Pentest Audit-Evidence Checklist

What your auditor expects in a pentest report

Senior-led penetration testing. Reports that close on first read.

What your auditor expects in a pentest report

A draft lead magnet for compliance-leader cold-outreach prospects.

This document is the canonical source for the PDF at `public/downloads/cyberguards-compliance-leader-audit-evidence-checklist.pdf`. Regenerate the PDF after edits with: `python3 scripts/generate-lead-magnet-pdfs.py`

Why this exists

The most expensive compliance failure of an audit cycle is not the control that fails — it is the control that fails on a *technicality* and generates a second cycle of evidence collection two weeks into field work. Pentest evidence is one of the most common technicality failures.

This is a 12-item checklist of what every pentest report needs to contain for your auditor to close the relevant controls on first read. Each item is mapped to which frameworks care about it, and to the common reason auditors send pentest reports back when the item is missing or wrong.

Use this checklist:

- Before you commission your next pentest, as part of vendor scoping.
- After you receive a draft pentest report, before you submit it as

audit evidence.

- As an internal audit-readiness review, to confirm your prior pentests

are still acceptable evidence under your current framework version.

The checklist covers SOC 2 (TSC 2017 and 2022), ISO 27001:2013 and 27001:2022, PCI DSS v3.2.1 and v4.0, and the HIPAA Security Rule.

The 12-item checklist

1. Scope statement that names every in-scope asset

What it is: An explicit list of the systems, hostnames, URLs, IP ranges, application names, and (where relevant) user roles that were tested. Not a generic "production environment" — an enumerable list.

Why auditors want it: They need to confirm that the in-audit-scope boundary is fully covered by the pentest scope. A gap between the two is the most common reason a pentest report gets bounced back.

Frameworks: SOC 2 (CC4.1), ISO 27001 (A.5.7 / A.5.30 / A.8.29 in the 2022 version), PCI DSS (Req 11.4), HIPAA (164.308(a)(8)).

Common bounce-back: "The pentest scope says 'web application' but your audit boundary includes the customer mobile app and the admin-console subdomain. Provide pentest evidence for those." Avoid by

naming every component, including mobile, admin consoles, and back-of- house tooling, in the SOW scope section.

2. Date the testing was performed (start and end)

What it is: The calendar dates the active testing window opened and closed, plus the report finalization date.

Why auditors want it: Most frameworks specify a pentest cadence relative to the audit period (e.g., PCI DSS requires testing at least annually and after any significant change). They need the dates to confirm the pentest falls inside the acceptable window.

Frameworks: SOC 2 (CC4.1, CC8.1), ISO 27001 (A.8.29), PCI DSS (Req 11.4.1, Req 11.4.4), HIPAA (164.308(a)(8)).

Common bounce-back: "The report is undated, or only the report publication date is shown. We need the testing window dates." Avoid by explicitly stating both the testing window and the report date in the cover page.

3. Methodology summary that names the standard followed

What it is: A short section that names the testing methodology — OWASP Web Security Testing Guide, PTES, NIST SP 800-115, OSSTMM, or a hybrid referencing one of these — and lists the categories of tests performed.

Why auditors want it: They need to confirm the testing was structured and repeatable, not ad-hoc. A named methodology lets them assess coverage against the framework's expectation.

Frameworks: SOC 2 (CC4.1), ISO 27001 (A.8.29), PCI DSS (Req 11.4.1 explicitly requires "industry-accepted approaches"), HIPAA (164.308(a)(8) references industry standards).

Common bounce-back: "The report shows findings but does not describe the methodology used to find them. Provide the testing approach." Avoid by including a one-page methodology summary in every report.

4. Severity rationale per finding (not just severity labels)

What it is: For every finding, a short paragraph explaining *why* the severity is what it is — usually based on a documented model like CVSS 3.1, or a custom model that names the inputs (exploitability, blast radius, data sensitivity, prerequisites).

Why auditors want it: Severity drives remediation prioritization, which feeds the risk-treatment evidence elsewhere in the audit. Auditors need the rationale to confirm the prioritization is defensible.

Frameworks: SOC 2 (CC7.1, CC8.1), ISO 27001 (A.5.27 incident management, A.8.8 management of technical vulnerabilities), PCI DSS (Req 6.3.1), HIPAA (164.308(a)(1)(ii)(A) risk analysis).

