Reliable Context

RELIABLE CONTEXT / RESEARCH & PRACTICAL GUIDANCE

Data Provenance Privacy: A Recipient Disclosure Test

> Reliable Context is a kernel of truth: a system of record for organizations that fuse private and public information and must defend every value they act on. One canonical record per real thing, a cited source and an observation date behind every value, unknowns declared rather than left blank, and any past state of the record reproducible.

Published by Reliable Context · 2026-10-22 · data provenance privacy · sensitive data provenance · provenance access control · provenance redaction · w3c prov · odrl · knowledge graph access control · context engineering · evidence disclosure

Original article: https://reliablecontext.com/articles/data-provenance-privacy-recipient-disclosure-test



In this article

1. Executive Summary
2. Introduction and Background
3. Define the disclosure boundary
4. What PROV-AQ and ODRL establish
5. A two-source example and the risk of inference
6. The recipient-by-component disclosure matrix
7. How to redact and test the delivered evidence
8. Policy changes, revocation, and export review
9. Data Analysis and Evidence
10. Implications and Future Directions
11. Conclusion

Executive Summary

Data provenance privacy is the problem of disclosing enough information to assess a value while protecting the sources, relationships, people, and activities behind it. The permission to receive a value does not establish permission to receive its complete evidence trail. The **World Wide Web Consortium (W3C)** published **Provenance Access and Query (PROV-AQ)** as a Working Group Note on **30 April 2013**. (www.w3.org) It distinguishes openly accessible resources from the sensitivity of their provenance. (www.w3.org) It identifies derivation traces and editing timestamps as disclosure channels. (www.w3.org) Research also separates protecting provenance from using provenance to authorize access to underlying data. (profsandhu.com)

The proposed **recipient disclosure test** produces a decision for a particular assertion, recipient, and export: release the complete authorized evidence; release a clearly qualified view; or withhold the assertion. It examines the value, citation, excerpt, source identifier, timestamps, and derivation path separately. This separation follows research treating data, attributes, and the causality graph as distinct security concerns. (www.usenix.org) A restricted view must say when its recipient cannot inspect enough evidence to verify the claim. Hiding a source name while preserving a revealing relationship does not settle the problem: relationships themselves disclose information about their endpoints. (www.usenix.org)

Open Digital Rights Language (ODRL) 2.2, a Recommendation dated **15 February 2018**, supplies policy vocabulary, including permissions, prohibitions, duties, constraints, and conflict handling. (www.w3.org) It supplies no evidence that this particular export was authorized and correctly rendered. That distinction is an analytical inference from its policy-expression purpose and the explicit separation of policy decisions from enforcement in **XACML**, the eXtensible Access Control Markup Language. (docs.oasis-open.org)

The worked example uses a **synthetic graph with 6 nodes and 8 directed edges**. Its complete view contains **3 paths** from the asserted value to a restricted source. The restricted projection contains **0 such paths**, but still needs an independent metadata-inference check. These are enumerated toy-graph results, not a security guarantee or performance benchmark. A published formal example demonstrates why withholding sometimes remains necessary: the public output can itself disclose a confidential dependency.

(www.usenix.org) The practical deliverables are a disclosure matrix, a reproducible path-count method, and a sign-off card. Permission changes require fresh output tests; even protected-document systems document circumstances in which changed permissions take effect only after an existing use license expires. (learn.microsoft.com)

Introduction and Background

Consider a hypothetical example: a reviewer receives a cited number. The number seems ordinary, but its footnote names an internal project, its source address identifies a confidential folder, and its derivation path reveals consultation of a restricted document. The review question is whether that recipient may receive the complete package, a limited package, or no assertion at all. **Sensitive data provenance** concerns this whole disclosure surface. The provenance record can be more sensitive than the value it describes. Braun, Shinnar, and Seltzer's provenance-security research distinguishes security settings for data and provenance. (www.usenix.org)

This report addresses context engineers, data stewards, and evidence export reviewers sharing individual claims across permission boundaries. Reliable Context publishes this analysis from the perspective of a system-of-record provider; the proposed test specifies observable outputs without assuming a provider's authorization mechanism. Its unit of analysis is an **asserted value with a specific meaning**, rather than an isolated scalar. A count without its population, exclusion rule, and observation point is a different review object from the same count with those qualifications. W3C's **Provenance Data Model (PROV-DM)** provides the underlying vocabulary of entities, activities, and people involved in producing information. (www.w3.org)

The scope includes confidential organizational knowledge and intellectual property. Auge, Scharlau, and Heuer explicitly frame provenance-query privacy in terms that include intellectual-property protection. (arxiv.org) Accordingly, the test asks what a recipient can learn about hidden sources and methods, alongside what the recipient can inspect about the exported claim.

