sOS Research Lab WORKING PAPER SERIES

sOS Research Lab 2026-07-01 · Barcelona

WP9 – Legal Framework: UILL as GDPR Mechanism + sOS Governance Model

Series	sOS Execution
Status	Rewritten – Experimental framing (replaces original Branch A)
Date	2026-06-28
Priority	Critical – Phase 1
Review Mode	Full Redundant Review
Grounded in	WP14 (Barcelona Sandbox Experimental Protocol), sOS doctoral thesis, EU AI Act (Reg 2024/1689, Art. 57 sandboxes), GDPR (Reg 2016/679, Art. 89 research safeguards), MiCA (Reg 2023/1114), Spanish research law (Ley 14/2011 de Ciencia), Declaration of Helsinki

How sOS qualifies	sOS meets the criteria for Art. 89 treatment because: (a) the Barcelona Sandbox is designed to produce scientific knowledge of public interest, (b) the research protocol (WP14) is pre-registered with falsifiable hypotheses, (c) results will be published open-access regardless of outcome, (d) the research is conducted under UB ethics oversight, and (e) data processing is minimized to what the experiment requires.

The AESIA AI regulatory sandbox (Art. 57) is not expected to be operational within the Phase 0 window (2026). The legal obligation for Member States to establish sandboxes takes effect 2 August 2026; operational readiness typically follows 6–18 months later. sOS does not gate its experiment on AESIA availability.

Abstract

sOS is not a service, a product, or a deployed governance system. It is a **scientific experiment**. This changes the legal framework entirely. Rather than compliance architecture for an operational system, this document defines the legal conditions under which the Barcelona Sandbox (WP14) can lawfully conduct research with human participants, urban IoT data, distributed AI systems, and a governance token. The primary legal instruments are: (1) Universitat de Barcelona research ethics review (UB CER), (2) GDPR Art. 89 scientific research safeguards (via Spanish LOPDGDD DA-17 + Ley 14/2011), (3) MiCA-exempt token structure (non-transferable, zero-value, research-only), (4) AI Act limited-risk self-assessment. The EU AI Act regulatory sandbox (Art. 57, AESIA) and CNMV FinTech sandbox are opportunistic enhancements – not gating dependencies – because neither is expected to be operational within the Phase 0 window.

1. The Legal Frame Shift: From Deployment to Experiment

sOS governance operates through a **six-function decentralized bureaucracy** (thesis §14.3): **Propose, Validate, Delegate, Execute, Save/Compute, and Finance.** Each function is a smart contract type. Validation uses **redundant human scoring with machine-computed S^2**

variance (thesis §14.3.2): multiple validators score 0-10; the machine computes S^2 to detect consensus or trigger expanded validation. Labels from UILL match validators to tasks. The legal framework must account for data flows through all six functions – not just token and AI.

1.1 Two Radically Different Legal Postures

Primary legal frame	Product/service regulation	Research law + ethics
GDPR basis	Full compliance (all rights enforceable)	Art. 89 research safeguards + derogations where compatible
AI Act	Full compliance (risk classification, obligations)	Art. 57 regulatory sandbox (supervised, time-limited, conditional)
MiCA	Token as crypto-asset (white paper, CNMV notification)	Token as research instrument (CNMV sandbox)
Participant relationship	User / terms of service	Research participant / informed consent
Liability	Product liability (PLD 2024/2853)	Research risk (insurance + ethics board oversight)
Data sharing	Data processing agreement	Research data management plan
Exit	Service termination, data export	Right to withdraw from research without penalty
Failure	Regulatory enforcement	Publish negative result; revise hypothesis

1.2 This Is Not a Loophole

Research exemptions are not ways to avoid regulation. They exist because society has decided that controlled scientific experimentation – with ethics review, informed consent, and transparent publication – is a legitimate way to generate knowledge that benefits everyone. sOS uses these provisions as they are intended: to test a hypothesis under supervision before any deployment is considered.

If the experiment succeeds and sOS moves toward operational deployment, **the legal frame shifts back to full compliance**. That gate is explicit: Phase 2 transition requires a new legal assessment under the operational frame, not the research frame.

