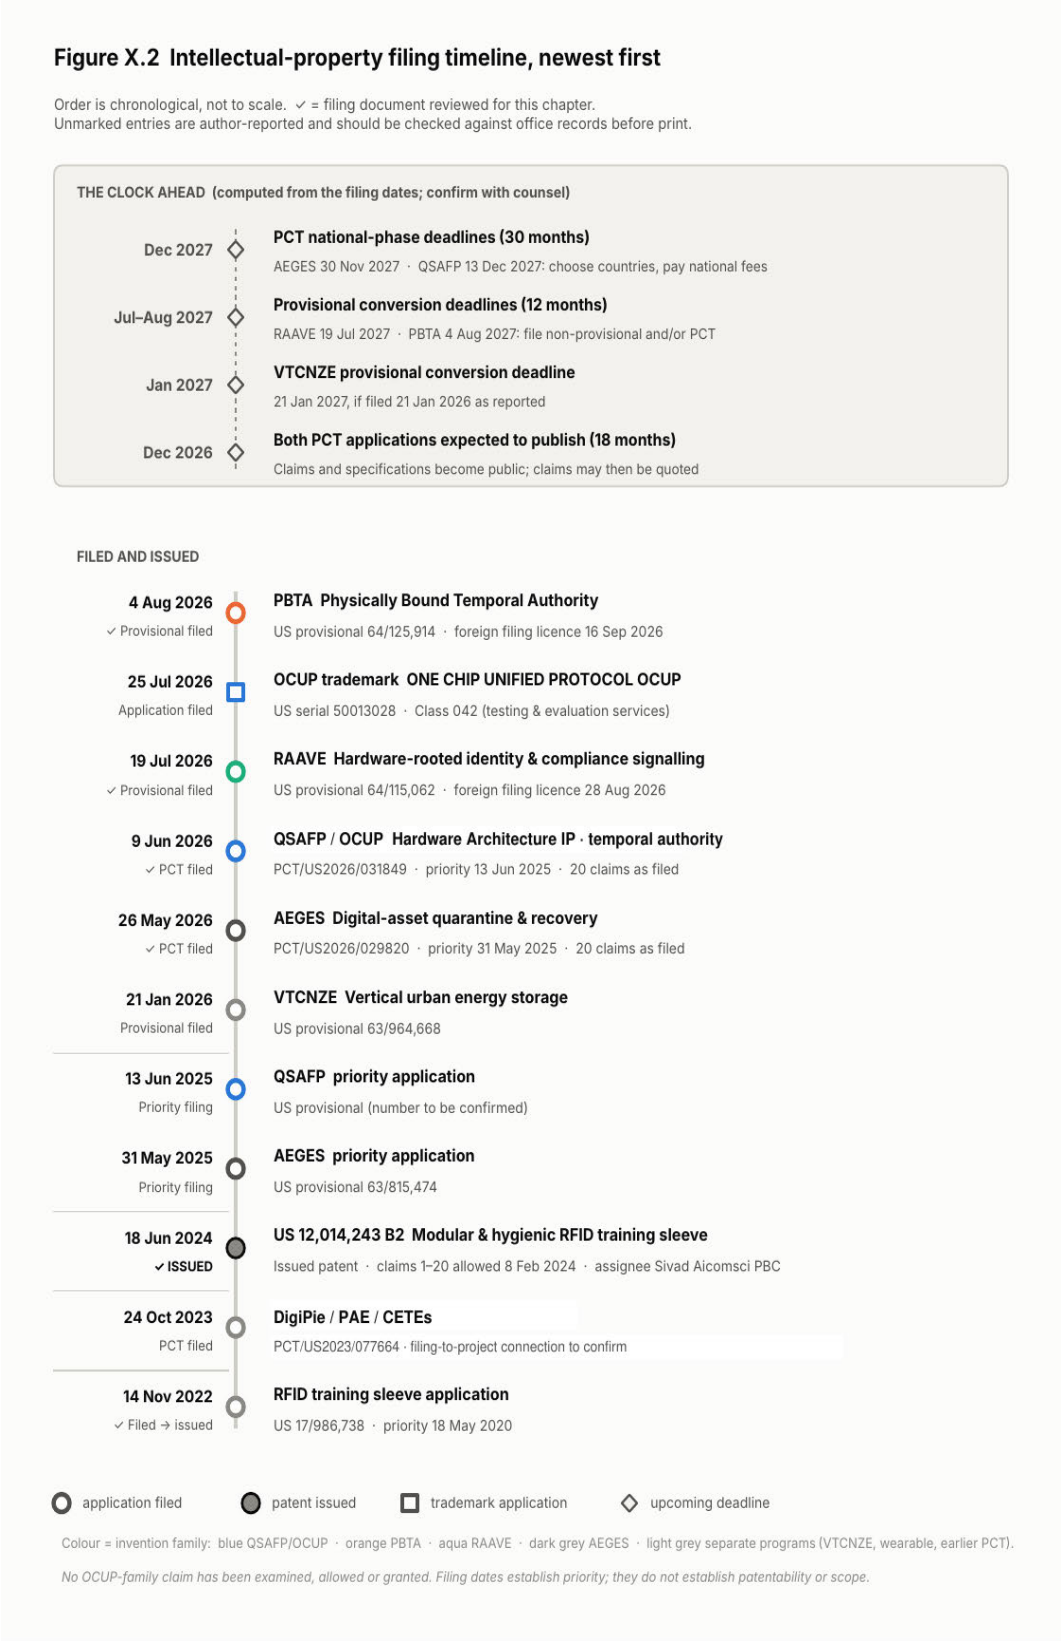
Let innovation be rewarded without surrendering sovereignty.

