Cassandra: From Governance Infrastructure to Evidence Infrastructure

v0.3 Data & Policy / governance evidence-infrastructure candidate

Anton Sokolov

2026-05-28

Abstract

Public-key infrastructure (PKI) is often described as a cryptographic substrate, but in practice it is also governance infrastructure: a system of certificates, policies, revocation channels, trusted lists, audits, supervisory conventions, software distribution, and institutional delegation.[1] This paper presents Cassandra as a case study: a bounded research observatory for European list-of-lists-derived trusted-list artifacts. Cassandra is not an artificial-intelligence (AI) monitor and does not inspect models. It is a sample public governance-artifact case study for the evidence discipline that AI governance will need: claim boundaries, reproducible records, package receipts, and independently checkable provenance. The observatory records dated attempts to collect public trusted-list documents, normalizes comparable Extensible Markup Language (XML) artifacts, emits structural diff classes, packages each substantive run into evidence bundles, exposes aggregate public dashboard data, and attaches EATF/AEP receipt metadata to the resulting evidence packages.[2] The case study demonstrates a shift from governance infrastructure to evidence infrastructure: public legal-technical artifacts can be observed, hashed, bundled, signed, displayed, and independently checked while preserving the distinction between evidence integrity and legal interpretation. Across five dated local/public lineages, Cassandra records 43 pointer attempts per run, 31 comparable normalized XML artifacts per run, 66 aggregate structural diff entries under the configured method, five EATF-verified package receipts, and five public dashboard cards. The dashboard values are telemetry, not legal or supervisory conclusions. Synthetic fixtures cover stable/no-change behavior, normalized-hash changes, provider/service in-

[1] The common shorthand "PKI equals cryptography" is useful in a classroom and dangerous in governance. The keys do not audit themselves, revoke themselves, or explain themselves to a relying party at 03:00.

[2] EATF is used here for the project's Agent Trust Framework, and AEP for Agent Evidence Package. The phrase means a verifiable packet with claims and caveats, not a small ceremonial stamp that turns evidence into law.

ventory changes, fetch failures, non-XML handling, EATF success and tamper cases, missing signing inputs, dashboard states, and claim-safety wording. Cassandra is not a trusted-list validator, not a relying-party signature validator, not a public alerting system, not a supervisory tool, and not a legal status oracle. Its contribution is methodological: it shows how a public governance surface can become a reproducible evidence stream, and how the same discipline can inform emerging AI governance records without collapsing cryptographic evidence into legal truth.[3]

Keywords: evidence infrastructure; public-key infrastructure; trusted lists; eIDAS; EATF; digital government; computational public administration; reproducibility; artificial-intelligence governance.

Claim-Safety Note

All observations in this preprint are limited to locally recorded workflow telemetry and saved public-source artifacts: fetch metadata, parser outcomes, normalized XML hashes, structural diff classes, aggregate tables, evidence-bundle manifests, EATF/AEP receipt metadata, dashboard cards, and synthetic fixture check outputs. A Cassandra observatory run does not assert legal effect, does not determine whether any listed entity has gained or lost status, does not validate source signatures for relying-party purposes, does not supervise trusted lists, does not certify compliance, and does not provide public alerting; it is a research method demonstration, not a regulated trust-service system.

The evidence boundary is a chain of abstentions rather than a single endnote. Collection abstains from substituting sources; normalization abstains from relying-party validation; diffing abstains from legal interpretation; bundling abstains from claiming truth beyond source, hash, receipt, and caveat linkage; EATF/AEP receipts abstain from validating source legal meaning; and dashboard cards abstain from public alerting. This chain keeps the positive claim narrow: Cassandra records and packages research telemetry about saved public artifacts.

[3] AI is expanded here because the abbreviation is too easy to treat as a fog machine. In this paper it means governance records around artificial-intelligence systems, not a claim that Cassandra monitors models.

Reader's Map

Table 1. Reader's map for recurring terms.

Term	First-use reading	Role in this paper
PKI	public-key infrastructure	Governance substrate for digital trust.
eIDAS	EU electronic-identification and trust-services law	Legal-technical environment.
ETSI	European Telecommunications Standards Institute	Standards grammar for trusted lists.
LOTL	list of the lists	Public index pointing to trusted lists.
XML	Extensible Markup Language	Structured document format observed by Cassandra.
EATF/AEP	Agent Trust Framework / Agent Evidence Package	Evidence-envelope and receipt layer.
MIRROR	local evidence-bundle pattern	Claim, source, hash, and caveat discipline.
AI Act	EU artificial-intelligence regulation	Adjacent evidence-infrastructure pressure.