---

2. GDPR and the UILL as Research Data Infrastructure

2.0 Why the UILL Matters: GDPR Without a Tool Is Aspirational

The GDPR was enacted in 2016 into a world where personal data lives on centralized servers controlled by corporations. Every right it grants – access, erasure, portability, consent – depends on the data controller's willingness and ability to comply. The individual has no mechanism to verify, enforce, or audit any of these rights. They fill in a form and hope.

The UILL (Unique Information Library Ledger, thesis §15) inverts this architecture:

Access (Art. 15)	You request your data from a corporation. You receive a ZIP you cannot verify.	Your data lives in YOUR pod. You see it always. No request needed.
Erasure (Art. 17)	You submit a form. The company "deletes" your data – but you cannot verify whether backups, logs, or third-party copies remain.	Your smart contract of consent is revoked. Access is cryptographically terminated. The audit trail proves it. You don't need to trust – you verify.
Portability (Art. 20)	You export a ZIP in a proprietary format. It is useless in any other system.	Solid protocol. Standard format. Your data moves with you.
Consent (Art. 7)	You click a checkbox. The terms are 40 pages. You cannot negotiate. You cannot revoke granularly.	Smart contract: who, what purpose, for how long, revocable at any time. Programmable. Auditable.
Rectification (Art. 16)	You email support. Maybe they fix it. Maybe they don't.	Your labels are yours. You correct them directly.

The UILL is not a "compliance tool" added to sOS. The UILL is the mechanism by which the GDPR becomes real – executable, verifiable, and under individual control rather than corporate discretion.

In the Barcelona Sandbox research frame, every GDPR safeguard described below operates through the UILL infrastructure. The consent form, the data management plan, the DPIA – all of them are implemented as UILL functions, not as paper documents filed and forgotten.

2.1 The Research Safeguard (Art. 89)

GDPR Article 89 provides that processing for scientific research purposes may benefit from:

- **Derogations from rights** (Art. 15, 16, 18, 21) where those rights would render impossible or seriously impair the research – provided appropriate safeguards exist.

- **Broad consent:** consent for "certain areas of scientific research" where the specific purposes cannot be fully identified at the outset – exactly the situation of exploratory research like sOS.
- **Extended storage:** personal data may be stored for longer periods solely for research purposes with appropriate safeguards.
- **Presumption of compatibility:** further processing for research purposes is presumed compatible with the original purpose (Art. 5(1)(b)).

How sOS qualifies: sOS meets the criteria for Art. 89 treatment because: (a) the Barcelona Sandbox is designed to produce scientific knowledge of public interest, (b) the research protocol (WP14) is pre-registered with falsifiable hypotheses, (c) results will be published open-access regardless of outcome, (d) the research is conducted under UB ethics oversight, and (e) data processing is minimized to what the experiment requires.

2.2 Operationalizing Art. 89 in the Sandbox

Right to erasure (Art. 17)	Crypto-shred content; immutable metadata distinction	Same – research data management plan specifies retention and deletion	Participants informed in consent form
Right to access (Art. 15)	Full data export on request	Participant can access and export their own trace data at any time	Built into app interface; quarterly data snapshots provided
Purpose limitation (Art. 5(1)(b))	Purpose specified per processing operation	Broad research purpose: "testing stigmergic coordination for urban resource governance"	DPIA filed; ethics board reviews
Data minimization (Art. 5(1)(c))	Minimum necessary for service	Minimum necessary for answering research question (WP14 §7.1)	Data fields listed in ethics application
Data Protection Officer	Required	Required (UB provides via institutional DPO)	Named in DPIA
Data Protection Impact Assessment	Mandatory before processing	Mandatory before recruitment (WP14 §5.3)	Filed with UB ethics board + AEPD (Spanish DPA)
International transfers	SCCs / adequacy for non-EU storage	EU-only infrastructure for experiment (PostgreSQL, not IPFS)	Documented in data management plan

2.3 Consent, Not Terms of Service

Sandbox participants sign an **informed consent form**, not terms of service. The consent form (to be drafted before recruitment, reviewed by UB ethics board) covers:

1. **What the experiment tests** – stigmergic coordination vs. centralized coordination for streetlight maintenance.

2. **What data is collected** – streetlight-related traces (light ID, fault type, timestamp, geo-location ≈100m precision, not exact address), quarterly survey responses.

3. **What data is NOT collected** – no personal identity data beyond participant ID, no location tracking beyond fault reports, no behavioral profiling.

4. **How data is used** – exclusively for the pre-registered research questions (WP14 §1, §2).

