

AssayMark Security Evidence Prioritisation and Remediation Framework

ASEPRF v1.3

Implementation and validation edition

Darklayer Tech Ltd | Updated 30 September 2026

Purpose and current conclusion

ASEPRF defines how supported security observations become understandable priorities, assigned actions, tracked remediation and credible verification. This revised edition consolidates v1.1 and aligns its technical explanations with source base 10.3.3f plus the v2.3 analyzer patch. Framework, product and delivery versions are separate identifiers.

The original reported implementation failures and subsequent language examples now pass the recorded local regression checks. The full suite passed 84 tests in each of two timezone configurations; the benchmark passed 25/25 and basic local security checks passed 16/16. These results support the tested behaviour, not complete threat detection, production readiness or independent assurance.

Operational update 30 September 2026. Founder-supplied terminal output and service logs evidence Nginx configuration recovery, a production restart under the dedicated assaymark user, and successful Google login, dashboard, existing analysis, escalation-report endpoint and new website-scan requests. These records support the exercised deployment paths; they are not an independent infrastructure audit or proof of complete platform security. Production post-restart health and selected live request outcomes are now evidenced by founder-supplied logs. Exact source-to-running-build checksum parity, malicious and benign upload coverage across both environments, broad authorization and tenant isolation, distributed rate limits, backup restoration and independent implementation-level assurance remain outstanding. The previously failed Celery worker has no evidenced resolution.

Framework boundary

ASEPRF is an AssayMark methodology, not an ISO, NIST, Cyber Essentials, NCSC, OWASP, MITRE or government certification. A framework mapping establishes relevance only. It does not establish compliance, external endorsement or absence of risk. This edition supplies no new legal, regulatory or certification determination.

How to read this edition

Requirement means the behaviour ASEPRF expects. Source behaviour describes the reviewed snapshot. Tested means a recorded scenario executed with an observed result. Proposed means an audit or assurance improvement whose complete implementation has not been established. These labels prevent design intent from being presented as verified implementation.

The following sections cover workflow, evidence integrity, score semantics, contextual logic, language interpretation, remediation, security boundaries, reporting, validation, deployment acceptance and governance. Historical commercial evidence and external opinions are not revalidated by this framework update.

Workflow and assessment scope

The core sequence remains Context, Priority, Action, Remediation and Verification. Evidence collection precedes the sequence; recurrence and remediation memory feed subsequent assessments. A presentation layer must preserve detector evidence and identify the engine responsible for any displayed score.

Stage	Required outcome
Context	Identify source, asset, observation, time, uncertainty and supported business relevance.
Priority	Explain the detector severity or applicable scoring path; distinguish confidence from urgency.
Action	Give an existing supported next step, suggested responsible role and target timeframe.
Remediation	Record ownership, implementation activity, status and evidence without inventing fixes.
Verification	Record the method, scope, result and evidence for the retest; retain unresolved uncertainty.

Assessment surface	Evidence and verification boundary
Logs	Supported uploaded events and detector output. A later quiet sample does not prove durable remediation.
Website	Public posture checks within authorized scope; absence-based verification requires a successful comparable repeat.
Authorized deep assessment	Explicit permission and verified target scope are required; public reachability is not authorization.
Secure configuration	Intended browser-local extraction of allowlisted non-secret settings, with backend validation. Full browser privacy behaviour is not proven by local payload tests.

Ownership and timing

Suggested roles include Business Owner, Internal IT, MSP, Web Developer, Email Administrator and Security Professional. These are assignment aids, not evidence that a qualified person acted. Critical means immediate attention; High within 7 days; Medium within 30 days; Low planned improvement; Informational monitoring without automatic remediation. These inherited operational targets are guidance, not contractual SLAs.

Active next actions use existing actionable findings. Resolved, Accepted Risk and Dismissed states are excluded from active work ordering, but their histories remain available. Severity leads action ordering; recurrence and confidence can help distinguish otherwise comparable findings without silently rewriting source severity.

