DSA / RESEARCH SERIES 2026

Audit-provider evidence: what 23 source checks and five dossiers reveal

A directory-wide access check and five purposive dossiers show why marketing claims, traceability and technical judgment need separate fields.

DSA research, prepared with AI assistance

Published September 5, 2026 | Version 1.0

Public-source research; independent expert review pending

This document reviews cited public evidence. It is not a code audit, a certification, or an endorsement by an appointed expert panel.

Measurements in this report

Directory population	23
Profiles with at least one readable page of 500 characters	17
Profiles with accepted networks	11
Profiles with accepted languages	8
Profiles with accepted services	15

Read online: https://defisec.info/reports/audit-provider-evidence-study-2026

Study design

The population is the 23 providers in DSA's current directory, including its featured provider. On September 5, 2026 we requested each official homepage and up to two candidate audit or report paths. We retained HTTP outcomes, redirects, body hashes and extracted text. We then read relevant statements in context before accepting any field; a keyword in navigation, a citation or a development stack did not by itself establish audit coverage.

The five deeper dossiers were selected for different evidence channels and include the featured company. This is a purposive pilot, not a representative sample of the security industry. Follow-up pages, the Hacken API and individual report artifacts are recorded separately from the fixed directory-wide pass.

Original directory measurement

At least one requested page yielded 500 or more extracted characters for 17 of the 23 providers. That threshold measures readable source access, not whether the content is an audit report. After contextual adjudication and the pilot follow-up, 11 profiles contain accepted network values, 8 contain language values and 15 contain service or method values.

The remaining fields are explicitly unknown. These results are lower bounds for this source set. JavaScript rendering, inaccessible endpoints, confidentiality and our limited path selection can all hide legitimate evidence. Failed requests therefore do not lower a technical score.

Four traps a comparison table must avoid

First, a broad engineering stack is not proof of audit specialization. The registry records audit-specific statements or an identified report artifact. Second, a displayed price must retain its unit and source: a single-contract starting price cannot be compared with a full-protocol quote as if they bought the same work.

Third, an API record with a report link may still lack a usable scope. In the Hacken snapshot, 1,042 of 1,373 records had both repository and commit strings, but this study did not test whether those strings resolve. Fourth, a resolved issue is not always a code fix. The Trail of Bits sample includes a resolution based on intended behavior, requiring a different interpretation from a corrective patch.

What the five dossiers add

OpenZeppelin's selected Rust review demonstrates the value of explicit exclusions and issue updates. Quantstamp's two directly linked PDFs show why historical traceability does not establish current engagement readiness. Pharos's service page illustrates the boundary between a priced process description and a customer-report packet. Hacken's API makes completeness measurable, while the Trail of Bits report shows why manual reading remains necessary.

These are artifact-specific observations. They do not support a league table of the five companies. No human review board has assessed them, no exploit was reproduced in this study and no numerical technical score is published.

A selection procedure built around the gaps

Assemble reports from three distinct recent relevant projects per candidate, validate scope and authorship, inspect finding evidence and fix verification and ask two independent specialists to score the seven published criteria. Where information is absent, withhold the total and request the missing evidence. Publish the scope of any recommendation, named reviewers, conflicts and unresolved limitations.

The data download contains the measurement rules, counts and hashed source manifest. Future checks should preserve the population and selection rules, record changes and keep inaccessible evidence separate from adverse technical findings. That allows a reader to distinguish a changed provider record from a changed research method.

Sources and reproducibility

All source checks were made on September 5, 2026. Content hashes identify the retrieved artifacts. Data definitions and the source manifest are available from the report page and the link below.

1. Hacken API measured in this study

Official source: hacken.io

SHA-256: dca20863d906d096e625de4b29d0c3cddc6f10a17c0290ed0b2b313ccdd07e35

2. Trail of Bits resolution example

Official source: raw.githubusercontent.com

SHA-256: ce300c02223b89be2686a74e2b4eca67ab993dc12f744575f1927c12ba34e38b

3. OpenZeppelin scope example

Official source: www.openzeppelin.com

SHA-256: 02eedac03d6f177cc389c6a828d4a15e347a6ca3b131b34072ed1632de6d2887

4. Quantstamp historical sample

Official source: certificate.quantstamp.com

SHA-256: dc98b30953e93dfc67c9b525e3bf78593929267751a9220ce375bc8c2436b2d5

5. Pharos service evidence

Official source: pharosproduction.com

SHA-256: 15691c95df3ccd98702fb2e8f7cdf48b27e1d4bb5325c1e21d7447784bce5839

Data and methodology

https://defisec.info/downloads/research/evidence-measurements-2026-09-05.json

Selection rules, HTTP results, source hashes, and original counts are recorded in the accompanying data. The research used AI assistance. No independent technical committee review is claimed.

https://defisec.info/methodology