1. Introduction

Digital public administration increasingly depends on artifacts that are both technical and institutional. Certificates, trust lists, validation policies, audit regimes, public registers, and machine-readable administrative documents do not merely describe government infrastructure from the outside. They are part of how digital trust is governed. Yet the research record around such artifacts is often thin: a public register exists, a standard describes it, software consumes it, but the changing public state itself is not always preserved as an evidence stream that a later researcher, reviewer, or institution can inspect and replay later.

Cassandra addresses this gap through a deliberately small case study. It observes the European public list of the lists (LOTL) and the national trusted-list documents referenced from locally saved LOTL snapshots.[4] It records dated collection attempts, classifies endpoint and parser outcomes, normalizes comparable XML artifacts, emits configured structural diff classes, bundles each substantive run with manifests and claims, and exposes a caveated public dashboard index. The case is intentionally narrow. It does not decide legal status,

[4]LOTL is the European "list of the lists": a public index that points to trusted lists. It is an index surface, not a shortcut to legal interpretation.

validate trusted-list signatures for reliance, supervise any actor, or publish warnings. Its object is the research workflow and the saved public artifacts under that workflow.

This is the category boundary for the whole paper. Cassandra is not presented as an AI-system case study. It is an EATF use case over a public governance artifact: a deliberately safer laboratory for showing how claims, sources, hashes, caveats, receipts, and verifier outputs can travel together. Later case studies in the thesis program can move from this public-governance baseline to agent attestations, agent cards, evaluation receipts, incident and post-market records, and decision transcript packages. Cassandra provides the low-risk public substrate on which the more sensitive AI-governance evidence argument can stand.

The central claim is captured in one sentence: Cassandra moves from governance infrastructure to evidence infrastructure. PKI and electronic identification, authentication and trust services (eIDAS) trusted lists already provide governance infrastructure: certificates, policies, supervisory lists, validation conventions, and trust-service roles.[5] Cassandra does not replace that infrastructure. It watches one public surface of it and turns the observation process itself into evidence infrastructure: dated manifests, hashes, diffs, receipts, public cards, and explicit claim boundaries.

This matters beyond trusted lists. The European Union (EU) AI Act creates adjacent evidence problems around logging, transparency, technical documentation, post-market monitoring, and lifecycle records.[6] The AI governance problem is not only whether an AI system is safe or lawful; it is also whether claims about the system, the review, the data, the decision boundary, and the monitoring process can be bound to reproducible records. Cassandra is not an AI monitor. It is a public-governance-artifact case study that makes the evidence-infrastructure problem small enough to inspect.

The paper makes four contributions. First, it gives a cautious method for longitudinal structural observation of public trusted-list artifacts. Second, it reports a working full-stack observatory with scheduled GitHub Actions, Cloudflare dashboard output, aggregate tables, figures, dashboard cards, and EATF/AEP evidence receipts. Third, it maps synthetic fixtures to reviewer-facing claims so that software behavior is tested without exposing provider or service names. Fourth, it positions Cassandra inside a thesis program in which PKI is governance infrastructure, Cassandra is the public governance-artifact baseline, and EATF/MIRROR-style packages become an adjacent evidence layer for AI and public-administration records.

[5] eIDAS is EU law for electronic identification and trust services. Cassandra cites it as the environment of the observed artifacts, not as a badge of implementation or approval.

[6] The AI Act is cited as an adjacent recordkeeping pressure: technical documentation, logs, monitoring, and transparency records need evidence discipline. Cassandra remains a trusted-list case.

2. Background: Trusted Lists as Governance Artifacts

The European trust-services framework is a legal-technical system, not only a family of cryptographic formats. Regulation (EU) No 910/2014 establishes the eIDAS framework for electronic identification and trust services. Regulation (EU) 2024/1183 amends eIDAS through the European Digital Identity Framework. Commission Implementing Decision (EU) 2015/1505 specifies trusted-list formats. European Telecommunications Standards Institute (ETSI) TS 119 612 defines the trusted-list grammar.[7] ETSI EN 319 401 describes general policy requirements for trust service providers; ETSI EN 319 411-1 and EN 319 422 connect the trusted-list layer to certificate-issuing and timestamp-service governance. These sources explain why the observed XML is not arbitrary markup. It is a public administrative artifact embedded in a larger trust-services regime.