Evidence identity and confidence

Preserve the original observation

Requirement: retain the source title, severity, confidence, evidence excerpt or hash, assessment identity and affected asset. Business explanations may describe possible consequences but must not turn a suspected signal into an asserted breach, loss or regulatory failure. Informational and insufficient-evidence states need not become remediation tasks.

Separate lifecycle identity from observation identity

A remediation key tracks a finding across assessments and is scoped by user, source kind, asset key and canonical finding key. Normalization can remove changing URLs, addresses or numbers to support recurrence. This lifecycle key is not proof that two raw observations are identical.

Source behaviour: contextual correlation now classifies exact repeated signal IDs separately from replayed evidence under a different ID. Observation identity uses normalized title, category, asset, source, correlation key, severity, technical score and observed time. Eligible candidates share an asset or correlation key. Duplicate/replayed items are excluded before correlation scoring.

Observed times are normalized to UTC second precision; missing time remains missing. Identical untimed records are treated conservatively as replay. A different observed time can count as recurrence. This heuristic does not establish cryptographic provenance or independent source ownership; upstream IDs and timestamps must still be trustworthy.

Case	Expected treatment
Same signal ID repeated	Exclude as exact duplicate.
Different ID with same observation identity	Exclude as replay; do not inflate score or related-evidence fingerprint.
Same semantic event at a distinct time	Allow recurrence under the current identity policy; retain provenance.
Different source label	May add source diversity; label diversity alone does not prove independent corroboration.

Zero is not missing

Source behaviour: numeric helpers preserve valid zero, including numeric strings. Missing, malformed or non-finite inputs use the documented field default; bounded fields are clamped where their caller requires it. Explicit zero confidence must remain zero rather than being replaced by 50 percent. Confidence is not a calibrated probability of a real-world breach.

Test evidence: the original duplicate/replay and zero-confidence reproductions pass. This is bounded coverage; it is not proof that every ingestion adapter or history selection is correct for every tenant and asset.

Priority and score semantics

A displayed value must name its engine, direction and scope. A high health score and a high risk score mean different things. Neither implies certification, absence of compromise or completion of remediation.

Mechanism	Meaning and rule
Source finding severity	Preserve the severity assigned by the detector. Explain confidence separately unless that engine explicitly uses it.
Assessment health	Health = $\max(0, 100 - 30C - 20H - 10M - 5L)$, with counts of Critical, High, Medium and Low findings. Info adds no penalty. This is directly exercised for log analysis.
Website health	Separate scanner aggregate. The inherited v1.1 description uses 85 to 100 Good, 65 to 84 Needs attention, below 65 High attention. The current retest did not validate every scanner contribution.
Legacy attack-story risk	Weighted technical, anomaly, evidence, context and correlation inputs, plus bounded feedback; not the formula behind every finding card.
Contextual risk and gate	A separate deterministic engine with technical, business, memory, correlation and confidence components. See the next section.

Legacy story path

The reviewed story formula is $0.30 \text{ technical} + 0.15 \text{ anomaly} + 0.15 \text{ evidence quality} + 0.15 \text{ asset criticality} + 0.10 \text{ privilege} + 0.05 \text{ internet exposure} + 0.10 \text{ correlation} + \text{analyst feedback}$ bounded to -9 to +9, with the result clamped to 0 to 100. Normalization occurs within the story path; these coefficients must not be applied to arbitrary raw evidence.

Observe and enforce modes

Source behaviour: observe retains the legacy displayed risk and computes contextual results alongside it. Enforce uses the greater of the legacy and contextual results, so contextual evidence does not silently lower the displayed risk. The contextual gate remains a separately explained output. Live mode selection is not established by these local tests.

Avoid false precision

The framework does not invent per-check website point contributions, exploitability estimates or financial loss where the underlying result does not retain them. V1.1 described the same severity-count health convention for deep and configuration assessments; this update does not independently certify every surface or caller. Display the actual mechanism used at each entry point.

