Where AI Becomes Visible

An Evidence-Infrastructure Ecosystem for the AI Trust Stack

Anton Sokolov / Tyche Institute

2026-06-14

Abstract

The governance of artificial intelligence keeps running into the same wall: the duties are written as outcomes, but the evidence that would let anyone check them is scattered, partial, and arrives in the wrong state for comparison. This article presents the Tyche AI trust-stack evidence ecosystem — a connected program of observatories, corpora, claim gates, and article packets that map *where* AI systems, capital, compute, patents, standards, public institutions, and legal decisions become visible, and *which gates must close* before stronger ownership, compliance, deployment, safety, or crisis claims become defensible. The program's organising idea is a discipline rather than a dataset: every observation keeps its own denominator and proof ceiling, and a shared corpus — the Thesis Atlas — links them as evidence objects without fusing them into one master claim. We describe the six contours of the ecosystem, the connective role of the Atlas, and a portfolio of seven companion studies that instantiate the pattern across public-sector visibility, the private capital-and-patent core, standards and evidence carriers, legal-temporal accountability, and the research operation itself. The recurring result is bounded and, we argue, more useful for being bounded: public evidence can build a defensible visibility map of the AI trust stack well before it can answer who owns it.

1. The Visibility Problem

Ask a sharp question about artificial intelligence — who controls the trust stack, is the boom a bubble, did regulation lag the harm, is this system compliant — and the honest first answer is almost always the same: we cannot see well enough to say. The evidence exists, but it is fragmented across public filings, patents, standards drafts, procurement notices, algorithm registers, court decisions, and incident reports, and each of those surfaces shows a different, partial slice. A system can be visible in a procurement notice and invisible in a registry; a patent can be counted but not attributed; a policy can change without any preserved prior version to diff; a decision can be public but unquotable.

The tempting move is to paper over the gaps — to read a patent count as an ownership map, a market concentration as a crisis, a standards citation as compliance. The Tyche program takes the opposite stance. It treats the visibility gap itself as the object of study, builds instruments that make each partial surface legible on its own terms, and refuses to upgrade a visibility signal into a stronger claim until a named gate closes. The payoff is a research program that can be ambitious about its questions and disciplined about its answers at the same time.

2. The Top-Level Frame

Public-sector AI was the program's first published bridge, because public procurement, registries, and official inventories happen to expose AI systems early and cleanly. But public-sector AI is one evidence surface, not the scope of the program. The right top-level frame is wider:

> Tyche is building an AI trust-stack evidence ecosystem: a set of observatories, corpora, diagrams, claim gates, and article packets that show where AI systems, capital, compute, patents, standards, public institutions, legal decisions, and evidence carriers become visible before stronger ownership, compliance, deployment, safety, or crisis claims become safe.

The ecosystem has six contours: public AI visibility; the private capital-and-patent core; patent topology and intellectual-property evidence; standards, wallets, and evidence carriers; legal-temporal accountability; and the research-factory control that turns all of it into governed output. The remaining sections walk the connective layer and then the six contours as a portfolio.

3. Atlas: The Connective Layer

The instrument that lets these surfaces connect without collapsing is the Tyche Thesis Atlas, a versioned, metadata-first research corpus. The Atlas records source families — legal/regulatory, standards, sandbox and case material, public algorithm registers, procurement, patents, literature, and internal claim-support artifacts — and converts them into analytical layers: a law-to-evidence matrix, a standards-to-evidence matrix, an actor map, gap analysis, and paper-readiness scoring. Its current freeze holds 422,048 raw history records reduced to 155,789 canonical public-metadata records; the active dedupe authority reports 424,592 raw and 162,994 canonical, the difference being lineage rather than error.

The Atlas's discipline is what makes it connective. It represents a patent, a procurement notice, a standards fragment, or a legal decision as an *evidence object* with a denominator and a proof ceiling, and it links objects across surfaces without letting any object inherit another's strongest claim. A patent row does not become an ownership finding; a procurement row does not become deployment proof; a standards row does not become certification; a legal row does not become legal truth. That single rule — link, but never fuse — is the backbone of the whole ecosystem.