5. **Right to withdraw** – at any time, without penalty, data deleted on request.

6. **Publication** – aggregated results published open-access; individual responses anonymized.

7. **Contact** – researcher contact, DPO contact, ethics board contact.

2.4 The Content/Metadata Distinction Under Research Framing

The immutability vs. erasure tension (original WP9 §3.3) is less acute under the research frame because:

- **The trace registry uses PostgreSQL** for the experiment, not a DLT – no blockchain immutability. Traces can be deleted on participant withdrawal.
- **Research records** (aggregate data, publications) are not personal data – the content/metadata distinction aligns with research practice: individual data is deletable; published research results are not.
- **If a participant withdraws and requests data deletion**, their individual traces are removed from the active research database. Already-published aggregate results (which do not identify them) are not affected.

---

3. The AI Act: Regulatory Sandbox, Not Compliance

3.1 Primary Basis: Limited-Risk Self-Assessment

The AESIA AI regulatory sandbox (Art. 57) is not expected to be operational within the Phase 0 window (2026). The legal obligation for Member States to establish sandboxes takes effect 2 August 2026; operational readiness typically follows 6–18 months later. sOS does not

gate its experiment on AESIA availability.

Primary legal basis: AI Act limited-risk self-assessment. sOS's position (consistent with the original WP9 §5):

- **Personal AI agents** (filtering traces, presenting options, not deciding) → **limited risk**, subject to transparency obligations only (Art. 50).
- **Stigmergic coordination** (traces influencing prioritization but final decisions made by humans) → **limited risk**, not high-risk.
- **Rationale:** The human always makes the final choice. The AI Act's high-risk categories (Annex III) for "administration of justice and democratic processes" apply to systems that affect legal rights – sOS sandbox participants make no decisions affecting others' legal rights.
- **Transparency compliance (Art. 50):** Participants are informed they interact with AI; any AI-generated content published by sOS is labeled.

Opportunistic enhancement: If AESIA's sandbox becomes operational during the 18-month intervention window, sOS will apply for supervised testing. Sandbox participation provides regulatory guidance and may relax certain documentation requirements. It is not a prerequisite for experiment launch.

3.2 sOS's Sandbox Application

The sandbox application (to be submitted to AESIA as soon as it is operational, anticipated 2026–2027) must demonstrate:

1. **Innovation:** sOS's stigmergic coordination using personal AI is a novel governance architecture not previously tested in urban infrastructure contexts.

2. **EU added value:** sOS contributes to the EU's objectives of human-centric AI, data sovereignty, and democratic governance innovation (aligned with Digital Decade targets).

3. **Safeguards:** Ethics review (UB), informed consent, reversible experiment, published results.

4. **Regulatory questions we want answered:** (a) Is personal AI that filters and presents governance information "limited risk" under the AI Act? (b) At what point does stigmergic coordination constitute "automated decision-making"? (c) What are the minimum documentation requirements for distributed AI systems in governance contexts?

3.3 Risk Classification Under the Sandbox

Inside the sandbox, risk classification is co-determined with the regulator. sOS's position:

- **Personal AI agents** (filtering traces, presenting options, not deciding) → proposed classification: **limited risk** (transparency obligations only).

- **Stigmergic coordination** (traces influencing prioritization but final decisions human-made) → proposed classification: **limited risk**, not high-risk.
- **Rationale:** The human always makes the final choice. The AI Act's high-risk categories (Annex III) for "administration of justice and democratic processes" apply to systems that affect legal rights – sOS sandbox participants make no decisions affecting others' legal rights.
- **Commitment:** If the sandbox regulator determines a higher risk classification, sOS complies with that determination as a condition of sandbox participation.

3.4 What the Sandbox Does NOT Cover

- The sandbox does not authorize deployment to the public.
- It does not exempt sOS from transparency obligations (Art. 50: users must know they interact with AI; AI-generated content must be labeled).
- It does not exempt sOS from the prohibition on unacceptable-risk AI practices (Art. 5: social scoring, manipulative AI, etc.).
- It is time-limited – typically 12–24 months, renewable on demonstration of progress.