Sources were checked on October 22, 2026. The report separates standards, historical research, documented control behavior, and its own proposed evaluation method. ProvAbs provides research on partial disclosure through abstraction; it does not validate the worksheet proposed here. (arxiv.org) Documented database controls likewise illustrate questions to test rather than establish protection for a complete evidence export. For example, BigQuery requires both dataset and policy-tag permission for protected columns, while Lake Formation documents filtering at column, row, and cell levels. (docs.cloud.google.com) (docs.aws.amazon.com) Neither statement determines what a rendered citation, copied excerpt, or exported graph reveals. That determination requires inspection of the actual recipient-visible output.

Define the disclosure boundary

The assertion and its evidence components

Data provenance describes production history. Within the proposed test, an **asserted value** is a value together with the predicate, scope, units where applicable, and observation context that give it meaning. An **evidence excerpt** is the exact source passage offered in support. A **source identifier** is the locator or

reference used to distinguish the originating material. A **provenance path** is a sequence of relationships connecting the assertion to that material. These are operational definitions for the worksheet, organized around the entities, activities, and responsible agents represented in PROV-DM. (www.w3.org)

A **recipient** is the person or service receiving a particular representation. A **permission boundary** separates information that recipient may receive from information the recipient may not receive under the applicable policy. Clearance is an input to that decision, not a substitute for evaluating the requested action. The **National Institute of Standards and Technology (NIST)** definition of **attribute-based access control (ABAC)** evaluates attributes associated with subjects, objects, operations, and sometimes the environment against policy. (csrc.nist.gov)

The reviewer should inventory the package before deciding whether redaction is sufficient:

- **Value:** The assertion, its units, scope, and qualifications.
- **Citation:** The displayed title, author, reference label, and footnote text.
- **Excerpt:** The quoted passage and any surrounding context delivered with it.
- **Identifier:** The source's Uniform Resource Identifier (URI), document key, or attachment name.
- **Time:** Observation, publication, retrieval, and activity timestamps where present.
- **Path:** Nodes, relationship labels, and intermediate activities included in the lineage.

This inventory extends beyond a visual footnote. OWASP, the Open Worldwide Application Security Project, recommends incorporating static resources into access-control policies. The proposed review therefore includes downloadable attachments and other delivered resources. (cheatsheetseries.owasp.org)

Authorization and inspectability are separate judgments

The **authorization judgment** asks whether disclosure is permitted. The **inspectability judgment** asks whether the visible evidence lets this recipient assess the assertion's actual meaning. The proposed test requires both judgments to be recorded. A source can be authorized but inadequate; an adequate source can be unauthorized.

This distinction matters because provenance protection is not achieved merely by denying direct access. Chong's analysis shows that apparently public data can reveal confidential provenance. (www.usenix.org) Similarly, PROVANON describes the receiving environment through agents, supplementary data, infrastructure, and governance. (hummedia.manchester.ac.uk) A clean export must therefore be evaluated against the recipient's relevant supplementary knowledge, with uncertainty recorded rather than assumed away.

What PROV-AQ and ODRL establish

A leakage warning, not a disclosure protocol

PROV-AQ describes access to and discovery of provenance. Its publication status matters: it is a **Working Group Note**, not a Recommendation. (www.w3.org) Its security discussion warns that “derivation traces might indicate access to sensitive materials” and recommends assessing sensitivity and applying appropriate access

controls. (www.w3.org) This supports reviewing source existence, relationships, and metadata alongside content. It does not specify the recipient disclosure worksheet proposed here.

The report's practical reading is conservative: exposing a value is a separate decision from exposing its evidence. The warning supplies a reason to test the package, rather than permission to publish a reduced graph under any particular policy.