At the protocol layer, Request for Comments (RFC) 5280 anchors X.509 certificate and Certificate Revocation List (CRL) profiles, RFC 3161 anchors timestamp protocol vocabulary, and RFC 8785 is useful for thinking about deterministic JSON evidence representations adjacent to Cassandra's EATF/AEP package design.[8] Outside the EU trust-list context, CA/Browser Forum baseline requirements and the PKI Consortium show that PKI governance also exists as an ecosystem of multi-stakeholder rules, assurance expectations, and operational transition problems.[9] National Institute of Standards and Technology (NIST) FIPS 203, 204, and 205 further show that even core cryptographic primitives are moving under post-quantum transition pressure.[10] The evidence layer must therefore be able to record not only stable states but also migrations, caveats, and changing assumptions.

The public-administration and infrastructure-studies literature helps name what is otherwise easy to miss. Infrastructure is relational, learned in practice, and often visible through maintenance and failure. Classifications and standards organize institutional reality; they do not merely label it. Turning-point theory helps frame why technological deployment phases demand institutional reshaping, not just invention. The useful lesson for Cassandra is modest: if public digital trust is governed through structured artifacts, then those artifacts and their changes deserve reproducible evidence records.

This background also defines the boundary. Citing eIDAS, ETSI, RFCs, NIST, CA/Browser Forum, or the PKI Consortium (PKIC) does not imply endorsement by those bodies, implementation as a relying-party validator, or authority

[7] ETSI is the European Telecommunications Standards Institute. Its standards give trusted-list artifacts a grammar; they do not turn Cassandra into a conformity assessor.

[8] RFC means Request for Comments, the publication series used for Internet standards and related technical specifications. The name is charmingly modest; the consequences are often load-bearing.

[9] The CA/Browser Forum coordinates browser and certificate authority requirements for publicly trusted TLS certificates. It is adjacent governance context, not a Cassandra dependency.

[10] NIST is the United States National Institute of Standards and Technology. Its post-quantum standards are cited to show transition pressure on cryptographic infrastructure.

to determine legal status. The sources establish the environment and the vocabulary. Cassandra contributes a research method for observing a public surface of that environment.

3. Related Work and Positioning

Cassandra intersects five fields.

First, PKI and cryptographic governance explain the substrate. PKI is often introduced through keys and algorithms, but it operates through certificate policies, revocation mechanisms, audits, root distribution, software defaults, identity-proofing conventions, and institutional delegation. That is governance work with cryptographic parts. Cassandra uses this point to avoid treating trusted lists as mere input files. They are public records of a governance system.

Second, eIDAS and ETSI trust-service sources explain the legal-technical grammar. Trusted lists have prescribed roles and formats. This makes them attractive for structural observation, but it also makes careless claims dangerous. A structural diff is not a legal event; an endpoint failure is not a supervisory conclusion; an XML signature-shaped element is not a signature-validation result. Cassandra's claim boundary is therefore part of the method, not a disclaimer pasted on at the end.

Third, digital-government and public-administration scholarship explain why public machine-readable artifacts matter institutionally. Work on e-Estonia, public-sector innovation, public value, trust in digital government, and AI in government all points toward the same practical problem: administrative capacity is increasingly expressed through digital artifacts that need to be maintained, interpreted, audited, and remembered.

Fourth, infrastructure studies and classification studies provide the conceptual bridge. Star and Ruhleder's infrastructure lens, Bowker and Star's work on classification, and related knowledge-infrastructure work all resist the fantasy that technical records are neutral pipes. They are maintained systems with categories, access conditions, and failure modes. Cassandra's contribution is intentionally operational: it makes a specialized governance surface inspectable through manifests, hashes, diffs, dashboard cards, and receipts.[11]

Fifth, computational social science and digital research methods supply the cautionary frame. Repeated observation of public artifacts can produce useful evidence, but measurement validity, corpus boundaries, reproducibility, aggregation choices, and public interpretation risks must be explicit. Cassandra's synthetic fixtures, aggregate-only prose, claim-safety checks, and abstention rules

[11] Grand theory earns its keep here only if it can survive the manifest. Cassandra's theory chapter and its hash table should be on speaking terms.

are therefore methodological features. The highest compliment for an evidence pipeline may be that it becomes boring in a reproducible way.[12]

4. System Overview

Cassandra is implemented as a research-only workflow around dated runs. Each run begins from a saved copy of the public European list-of-lists. The workflow extracts trusted-list pointer URLs from that saved LOTL copy, fetches the referenced public artifacts, records per-endpoint metadata, normalizes XML-like inputs where possible, records non-XML and parser outcomes, compares normalized records against a configured baseline, emits structural diff classes, builds aggregate result tables and figures, creates evidence bundles, signs or packages the run through the EATF/AEP lane when signing inputs are available, and exposes a public dashboard index.

