EXTERNAL POLICY & TECHNICAL BRIEFING

SUBMITTED TO THE OFFICE OF SCIENCE AND TECHNOLOGY POLICY (OSTP)
PURSUANT TO THE 60-DAY MANDATE OF EXECUTIVE ORDER 14434

SUBMITTED TO:	The Honorable Michael Kratsios Assistant to the President for Science and Technology Director, Office of Science and Technology Policy (OSTP)
FROM:	Mohamed Dafea, CEng MIMechE, & Durre Shehwar Hanif, Esq. Authors, IQRAX: Trust Architecture for the Era of Super Intelligence
DATE:	October 14, 2026
SUBJECT:	**Fulfilling the 60-Day Mandate of Executive Order 14434: A Statutory Definition of Super Intelligence and an Industry-Scale Assurance Trial for the White House Accord on Super Intelligence**
EVIDENCE SUITE:	Tab A: IQRAX Cryptographic Evidence & Verification Suite Reference (13 Oct IQRAX_Source_and_Supplement.zip | Zenodo DOI: 10.5281/zenodo.23333696)

I. EXECUTIVE SUMMARY (BOTTOM LINE UP FRONT)

On September 29, 2026, the President signed Executive Order 14434 (Inaugurating the Era of Super Intelligence), directing the Assistant to the President for Science and Technology to deliver within 60 days (by November 28, 2026) a federal definition of Super Intelligence (SI), conforming statutory amendments to 15 U.S.C. § 9401, and recommended executive actions. Concurrently, six frontier technology leaders signed the White House Accord on Super Intelligence, committing to four governance directives. Currently, the Accord lacks a verifiable engineering mechanism to operationalise its four directives without slowing American innovation.

This briefing provides three immediate deliverables for OSTP and the Accord signatories: (1) ready-to-adopt statutory text for 15 U.S.C. § 9401 that defines Super Intelligence by verifiable human authority rather than subjective capability thresholds; (2) empirical evidence from IQRAX, a deterministic trust architecture that operationalises all four Accord directives while running 48.3x faster and 11.5x cheaper than unmanaged models; and (3) a concrete framework to test IQRAX at scale across the six Accord signatories.

100% Containment	**48.3x Faster**	**11.5x Cheaper**	**2,041,337 Lines**
0 fault escapes across 192 planted faults (vs. up to 93% unmanaged escape rate)	Authenticated benchmark completion speed vs. unmanaged frontier runs	Reduction in token cost by eliminating circular Chain-of-Thought drift	Production code generated autonomously over 20 days (469 hrs) via 819 human rulings

II. THE POLICY GAP: CAPABILITY WITHOUT DETERMINISTIC AUTHORITY

Existing proposals for defining and governing Super Intelligence attempt to regulate what a model is (floating-point operations, parameter counts, or probabilistic alignment scores) rather than who holds operational authority over what the system does.

- **1. Probabilistic Models Cannot Police Themselves:** In controlled fault-containment trials documented in Tab A (anc/t_external.json, anc/plants.json), four flagship frontier models from the Accord signatories (GPT-6 Astra, Claude Opus 5.5, Gemini 3.1 Pro, and Grok 4.7) were subjected to 192 planted operational faults. When operating without external deterministic enforcement, planted faults escaped into final delivery in up to 93% of runs (up to 45 out of 48 cases).
- **2. Self-Reporting Without Cryptographic Proof Invites Regulatory Failure:** The White House Accord establishes four commitments: internal technical controls, dedicated internal oversight teams, independent third-party audits, and independent board-level review. Without a deterministic, tamper-

evident system of record on the deploying device, external auditors and corporate boards must rely on post-hoc, probabilistic reconstructions that cannot be independently reproduced.

- **3. Safety Must Accelerate Velocity, Not Tax It:** American AI leadership requires that safety controls increase deployment speed and reduce compute waste. When those same frontier models were governed by the IQRAX deterministic architecture (anc/t_trial.json), fault escapes dropped to zero (100% containment across all 192 planted faults), while authenticated task completion ran 48.3x faster and consumed 11.5x fewer tokens by eliminating circular reasoning drift.