Policy vocabulary and conflict handling

ODRL represents usage policies concerning assets and actions. A permission applies when its constraints are satisfied and its duties fulfilled; a duty is an obligation to exercise an action. (www.w3.org) A prohibition expresses a disallowed action under its applicable conditions. The model can describe the policy attached to evidence, but the export reviewer still needs the applicable policy and an observable result. (www.w3.org)

For conflicting permissions and prohibitions, ODRL defines **perm**, **prohibit**, and **invalid** strategies. The default is **invalid**: the policy is void when a conflict is detected. These are model semantics, not a universal instruction that permission always wins or denial always wins across all systems. (www.w3.org) The reviewer should record the applicable conflict interpretation rather than infer it from the presence of a policy file.

The distinction between expression and enforcement is particularly clear in XACML. Its policy decision point evaluates policy; its policy enforcement point requests and enforces authorization decisions. (docs.oasis-open.org) From that separation, this report infers that an ODRL statement alone cannot demonstrate enforcement of an evidence export. ODRL terminology should not be treated as a licence to disclose restricted sources.

The European Commission's **Simpl** ecosystem page explicitly refers to ODRL for expressing data usage policies. That is evidence of the reference only. It is not an endorsement of this report's test, a recipient-clearance guarantee, or a validation result. (simpl-programme.ec.europa.eu)

What the reviewer must establish

Translate the applicable policy into observable expectations:

- **Policy scope:** Which assertion and evidence components does the rule cover?
- **Recipient context:** Which identity, role, purpose, or environment is relevant?
- **Requested action:** Is the request to inspect, download, forward, or otherwise use material?
- **Conflict treatment:** Which applicable interpretation resolves overlapping rules?
- **Decision evidence:** Who approved the interpretation and which version was applied?

These are proposed review fields. NIST's ABAC definition supports treating authorization as a context-dependent evaluation, rather than a property permanently attached to a citation. (csrc.nist.gov)

A two-source example and the risk of inference

Synthetic example: a public count and a restricted scope rule

The following is a **hypothetical example**, constructed solely for this report. The proposed assertion **V** is: **37 entries satisfy the release scope at the stated cut-off**. Public source **P** lists the entries. Restricted source **R** supplies the release-scope rule. The public list supports the count of listed entries, but cannot by itself justify the additional assertion that those entries satisfy the restricted rule.

The graph also contains activities **A**, **B**, and **C**, representing counting, reconciliation, and a supporting check. For this calculation, each arrow points from a result toward a dependency. The edge list is **V→A, V→B, A→P, A→R, B→R, B→C, C→R, C→P**. These labels and relationships are synthetic; they make no claim about a production implementation.

The complete view exposes the restricted source through several paths. Removing the direct connection from **B** to **R** would leave other connections intact. This illustrates the graph-structural concern described by Braun and colleagues: relationships disclose information about connected entities. (www.usenix.org) A path labeled merely “scope check” can also be revealing if the recipient knows that such checks occur only when a particular confidential source exists. That is an inference challenge, even when the source text is absent, consistent with Chong's analysis of public information disclosing private dependencies. (www.usenix.org)

Redaction must preserve honest meaning

Provenance redaction access control concerns both removal and the meaning of what remains. Cadenhead, Khadilkar, Kantarcioglu, and Thuraisingham's **2011** research proposed graph-grammar transformations for provenance redaction. It is a historical research example, with a prototype, rather than proof of any contemporary system's behavior. (www.profsandhu.com)

ProvAbs examines grouping graph nodes into abstractions. Its rewriting aims to preserve graph validity; arbitrary deletion can leave disconnected fragments. (arxiv.org) The proposed test therefore prohibits repairing a broken export by inventing a direct source relationship that the original provenance does not support.

For the restricted recipient, the visible evidence can support a narrower statement: “The public register lists 37 entries.” That is a **different assertion**, even though the number is unchanged. The export should identify the change of scope and retain a distinct review outcome for the original claim. If the required decision depends on the release-scope rule, the original assertion is **unverifiable to that recipient**. Where recipient inspection of that basis is required, withhold the original assertion.