Common bounce-back: "Why is this finding rated High and that finding rated Medium? We cannot assess your remediation prioritization without the rationale." Avoid by including a "Severity rationale" paragraph on every finding, not just a severity label.

5. Reproduction steps a non-tester can replay

What it is: Numbered steps that a developer or compliance lead can follow to confirm the finding without the original tester present.

Why auditors want it: They occasionally request reproduction of findings during field work, particularly when the remediation status is "fixed" and they want to confirm the fix is effective.

Frameworks: SOC 2 (CC8.1), ISO 27001 (A.8.8), PCI DSS (Req 6.3), HIPAA (164.308(a)(1)(ii)(B) risk management).

Common bounce-back: "We tried to reproduce this finding to confirm remediation and could not follow the report. Provide clearer reproduction." Avoid by writing each finding's reproduction as if the reader does not have access to your testing tools or environment.

6. Working proof (screenshot, captured request/response, or data sample)

What it is: An artifact captured during the test that demonstrates the finding works. Not a theoretical "this could happen" — a captured "this did happen" proof.

Why auditors want it: Distinguishes a real finding from a scanner false-positive. Frameworks that require remediation tracking (PCI DSS in particular) penalize false-positive-inflated finding lists.

Frameworks: SOC 2 (CC4.1, CC7.1), ISO 27001 (A.8.8), PCI DSS (Req 11.4.4 explicit on validated findings), HIPAA (164.308(a)(8)).

Common bounce-back: "These look like scanner output. Where is the manual validation?" Avoid by including a captured artifact per finding, and stating somewhere in the report that scanner-only findings without manual validation have been removed.

7. Control reference per finding (where applicable)

What it is: The framework control ID that the finding most directly relates to, in the title or first line of the finding (not buried in an appendix).

Why auditors want it: They are testing controls. A finding mapped to a control makes their job mechanical. A finding *not* mapped to a control forces them to do the mapping themselves, which is slow and error-prone and often results in mappings their assessment program doesn't accept.

Frameworks: SOC 2 (per Trust Services Criteria — CC, A, C, P, PI domains), ISO 27001 (Annex A controls in the relevant version), PCI DSS (per requirement), HIPAA (per Security Rule safeguard).

Common bounce-back: "Which control does this finding implicate? We cannot use this report as evidence without the mapping." Avoid by having your pentest vendor map findings to your specific framework version, in the finding title line, not in a back-of-report appendix.

8. Remediation recommendation per finding

What it is: A specific, actionable remediation. Not "patch the underlying issue" — a specific patch, configuration change, or code change with enough detail that an engineering team can scope the work.

Why auditors want it: Frameworks expect remediation to be tracked through to closure. The recommendation is the starting point for that tracking.

Frameworks: SOC 2 (CC7.1, CC8.1), ISO 27001 (A.8.8, A.5.27), PCI DSS (Req 6.3.1), HIPAA (164.308(a)(1)(ii)(B)).

Common bounce-back: "What is the remediation plan for this finding? 'Generic best practice' is not a remediation." Avoid by requiring specific, scoped remediation recommendations from your pentest vendor.

9. Remediation status updated after fixes

What it is: After your engineering team ships the remediation, the report is updated to reflect each finding's current status (Open, Remediated, Risk Accepted, Compensating Control In Place). Either the original report is updated in place, or a remediation appendix / follow-up letter is issued.

Why auditors want it: They want the report version they read to reflect the post-fix state, not the test-day state. If the report says "Open" but the control is actually fixed, the auditor sees an unmitigated finding and assumes the worst.

Frameworks: SOC 2 (CC7.1, CC8.1), ISO 27001 (A.8.8), PCI DSS (Req 6.3.1 — remediation must be tracked), HIPAA (164.308(a)(1)(ii)(B)).

Common bounce-back: "Your report shows 12 open critical findings. Are these still open?" Avoid by requiring a retest + report update from your pentest vendor before audit field work begins.

10. Retest of remediated findings

What it is: After your engineering team ships the fixes, the pentest vendor retests each remediated finding and confirms (or rejects) the fix in writing.

Why auditors want it: A self-reported fix is weaker evidence than a fix retested by the original tester. For Critical and High findings, some auditors require retest evidence specifically.

