



AssayMark Security Evidence, Prioritisation and Remediation Framework

ASEPRF v1.1

Actionability & Verification Edition

Core workflow

CONTEXT → PRIORITY → ACTION → REMEDIATION →
VERIFICATION

Framework update date: 16 September 2026

Product alignment: AssayMark V10.3.1

Status and disclaimer. ASEPRF is an original AssayMark assessment methodology. It is not an ISO, NIST, Cyber Essentials, NCSC, OWASP, MITRE or government certification, and no external endorsement is claimed. It is intended to make supported security findings more evidence-based, explainable and actionable. Framework mappings are shown only where an existing validated relationship is available; mapping does not by itself establish certification or compliance. The framework does not guarantee that a system is secure and does not replace appropriate professional, legal, regulatory, penetration-testing or incident-response judgement.

Contents

1	Executive Summary	4
2	Purpose, Scope and Design Principles	5
2.1	Purpose	5
2.2	Systems and assessment surfaces in scope	5
2.3	Out of scope	6
3	Core Methodology	7
3.1	The five-stage ASEPRF workflow	7
3.1.1	Context	7
3.1.2	Priority	7
3.1.3	Action	7
3.1.4	Remediation	7
3.1.5	Verification	7
4	Evidence Model and Finding Integrity	8
4.1	Evidence preservation	8
4.2	Canonical finding identity	8
4.3	Actionable vs non-actionable observations	8
5	Privacy and Trust Boundaries	9
5.1	Zero-secret configuration handling	9
5.2	Authorisation boundary	9
6	Business View and Technical View	10
6.1	Why two views exist	10
6.2	Business View	10
6.3	Technical View	10
6.4	Session preference	11
7	Business Impact Context	12
7.1	Impact categories	12
7.2	Example: weak email-domain protection	12
7.3	Example: authentication finding	12
8	Priority and Score Transparency	13
8.1	Priority first, score second	13
8.2	Finding priority preservation	13
8.3	Log, Deep Assessment and Secure Configuration health score	13
8.4	Website assessment health	13
8.5	V8 attack-story legacy risk formula	14
8.6	Deterministic contextual production-gate engine	14
8.7	Observe and enforce discipline	15
9	Security Memory and Longitudinal Context	16
9.1	Purpose	16
9.2	Why recurrence matters	16
9.3	False-positive history	16

10 Owner, Timeframe and Next Best Action	17
10.1 Recommended timeframe	17
10.2 Suggested responsible roles	17
10.3 Your next three actions	17
11 Remediation Model	18
11.1 Remediation record	18
11.2 Append-only remediation events	18
11.3 Status model	18
11.4 Resolved vs verified	19
12 Verification Model	20
12.1 Automatic re-test eligibility	20
12.2 Verification pending behaviour	20
12.3 Manual verification	20
12.4 Recurrence after resolution	20
13 Framework and Control Mapping Discipline	21
13.1 General rule	21
13.2 Cyber Essentials alignment	21
13.3 Other mappings	21
13.4 ASEPRF control catalogue status	21
14 Advisory Machine Learning and Human Judgement	23
14.1 Isolation Forest role	23
14.2 Why ML is bounded	23
14.3 Human validation	23
15 Assessment Summary Requirements	24
15.1 Overall posture	24
15.2 Remediation progress semantics	24
15.3 Framework coverage semantics	24
16 Edge States and Honest Failure Modes	25
17 Roles, Permissions and Security	26
17.1 User scoping	26
17.2 Existing platform protections	26
17.3 Sensitive evidence	26
18 Reporting Standard	27
19 Worked Example – Email Domain Protection	28
19.1 Business View example	28
19.2 Technical View example	28
20 Implementation Alignment – AssayMark V10.3.1	29
20.1 What V10.3 already implements	29
20.2 Known limitations	29
21 Acceptance Criteria	31
22 Governance and Versioning	32
22.1 Versioning rule	32
22.2 Change log	32

A	Reference Tables	33
A.1	Priority and timeframe	33
A.2	Current external control areas exposed for raw findings	33
B	Methodology Summary	34

Chapter 1

Executive Summary

The AssayMark Security Evidence, Prioritisation and Remediation Framework (**ASEPRF**) defines how AssayMark should move from a security observation to an auditable remediation decision. The framework is designed for small and medium-sized organisations that may have real digital risk but do not operate a dedicated security team.

The framework deliberately rejects a weak product pattern:

SCAN → SCORE → REPORT

A long list of findings or an unexplained score does not answer the operational questions an SME actually has. ASEPRF instead requires supported findings to be presented through the following sequence:

CONTEXT → PRIORITY → ACTION → REMEDIATION → VERIFICATION

ASEPRF outcome test. A useful finding should allow the user to answer: What is wrong? Why does it matter? Why did it receive this priority? What evidence supports it? What should be done? Who should do it? When should it be done? How is it fixed? How is the fix verified? Which relevant control area does it relate to, where a validated mapping exists?