The recipient-by-component disclosure matrix

The proposed **data provenance access control** test compares a complete authorized recipient, a restricted recipient authorized for the public evidence, and a recipient authorized for neither package. These classes are test inputs, not inferred user identities. OWASP's recommendation to deny by default supports making unestablished permission an explicit negative case. (cheatsheetseries.owasp.org)

Table 1 specifies expected outputs for the synthetic example. **Visible** means inspectable in the delivered representation. **Hidden** means omitted, including revealing substitutes. **Withheld** means the assertion is not supplied.

Component	Full-view recipient	Restricted-view recipient	No-access recipient
Asserted value	Original scoped assertion visible.	Number visible only with the original assertion marked unverifiable to this recipient; otherwise withhold.	Withheld.
Citation	Authorized public and restricted references visible.	Public reference visible; restricted reference and revealing title hidden.	Hidden.
Evidence excerpt	Relevant authorized excerpts visible.	Public excerpt visible; restricted excerpt hidden.	Hidden.
Source URI	Authorized locators visible.	Public locator visible; restricted locator and identifying replacement hidden.	Hidden.
Timestamps	Authorized source and activity times visible.	Public-source time visible if authorized; identifying restricted-activity times hidden or explicitly generalized.	Hidden.
Derivation path	Complete authorized graph visible.	Public-supported path visible; restricted dependencies and revealing intermediate labels hidden.	Hidden.
Evidence status	Recipient can inspect the offered basis, subject to ordinary evidence review.	Limited view explicitly states the original scope cannot be checked; a narrower claim requires separate identification.	Generic non-disclosure outcome without a restricted-source inventory.

The matrix does not classify a restricted view as fully evidenced. Its central safeguard is the separation between an available number and an inspectable basis. A source locator that returns an access error can itself disclose existence. **Hypertext Transfer Protocol (HTTP)** semantics recognize this distinction: **Request for Comments (RFC) 9110** permits a server to use **404** when it is unwilling to disclose that a resource exists, while **403** indicates refusal to fulfill an understood request. (www.rfc-editor.org)

Expected outcomes and acceptance rules

The recipient's evidence review should record a deterministic outcome for the intended use:

- **Complete release:** Every delivered component is authorized and the required basis is inspectable.
- **Limited release:** The allowed view states what evidence is unavailable and what cannot be verified.
- **Narrower assertion:** Any revised claim is explicitly identified and assessed against its visible evidence.
- **Withholding:** The original assertion is absent when its required basis cannot be disclosed safely.
- **Generic explanation:** A denial reason does not enumerate hidden sources unless that disclosure is authorized.

Avoid equating a valid container with a complete evidentiary basis. The **Coalition for Content Provenance and Authenticity (C2PA)** specification, for example, does not require inability to retrieve or validate external assertion data to reject a claim in standard validation. That documented behavior concerns its validation

procedure, not the sufficiency of evidence for this recipient's decision. (spec.c2pa.org) The worksheet asks the latter question explicitly.

How to redact and test the delivered evidence

Inspect every representation the recipient receives

Review the actual file or response under the recipient's permissions. OWASP recommends checking permissions on every request. (cheatsheetseries.owasp.org) For this test, opening an authorized preview is insufficient when the recipient also receives an attachment, machine-readable export, link preview, or cached representation.

Documented controls illustrate why representative permissions matter. PostgreSQL normally permits table owners to bypass row security; its default-deny behavior applies when row security is enabled without an applicable policy. (www.postgresql.org) Thus an export inspected only through a privileged account does not establish what a restricted recipient sees.

A practical **citation metadata disclosure risks** checklist should cover:

- **Titles:** Internal project names, restricted document classes, and descriptive filenames.
- **Locators:** URI paths, query parameters, folder names, and document keys.
- **Excerpts:** Headers, annotations, and surrounding text copied with the relevant passage.
- **Relationships:** Hidden endpoints exposed by named activities or dangling references.
- **Time patterns:** Exact activity times that distinguish a confidential workflow.
- **Counts:** Omitted-reference totals or graph sizes that reveal a hidden source inventory.
- **Stable tokens:** Repeated identifiers that let recipients link separate restricted events.
- **Alternate outputs:** Footnotes, attachments, structured fields, previews, and error responses.

