



PUBLIC METHODOLOGY WHITE PAPER

V2P Source-Code Audit and Verification Methodology

A practical standard for turning source evidence into a scoped, publicly verifiable result.

Vibe to Production (V2P) connects two needs that are often separated: engineering teams need concrete evidence they can use to improve source code, while customers, buyers, partners, and procurement teams need a clear way to understand what was assessed and what the result means.

V2P assesses one submitted source snapshot across four pillars and 20 weighted rubric categories. It combines deterministic, non-AI static analysis with contextual review by AI specialist agents. The evidence is normalized into a governed scoring process, returned privately for remediation, and—when eligible—expressed as a scoped certificate and clickable trust seal.

4 pillars

Security, Performance, UX & Reliability, and Code Quality

20 categories


A weighted source-code rubric totaling 100 points

Scoped result

Revision, scope, date, audit version, and fingerprint

IMPORTANT SCOPE

V2P is a source-code assessment—not a guarantee of production readiness. It does not prove runtime behavior, infrastructure configuration, deployed-code identity, repository or domain ownership, legal compliance, or the absence of vulnerabilities.



From one source snapshot to a controlled public status

The V2P program follows five stages. Each stage answers a different question and keeps the evidence boundary explicit.

1. Define the source

Connect a GitHub repository and select the branch, revision, or source snapshot to assess. The result refers to that submitted state—not every past or future version of the application.

2. Analyze eligible files

Deterministic scanners look for known programmatic patterns. AI specialist agents examine relevant source across Security, Performance, UX & Reliability, and Code Quality.

3. Review private evidence

The report connects outcomes to categories and, when available, to source files, excerpts, severity, confidence, and explanatory context.

4. Remediate and re-audit

V2P can provide prompts, suggested patches, and repair guidance. The team reviews, applies, builds, tests, and approves changes through its own controls, then submits a fresh snapshot.

5. Issue an eligible result

When the score and applicable gates qualify, the team can issue a certificate and seal tied to the stated scope, revision, date, audit version, and fingerprint.

Private by default	Scoped when published
Source files, detailed findings, excerpts, remediation notes, and internal workflow stay private.	The public record can show score, status, scope, issue date, audit version, and fingerprint.

Two analysis layers. One governed result.

Deterministic static analysis

Non-AI programmatic checks identify known source patterns, including potential exposed secrets. Where relevant and available, specialist tools can inspect Python security patterns and declared dependency vulnerabilities.

AI-assisted specialist review

Relevant files are reviewed in context by specialists aligned to the four pillars. The proprietary implementation governs classification, routing, analysis, normalization, and synthesis.

Neither layer independently issues a public status. Findings are mapped into the rubric, normalized, and evaluated against deterministic scoring and eligibility rules. AI-generated recommendations support engineering review; they do not autonomously decide what is published.

THE FOUR-PILLAR, 20-CATEGORY RUBRIC

Pillar	Weight	Public rubric categories
Security	35 / 100	Secrets & API Key Management; Authentication & Session Management; Authorization & Access Control; Input Validation & Injection Prevention; Data Exposure & Privacy.
Performance	20 / 100	Database & Query Performance; Frontend Performance; Caching & Data Fetching; API & Backend Performance.
UX & Reliability	22 / 100	Error Handling & Resilience; User Experience & Accessibility; Form & Input Handling; Responsive Design & Mobile; Loading & Empty States.
Code Quality	23 / 100	Code Organization & Architecture; Type Safety & Consistency; Dependency Management; Testing Coverage; Documentation & Readability; Vibe-Code Patterns.

This paper intentionally does not disclose internal prompts, model configuration, orchestration logic, scanner rules, or scoring implementation details.

A 100-point score with separate status gates

Weighted category outcomes combine into an overall score out of 100. Pillar rollups prevent a high overall score from hiding a materially weak domain. The public score band is necessary, but it is not sufficient on its own.

Below 75	Does Not Pass	No V2P verification status is issued.
75–89.9	Certified Site	The assessed site is Verified by V2P, subject to all gates.
90–100	Verified with Distinction	The assessed site is Verified with Distinction by V2P, subject to all gates.

ELIGIBILITY GATES

Coverage floor

Less than 50% eligible-source coverage cannot receive a public verification status.

Coverage cap

Less than 80% eligible-source coverage caps the result at Certified Site.

Critical Security gate

A critical Security finding results in Does Not Pass.

Pillar minimums

Verified with Distinction requires minimum performance in every pillar, not only a 90-plus overall score.

Why the gates matter

Confidence depends on both the observed quality and the breadth of the assessment. Passing language is therefore always tied to the stated scope, revision, coverage, and date.

A fixed certificate and a dynamic trust seal

Certificate

A fixed record of an issued result for a particular source snapshot. It preserves the score, status, scope, fingerprint, audit version, and issue date associated with that assessment.

Trust Seal

A clickable public surface designed to resolve to the most recently issued result for the relevant subject. A newer eligible re-audit can replace the prior displayed result.


WHAT 'DYNAMIC' MEANS

Dynamic does not mean that V2P continuously monitors a live production system. The public status changes when a newer eligible source snapshot is audited and a new result is issued. The issue date and fingerprint help a reader understand freshness and scope.

WHAT CAN BE PUBLIC

Public verification record	Private audit workspace
Overall score and public status Declared source scope Issue date and audit version Source fingerprint	Proprietary source files Detailed findings and excerpts Confidence and remediation notes Internal review workflow

The seal communicates that a stated source snapshot met the V2P source-audit standard at the displayed revision and scope. It does not prove that the assessed source is currently deployed, that the submitter owns the repository or domain, or that the application is free of vulnerabilities.



Use the result as a scoped decision tool

Before the audit

Define the repository, revision, scope, and exclusions. Select a meaningful snapshot and make sure the eligible-source coverage can support the decision you want to make.

During review

Prioritize severe and high-confidence findings, but also look for clusters that reveal an architectural or process weakness. Treat suggested fixes as inputs to engineering review, not as verified changes. Build and test changes through the team's normal controls.

When to re-audit

Re-audit after material source changes, major dependency updates, or significant remediation. A time-stamped result becomes less representative as the code diverges from the assessed snapshot.

What V2P complements

Use V2P alongside runtime observability, penetration testing, infrastructure review, compliance programs, deployment controls, and ownership verification. Each discipline answers a different question. V2P's role is to create a repeatable source-evidence loop and a scoped way to transfer the resulting status.

The practical next step

Choose one meaningful source snapshot, run the complete 4-pillar audit, review the private evidence, and decide what to repair before issuing any public result.

Start at your V2P sign-up page or review the current public methodology before connecting a repository.

Publication note

This white paper describes the public methodology at a high level and intentionally excludes trade secrets. Certification, security, and legal language should be human-reviewed before public deployment. Refresh this paper whenever rubric weights, status thresholds, issuance rules, or seal resolver behavior change.

Created by Human Co-Pilot