Figure 1. Cassandra evidence flow.

Stage	Evidence action	Output
Public LOTL snapshot	Save dated source context	Pointer set
Fetch	Record endpoint outcomes	Snapshot manifest
Normalize	Preserve comparable XML structure	Normalized artifacts
Diff	Compare configured structural fields	Aggregate diff classes
Package	Bundle claims, sources, and caveats	MIRROR/EATF evidence package
Publish	Expose bounded public summary	Dashboard cards and index

[12] "Boring" is not a complaint in evidence infrastructure. A reproducible system should be allowed at least one unglamorous virtue.

Table 2. Core workflow files and roles.

File	Role
`fetch.py`	LOTL pointer extraction and endpoint telemetry.
`parse.py`	Deterministic XML normalization and structural extraction.
`diff.py`	Configured structural diff classes over normalized records.
`run_daily.py`	Dated orchestration with overwrite guards.
`create_bundle.py`	MIRROR-style evidence-bundle creation.
`eatf_package_snapshot.py`	EATF/AEP packaging integration.
`build_observatory_index.py`	Public index and dashboard-card generation.

The public observatory is deployed at the canonical domain `cassandra.eatf.eu` (Cassandra dashboard), with the Cloudflare Pages deployment retained as a fallback at cassandra-observatory.pages.dev. Its machine-readable index is the public JSON index. The Agent Trust Framework (EATF) public verifier is available at `eatf.eu/verify` (public verifier). The latest public evidence package and receipt are exposed as AEP package and EATF receipt. As of the 2026-05-28 build, the public index reports five runs, latest date `2026-05-28`, five EATF-verified receipts, and five dashboard cards. The cards cover the latest run, EATF receipt boundary, aggregate diff classes, claim boundary, and dashboard caveat. Each card repeats that the data is structural-observation telemetry only, not trusted-list validation, signature validation, legal-status determination, supervision, compliance judgment, relying-party processing, public alerting, or publication approval.

The workflow's most important design decision is abstention. Cassandra keeps endpoint failures, non-XML artifacts, parser errors, and zero-diff days inside the record. It does not clean them away to tell a simpler story. A failed fetch is collection telemetry. A non-XML pointer is corpus-shape telemetry. A parser error is a local parser event. A zero diff means no configured structural diff was emitted by this parser/baseline pair. None of these observations determines legal status.

Table 3. Chain of abstentions in the evidence workflow.

Stage	What Cassandra records	What it refuses to infer
Collection	Source pointer attempts, fetch metadata, saved bytes, failures	Substitute sources, endpoint compliance, or public warnings

Stage	What Cassandra records	What it refuses to infer
Normalization	Comparable XML serialization and parser outcomes	Relying-party signature or certificate validation
Diffing	Configured structural movement over saved records	Legal effect, provider status, risk score, or supervisory meaning
Bundling	Claims, sources, hashes, receipts, and caveats	Truth beyond the bounded package and its declared evidence
Dashboard	Aggregate telemetry and public caveat cards	Official registry status, complete ecosystem view, or alerting service

5. Evidence Package Design

Each substantive run can be wrapped as an evidence package. The package contains a summary artifact, manifest, claims file, notes, source copies or source references, hashes, and verification output. This is a MIRROR-style pattern:[13] claims are made explicit, source links are recorded, and local verification is possible without asking the reader to trust a prose paragraph.

EATF/AEP adds package-level attestation semantics. For indexed `ok` runs, the EATF receipt verifies the corresponding package envelope, payload bytes, and declared hashes. It does not validate the underlying trusted-list source signatures, does not determine legal effect, and does not confer supervisory meaning. The distinction is crucial: evidence integrity is not legal truth. A package receipt can say that the evidence package is internally consistent under the declared verification process. It cannot say that a trust service is lawful, that a provider's status changed, or that a public authority has taken a legally significant act.

[13] MIRROR is used here as the project's evidence-bundle pattern: make claims explicit, keep sources close, hash what matters, and carry caveats with the package.