These are proposed adversarial review prompts. The underlying distinction between content, attributes, and graph structure comes from provenance-security research. (www.usenix.org) UK Anonymisation Network's PROVANON discussion also emphasizes supplementary data that could be integrated with the shared data. (ukanon.net) The reviewer should identify plausible linking information for this recipient, without claiming to enumerate everything the recipient knows.

Preserve the status of redacted information

How to redact sensitive source citations is a question about the whole evidence package. Remove unauthorized excerpts, locators, titles, timestamps, and relationships together where policy requires. Then test whether the remaining package still supports the assertion's meaning. BigQuery's documented masking options include substituting null, default, or hashed content; such a substitution is not itself a statement that the original evidence remains inspectable. (docs.cloud.google.com)

Redaction rules are also format-specific. C2PA permits specified assertion redaction methods but prohibits redacting its actions assertions because of their historical significance. (spec.c2pa.org) A reviewer should not generalize a redaction technique across formats without checking the applicable semantics.

For a **knowledge graph access control** review, preserve the distinction between a hidden relationship and a replacement abstraction. ProvAbs provides a formal research treatment of graph abstraction. (arxiv.org) The proposed acceptance rule is straightforward: disclose the loss of evidentiary scope; do not silently present a transformed graph as the complete original basis.

Policy changes, revocation, and export review

Recipient authorization for data provenance is a decision at a specified time and policy state. It should be re-tested when the recipient's permissions, the source classification, the sharing purpose, or the receiving environment changes. NIST Privacy Framework guidance recommends reassessment to verify that privacy capabilities and requirements remain fulfilled. (www.nist.gov) The UK Anonymisation Network similarly describes data situations as dynamic. (ukanon.net)

The proposed change tests concern observable outputs:

- **Recipient downgrade:** Repeat the export after removing access to the restricted source; compare every component with the expected limited or withheld result.
- **Recipient revocation:** Confirm fresh requests and fresh exports no longer disclose the formerly authorized package.
- **Source reclassification:** Re-test citations, excerpts, paths, and attachments after the source becomes restricted.
- **Policy replacement:** Record the new interpretation and regenerate the expected matrix before approving an export.
- **Previously delivered copy:** Identify the actual protection and access behavior of that copy instead of assuming a new policy retracts it.
- **Receiving-environment change:** Reassess linking risks when recipients gain supplementary data or the package moves to a different environment.

Changed access does not imply immediate withdrawal from every delivered representation. Microsoft documents use-license circumstances in which users who already opened encrypted content do not receive changed permissions until license expiry and reauthentication. (learn.microsoft.com) Microsoft also documents encryption persisting inside or outside the organization, a property that should be tested separately from when a new permission takes effect. (learn.microsoft.com)

Caching requires its own observation. RFC 9111 restricts shared-cache reuse of authenticated responses unless an appropriate response directive permits it; it also warns that **no-store** alone is not sufficient to ensure privacy. (www.rfc-editor.org) These are reasons to inspect delivered and retained representations, without inferring that one header satisfies the proposed disclosure test.

Product-specific historical access paths also matter. Google's column-security documentation discusses waiting for the configured time-travel window after sensitive-data cleanup before safely relaxing security. (docs.cloud.google.com) AWS documents that an all-rows filter can prevail over other concurrent filter expressions. (docs.aws.amazon.com) Neither behavior is a general provenance rule. Both illustrate why the test should use the effective recipient view rather than an isolated rule read in a configuration screen.

Data Analysis and Evidence

A reproducible path-count calculation

The quantitative core is a **synthetic evaluation**, not an empirical benchmark. Starting with the edge list above, enumerate every directed simple path from **V** to restricted source **R**. A simple path does not repeat a node. The complete graph yields **V→A→R**, **V→B→R**, and **V→B→C→R**. Its public-source paths are **V→A→P** and **V→B→C→P**.

The restricted projection retains only **V, A, P** and edges **V→A, A→P**. The no-access projection exports no graph. The counts were calculated by depth-first enumeration over the listed adjacency relationships. For each outgoing edge, extend the path unless its endpoint already occurs in that path; count a path when its endpoint is the target. This procedure can be repeated directly from the printed edge list.

Table 2 reports the resulting counts. A **zero restricted-path count** is necessary for this particular hidden-endpoint objective, but does not establish absence of indirect disclosure.