4. The Portfolio

The ecosystem is deliberately a portfolio of bounded studies rather than one sweeping paper. Each study owns one contour, keeps its own denominators, and reuses the others' grammar. Table 1 is the current portfolio.

Study	Contour	Strongest safe claim
From Procurement to Public Visibility	Public AI visibility (Registry, Procurement Gap, GAIA, PALLAS, LIMEN)	Public AI visibility is a chain of partial surfaces, not one register.
The Private Core of AI Trust Infrastructure	Private capital + patents (AION, Patent Topology, Trust Turning Point)	Public evidence maps visibility channels before it can answer who owns the stack.

Study	Contour	Strongest safe claim
Standards in the Gap	Standards + evidence carriers (SIGIL, EUDIW, Computational Trust, VESTA, RHR)	Carriers structure evidence and reduce ambiguity; they do not certify compliance.
When Evidence Arrives Late	Legal-temporal accountability (LIMEN, Temporal Commons, NOMAD, CLIO)	Governance evidence is often useful before it is comparable; type the row, name the gate.
Harvest Mode	Research-factory control (factory control, NIKA, sufficiency, papergrade)	A research operation needs governance for stopping, not only for starting.
Thesis Atlas Data/Method	The connective corpus	Governance evidence needs evidence objects, packages, and claim ladders.
Patent-Topology Pipeline (data + code)	Reproducible IP evidence	A repair-first pipeline can defensibly report scale without claiming ownership.

5. Public Visibility Is a Chain

The first contour shows that "public-sector AI visibility" is not a single register but a chain of partial surfaces — procurement traces, registry and source-accession rows, named system records, country-route evidence, and edge-case accountability panels. The companion study draws on a public-evidence edge-case atlas whose live panels are denominated explicitly (fifteen source-family rows, a twenty-one-lineage publication-safe evidence funnel, a twelve-row jurisdiction/language map), and it keeps each surface's missingness visible rather than smoothing it away. The blocked upgrades are exactly the tempting ones: deployment proof, registry or procurement completeness, country ranking, and public-sector prevalence.

6. The Private Core, Seen Through Public Evidence

The second contour asks the program's most quotable question — who owns the AI trust stack? — and answers it honestly: not yet, on public evidence. Three channels each expose a different slice. A financing-and-markets channel reads the AI-equity basket as concentrated and heterogeneous, with a cash-richer hyperscaler core and a transmission-sensitive supplier edge, and treats crisis as a conditional, staged transmission question rather than a settled fact. A patent channel turns a candidate flood into defensible evidence — 647,410 normalized candidate rows, 151,890 canonical records, a 495,520 duplicate-candidate burden — and reports its scale and a negative calibration result (automated substantive labelling fails at current precision) without claiming ownership, concentration, or freedom-to-operate. The patent data and code are openly archived (Zenodo `10.5281/zenodo.20644809`). A standards-and-infrastructure channel frames trust infrastructure as a turning-point problem. Together they map the private core's visibility channels and the gates — assignee normalization, family and citation closure, transmission denominators — that stand between the map and any ownership answer.

7. Carriers, Not Certificates

The third contour addresses the gap between a legal duty and an inspectable artifact. Duties are written as outcomes; an auditor can only inspect a carrier. The companion study assembles a duty-to-carrier chain across five observatories: a standards-fit ledger maps fifteen AI Act duty rows to candidate standards, gap scores, and checksummed artifacts (with an openly archived evidence package, Zenodo `10.5281/zenodo.20600721`); a specification-status ledger pins six wallet and credential specification families by status, date, and hash; a typed-verification layer and a computational-trust grammar supply receipts, traces, and fixtures; and a procurement observatory packages public AI/PQC/PKI signals with verifiable receipts. The bounded claim is that carriers can structure evidence and reduce ambiguity, but do not by themselves prove compliance, conformance, certification, or production assurance.

