

THE EVIDENCE AVALANCHE IS A DATA-MODELING FAILURE

If the assessment becomes less understandable every time another file is uploaded, you do not have an evidence-volume problem. You have lost the relationships that give the evidence meaning.

PRACTITIONER POSITION

PCI DSS does not prescribe an evidence graph, claims ledger, database, or W/ relationship model. It does require assessors to perform and document testing procedures, and PCI SSC FAQs 1566 and 1567 preserve the assessor's responsibility to evaluate whether evidence is specific to the assessment scope, directly related to the requirement, within the applicable timeframe and testing conditions, and sufficient to support the assessor's opinion. FAQ 1569 makes another relationship explicit when sampling is used: a selected item must be understood against the complete population and its variants. My practitioner position is that the so-called evidence avalanche is usually not caused by too little storage, too few portals, or too few screenshots. It is caused by treating evidence as a pile of files rather than a network of bounded claims, subjects, populations, periods, authoritative sources, collection methods, tests, contradictions, owners, reviewers, dispositions, and report conclusions. Better data modeling does not decide compliance. It makes the assessment story traversable enough for accountable human judgment.

THE FIELD PATTERN

WHAT THIS LOOKS LIKE IN PRACTICE

A program has 14,000 assessment files arranged by requirement number. The same access export appears in nine folders under seven filenames. One version covers production; another quietly excludes a support tenant. Screenshots show successful settings but not the complete population or the assessment period. A ticket explains an exception, but no record identifies which claim the exception changes. The assessor asks a narrow follow-up and receives another archive containing 600 files. Everyone is working hard. Nobody can reliably answer which evidence supports which bounded claim, which population was tested, which version is authoritative, which contradiction remains open, or which report sentence must change when one fact changes. The avalanche is not the file count. It is the absence of declared relationships.

SIGNALS DETECTED

WHAT SHOULD MAKE YOU STOP

- Files are organized primarily by requirement folder, while the subject, environment, population, period, and bounded claim remain implicit.
- The same artifact is copied into several folders, but its permitted uses, limitations, consumers, and expiration trigger do not travel with it.

- A screenshot, export, ticket, policy, interview, and observation are all labeled evidence without preserving how each was produced or what each can and cannot establish.
- The population is reconstructed after a sample is selected, or the sample cannot be traced back to all material variants in the candidate population.
- Evidence delivery, admission, testing, evaluation, disposition, and reporting are represented as one status such as uploaded, complete, or passed.
- A changed component, provider, account population, period, or control implementation does not identify the affected evidence records and conclusions.
- Contradictory artifacts are overwritten, renamed, or explained in chat instead of remaining connected to the claims they reopen.
- Evidence owners are known socially, but collection, source authority, review, decision, and remediation responsibilities are not separately attributable.
- A dashboard can show green while the underlying source, query, filters, exclusions, observation time, and represented population remain unknown.
- The assessment team measures file completion and request closure but cannot measure unresolved claims, unrepresented variants, stale relationships, or semantic clarification loops.

WHY IT MATTERS

THE COST OF LEAVING IT UNEXPLAINED

Start with claim identity. Give each bounded assertion a stable identifier and state the requirement or testing objective it informs, the exact subject and environment, the period, the responsible party, and the conclusion that remains to be made. A control label such as MFA or quarterly review is not yet a claim. The useful unit is closer to: the defined remote-access population for the named production environment used the specified authentication control during the applicable period, subject to the recorded exceptions.

Model subjects and populations separately. A system, service, account class, software component, endpoint, provider service, process, or location may be the subject. A population is the complete candidate set relevant to a test at a stated point or period, with inclusion rules, exclusions, variants, and an authoritative source. Evidence about one sampled item is not a population record, and a folder of examples does not define completeness.

Treat an evidence object as an attributable observation, not a magical file. Record its source system, native object or query, parameters and filters, collector, collection time, represented time, transformations, storage location, integrity reference where appropriate, sensitivity handling, limitations, and supersession relationship. The file may be the payload; the evidence record preserves its meaning.

Separate collection method from assessment test. API output, CLI output, screenshot, export, interview, observation, configuration review, log inspection, reperformance, and document inspection describe how information was obtained. The assessment test describes what the assessor did with relevant evidence and why the result supports or fails to support a bounded proposition. Automation may collect or normalize facts without inheriting the assessor's judgment.

Make relationships first-class. A claim may apply to many subjects; a population may contain many variants; one evidence object may support several claims within explicit limits; one test may consume several evidence objects; one contradiction may reopen multiple dispositions; one source change may invalidate downstream conclusions. A spreadsheet can represent these joins when disciplined. A graph can make them easier to traverse. Neither representation is a compliance result.

Preserve the decision chain. Delivery records that something arrived. Admission records that it is usable for a named purpose with known limitations. Evaluation records what testing found. Disposition records the accountable interpretation. Reporting records the bounded statement that may be made. Collapsing these into complete or passed creates false certainty and makes later challenge look like rework instead of assessment.

Design change propagation deliberately. When a component, population predicate, provider responsibility, control configuration, period, authoritative source, or contradiction changes, the model should identify potentially affected claims, evidence objects, tests, dispositions, report language, and owners. It should create a review queue, not automatically change an assessment result.