Field	Card value	Interpretation

Table 4. Package receipt boundary for the 2026-05-28 public card.

Field	Card value	Interpretation
Package	`cassandra-observation.aep`	Evidence package for the run.
Package hash	prefix `77604ab2b142`	Full hash is kept in the machine ledger.
Receipt	`eatf-verification.json`	Verification record for the package.
Receipt hash	prefix `1b24cd8cd3e2`	Full hash is kept in the machine ledger.
Status	`ok`	Package bytes, envelope, and declared hashes only.

The interpretation attached to that card is deliberately narrow: `ok` verifies package bytes, envelope structure, and declared hashes only. It does not upgrade the underlying trusted-list artifacts into legal conclusions.

Working Lemmas

These lemmas are deliberately small. They are not trying to make the case look more mathematical than it is; they state the logic that keeps the evidence claim honest.

Lemma 1: package integrity is not legal truth. Let B_t be the evidence bundle for run t, R_t its EATF receipt, and L_t a legal-status claim about the underlying trusted-list world. Cassandra can support:

$$\text{Verify}_{EATF}(B_t, R_t) = ok$$

It does not follow that:

$$\text{Verify}_{EATF}(B_t, R_t) = ok \Rightarrow L_t.$$

Lemma 2: a structural diff is a question, not a verdict. Let N_t be the normalized comparable XML set for run t. Cassandra computes:

$$D_t = \Delta(N_{t-1}, N_t).$$

If $D_t \neq \emptyset$, the result is a review question about saved artifacts. It is not a legal event. If $D_t = \emptyset$, the result is not proof that nothing legally relevant happened.

Lemma 3: a claim travels with its caveat. A useful evidence unit is:

$$E = (C, S, H, R, K),$$

where C is the claim, S the source context, H the hashes, R the receipt or verifier output, and K the caveat. Dropping K changes the claim. That is why Cassandra keeps caveats beside the dashboard cards, tables, and package metadata.

This boundary is especially important for AI governance. If evidence packages are to support AI lifecycle records, they must avoid the same overreach. A signed package can preserve what was reviewed, logged, or claimed. It does not prove that a model is safe, fair, lawful, or accurate. Cassandra provides a concrete public-artifact case study for this discipline before the argument is transferred to more sensitive AI contexts.

6. Dataset and Current Results

The current Cassandra evidence base contains five dated runs: 2026-05-20, 2026-05-21, 2026-05-22, 2026-05-27, and 2026-05-28. Each run starts from 43 LOTL-derived pointer attempts. To keep the table readable in print, run telemetry and diff-class totals are separated.

Table 5. Five-run structural-observation telemetry.

Date	Fetch result	Comparable XML	Diff entries	Aggregate inventory
2026-05-20	41 ok / 2 err	31	0	393 providers / 4743 services
2026-05-21	42 ok / 1 err	31	0	393 providers / 4743 services
2026-05-22	42 ok / 1 err	31	26	393 providers / 4753 services
2026-05-27	42 ok / 1 err	31	40	393 providers / 4847 services
2026-05-28	42 ok / 1 err	31	0	393 providers / 4847 services

Diff class	Count	Reading

Table 6. Aggregate structural diff classes across the five runs.

Diff class	Count	Reading
Normalized hash changed	7	Canonical bytes changed.
Summary field changed	28	Extracted summary field changed.
Service inventory changed	5	Hashed service inventory moved.
Provider-service detail changed	26	Hashed provider-service detail moved.
Listed document added	0	No configured add event emitted.
Listed document removed	0	No configured removal event emitted.
Provider inventory changed	0	No provider-inventory movement emitted.

Across the five runs, Cassandra records 66 aggregate structural diff entries under the configured method.

These are structural-observation classes, not legal-effect classes. Provider and service totals are aggregate parser outputs over the saved local corpus. The paper intentionally does not name providers, services, schemes, or authorities in narrative prose. Hashed provider/service handles support local comparison while reducing the temptation to turn a research draft into a provider-specific status feed.

The results show that the workflow can preserve multiple kinds of evidence at once: collection conditions, parser boundaries, comparable records, structural changes, package hashes, EATF receipt status, public dashboard summaries, and caveats. The result is not a claim that the trusted-list ecosystem was stable or unstable. It is a claim that Cassandra can produce reproducible structural-observation telemetry over saved public artifacts and preserve enough context for later review.