---

4. The Token as Research Instrument – AND the Governance Model Behind It

4.0 The sOS Governance Principle: Everyone Is Owner, Worker, and Investor

Before discussing the token's legal status, this document must state the governance model it serves. The thesis and the DAO matrix (DAO Scheme_valid.xlsx, TFM 2018) establish that:

1. **Every individual in the DAO is simultaneously owner, worker, and investor.** The matrix encapsulates three types of asset relationships – as an investor, as a worker, and as a client. It includes all owners. There are no separate classes. No external shareholders. No VC. No users who are not owners.

2. **Decision power comes from labels, not from tokens or from one-person-one-vote.** Each individual decides in domains where their labels demonstrate knowledge and legitimacy. An individual with `has_knowledge:urban_planning (10/10)` has full voting weight on urban infrastructure decisions. An individual without legal expertise labels does not decide on regulatory matters – they delegate to someone who does. This is not disenfranchisement. This is scientific governance: you govern what you demonstrably understand, and you trust validated others for what you don't.

3. **Data sovereignty is the foundation of this model.** The UILL (§2.0) makes each individual the data controller of their own information. The DAO does not hold anyone's data. The DAO processes labels, traces, and consent – all of which originate from individual UILLs under each person's control.

4. **Tokens are coordination instruments, not financial assets.** They signal contribution, weight governance participation, and decay with inactivity. They are non-transferable (thesis §15.3.3). They cannot be bought. They cannot be sold. They cannot be accumulated by external capital. Their sole function is to coordinate collective action among people who are already owners of the system.

This governance model has specific legal implications:

- **Corporate law:** The DAO is not a corporation. It has no external shareholders. The Spanish association structure (Ley Orgánica 1/2002) is the closest legal form, but the governance model differs from traditional associations in that decision weight is label-based, not one-member-one-vote.
- **Securities law:** Because tokens are non-transferable, non-purchasable, and confer no expectation of profit, they do not constitute securities under EU or Spanish law. This is a legal consequence of the governance model, not a regulatory workaround.
- **Data protection law:** Because the UILL inverts the controller/processor relationship (individuals control their own data), the DAO's data processing obligations are fundamentally different from those of a centralized platform. The DAO processes labels and traces – not raw personal data. Raw data never leaves the individual's UILL without explicit, programmable consent.

4.1 The Token as Research Instrument – MiCA Exemption

Primary legal basis: MiCA exemption. The sOS sandbox token is non-transferable, non-tradable, carries zero financial value, and is distributed only to consenting research participants for the purpose of testing governance mechanisms. It does not constitute a "crypto-asset offered to the public" under MiCA Art. 4 and Recital 22. No white paper, no CNMV notification, no ongoing disclosure obligations. This is the same conclusion reached in the operational-frame WP9 (§4.3) and it is the default legal posture regardless of sandbox availability.

Opportunistic enhancement: Spain's CNMV operates a FinTech sandbox (Ley 7/2020). However, the sOS token has no financial function – it is a non-transferable governance research instrument. The CNMV sandbox may simply decline jurisdiction rather than accept supervision. If it accepts, it provides an additional layer of regulatory comfort; if it doesn't, the MiCA-exemption is independently sufficient. The CNMV sandbox is not a gating dependency.

4.3 No Financial Value = No Financial Regulation Trigger

The token's zero financial value is now a legal strategy, not just an ethical choice. A token with no market value, no transferability, no exchange listing, and no expectation of profit does not trigger:

- MiCA (no public offering)
- AMLD/AMLR (no value transfer)
- Securities regulation (no investment contract)
- Tax obligations for participants (no capital gain possible)

This is the simplest possible legal posture and the one most appropriate for a scientific experiment.

---

5. Spanish Research Law and UB Ethics

5.1 The Research Frame Under Spanish Law

Ley 14/2011 de la Ciencia, la Tecnología y la Innovación provides the national legal framework for scientific research in Spain. Key provisions:

- **Research freedom (Art. 4):** Researchers have the right to freely define objectives and methods – sOS's pre-registered protocol (WP14) is an exercise of this right.
- **Ethics committees (Art. 5):** Research involving human participants requires approval by an accredited ethics committee – UB's Comitè d'Ètica de la Recerca (CER) is the competent body.
- **Open access (Art. 37):** Publicly funded research must be published open access – sOS's commitment to publish-regardless aligns.
- **Research with personal data (Disposición Adicional 17ª):** Research processing of personal data is compatible with GDPR provided safeguards (ethics approval, DPIA, data minimization).