Frameworks: SOC 2 (CC8.1), ISO 27001 (A.8.8), PCI DSS (Req 11.4.4 explicit), HIPAA (164.308(a)(8)).

Common bounce-back: "How was the remediation validated? Internal claim of fix is not sufficient for this control." Avoid by selecting a pentest vendor whose engagement includes a retest within the standard scope, not as a separate paid line item.

11. Tester identification (firm + role, not necessarily name)

What it is: The firm that performed the test and the role of the person who led the engagement (e.g., "Senior penetration tester, OSCP / OSCE / CRTP / OSWE certified"). Some auditors accept generic role attribution; others want the specific tester named.

Why auditors want it: They want to confirm the test was performed by a qualified third party, not by internal staff (which would be a self-assessment and a different category of evidence) and not by a junior on their first engagement.

Frameworks: SOC 2 (CC4.1 — qualified personnel performing assessments), ISO 27001 (A.8.29), PCI DSS (Req 11.4.4 — qualified personnel), HIPAA (164.308(a)(8)).

Common bounce-back: "Who performed this test? Was the tester qualified and independent?" Avoid by including a "Testing personnel" paragraph naming the firm and the lead tester's qualifications.

12. Independence statement (third-party assertion)

What it is: A short paragraph asserting that the pentest firm is organizationally independent of the audited entity, and that the testers who performed the engagement are not involved in the operation or development of the systems tested.

Why auditors want it: Independence is a hard requirement under most frameworks for the test to count as third-party evidence rather than self-assessment.

Frameworks: SOC 2 (CC4.1), ISO 27001 (A.5.35 / A.8.29), PCI DSS (Req 11.4.5 — internal pentests by qualified resource organizationally separated from those who manage the CDE), HIPAA (164.308(a)(8) generally, plus 164.308(a)(2) for security responsibility).

Common bounce-back: "Is this an internal or external assessment? We need an independence statement." Avoid by requiring the independence language in the SOW and confirming it appears in the final report.

Where pentest reports usually fail audit

Across the engagements we have run, the items that bounce most often are:

- **Item 7 — control reference per finding.** Most pentest firms map

findings to an appendix. Most auditors want them in the title line.

- **Item 9 — remediation status updated after fixes.** The original

test-day report is submitted as-is, and the auditor reads "Open" findings that have actually been fixed.

- **Item 10 — retest.** If retest is a paid line item the client

sometimes skips it, and the report goes to audit without validation.

If you commission pentests as a one-time engagement without a retest in scope, you are buying a snapshot of test-day state. Auditors will accept this for some frameworks and not others, and they will increasingly want the post-fix state on first read.

How CyberGuards engagements address each of the 12 items

Every CyberGuards engagement includes:

- **Scope statement** — written into the SOW, every in-scope asset

enumerated by name. SOW scope and final-report scope are identical.

- **Dates** — testing window dates on the cover, plus report finalization

date.

- **Methodology** — OWASP WSTG, PTES, and NIST SP 800-115 references in

the methodology section, with the specific test categories named.

- **Severity rationale** — every finding includes a "Severity rationale"

paragraph. We use a documented model based on CVSS 3.1, modified to include business-context inputs.

- **Reproduction** — every finding includes numbered reproduction steps.
- **Working proof** — every finding includes a captured request/response,

screenshot, or data sample as Appendix evidence.

- **Control reference** — every finding's title line includes the

relevant control reference for your specific framework version, confirmed on the scoping call before the SOW is signed.

- **Remediation recommendation** — every finding includes a specific,

paste-ready remediation snippet (code, configuration, or policy).

- **Remediation status** — the report is updated in place after your

engineering team ships fixes. The version your auditor reads reflects post-fix state.

- **Retest** — included in every engagement scope. Not a separate paid

line item. The retest is sequenced to land in the final week before your audit field work begins.

- **Tester identification** — the lead senior tester's qualifications

are named in the report.

- **Independence statement** — included in every report. We are an

independent third-party firm; testers who run an engagement are not involved in the operation of the systems tested.

If this is the audit-evidence shape you want from your next pentest, we'd like to walk your audit window.

Book a 30-minute scoping call: https://cyberguards.ai/contact/

Read the full For Compliance Leaders page: https://cyberguards.ai/for-compliance-leaders/