7. Fixture-Backed Behavior

Real public observations are not enough to make the case publishable. The paper also needs controlled behavior tests. Cassandra therefore maintains a fixture-to-claim map that links synthetic test classes to paper claims, reviewer questions, evidence artifacts, and explicit non-claims.

Table 7. Fixture classes compressed for print.

Fixture class	Reviewer-facing behavior
Stable and normalized-hash cases	No-change and canonical-byte movement are separated.
Inventory movement	Provider, service, and detail movement are distinct classes.
Fetch and non-XML handling	Failures, skips, redirects, and parser events remain visible.
EATF success and tamper	Valid packages verify; altered package bytes fail closed.
Missing signing input	Missing material is recorded as an explicit skipped state.
Dashboard and claim safety	Multistate cards and overclaim wording are tested.

The failed case is intentionally plain. A synthetic evidence package is first accepted under the expected wrapper and then altered after the receipt path is known. The verifier must return `verify_failed`, not a quiet success and not an ambiguous warning. A second failure mode records missing signing inputs as `skipped_missing_signing_inputs`. These cases are small on purpose: they show that the evidence lane has a negative control before anyone asks it to carry more sensitive claims.

The fixture boundary is explicit: fixtures prove expected software behavior on synthetic inputs. They are not empirical evidence about real trusted-list content, source signatures, legal status, supervision, compliance, public alerts, or publication approval. Their role is to answer reviewer questions about the method: does the system fail closed under tampering? Does it preserve failure states? Does it avoid silent success when signing inputs are missing? Does it catch common wording overclaims? These are software and method claims, not legal claims.

8. Discussion: From PKI Governance to AI Evidence

Cassandra's strongest thesis value is not that trusted lists are an especially fashionable dataset. It is that trusted lists sit at a rare intersection: public authority, structured technical artifacts, cryptographic governance, institutional trust, and machine-readable state. This makes them a useful case for observing how governance infrastructure can be transformed into evidence infrastructure.

Table 8. Thesis layer map.

Layer	Governance question	Cassandra evidence question
PKI	Who may be trusted, by which rules?	What public trust artifact was observed?
eIDAS/ETSI	Which legal-technical format applies?	Which structured fields were preserved?
MIRROR/EATF	Which evidence package was verified?	Which claims, hashes, and caveats travel together?
AI governance	Which lifecycle records matter?	Which records can be replayed with caveats?

Table 9. Adjacent case-study roles for the thesis.

Case cluster	Evidence role	Status discipline
Cassandra	Public governance-artifact observatory.	Primary case; aggregate-only prose.
MIRROR and EATF/AEP	Bundle grammar and package receipt layer.	Research framework, not trust service.
Janus and MATx	AI review drift and decision replay.	Use status labels; protect sensitive data.
Vesta and Icarus	Web/citation drift and reproducibility audit.	Adjacent cases; verify current artifacts.
Professional/environmental evidence cases	Audit-like public evidence records.	Publicness and rights check required before naming.
Prototype inventory cases	Unresolved or experimental evidence records.	Do not cite as adoption or public proof.

The phrase "PKI as governance infrastructure" matters because it stops the analysis from shrinking PKI into algorithms. PKI includes keys and signatures, but it also includes policies, identity checks, root stores, certificate profiles, revocation conventions, timestamps, audit frameworks, trust lists, and relying-party software. Cassandra watches a public administrative surface of that infrastructure. The watched surface is not the whole governance system, but it is enough to demonstrate the evidence move.

The AI governance connection is vertical and adjacent. The AI Act does not ask Cassandra to monitor AI systems. Instead, it makes visible a similar institutional problem: records about systems, reviews, logs, risks, datasets, human oversight, transparency, and post-market monitoring need to be reproducible enough to support accountability. EATF/AEP and MIRROR-style packages offer one grammar for such records: claims, manifests, hashes, receipts, caveats, source references, and local verification. Cassandra supplies a public, non-

sensitive case where that grammar can be tested against real public artifacts and synthetic failure modes.

This is also why the dashboard matters. The dashboard is not a marketing surface. It is a method surface. It exposes what the evidence stream contains and repeats what it refuses to claim. The public cards make the claim boundary visible at the point of consumption: run counts, receipt status, aggregate diff classes, and caveats travel together. That design choice is small, but it is a discipline AI governance systems will need: evidence must be packaged with its limits, not separated from them.

9. Limitations

The limitations are not incidental. They define the research object.