Synthetic recipient view	Visible nodes	Visible edges	Paths from V to R	Paths from V to P	Required interpretation
Full view	6	8	3	2	Restricted basis inspectable by the authorized recipient.
Restricted projection	3	2	0	1	Original scoped assertion unverifiable; public-count claim requires separate labeling.
No-access projection	0	0	0	0	Assertion withheld.

The same restricted-path count appears for the limited and withheld outputs, although their utility differs. Path suppression therefore cannot be the sole success metric. An additional test should ask whether a surviving title, timestamp, token, or combination of public facts identifies **R**. Chong's formal counterexample establishes that a secure provenance trace need not exist when the public output already reveals the dependency. (www.usenix.org)

What historical measurements do and do not measure

Cadenhead and colleagues varied graph size with a fixed **200-rule** set and, in another experiment, varied rules with a fixed **87-triple** graph. Their utility measure concerned original triples left unaltered. These are historical prototype settings, not evidence about modern export latency, recipient verification, or resistance to inference. (www.profsandhu.com)

Completeness also requires a defined meaning. **Supply-chain Levels for Software Artifacts (SLSA)** provenance describes where, when, and how software artifacts were produced; its documented resolved-dependency completeness is best effort at least through **Build L3 (level 3)**. (slsa.dev) An **in-toto** Statement binds an attestation to a subject and predicate type; subjects are matched by digest. These structural properties do not answer whether the recipient can inspect the basis of a business assertion. (github.com)

The proposed outcome record therefore keeps authorization, inspectability, path exposure, and residual inference risk separate. No aggregate score is offered: these dimensions have different failure meanings, and the synthetic calculation provides no basis for assigning weights.

Implications and Future Directions

A defensible export review needs a record of the decision and its limits. *Table 3* is a reusable **sign-off card** for the proposed test. It is an evaluation artifact, not a certification, compliance claim, or description of a vendor's internal authorization mechanism.

Field	Required entry	Review purpose
Assertion and version	Exact value, predicate, scope, and observation context.	Prevent a narrower public claim from being confused with the original assertion.
Policy owner	Named accountable role or responsible owner.	Identify who can resolve policy interpretation.
Policy version and date	Applied policy identifier, effective date, and review time.	Bind the result to the policy used.
Recipient class and context	Relevant permissions, purpose, and receiving environment.	Define the tested disclosure boundary.
Evidence visibility	Visible, hidden, or withheld status for every matrix component.	Make the delivered package reviewable.
Inspectability and decision	Complete, limited, narrower assertion, or withheld, with reason.	State whether the recipient can assess the claim.
Residual uncertainty	Known inference risks and limits of supplementary-knowledge assessment.	Avoid treating untested inference channels as resolved.
Output and change tests	Reviewed artifact reference and applicable re-test triggers.	Connect approval to an observable output.

The card distinguishes policy interpretation from output inspection, consistent with XACML's separation of decision and enforcement. (docs.oasis-open.org) It also makes later reassessment possible without implying that a past approval applies indefinitely. (www.nist.gov)

For **protecting confidential data lineage**, future evaluation should examine compositions: multiple individually acceptable exports may expose a relationship when combined. The receiving-environment perspective provides a basis for selecting supplementary information to test. (ukanon.net) SLSA offers a relevant design caution in its own domain: ignoring an extension should not turn a denial into an allowance. That is a useful analogy for reviewing missing metadata, not a claim that SLSA governs this worksheet. (slsa.dev)

The article's scope is the evidence package, so it does not compare vendors or infer a publisher's implementation. Choosing a database authorization feature requires its own evaluation, including documented limits such as Lake Formation's cell-security exclusions for nested columns, views, and resource links. (docs.aws.amazon.com) The recipient test remains a review of delivered information and declared uncertainty.

Conclusion

Sharing cited values without exposing sources requires an explicit decision about both permission and evidence. The value, displayed citation, excerpt, source locator, timestamps, and derivation path are separate disclosure components. A recipient may be allowed to receive a count while lacking permission to inspect the rule that makes the count relevant. Provenance-security research supplies the reason to inspect relationships as well as source content. (www.usenix.org)

The proposed recipient disclosure test turns that distinction into observable outputs. A complete release provides the authorized basis. A limited release states exactly which claim the visible evidence supports and what remains unverifiable to the recipient. When a decision requires restricted evidence that cannot safely be shown, withholding the original assertion is a valid result. Changing the wording to a narrower public claim must be visible as a change of assertion.