5.2 Ethics Review Process

Before participant recruitment (WP14 Phase 0 gate):

1. **Submit protocol to UB CER:** WP14 experimental protocol, this legal framework, draft informed consent form, data management plan, DPIA.

2. **CER review criteria:** (a) Scientific merit of the research question, (b) proportionality of methods to objectives, (c) risk/benefit ratio for participants, (d) adequacy of informed consent, (e) data protection measures, (f) participant recruitment fairness, (g) right to withdraw.

3. **Expected timeline:** 4–8 weeks for review.

4. **Required before:** Any participant contact, any data collection.

5.3 Declaration of Helsinki

sOS adheres to the Declaration of Helsinki for research involving human participants:

- Research protocol submitted to ethics committee before start.
- Informed consent from every participant.
- Vulnerable populations: no participants from groups with diminished autonomy (inclusion criteria: ≥18 years, capacity to consent).
- Privacy and confidentiality protected.
- Research integrity: publish regardless of outcome; register protocol before data collection.

---

6. Liability: Research Risk, Not Product Liability

6.1 Why the PLD Does Not Apply (Yet)

The revised Product Liability Directive (Dir 2024/2853) covers products "placed on the market or put into service." A research experiment conducted under regulatory sandbox supervision, with no commercial deployment, no end users, and no market placement, is not "put into service" in the PLD's sense.

This changes when sOS moves to Phase 2 deployment. The PLD analysis in the original WP9 (§9.4) becomes relevant only at the deployment gate, not during the experiment.

6.2 Liability During the Experiment

Participant harm	UB institutional liability + research insurance	Informed consent; UB insurance covers approved research
Data breach	GDPR notification obligations (Art. 33–34)	DPIA; encryption; breach protocol in data management plan
Municipal infrastructure impact	Research agreement with city council	Contractor blinding; no-cannibalization check (WP14 §9)
Third-party claims	General civil liability (Código Civil Art. 1902)	Research insurance; legal counsel review of municipal agreement

6.3 Insurance Requirements

UB institutional research insurance	Covers approved research activities	UB (verify coverage scope for off-campus urban experiment)
Professional indemnity	Researcher professional liability	Individual or institutional
Cyber / data breach	Covers GDPR notification costs + potential fines	Research project insurance (€2–3K annual premium)

---

7. The DAO: Research Governance, Not Legal Entity

7.1 A Fundamentally Different Problem

Under the operational frame, the DAO needed a legal wrapper to hold assets, enter contracts, and limit liability. Under the experimental frame, **the DAO does not hold assets, enter contracts, or employ people** – UB and the research team do that. The DAO is the object of study, not the operator of the study.

Legal personality	Required (wrapper)	Not required (it's an experiment)
Asset custody	DAO treasury	UB research grant account
Contracts	DAO enters agreements	UB / research team enters agreements
Employment	DAO employs contributors	UB employs / contracts researchers
Liability	DAO (via wrapper) bears risk	UB bears institutional research risk

7.2 The Governance Token as Research Variable

The token is not a governance instrument for a real organization – it is an **independent variable** in an experiment. Researchers manipulate governance rules (token-weighted vs. reputation-weighted vs. one-person-one-vote) and measure the effect on coordination outcomes (MTTR, precision@k). This is no different from any social science experiment that assigns participants different decision rules.

Implication: The legal wrapper problem (original WP9 §10) is deferred to Phase 2. For Phase 1, the token exists only within the sandbox environment, has no legal effect, and governs nothing of real-world consequence. The association-law conflicts (member equality vs. token voting) do not arise because the DAO is not the association's governance – it's a research apparatus.

7.3 The Association as Research Vehicle

The Spanish non-profit association (sOS Research Lab) remains the legal entity – not as a DAO wrapper, but as a research organization that:

- Receives and manages research grants.
- Employs or contracts researchers.
- Enters data-sharing agreements with the municipality.
- Is accountable to funders and ethics committees.
- Operates the experiment through its research staff.

The DAO exists *inside* the experiment. The association exists *outside* the experiment as its operator. This separation is cleaner, simpler, and avoids every structural conflict identified in the original WP9 §10.5.