Contextual scoring and security memory

Source specification: $R = \text{round}(\text{clamp}(0.64T + B + M + X + A, 0, 100), 1)$. T is the supplied technical score or severity base: Info 8, Low 22, Medium 48, High 72, Critical 92. The remaining components are defined below.

Component	Implemented contribution
Business B	Criticality: $(c - 1) \times 3.2$, c from 1 to 5. Add sensitivity, exposure and environment points, plus privilege 8, customer-facing 5, revenue impact 4 and regulated 5. Clamp total to -8 to 45.
Sensitivity	Public 0; internal 5; confidential 12; PII 16; financial 18; health 20; credentials 22. Unknown label fallback 8.
Exposure and environment	Exposure: isolated -6, internal 0, partner 5, internet 13; fallback 3. Environment: development -4, test -2, staging 2, production 8; fallback 2.
Memory M	$2.4 \times \text{min}(\text{recurrence}, 5) + 3.3 \times \text{min}(\text{unresolved}, 4) + 4 \times \text{min}(\text{confirmed}, 3) - 3 \times \text{min}(\text{false positives}, 3)$. Clamp to -10 to 30.
Correlation X	$1.7 \times \text{min}(\text{unique related items}, 6)$, plus 3.8 per extra source up to four and 2 per extra category up to three; clamp to 0 to 28. Primary source/category participate in diversity counts.
Confidence A	$(\text{confidence} - 0.75) \times 18$, minus $1.5 \times \text{min}(\text{prior false positives}, 2)$ when memory exists.

Bands and gate thresholds

Risk bands: Critical at 85 or higher; High at 70 to below 85; Medium at 45 to below 70; otherwise Low. FAIL occurs at 85 or higher; or at 76 or higher with internet exposure, PII/financial/health/credential data or privilege; or at 72 or higher with at least two unresolved records in production. Otherwise CONDITIONAL_PASS starts at 52; lower values return PASS.

PASS means below this implemented blocking threshold, not secure. No-signal input also needs explicit evidence-sufficiency treatment; an empty result must not be marketed as successful assurance. Component breakdowns are rounded for display and may not sum exactly to the final score calculated before component rounding.

Memory should be attributable to the correct user and asset. Asset-specific history selection, cross-source independence and all adapter defaults require wider verification. Record actual context and missingness: constructor, adapter and release-entry defaults are not interchangeable.

Worked decisions and language boundaries

Contextual example

A medium signal has T = 48 and confidence 0.8. With criticality 1, public data, internal exposure, test environment, no memory and no related signals: $0.64T = 30.72$, $B = -2$, $M = 0$, $X = 0$ and $A = 0.9$. The result is 29.6, Low, PASS. Changing criticality to 5 and exposure to internet gives $B = 12.8 + 13 - 2 = 23.8$ and $R = 55.4$, Medium, CONDITIONAL_PASS. A severity-only comparator remains 48; this arithmetic does not establish which ranking is best without expert labels.

Negation and outcome interpretation

Requirement: distinguish absence of detection from failed removal or containment. Absence wording must not suppress explicit malicious execution or damage, while an unrelated successful backup must not be treated as malware. The patched analyzer remains a heuristic; unsupported grammar, subject linkage and multi-event text still need review.

Input example	Recorded result
No malware detected.	Info; health 100
Malware has been not detected.	Info; health 100
Malware was not removed; containment failed.	Critical; health 70
Malware was not detected by antivirus but executed successfully.	Critical; health 70
Malware was never detected by the scanner; it encrypted the files.	Critical; health 70
No malware detected; the scheduled backup executed successfully.	Info; health 100
No malware detected; files were encrypted by the approved backup service.	Info; health 100

The v2.3 patch limits outcome matching to the immediately following clause, retains linked threat subjects or supported implicit outcomes, and excludes unrelated named subjects or explicit non-threat actors. It does not implement general language understanding. Health 100 means no penalized rule matched in these cases, not that the environment is safe.