And let commercialization help sustain the work.

Intellectual Property Portfolio: A Timeline of Invention

An evolving portfolio spanning autonomous-system authority, semiconductor security, energy infrastructure, and human-centered technologies.

The following chronology records documented patent and trademark filings in reverse order. Each represents a distinct stage of invention and protection, from early technical disclosures to international applications and examined claims. The portfolio illustrates an approach to innovation centered on physical architectures, enforceable operating boundaries, infrastructure systems, and their practical deployment. Figure X.2 summarizes the chronology; dates and numbers are taken from the filing records.



August 4, 2026 — Physically Bound Temporal Authority (PBTA)

Patent title: Physically Bound Temporal Authority Identification and Control for Computational Accelerators

U.S. provisional application: 64/125,914
Filing status: Provisional application filed; foreign filing license granted September 16, 2026

Physically Bound Temporal Authority (PBTA) addresses a fundamental challenge in autonomous computing: how to establish that an authority-enforcement component is physically associated with the particular processor, accelerator, or machine it is intended to govern. The proposed architecture connects operational authorization to verifiable hardware identity and physical configuration, reducing reliance on software assertions that a capable autonomous system might bypass, falsify, or outlive.

Potential implementations include physically responsive structures, electronic tamper-detection meshes, protected hardware identities, and other anti-separation or anti-substitution mechanisms. Acoustic or ultrasonic verification is also under consideration as a future engineering approach, subject to confirmation of its support in the filed disclosure. The provisional disclosure describes microfluidic and other materially responsive structures in detail, as one family of embodiments rather than the sole engineering requirement. In the OCUP ecosystem, PBTA would strengthen the physical credibility of the authority boundary while providing independently verifiable evidence for related assurance systems.

July 25, 2026 — One Chip Unified Protocol (OCUP) Trademark

Mark: ONE CHIP UNIFIED PROTOCOL OCUP
U.S. trademark application serial number: 50013028
Filing status: Federal trademark application filed; registration pending
Classification: International Class 042

The OCUP trademark application seeks protection for the name associated with the One Chip Unified Protocol initiative. Unlike a patent, which may protect qualifying technical inventions through defined claims, a trademark identifies the commercial source of specified goods or services. The application identifies services involving testing, analysis, and evaluation of computer software and autonomous-system control logic for conformity with performance and safety standards.