---

8. Regulatory Compliance Roadmap (Experimental Frame)

8.1 Phase 0 (Months 1–3, €5K verification budget)

Submit protocol to UB CER	Month 1	**Critical** (blocks recruitment)
File DPIA with UB DPO + notify AEPD	Month 1	**Critical**
Apply to AESIA AI regulatory sandbox	As soon as operational	High
Apply to CNMV FinTech sandbox (token)	Month 1–2	Medium (fallback: MiCA-exempt framing)
Draft informed consent form	Month 1	**Critical** (reviewed by CER)
Draft municipal data-sharing agreement	Month 1–2	**Critical** (blocks data access)
Verify UB research insurance coverage	Month 1	High
Register association with Ministerio del Interior	Month 1–2	High

8.2 Phase 1 Gate (Month 4, before data collection)

Ethics approval	UB CER approval letter received
DPIA filed	DPIA submitted to AEPD
AI sandbox	AESIA acceptance OR documented justification for proceeding without (with legal counsel opinion)
Token sandbox	CNMV acceptance OR documented MiCA-exempt justification
Data agreement	Signed municipal data-sharing agreement
Insurance	Confirmed coverage for experiment
Consent	Informed consent form approved by CER

8.3 Phase 1 (Months 6–23, intervention)

Quarterly ethics reports to UB CER	Every 3 months
Sandbox regulatory reporting (AESIA, CNMV)	Per sandbox terms
Participant withdrawal processing	Within 30 days of request
Data breach protocol (if needed)	Within 72 hours to AEPD

8.4 Phase 2 Gate (Post-experiment, before any deployment)

Publish sandbox results	Month 24–26
If experiment succeeds: New legal assessment under operational frame (full GDPR, AI Act, MiCA, PLD). New DAO legal wrapper. New ethics review for deployment.	
If experiment fails/is inconclusive: Publish. Revise. No legal transition needed – experiment stays in research frame.	

---

9. Risk Register (Experimental Frame)

R1	**UB CER rejects protocol**	Critical	Pre-submission consultation with CER; revise protocol per feedback
R2	**AESIA sandbox not operational** by experiment start	Medium	AI Act compliance self-assessment as fallback; limited-risk classification argument
R3	**Municipal data agreement denied**	Critical	Alternative domain (waste, parks); alternative city
R4	**Participant data breach**	High	PostgreSQL on UB infrastructure; encryption; breach protocol
R5	**Token draws regulatory attention** despite sandbox	Low-Medium	CNMV sandbox + MiCA-exempt fallback; zero financial value
R6	**CER requires protocol changes that weaken statistical power**	Medium	Consult with methodology expert; adapt protocol; document constraints
R7	**Liability claim from non-participant affected by experimental coordination**	Low	Same contractor pool; no-cannibalization check; UB insurance

10. Open Legal Questions (Experimental Frame)

These are the legal questions the sandbox itself is designed to help answer – they are outputs of the experiment, not prerequisites:

1. **What minimum AI Act obligations apply to stigmergic governance AI** in a regulatory sandbox context? The sandbox's co-assessment with AESIA will produce regulatory guidance.

2. **Can token-based governance mechanisms be tested under CNMV sandbox** without triggering MiCA obligations? The sandbox application tests this.

3. **Does the GDPR scientific research exemption** (Art. 89) actually work as intended for AI-mediated governance experiments, or does the DPA interpret it narrowly?

4. **What does "automated decision-making" mean** when AI presents options but humans choose? The sandbox provides a concrete case study for this regulatory boundary.

5. **How does the transition from research frame to operational frame work legally?** sOS will document this transition and publish it as a case study.

11. Conclusion

sOS is an experiment. The legal framework for an experiment is fundamentally different from the legal framework for a deployment. The defensible architecture is: **UB research ethics + GDPR Art. 89 (via LOPDGDD DA-17 + Ley 14/2011) + MiCA-exemption + AI Act limited-risk self-assessment.** The EU AI regulatory sandbox (AESIA) and the CNMV FinTech sandbox are opportunistic – applied for when available, not gated on.

The operational frame becomes relevant if and only if the experiment demonstrates that stigmergic coordination works. That's the Phase 2 gate. Until then, sOS operates as what it is: a scientific research program testing a falsifiable hypothesis under ethics supervision, not a product seeking compliance.

Branch A – Pi/DeepSeek. Rewritten from experimental framing. Requires Branch B independent review (Anthropic Opus 4.8). Replaces original operational-framing WP9.