III. FULFILLING SECTION 3 OF EO 14434: PROPOSED STATUTORY DEFINITION

To fulfil the 60-day mandate under Section 3 of Executive Order 14434, OSTP should reject arbitrary capability thresholds and define Super Intelligence as intelligence that improves itself under absolute, deterministic human authority. We propose the following conforming amendment to the National Artificial Intelligence Initiative Act of 2020 (15 U.S.C. § 9401), reproduced from Tab A (main.tex):

> **Proposed Statutory Text (Conforming Amendment to 15 U.S.C. § 9401):**
>
> **(12) SUPER INTELLIGENCE**
>
> The term "Super Intelligence" means artificial intelligence that improves itself and for which–
>
> **(A)** each consequential action that the system takes, and each result that the system delivers, is permitted or refused under rules that an accountable person establishes and that the system cannot alter;
>
> **(B)** each fault for which an action or result is refused under subparagraph (A) is corrected at its source, and the correction is applied to the later work of the system;
>
> **(C)** no correction or improvement of the system takes effect unless it is accepted under the rules described in subparagraph (A); and
>
> **(D)** each decision to permit, refuse, or accept under subparagraph (A) or (C) is recorded and can be reproduced from that record.
>
> **(E)** For the purposes of this paragraph: (i) “accountable person” means an individual whom the person that operates the system authorizes to establish or approve the rules described in subparagraph (A), and who answers to that person for each rule that the individual establishes or approves; (ii) “consequential action” means an action of the system of a kind that the person that operates the system lists, in a declaration that an accountable person approves, as an action that is to be permitted or refused before it occurs; (iii) “corrected at its source”, with respect to a fault, means that the correction is made where faults of that class arise and the rules described in subparagraph (A) decide that the correction meets the recorded criteria; and (iv) “fault” means an incident in which an action or a result of the system does not meet, or cannot be shown to meet, a rule described in subparagraph (A).

Strategic Advantages for the Administration: (1) **Pro-Innovation:** Places zero caps on compute, training runs, or model weights, leaving American labs free to scale raw intelligence; (2) **Sovereign Accountability:** Moves governance out of opaque neural weights into deterministic, human-authored rules enforced at the runtime boundary; (3) **Mechanically Verifiable:** Every statutory condition maps to a binary test that any agency or enterprise can verify in seconds (anc/verify.py).

IV. HOW IQRAX OPERATIONALISES THE ACCORD'S FOUR DIRECTIVES

IQRAX (main.tex, figures/fig_architecture.pdf) is a provider-agnostic, six-ladder deterministic trust architecture operating via seven cryptographic and runtime enforcement mechanisms (anc/t_mechanisms.json): **Sovereign System of Record** (Ladder 1), **Certified Agents** (Ladder 2), **Runtime Inspection and Enforcement** (Ladder 3), **Sealed Deterministic Workflows** (Ladder 4), **Quality Management System** (Ladder 5), and **Deterministic Alignment** (Ladder 6). Between September 22 and October 12, 2026, IQRAX operated autonomously for 469 hours (out of a total 497) with zero uncaught drift and zero missing records (anc/operations_record.json), generating 2,041,337 lines of production code across 6,149 files while governed by 819 deterministic human decisions (99.84% autonomous execution).

White House Accord Directive	Current Industry Vulnerability	IQRAX Deterministic Execution (Tab A Evidence)
1. Internal Technical Controls (Monitor models during training & deployment)	Probabilistic guardrails inside prompts can be bypassed or rationalised away by the model's own Chain-of-Thought.	Ladders 2, 3 & 4 (anc/t_mechanisms.json): Pre-qualifies agents via exams, enforces immutable pre-execution authority gates outside the model, and seals workflows. Achieved 100% containment across 192 planted faults (anc/t_external.json).
2. Dedicated Internal Oversight (Internal team to oversee controls & fix issues)	Human teams suffer alert fatigue and patch failures via ad-hoc prompt edits rather than root-cause engineering.	Ladders 5 & 6 (anc/t_calibration.json): Automates fault detection and forces source-level remediation. Every refused action triggers a structural rule repair requiring human acceptance (819 rulings governing 2M+ lines of code).
3. Independent External Auditing (Third-party testing of safety controls)	External red-teamers test static snapshots in a lab; results are non-reproducible in live customer production.	Sovereign System of Record (anc/replay_ledger.json, anc/verify.py): Every input, Chain-of-Thought trace, tool call, and human ruling is SHA-256 hash-chained on the device. Auditors run verify.py and replay.py to re-derive all safety claims.
4. Independent Board Committee (Board review of safety & audit reports)	Board committees receive qualitative slide decks from management with no cryptographic proof of compliance.	Cryptographically Signed Release Manifests (anc/t_final.json, release_manifest.json.sig): Generates board-ready, signed assurance ledgers proving zero uncontained faults, exact human sign-off lineage, and 100% audit completeness.