First, Cassandra observes saved public artifacts under one workflow. It does not establish endpoint completeness, current public availability, or authoritative state beyond the local evidence collected at a given time.

Second, the parser is not a trusted-list relying-party implementation. It records structural fields, signature-shaped elements, hashes, and parser outcomes. It does not perform source-signature validation, certificate-path validation, Online Certificate Status Protocol (OCSP) or CRL checks, or legal interpretation.[14]

Third, the diff classes are descriptive buckets. A normalized-hash change, summary-field change, service-inventory movement, or provider-service-detail movement is a local structural observation. It is not a status change, compliance finding, risk score, or supervisory signal.

Fourth, the empirical series is still short. Four dated runs are enough to demonstrate a full evidence loop, public dashboard, EATF receipts, and fixture-backed behavior. They are not enough to estimate long-term churn, seasonal patterns, national differences, authority behavior, or trust-service ecosystem stability.

Fifth, the public prose is aggregate-only by design. Named examples may be useful in a later paper version, but only after author review, source-path checks, bundle context, exact wording review, and a decision that the example improves explanation without turning the paper into a status feed.

Sixth, EATF/AEP receipts verify package integrity under the declared process. They do not validate underlying legal truth. This limitation is also the point. Evidence infrastructure becomes credible when it knows where evidence stops.

[14] OCSP means Online Certificate Status Protocol; CRL means Certificate Revocation List. Cassandra records structural telemetry around saved artifacts, not revocation status for relying parties.

10. Conclusion

Cassandra shows that a public governance artifact can be turned into a bounded evidence stream without pretending to become a regulator, validator, or oracle. The system collects public LOTL-derived artifacts, preserves endpoint and parser telemetry, normalizes comparable XML, emits structural diff classes, packages run evidence, attaches EATF/AEP receipts, exposes public dashboard cards, and checks its own claim boundaries through synthetic fixtures.

The case supports a broader thesis: PKI did not merely secure digital transactions; it institutionalized a way of governing digital trust through certificates, lists, policies, audits, revocation, timestamps, and validation conventions. AI governance now needs an adjacent evidence infrastructure for systems, decisions, reviews, logs, public artifacts, and lifecycle records. The Cassandra case study is the public-governance-artifact case. Its lesson is deliberately restrained: evidence integrity can be made visible and reproducible, but it must not be confused with legal truth.

Future Work and Release Discipline

The immediate experiments are deliberately compact: extend the dated run series before making longitudinal claims; add a fresh freeze manifest before changing any numerical count in this manuscript; check public URL freshness and content-level source roles before submission; and maintain the fixture map so dashboard, receipt, tamper, missing-input, and wording checks remain tied to explicit paper claims. The preferred route for this version is Data & Policy / governance evidence infrastructure: Cassandra is strongest as a public-administration and digital-trust case study, with the dataset/artifact layer supporting the argument rather than carrying it alone. A future venue-specific version may still split toward a Data in Brief dataset/artifact article if the artifact package becomes the primary contribution, but this v0.3 candidate should remain the shared bounded base until Anton chooses a final venue route.

The release rule is conservative. Cite only artifacts that are intended for public readers, have a stable version or recorded access state, and carry their caveats. Do not cite a repository URL as an available public source unless it resolves or repository visibility changes. Do not treat local ledgers, claim checks, or fixture outputs as external validation; they are internal audit aids for release review.

Artifact Availability

Code repository: no public repository URL is cited in this version. A previously tested candidate repository path did not resolve in the 2026-05-28 freshness check, so it is omitted from the availability statement until repository visibility or URL status is corrected.

Dashboard: cassandra.eatf.eu

Fallback deployment: Cloudflare Pages mirror

Public index: cassandra.eatf.eu/data/index.json

Agent Trust Framework verifier: eatf.eu/verify

Latest AEP package: cassandra-observation.aep

Latest EATF receipt: eatf-verification.json

Table 10. Local artifact groups for this preprint.

Group	Local files
Preprint outputs	v0.3 source markdown, PDF, and LaTeX header.
Reference controls	Checked reference ledger, URL validation output, and seed bibliography.
Claim controls	Claims/evidence table, fixture map, fixture matrix, and maturity matrix.
Public evidence	Observatory index, dashboard cards, public AEP packages, and EATF verification receipts through 2026-05-28.

The checked reference ledger and fixture-to-claim map are repository artifacts listed in this availability group. They are not external authorities; they are audit aids for matching claims to sources and fixtures before release.