The synthetic path calculation measures one precise property: whether exported dependencies still reach the restricted source. It does not measure every inference available to a recipient. That limitation follows the literature's broader warning that public information can disclose confidential provenance. (www.usenix.org)

A completed review should retain its policy owner, applied policy state, recipient context, component visibility, inspectability decision, and residual uncertainty. Repeating the test after relevant changes preserves the distinction between a prior authorized view and a newly authorized export. The useful outcome is a claim whose evidentiary status is clear to its actual recipient, including an explicit limited view or a reasoned decision to withhold it.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. https://www.w3.org/TR/prov-aq
2. https://profsandhu.com/confrnc/misconf/pst12_pbac_model_preprint.pdf
3. https://www.usenix.org/legacy/event/hotsec08/tech/full_papers/braun/braun_html/index.html
4. https://www.w3.org/TR/odrl-model
5. https://docs.oasis-open.org/xacml/3.0/xacml-3.0-core-spec-os-en.html
6. https://www.usenix.org/legacy/event/tapp09/tech/full_papers/chong/chong_html/index.html
7. https://learn.microsoft.com/en-us/purview/encryption-sensitivity-labels
8. https://www.w3.org/TR/prov-dm
9. https://arxiv.org/abs/2101.04432
10. https://arxiv.org/abs/1406.1998
11. https://docs.cloud.google.com/bigquery/docs/column-level-security-intro
12. https://docs.aws.amazon.com/lake-formation/latest/dg/data-filtering-notes.html
13. https://csrc.nist.gov/pubs/sp/800/162/upd2/final
14. https://cheatsheetseries.owasp.org/cheatsheets/Authorization_Cheat_Sheet.html
15. https://hummedia.manchester.ac.uk/institutes/cmist/archive-publications/reports/2022-01-PROVANON-report.pdf

16. https://simpl-programme.ec.europa.eu/programme/ecosystem
17. https://www.profsandhu.com/provsec_publications/2011sacmat-transform-provenance.pdf
18. https://arxiv.org/html/1406.1998v1
19. https://www.rfc-editor.org/rfc/rfc9110.html
20. https://spec.c2pa.org/specifications/specifications/2.4/specs/C2PA_Specification.html
21. https://www.postgresql.org/docs/current/ddl-rowsecurity.html
22. https://ukanon.net/the-provanon-project-2
23. https://www.nist.gov/privacy-framework/using-privacy-framework-11
24. https://www.rfc-editor.org/rfc/rfc9111.html
25. https://slsa.dev/spec/v1.2/build-provenance
26. https://github.com/in-toto/attestation/blob/main/spec/v1/statement.md

THE PUBLISHER

About Reliable Context

Reliable Context is a system of record for organizations that fuse private and public information and must defend every value they act on.

The guarantees

One canonical record per real thing: duplicates resolve by rule, ambiguity is shown rather than hidden, and related organizations relate without ever being merged. A verbatim, cited source and an observation date sit behind every value. Unknowns are declared as not found, blocked or not applicable with a reason, never left as an empty cell. Any past state of the record is reproducible, together with what changed since.

Where it runs

It runs inside the customer's perimeter, one database per organization. Permitted sources, model routing, reader clearance and retention are set as the customer's own policies, and every value is exportable with its chain of custody for an auditor or regulator.

What it is not

It is not a vector database or a RAG framework, not a data catalog, MDM or data-quality suite, not a CRM or data-enrichment vendor, and not a chatbot or AI assistant. Where a reader needs one of those, the honest answer is to say so and point them there.

Visit

See Reliable Context.

Reliable Context is in production use in two sectors; more than six thousand of its identity decisions were made by rule and none has been reversed. Articles published here explain general practice, cite their primary and academic sources, and name the jurisdiction of any regulatory rule they describe. They are not legal, regulatory or compliance advice for a specific organization, and requirements change.

Website: https://reliablecontext.com

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.