Measure semantic debt, not only collection volume. Useful measures include evidence objects with no bounded claim, claims with no current evidence relationship, samples with no declared population, copied artifacts with divergent metadata, contradictions without dispositions, expired periods, and clarification loops caused by missing semantics. File count alone tells you almost nothing about assessment readiness.

Build the smallest viable relationship model before buying another portal. A claims ledger, subject inventory, population register, evidence record, test record, contradiction record, and disposition record can begin as controlled tables with stable identifiers and explicit joins. Maturity comes from governed meaning, not from calling an attachment repository a continuous-compliance platform.

Keep the authority boundary visible. The entity owns its facts, implementation records, explanations, and assertions. Tools can collect, relate, compare, and reopen. The assessor performs the required testing and records professional judgment. The compliance-accepting entity determines acceptance under its program. A beautiful graph can make weak evidence easier to see; it cannot make the evidence sufficient by itself.

QUESTIONS

QUESTIONS THAT EXPOSE THE WACK

- 01 What bounded claim does each high-value evidence object inform, and what conclusion remains outside the object?
- 02 Can every selected item be traced to a declared population, its authoritative source, inclusion rules, exclusions, and material variants?
- 03 Which timestamps describe source state, represented period, observation, collection, submission, review, and disposition?
- 04 Which object is authoritative, which is a derivative, and which transformations or filters produced the derivative?
- 05 Where are limitations and non-claims recorded so evidence cannot be stretched into a different subject, population, period, or requirement?

- 06 How are delivery, admission, evaluation, disposition, and report conclusion represented as different events?
- 07 What happens to downstream claims and conclusions when a component, provider responsibility, population, period, or source changes?
- 08 Can contradictions coexist without being overwritten, and can the affected claims be reopened deliberately?
- 09 Which relationships depend on human professional judgment, and is the accountable decision maker attributable?
- 10 Could a cold reviewer traverse one report statement back through disposition, test, evidence, population, subject, source, and owner without oral history?

THE CLUSTER OF WACK

FIND / MAP / EXPLAIN / REDUCE / PROVE

ESCAPE THE WACK

01 / FIND

Choose one evidence-heavy requirement family and inventory the claims, subjects, populations, evidence objects, tests, contradictions, dispositions, and report statements that currently exist. Mark copied files, missing identifiers, and relationships carried only in people's memory.

02 / MAP

Assign stable identifiers and create explicit joins from claim to subject, population, evidence, test, disposition, and report conclusion. Record period, source authority, variants, owners, limitations, non-claims, and reopening triggers at the relationship where they matter.

03 / EXPLAIN

Walk one conclusion backward and one changed fact forward. A cold reviewer should be able to see why each object is relevant, what it does not prove, who evaluated it, and which conclusions would reopen if it changed.

04 / REDUCE

Remove uncontrolled copies, duplicate collection, ambiguous folder semantics, binary completion statuses, orphan artifacts, and manually retyped relationships. Preserve required artifacts, original evidence, assessment procedures, history, contradictions, and accountable decisions.

05 / PROVE

Run a fictional or safely sanitized change drill. Change one population predicate or source fact, identify every affected claim and evidence relationship, route the items to human review, and demonstrate that no tool silently converts propagation into an assessment result.

THE FIELD LAW

THE LINE WORTH REMEMBERING

EVIDENCE BECOMES AN AVALANCHE WHEN ITS VOLUME GROWS FASTER THAN THE RELATIONSHIPS THAT PRESERVE ITS MEANING.

Build the Evidence Relationship Record for one requirement family. Trace one conclusion backward, change one fact, and prove that the right claims reopen without pretending the model can decide compliance.

TECHNICAL GROUNDING

START WITH THE SOURCE MATERIAL

These references ground the practitioner interpretation. They do not turn the metaphor or method into official PCI DSS terminology.

- WACK-SOURCE-0001 / PCI DSS v4.0.1 / PCI Security Standards Council / [Open source](#)

- WACK-SOURCE-0080 / FAQ 1566: Can a QSA ask an auditor from the same company to collect evidence for a PCI DSS assessment? / PCI Security Standards Council / [Open source](#)

- WACK-SOURCE-0081 / FAQ 1567: Can a QSA rely on results from a non-PCI DSS assessment? / PCI Security Standards Council / [Open source](#)

- WACK-SOURCE-0118 / FAQ 1569: Is sampling allowed in PCI DSS v4.x? / PCI Security Standards Council / [Open source](#)

- WACK-SOURCE-0124 / 2026 North America Community Meeting Agenda Highlights / PCI Security Standards Council / [Open source](#)

AUTHORITY CLASSIFICATION / PCI_SSC_SOURCE / ASSESSOR_JUDGMENT / PRACTITIONER_POSITION / PRACTITIONER_METHOD / AI_ASSISTED

Practitioner reviewed by Scott Norton on 2026-09-17 against canonical digest

cb6548a8f96948dfbc04c2c3031421b2a814828f7e1e8790621f41eaa29df2f8. Educational content, not a compliance determination. Independent practitioner resource; not affiliated with or endorsed by PCI SSC. Authorship and authority classification do not determine legal intellectual-property ownership.