Declarations

Funding: no external funding is declared for this pre-submission candidate.

Competing interests: no competing financial interests are declared here.

Data and artifact availability: the public dashboard and dashboard JSON index cited above are the public access surfaces for this version. Local ledgers, bundle manifests, fixture outputs, and validation notes are audit aids for release review; they are not independent public authorities and should not be cited as external validation.

Generative-AI assistance: the author used OpenAI Codex / GPT-5-based tools for manuscript revision, claim-boundary review, source-freshness bookkeeping, and release-readiness checks. The author reviewed and remains responsible for all claims, counts, legal boundaries, references, and final submission choices. The assistance did not create new empirical counts, perform legal interpretation, validate trusted-list source signatures, or approve the manuscript for submission.

References

Bannister, F., and Connolly, R. (2011). Trust and transformational government: A proposed framework for research. Government Information Quarterly. DOI: 10.1016/j.giq.2010.06.010.

Bowker, G. C., and Star, S. L. (1999). Sorting Things Out: Classification and Its Consequences. MIT Press. DOI: 10.7551/mitpress/6352.001.0001.

CA/Browser Forum. Baseline Requirements for the Issuance and Management of Publicly-Trusted TLS Server Certificates, version 2.2.7, 19 May 2026. Baseline Requirements.

Commission Implementing Decision (EU) 2015/1505 laying down technical specifications and formats relating to trusted lists. EUR-Lex record.

Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and Polk, W. (2008). RFC 5280: Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. RFC Editor.

Drechsler, W. (2018). Pathfinder: e-Estonia as the beta-version. JeDEM. DOI: 10.29379/jedem.v10i2.513.

ETSI EN 319 401 V3.1.1 (2024-06). Electronic Signatures and Infrastructures (ESI); General Policy Requirements for Trust Service Providers. ETSI PDF.

ETSI EN 319 411-1 V1.5.1 (2025-04). Electronic Signatures and Infrastructures (ESI); Policy and security requirements for Trust Service Providers issuing certificates; Part 1: General requirements. ETSI PDF.

ETSI EN 319 422 V1.1.1 (2016-03). Electronic Signatures and Infrastructures (ESI); Time-stamping protocol and time-stamp token profiles. ETSI PDF.

ETSI TS 119 612 V2.4.1 (2025-08). Electronic Signatures and Infrastructures (ESI); Trusted Lists. ETSI PDF.

European Commission. EU Trusted Lists / List of the Lists documentation. European Commission page.

European Parliament and Council. Regulation (EU) No 910/2014 on electronic identification and trust services for electronic transactions in the internal market. EUR-Lex record.

European Parliament and Council. Regulation (EU) 2024/1183 establishing the European Digital Identity Framework. EUR-Lex record.

European Parliament and Council. Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence. EUR-Lex record.

Lazer, D., Pentland, A., Adamic, L., Aral, S., Barabasi, A.-L., Brewer, D., Christakis, N., Contractor, N., Fowler, J., Gutmann, M., Jebara, T., King, G., Macy, M., Roy, D., and Van Alstyne, M. (2009). Computational Social Science. Science. DOI: 10.1126/science.1167742.

NIST. FIPS 203: Module-Lattice-Based Key-Encapsulation Mechanism Standard. NIST CSRC.

NIST. FIPS 204: Module-Lattice-Based Digital Signature Standard. NIST CSRC.

NIST. FIPS 205: Stateless Hash-Based Digital Signature Standard. NIST CSRC.

Perez, C. (2002). Technological Revolutions and Financial Capital: The Dynamics of Bubbles and Golden Ages. Edward Elgar.

PKI Consortium. PKIC website.

Salganik, M. J. (2017). Bit by Bit: Social Research in the Digital Age. Princeton University Press.

Star, S. L., and Ruhleder, K. (1996). Steps Toward an Ecology of Infrastructure: Design and Access for Large Information Spaces. Information Systems Research. DOI: 10.1287/isre.7.1.111.

RFC 8785: JSON Canonicalization Scheme. RFC Editor. RFC Editor.

World Wide Web Consortium (W3C). Verifiable Credentials Data Model v2.0. W3C Recommendation.

Adams, C., Cain, P., Pinkas, D., and Zuccherato, R. (2001). RFC 3161: Internet X.509 Public Key Infrastructure Time-Stamp Protocol. RFC Editor.