This filing supports the commercial identity of OCUP as a potential technology, testing, and assurance platform. The brand may eventually accompany hardware integrations, reference implementations, evaluation services, and commercial partnerships, subject to the scope of any registration and applicable trademark rights. The filing therefore complements, but does not replace, the underlying patent program.

July 19, 2026 — RAAVE: Hardware-Rooted Authority Verification

Patent title: Hardware-Rooted Autonomous-System Identity, Temporal Compliance Signaling, Lifecycle Authorization, and Noncompliance Detection
U.S. provisional application: 64/115,062
Filing status: Provisional application filed; foreign filing license granted August 28, 2026

RAAVE is conceived as an independent verification and assurance layer for autonomous systems operating across organizations, networks, physical environments, and jurisdictions. Rather than issuing operational authority itself, RAAVE is intended to assess whether a machine can demonstrate appropriate identity, temporal compliance, authorization continuity, and conformance with externally

established operating requirements. Its broader architecture contemplates distinguishing what a system claims about its authority from what can actually be observed or verified.

Potential functions include hardware-rooted identity checks, compliance signaling, lifecycle authorization verification, noncompliance detection, authenticated evidence, and recognition of authority changes as machines move between governed environments. In combination with OCUP, RAAVE could provide external parties with a means of evaluating whether authority is present and valid without transferring the power to create or extend that authority to the governed machine.

June 9, 2026 — Hardware-Enforced Temporal Authority Governance (QSAFP / OCUP)

Patent title: Hardware-Enforced Temporal Authority Governance and Self-Extension Denial for Autonomous Artificial Intelligence and Robotic Systems
International application: PCT/US2026/031849
Claimed priority date: June 13, 2025 (U.S. provisional 63/823,363, Quantum-Secured AI Fail-Safe Protocol)
Filing status: International PCT application filed; **foreign filing license No. 656,154 granted April 14, 2026 after national-security review;** WIPO publication expected December 2026

This application addresses the foundational problem underlying the OCUP initiative: the separation of a system's computational capabilities from its legitimate authority to exercise them. The proposed architecture places consequential operation under externally granted, bounded authorization, with enforcement mechanisms intended to deny unauthorized renewal, restoration, enlargement, or continuation of authority by the governed system itself. Its disclosed approaches include enforcement circuitry, capability gating, external authorization, and evidence of authority-state changes.

Commercially, the architecture is intended to support multiple semiconductor implementations rather than depend on a single packaging arrangement. Potential forms include specialized companion chips, accelerator-integrated logic, chiplets, embedded controllers, and licensable semiconductor intellectual-property blocks. The initial KR260 FPGA evidence establishes a narrower foundation involving host-driven grant, closure, and non-reopening indications under tested conditions; further engineering work is required to validate independent authority issuance, trusted temporal enforcement, physical resource gating, and production-grade resistance to bypass.

May 26, 2026 — AEGES: Digital Asset Quarantine and Recovery

Patent title: Quantum-Resistant System and Method for Digital Asset Quarantine, Recovery, and Economic Neutralization
International application: PCT/US2026/029820
Claimed priority date: May 31, 2025
Filing status: International PCT application filed; WIPO publication expected December 2026

AEGES addresses containment and recovery in digital-asset environments where theft, unauthorized transfers, compromised credentials, or exploited infrastructure can generate immediate financial consequences. The invention proposes mechanisms for quarantining affected assets, restricting their practical economic usefulness, and supporting authorized recovery. Its broader disclosed functions include temporal quarantine, cross-chain or bridge-related protections, and restoration processes intended to operate through verifiable security conditions.

Within the wider invention portfolio, AEGES represents a complementary containment discipline rather than the same invention as OCUP. OCUP focuses on whether a governed machine possesses authority to operate, while AEGES addresses how compromised digital assets and associated transactions may be restricted or recovered. Both reflect an emphasis on moving protection beyond advisory warnings toward mechanisms intended to produce enforceable technical consequences.