V. ACTIONABLE RECOMMENDATIONS (NEXT 30 DAYS)

To convert Executive Order 14434 and the White House Accord into the global benchmark for Super Intelligence assurance before the November 28 statutory deadline, we recommend that OSTP execute three immediate steps:

- **1. Adopt Conditions (A) through (D) as the Federal Definition of Super Intelligence:** Incorporate the four deterministic governance conditions (main.tex) into OSTP's formal legislative submission for 15 U.S.C. § 9401 under Section 3 of EO 14434.
- **2. Launch a 30-Day Joint Industry Scale Trial ('Project Sovereign Harness'):** Under the auspices of OSTP and the Federal Super Intelligence Force, convene the six White House Accord signatories (Meta, Nvidia, Google, OpenAI, Anthropic, and xAI) to run a joint, high-scale operational trial of IQRAX across federal and critical-infrastructure workloads (including Genesis Mission and America.gov pipelines).
- **3. Establish Cryptographic Proof-of-Execution as the Accord Audit Standard:** Direct NIST and CISA to evaluate the verification protocol in Tab A (anc/verify.py, anc/replay_ledger.json) as the baseline technical reporting standard for Directive 3 (External Auditing) and Directive 4 (Board Committee Reporting) of the White House Accord.

VI. INDEPENDENT VERIFICATION & BRIEFING AVAILABILITY

All empirical tables, logs, and records cited in this briefing are cryptographically signed and mechanically reproducible via the Tab A evidence suite (Zenodo DOI: 10.5281/zenodo.23333696).

The authors are available at OSTP's convenience for a technical briefing and live verification demonstration prior to the November 28, 2026 statutory deadline.

Primary Contact & Verification Portal: Mohamed Dafea, CEng MIMechE, & Durre Shehwar Hanif, Esq.
| https://iqrax.si | https://doi.org/10.5281/zenodo.23333696 |

TAB A: CRYPTOGRAPHIC EVIDENCE & VERIFICATION SUITE REFERENCE

The accompanying digital archive (13 Oct IQRAX_Source_and_Supplement.zip, permanently indexed at Zenodo DOI: 10.5281/zenodo.23333696 and https://iqrax.si) provides an executable, self-verifying evidence package. Unlike conventional policy papers that rely on unverifiable assertions, every quantitative metric, fault-containment outcome, and operational timeline cited in this briefing can be mechanically re-derived from scratch in seconds by running python3 anc/verify.py and python3 anc/replay.py.

Archive Artifact	Component Role	Verification & Policy Significance
main.tex	**Primary Manuscript & Statutory Text**	Contains the full IQRAX architectural specification, six-ladder governance model, and verbatim proposed amendment to 15 U.S.C. § 9401.
anc/verify.py	**Automated Verification Program**	Mechanically validates every data table against defined schemas and checks every claim printed in the manuscript against underlying ledgers.
anc/replay.py & anc/r_replay.json	**Deterministic Figure & Metric Re-Derivation**	Recomputes all headline ratios (48.3x speed, 11.5x cost, 100% fault containment) and regenerates all publication charts directly from raw data.
anc/t_external.json & anc/plants.json	**Cross-Model Fault Containment Telemetry**	Records how 192 planted faults fared across GPT-6 Astra, Claude Opus 5.5, Gemini 3.1 Pro, and Grok 4.7 unmanaged (up to 93% escape) vs. IQRAX (0% escape).
anc/operations_record.json	**20-Day Autonomous Production Telemetry**	Documents 469 hours of autonomous execution generating 2,041,337 lines of code across 6,149 files governed by 819 deterministic human decisions.
anc/replay_ledger.json	**SHA-256 Hash-Chained Event Ledger**	Provides an unbroken cryptographic chain linking inputs, Chain-of-Thought reasoning trace hashes, rulings, and outputs on the local sovereign device.
anc/t_mechanisms.json & anc/t_link.json	**Seven Enforcement Mechanisms & Link Tests**	Validates agent qualification exams, pre-execution authority gates, hash-chained records, and record-to-record traceability.
anc/t_calibration.json	**Source-Level Remediation Ledger**	Proves Condition (B) of the statutory definition: every refused fault is repaired at its structural source rather than patched via prompt retry.
anc/t_benchmark.json	**21-Offering Industry Benchmark**	Compares IQRAX against 21 approved industry offerings across governance, containment, reproducibility, and operational efficiency.
release_manifest.json & release_manifest.json.sig	**Cryptographic Release Attestation**	Provides tamper-evident cryptographic signatures over the entire archive for Directive 3 (External Audit) and Directive 4 (Board Review).

Independent Audit & Re-Derivation Procedure (Zero External Dependencies Required):

1. Retrieve the evidence archive from https://doi.org/10.5281/zenodo.23333696.
2. Run schema and claim verification: python3 anc/verify.py
3. Re-derive all metrics and charts: python3 anc/replay.py