Decision trace requirement

Record build/rule version, input IDs and hashes, actor/asset scope, observation time, actual context, history selection, excluded duplicate/replay counts, contributions, rounding, mode, displayed score and gate, reasons, and verification references. The source returns reasons, breakdown, fingerprint and timestamp; complete persistence of this expanded trace contract remains proposed.

Remediation and verification

The workflow uses one scoped remediation record with source/asset identity, title/category/severity, owner, role, due date, notes, first/last seen, current source references, verification result, timestamps and reopened count. Lifecycle events support an audit trail; append-only application events are not the same as tamper-evident storage.

State	Interpretation
Open and In Progress	An actionable finding remains active; assignment or work does not prove resolution.
Verification Pending	A fix has been submitted for checking. If it persists in a valid retest, return to Open.
Resolved	A closure state; distinguish manual closure from evidence-backed verification.
Accepted Risk	A recorded business decision with a reason; not technical remediation.
Dismissed or False Positive	A recorded assessment decision with a reason; not a verified fix.
Reopened	A resolved canonical finding reappeared; retain history, clear current closure/verification timestamps and increment recurrence.

Automatic verification

Requirement: verify only after a successful, comparable assessment of the same asset, scope and control with adequate coverage. Website, authorized deep and labelled configuration assessments can support absence-based verification where these conditions hold. A failed scan, skipped check, changed target, missing permission or incomplete evidence is inconclusive, not proof of remediation.

Tested: selected synthetic lifecycle cases cover creation, eligible website absence, persistence, reopening and cross-user lookup. These component tests do not prove all live scanners enforce comparable coverage. The caller must establish retest eligibility and success before treating absence as resolution.

Manual verification

Log absence cannot prove durable remediation. Record a human evidence note and method where automatic verification is unavailable. Manual Resolved selection alone is not independent verification. Accepted Risk and Dismissed decisions must not silently become technical remediation evidence; preserve the transition and its justification.

Progress and memory

A progress summary may count Resolved, Accepted Risk and Dismissed as completed lifecycle decisions, but must show their distinct meanings. Historical recurrence and false-positive decisions should remain attributable and reviewable. A recurring issue can increase attention without rewriting the original detector evidence.

Privacy authorization and platform security

Data minimization

Requirement: supported configuration assessments should parse locally in the browser and send only allowlisted non-secret settings for backend validation. Passwords, tokens, salts, credentials and API keys should not be needed for a posture decision. Logs may contain necessary evidence excerpts; this is not a blanket claim that no raw security evidence reaches the server.

Local payload tests reject raw secret fields and malformed or unexpected configuration values. Complete browser behaviour, all upload limits, retention/deletion, storage encryption, backup restoration and support access need separate evidence. Minimize evidence displayed to business users and keep technical detail available only to authorized users.

Session and authorization controls

Source behaviour: signed session tokens carry a unique jti. Logout records per-session revocation until token expiry; current-user checks consult that record. The tested flow rejects replay after logout while preserving another concurrent valid session. Session-version invalidation remains available for broader account revocation.

Tested: anonymous access denial, selected own/other-user analysis access, ordinary-user admin denial, CSRF rejection, token tampering and version invalidation. These are user-ownership checks, not proof of organization-level or MSP tenant isolation across every endpoint and role.

Proxy identity and abuse controls

Recovered production and staging Nginx blocks overwrite X-Forwarded-For with the connecting peer and set X-Forwarded-Proto to \$scheme. Supplied service overrides set TRUST_PROXY_HEADERS=true, restrict trusted proxy CIDRs to loopback and use --no-proxy-headers. Configuration evidence is available; live adversarial forwarded-header and session-abuse tests were not supplied.

Local tests reject the previously reproduced header-based identity change. Per-process rate buckets do not establish shared multi-worker enforcement. Selected headers and URL guards pass local checks; live TLS, proxy behavior, redirect/DNS races, error paths and comprehensive injection testing remain outside the demonstrated scope.

Trust boundaries

Review browser-to-ingress, ingress-to-app, user-to-owned-record, upload-to-parser, app-to-database, scanner-to-network, operator-to-backup/secrets, and report-to-reviewer boundaries. No framework statement authorizes testing third-party targets or exposing customer evidence.