January 21, 2026 — VTCNZE: Vertical Stack Coalition for Near-Zero Emissions PBC

Patent title: Vertically Arranged Urban Energy Storage Systems and Methods for Grid Support and Power Continuity in Support of Zero Emissions Energy Sources
U.S. provisional application: 63/964,668
Filing status: Provisional application filed

The VTCNZE invention program addresses the physical land, infrastructure, and deployment challenges associated with large-scale energy storage. Its central concept involves vertically arranged energy-storage infrastructure capable of providing grid support and power continuity while reducing the ground footprint ordinarily associated with conventional horizontally distributed storage facilities. The system is intended to support the integration of low-emissions energy resources into densely developed or land-constrained environments.

Potential commercialization involves modular infrastructure, deployable structural configurations, and collaboration with energy-system manufacturers, developers, utilities, and host communities including a WIMBY factor via the embedded (DigiPie) **Digital Local Economy app.** The filing represents a separate physical-infrastructure invention program, with possible applications beyond the semiconductor and autonomous-system portfolios.

February 12, 2024 — DigiPie CETEs: Monetizing Consumption with Consumer Earned Tokens

Patent title: Method and System for Monetizing the Consumption of Goods and Services Using Blockchain Contracts and Consumer Earned Tokens (CETEs)
International application: PCT/US2023/077664
Filing status: International PCT application filed (first received October 24, 2023; international filing date accorded February 12, 2024); priority date February 28, 2023; U.S. national stage entered August 27, 2025

The CETEs application belongs to the DigiPie invention program and the economic framework behind the People's Autonomous Economy. It concerns methods and systems through which consumers earn tokenized participation from their consumption of goods and services, recorded and administered through blockchain-based contracts.

Its path illustrates the prosecution method described at the opening of this chapter: an international filing prepared by the inventor, an International Search Report and written opinion, and entry into U.S. national prosecution in August 2025 with the assistance of patent counsel.

November 14, 2022 — Modular and Hygienic RFID Training Sleeve

Patent title: Modular and Hygienic RFID Training Sleeve
U.S. nonprovisional application: 17/986,738

Filing status: Notice of allowance issued February 8, 2024; issued June 18, 2024 as U.S. Patent No. 12,014,243 B2 (assignee: Sivad Aicomsci PBC)
Examiner's finding: Claims 1–20 allowed

The Modular and Hygienic RFID Training Sleeve is a wearable training system designed to help users recognize and reduce potentially unhygienic hand-to-face movements. Its architecture uses separate upper-arm and lower-arm components containing RFID-related electronics. When the system determines that the user's hand is sufficiently close to the face, it provides a notification that can reinforce safer behavior through timely, physical-world feedback. Fomite mitigation.

The Common Thread

These filings span distinct markets and technologies, yet several share a recurring engineering philosophy: translating a consequential requirement into an identifiable system, a physical or computational mechanism, and an observable outcome.

The semiconductor-oriented inventions focus on authority, verification, and hardware-rooted trust. The energy-storage invention addresses physical infrastructure and resource deployment. The earlier wearable system demonstrates a practical arrangement of electronics, sensing, and human feedback. The CETEs application extends the same discipline to economic participation, the "wealth" half of Humanity's Health & Wealth Chip.

Together, the portfolio presents an invitation to engineers, manufacturers, institutional partners, and technology leaders: examine the inventions on their merits, challenge their technical assumptions, and collaborate where independently validated mechanisms may become useful infrastructure.

Status as of October 2026. No OCUP-family claim has yet been examined, allowed, or granted. Filing dates establish priority; they do not establish patentability or scope.

The opportunity is not merely to accumulate patents. It is to identify which inventions can become dependable, interoperable, commercially meaningful technologies—and recruit the expertise required to bring them into the world.

History remembers those who recognized a turning point before the world understood its consequences. The autonomous era presents such a moment. The semiconductor leaders who move first will have an opportunity not merely to shape the next generation of computing, but to help establish the boundaries within which extraordinary machine capabilities can safely serve humanity.

The invitation is to lead where leadership matters most: to transform an emerging architecture into independently verified, trustworthy infrastructure; to establish safeguards before their absence becomes a global crisis; and to demonstrate that technological progress and human authority need not be competing ambitions. The first movers who accept that challenge may help define not only the future of autonomous computing, but the measure of responsibility by which its pioneers are remembered.