8. When Evidence Arrives Late

The fourth contour treats a temporal problem: AI-governance evidence is frequently visible in the wrong state for comparison. A row-state ladder — visible, source-family coded, date-role coded, claim-support linked, quote/denominator ready — runs across an edge-case atlas, a regulatory-lag observatory that types incident and response dates and defaults causal language to "not claimed," a normative-diff observatory that demands segment-pair evidence rather than a loose "the policy changed," and a legal-decision workbench that tiers six European decisions by anchor quality and quotation readiness. The discipline is that a row can be evidence-bearing while held below truth, prevalence, causation, legal effect, or compliance — and that visibility must stay separate from substantive use.

9. Governing the Operation Itself

The fifth contour turns the lens inward. A research operation that can spawn many parallel investigations needs governance for *stopping* — for knowing when to stop crawling and start harvesting. The companion study (a public-safe derivative, produced after a privacy/security scrub) describes a harvest-mode control loop: a readiness score sorts projects into classes, a gate-type taxonomy names each candidate's actual blocker, a hostile-reviewer pass holds claims to ceiling, and a route-control discipline keeps packages honest over time. It deliberately publishes only the governance pattern, not the operation it governs — and the scrub boundary that makes that possible is treated as part of the method.

10. The Common Grammar

What lets seven studies share logic without sharing a denominator is a common metadata grammar. Every node in the ecosystem carries the same minimum fields: the owner surface that holds the row, the evidence unit actually measured, the exact denominator, the source family and authority class, the proof ceiling (the strongest safe claim), the blocked reading (the tempting stronger claim that must not be inferred), the package route with source path and checksum, and the article use. The grammar is the real connective tissue. It lets the public-sector study, the private-core study, the standards study, the legal-temporal study, and the operations study reuse each other's structure while each keeps its own population.

11. Claim Ceilings for the Whole Ecosystem

Stated once, for the program as a whole. The ecosystem can safely claim that Tyche has multiple article-ready observatory surfaces; that the Atlas can connect them through evidence objects and claim-support gates; that public, private-core, patent, standards, legal-temporal, and operational evidence are distinct but connectable; and that the right production mode is harvest — audit, assemble, review, and record. The ecosystem must not claim who owns the AI trust stack; ownership concentration, market control, licensing, or citation influence; public-sector prevalence, deployment proof, or registry/procurement completeness; compliance, certification, conformance, trust-service qualification, official audit, or legal sufficiency; safety, incident truth, causation, representative regulatory lag, or legal truth; or any private operational or security posture. "Who owns the AI trust stack?" stays a research-program hook: the evidence can show why the question is hard, which channels are available, and which gates must close before an answer is defensible.

12. Conclusion

The Tyche ecosystem is not "public-sector AI," and it is not one mega-dataset, one dashboard, or one sweeping claim. It is an AI trust-stack evidence program: a connected set of instruments — with diagrams, denominators, proof ceilings, and article routes — built around a single corpus that links evidence objects without fusing them. Its contribution is to make the AI trust stack *visible* and to make the limits of that visibility *explicit*, so that the strong questions stay on the cover and the honest answers stay in the results. The portfolio above is the first full turn of that program; each study is now ready to stand on its own, and each points back to the same map.

Appendix. Companion Studies and Archives

The studies referenced above are available as internal article-facing packets in the Tyche research vault: the public-visibility bridge and the private-core, standards, legal-temporal, and harvest-mode bridges under `papers/observatory-bridges-2026-06-14/`; the ecosystem overview and this flagship article under `papers/tyche-ai-trust-stack-ecosystem-2026-06-14/`; and the Atlas data/method refresh under `papers/thesis-atlas-data-method-2026-06-14/`. Two companion data-and-code deposits are openly archived on Zenodo: the Patent-Topology pipeline (`10.5281/zenodo.20644809`) and the SIGIL standards-fit evidence package (`10.5281/zenodo.20600721`). Each packet carries its own verification ledger with SHA-256 checksums and explicit claim ceilings.