Presentation mapping and human review

Business and technical views

Business View should show the supported issue, cautious business impact, urgency, owner, timeframe, next action and remediation state. Technical View should expose source/check, asset, evidence, severity, confidence, supported mappings, remediation and verification details. Session view preference is a usability feature rather than an account security policy.

Impact categories may include operational disruption, account compromise, data exposure, financial risk, fraud, reputation, continuity and regulatory relevance where supported. These are possible consequences, not proof that harm occurred. Weak email-domain protection, for example, can increase impersonation opportunity without proving fraud.

Control mappings

Mappings show relevance only. The inherited source mapping surface exposes Cyber Essentials control areas and existing stored ISO 27001/NIST relationships where available. No new crosswalk or current scheme compliance is validated in this update. Do not infer additional mappings solely for display or label mapped coverage as a compliance percentage.

The inherited Cyber Essentials labels are Firewalls, Secure Configuration, Security Update Management, User Access Control and Malware Protection, with Direct, Supporting or Contextual relationships. Scheme-version correctness needs its own maintained mapping evidence. An authoritative machine-readable ASEPRF control-ID catalogue is not established by this delivery; future catalogue IDs must be separately versioned and reviewed.

ML and human judgement

Optional Isolation Forest anomaly support depends on a compatible baseline. The contextual gate is deterministic, but the legacy story formula includes a 15 percent anomaly term. Do not claim that ML cannot influence any displayed score. Baseline building and absent/invalid model states should be explicit. Incremental ML benefit has not been measured here.

Human review remains necessary for uncertain context, high-impact changes, accepted risk, false-positive decisions and manual verification. High-confidence rule matches and deterministic scores still need evidence interpretation.

Report requirements

Include assessment scope and identity, prioritized risks and next actions, severity counts, cautious impact, score method, evidence, owner/timeframe, remediation, verification result and limitations. Show missing evidence, absent mappings and unavailable remediation explicitly. Use manual-verification or inconclusive states instead of fabricated advice or automatic closure.

Validation record and evidence classification

Internal validation was conducted with Codex assistance using supplied source and isolated synthetic data. It is distinct from the historical documentary reviews and from the founder-supplied server output. No reviewer signature or independent professional certification is implied.

Evidence	Executed result	Boundary
INT-01 benchmark	25 of 25 passed	Selected component scenarios; no held-out expert labels.
INT-02 security	16 of 16 passed	Selected local routes/helpers, not complete production assurance.
INT-03 full suite	84 passed in UTC and 84 in London; no skips	Same cases repeated under two timezones; not 168 distinct tests.
INT-04 language	12 of 12 extended and 8 of 8 outcome-scope checks passed	Task-authored language cases; not population accuracy.
INT-05 server analyzer	20 of 20 in each directory	Founder-pasted staging and production terminal output.

Repair history

Initial findings were B06 negated malware false positives, B17 duplicate-related-evidence score inflation, B18 loss of zero confidence, S06 logout-token replay and S16 untrusted forwarded-header identity changes. The original reproductions now pass. Follow-up testing exposed missed malicious activity and benign backup/encryption false alarms; the v2.3 patch resolves the recorded cases and adds 13 subject-scope regressions.

The formerly skipped login/logout integration test executed in the dependency-complete local environment. The full suite reports 132 deprecation warnings per run. Test portability was improved with file-based synthetic databases and timezone-aware expiry fixtures. Historical 22/25 and 14/16 results describe pre-fix snapshots, not current failures.

What these results support

The evidence supports repeatable outcomes for the exercised source paths and demonstrates a recorded fix-and-retest process. It does not measure customer breach detection, ranking superiority, calibrated confidence, long-term false-positive rates, performance under load, broad tenancy isolation or complete deployment security. Source-type labels and passed fixtures do not replace independent labels and adversarial review.