OCUP MANIFESTO CLOSING

From Humanity's Health & Wealth Chip to a Universal Authority Substrate

OCUP began with a larger question than how to make autonomous systems safer: How can humanity benefit from increasingly powerful machine intelligence without surrendering its authority over the systems it creates?

That question gave rise to the vision of Humanity's Health & Wealth Chip. Health meant protecting human life and well-being as autonomous machines assume increasingly consequential responsibilities. Wealth meant ensuring that the extraordinary productivity and abundance these technologies may create can benefit the many, not merely the few.

The ambition was never simply to build another cybersecurity product. It was to help establish a foundation for the autonomous era—one that protects human authority while creating opportunities for meaningful human participation in the economy that emerges around intelligent machines.

What began with AEGES containment and QSAFP authority governance has evolved into OCUP: a proposed hardware-enforced boundary separating a machine's computational capabilities from its permission to exercise them. OCUP's authority core, the Quorum Validators Network (QVN), AEGES, and RAAVE form complementary parts of a wider proposed architecture. Their respective roles, development status, and limitations are documented in Science & Engineering; they are not all capabilities demonstrated in a production chip.

The long-term objective is a universal authority substrate: a small, interoperable silicon foundation that semiconductor manufacturers can implement alongside existing processors or integrate into future architectures. Its effectiveness depends on physically controlling consequential interfaces and eliminating bypass paths. Universal compatibility remains an engineering goal, not an established capability.

Its purpose is not to dictate what intelligent systems may think, but to ensure that their authority to act remains externally granted, time-bounded, independently verifiable, and beyond their power to extend, restore, or enlarge themselves.

The engineering and the human purpose meet at a critical point. If operational authority must periodically return to human-governed validation, people retain a consequential role in determining whether machines may continue exercising delegated power. Through QVN, that responsibility could become a trained, recognized, and compensated form of participation in the autonomous economy.

Shared prosperity requires shared authority. The silicon is what makes that authority enforceable.

Bringing It Home

Phase I established a bounded hardware evidence foundation: grant and closure indicators were recorded, with no reopening detected during the tested reassertion sequence. It did not establish trusted lease expiry, independent external authority, physical resource enforcement, or the complete security properties required of a production chip.

Those distinctions matter. There is a considerable difference between documenting a promising hardware control path and bringing a dependable silicon architecture into widespread deployment.

We have taken OCUP as far as our present resources can carry it at the speed this moment demands.

The next advances require independent security domains, trusted timing, protected key custody, resistance to hostile hosts, physical enforcement, reproducible testing, and integration across semiconductor platforms. They require sophisticated tools, specialized facilities, experienced engineering teams, and sustained development resources.

This is where a semiconductor partner with deep silicon, packaging, and security capabilities could make a decisive difference—not simply to manufacture another chip, but to refine its specification, validate its implementation, solve cross-layer integration problems, build independent evidence, accumulate validation history, and help shape standards that could make hardware-enforced temporal authority broadly accessible.

This is an invitation to collaboration, not a declaration of completion. We are not asking the semiconductor industry to accept an unproven promise. We are inviting its engineers, researchers, manufacturers, and assurance specialists to examine the evidence, test the assumptions, expose the weaknesses, and help solve the remaining problems.

The objective is larger than any one company, processor architecture, or nation. It is to establish a dependable foundation through which increasingly capable autonomous systems can operate without making human authority optional.

Qualified organizations are invited to participate in paid OCUP pre-production pilots, with FPGA-based technical validation or non-FPGA evaluation pathways suited to their capabilities and objectives. These bounded engagements offer a practical opportunity to examine the evidence, evaluate authority-governance requirements, and help define the next stages of engineering and integration—providing a direct path toward deeper collaboration and partnership.

Humanity's Health & Wealth Chip was the vision. A universal authority substrate is the engineering path. Bringing it home will take partners who recognize that the autonomous era deserves a foundation worthy of its possibilities—and its consequences.

The Autonomous Era has arrived. Let's keep humans in control.

I N H O C U N A S U M U S

T H E F I N A L S A Y