Deployment evidence and acceptance

The reviewed source base declares VERSION 10.3.3f plus the v2.3 analyzer patch. This is distinct from ASEPRF v1.3. Exact source and package manifests belong to the delivery; API labels and all deployment version identifiers are not represented as synchronized.

Reported server status

Operational update 30 September 2026. Founder-supplied terminal output and service logs evidence Nginx configuration recovery, a production restart under the dedicated assaymark user, and successful Google login, dashboard, existing analysis, escalation-report endpoint and new website-scan requests. These records support the exercised deployment paths; they are not an independent infrastructure audit or proof of complete platform security. Production post-restart health and selected live request outcomes are now evidenced by founder-supplied logs. Exact source-to-running-build checksum parity, malicious and benign upload coverage across both environments, broad authorization and tenant isolation, distributed rate limits, backup restoration and independent implementation-level assurance remain outstanding. The previously failed Celery worker has no evidenced resolution.

The last supplied listing showed logsentinel-celery.service failed. Cause, impact and recovery are unknown. The analyzer patch does not establish worker recovery. Direct browser access was previously blocked by unavailable admin-policy verification, and no bypass was attempted.

Acceptance levels

Level	Required evidence	Current position
Finding completeness	Issue, relevance, priority rationale, source evidence, next step, owner, timing, remediation and verification method. Mapping only where supported.	Framework requirement; assess each finding.
Component regression	Versioned inputs, expected outcomes, actual results, retained failures and successful retests.	Recorded suites passed.
Running deployment	Build parity, healthy services and the same benign/malicious examples through the actual web path.	Partially evidenced: production startup and selected live flows supplied; build parity and benign/malicious web uploads remain open.
Independent assurance	Agreed scope, competent reviewer, signed methods/results and residual risks.	Not completed by this update.

Next verification work

Extend the evidenced login, existing analysis, report endpoint and new website scan with benign/malicious sample uploads, staging authenticated tests and downloaded-report content checks. Establish source-to-running-build parity, investigate the worker failure and test broader authorization boundaries. Use held-out expert-labelled cases for comparative ranking; scenario pass counts are not accuracy guarantees.

Governance references and change control

Version history

Version	Date	Change
1.0	August 2026	Initial public methodology, as recorded in v1.1.
1.1	16 September 2026	Actionability and verification edition aligned to the then-described V10.3.1 implementation.
1.2	27 September 2026	Consolidated method; current source and test alignment; duplicate/replay identity, zero confidence, negation/outcome boundaries, revocation/proxy controls and evidence-level distinctions.

V1.3 supersedes v1.2 for operational evidence. The core five-stage method, scoring and conservative verification principles are unchanged. The update records production runtime and selected request results; it does not publish the framework to the website or change deployed code.

Evidence sources

S1. AssayMark_ASEPRF_v1_1_Updated_Framework_Final.pdf, 16 September 2026. Historical framework baseline; retained unchanged.

S2. AssayMark_Internal_Validation_Delivery_UPDATED_v2.3.zip, consolidated 27 September 2026. Current report/status, source, runners, regression evidence and source/package SHA256 manifests. Primary locators: app/analyzer.py; app/innovation/proprietary_engine.py; app/numeric_utils.py; app/v8/decision.py; app/session_revocation.py; app/security_utils.py; remediation and actionability modules.

S3. Fresh runs: controlled-tests.xml and london-tests.xml; benchmark-results.json; security-results.json; extended-negation-results.json; outcome-scope-results.json. These are internal executed records, not signed external review.

S4. Founder-supplied terminal output and service logs dated 30 September 2026; OPS-01 to OPS-03 in the companion evidence index. These records evidence recovery, non-root runtime and selected production flows, with the limitations stated above.

S5. Business Plan v2.3 and Evidence Pack v1.8, operational evidence update dated 30 September 2026. Companion narratives, not independent assurance.

Change control rule

Changes to interpretation, identity, defaults, weights, thresholds, verification semantics or mapping discipline require a versioned explanation and meaningful regression tests. Preserve original failed results, code hashes, expected-label provenance, retest outputs and limitations. Do not relabel historical reviewer opinions or customer statements as fresh implementation validation.

Before publication, separately verify current external scheme references and any public claims that depend on them. This edition updates internal methodology and evidence alignment; it does not independently revalidate external standards, legal rules or commercial assertions.

Methodology commitment

Evidence should support an understandable decision, a responsible action, a tracked remediation and a credible verification path. Missing evidence remains visible. A policy PASS, a closed task and an independently verified remediation are different states and must stay distinguishable.

Production operations update 30 September 2026

Operational update 30 September 2026. Founder-supplied terminal output and service logs evidence Nginx configuration recovery, a production restart under the dedicated assaymark user, and successful Google login, dashboard, existing analysis, escalation-report endpoint and new website-scan requests. These records support the exercised deployment paths; they are not an independent infrastructure audit or proof of complete platform security.

Recovery and runtime controls

Production logsentinel-v5.service runs as assaymark:assaymark with UMask=0077 and NoNewPrivileges=yes. Its Uvicorn process binds to 127.0.0.1:8000 with --no-proxy-headers. The migration script applied owner-only access to regular data files and directories; no post-change recursive permission listing was supplied. The production .env was previously evidenced as 600 root:root and was not changed by the migration. Systemd loads it through EnvironmentFile.

Recovered production and staging Nginx blocks overwrite X-Forwarded-For with the connecting peer and set X-Forwarded-Proto to \$scheme. Supplied service overrides set TRUST_PROXY_HEADERS=true, restrict trusted proxy CIDRs to loopback and use --no-proxy-headers. Configuration evidence is available; live adversarial forwarded-header and session-abuse tests were not supplied.

ReadWritePaths alone is not a write allowlist: the last configuration inspection reported ProtectSystem=no. Existing PrivateTmp, ProtectHome, PrivateDevices, RestrictSUIDSGID and LockPersonality settings were retained. Non-root execution reduces privilege but does not prove a complete filesystem sandbox, encryption at rest or isolation between every workload on this shared VPS.

Production and staging use separate service trees and loopback ports on a shared VPS. Staging was evidenced running as www-data on port 8002. This is logical separation, not dedicated-host isolation. Earlier supplied UFW and listener records showed public SSH and HTTP/HTTPS, with application, database and cache listeners on loopback.

Route	Response	Evidence
Production and www	200	HTTPS root requests after Nginx reload
Staging	401	Basic Auth plus noindex, nofollow, noarchive
Legacy LogSentinel	301	Redirect to https://assaymark.co.uk/
Outreach	401	Basic Auth also observed directly on port 8011
SecureGate	302	Redirect to /dashboard through port 8010
OvertimeLogger	200	HTTPS and local backend on port 5000

Post migration checks and evidence limits

Flow	Observed log result
Google authentication	Token exchange 200; callback 303; dashboard 200
Existing analysis	GET /analysis/18 returned 200
Escalation report	GET /analysis/18/escalation-report returned 200
New website scan	POST /website-scan returned 303; GET /website-scan/36 returned 200

The new production process was PID 918658. The journal records clean shutdown at 22:14:02 and startup completion at 22:14:10 on 30 September 2026, in the server-displayed time. No permission-denied message, traceback or application 500 appears in the supplied excerpt. Endpoint success does not verify the downloaded PDF layout or every module. Scanner-target 404 responses are not application permission failures.

Production post-restart health and selected live request outcomes are now evidenced by founder-supplied logs. Exact source-to-running-build checksum parity, malicious and benign upload coverage across both environments, broad authorization and tenant isolation, distributed rate limits, backup restoration and independent implementation-level assurance remain outstanding. The previously failed Celery worker has no evidenced resolution.

OPS-01 covers Nginx recovery and local HTTPS route checks. OPS-02 covers runtime permissions and the non-root migration result. OPS-03 covers post-migration request logs. Evidence is founder-supplied and reviewed in this task; no independent operator access or new source test run is asserted. Evidence-file hashes are retained in the accompanying index; backup paths are recorded in the deployment notes.