Version 1.1 updates the framework to align with AssayMark V10.3.1. The major change is that remediation and verification are now first-class parts of the methodology rather than being treated as report afterthoughts. V1.1 also formalises Business View and Technical View, priority transparency, owner/timeframe guidance, next-best-action ordering, lifecycle edge states and verification boundaries.

The update does **not** replace AssayMark's established detection or scoring engines. Where a score is displayed, the real mechanism used by the relevant assessment surface must be explained. The presentation layer must not fabricate factors such as exploitability, financial impact or control failure when the underlying detector did not calculate them.

Purpose, Scope and Design Principles

2.1 Purpose

ASEPRF is intended to provide a repeatable method for converting supported cybersecurity evidence into a decision and remediation workflow that can be understood by both business and technical users.

Its objectives are to:

- preserve the original evidence and severity assigned by the relevant detector or assessment workflow;
- add cautious business context without pretending that a finding proves a breach, loss or regulatory failure;
- make priority reasoning visible rather than hiding it behind a generic score;
- make ownership and target timeframe explicit;
- connect each actionable finding to remediation guidance and a verification path;
- keep historical recurrence and remediation state available through Security Memory and lifecycle records;
- distinguish automatic verification from manual verification;
- use framework/control mappings conservatively and only where a legitimate mapping exists;
- support non-technical business decision makers and technical implementers without duplicating the same paragraph under different headings.

2.2 Systems and assessment surfaces in scope

The current AssayMark implementation applies ASEPRF concepts across four main assessment surfaces:

Surface	Primary evidence	ASEPRF role
Log Analysis	Uploaded supported log evidence and detector output	Finding interpretation, contextual priority, Security Memory, remediation and manual verification discipline.
Website Security Assessment	Public website/security posture checks	Exposed controls, web hardening, actionable remediation and repeat-scan verification where applicable.
Authorised Deep Assessment	Explicitly authorised and ownership-verified target checks	Higher-context technical findings with bounded reassessment of the same verified target.
Secure Configuration Assessment	Browser-local extraction of allow-listed non-secret configuration settings	Privacy-aware configuration evidence, remediation and repeat labelled assessment verification.

2.3 Out of scope

ASEPRF is not by itself:

- a penetration-test methodology for unrestricted offensive testing;
- a guarantee that a system is secure;
- an automated legal or regulatory opinion;
- a Cyber Essentials certification mechanism;
- a substitute for full SOC, SIEM, EDR/XDR, incident-response or specialist security services;
- permission to test systems without authorisation;
- a justification for inventing framework mappings or unsupported risk factors.

Core Methodology

3.1 The five-stage ASEPRF workflow



Figure 3.1: ASEPRF actionability sequence.

3.1.1 Context

Context answers what the finding means in the environment where it was observed. Depending on the assessment surface, context can include the affected asset, category, business impact language, asset criticality, privilege, internet exposure, historical recurrence, related evidence and existing controls.

3.1.2 Priority

Priority answers how urgently the issue should be handled and why. Priority must be traceable to the existing detector severity or to the documented contextual engine when that engine is the relevant source. ASEPRF does not permit a presentation layer to silently change severity.

3.1.3 Action

Action answers what the organisation should do next. An actionable finding should identify a concise next step, a sensible responsible role and an appropriate target timeframe.

3.1.4 Remediation

Remediation contains the implementation guidance required to reduce or address the issue. Where the source detector already supplies remediation text, the framework reuses it rather than inventing new technical steps.

3.1.5 Verification

Verification records whether the remediation was actually checked. Automatic re-test is used only where the underlying assessment can legitimately re-check the same asset/control. Where that is not possible, manual verification evidence is required.

Evidence Model and Finding Integrity

4.1 Evidence preservation

ASEPRF separates **finding generation** from **finding presentation**. The V10.3 actionability layer is designed to preserve the title, severity, confidence and evidence generated by existing detectors.

This distinction matters because it prevents a usability layer from quietly becoming a second risk engine. The Business View may translate the likely business significance into cautious language, but it must not rewrite the underlying technical evidence.

4.2 Canonical finding identity

Where the remediation workflow tracks a finding across assessments, a stable canonical key is derived from the finding title/category with selected normalisation. Dynamic URLs, IP addresses and numbers are normalised so that recurrence can be tracked without requiring byte-for-byte identical text.

A remediation record is scoped by user, source kind, asset key and canonical finding key. This supports longitudinal tracking while avoiding a global cross-customer finding identity.

4.3 Actionable vs non-actionable observations

Informational/pass observations are retained where useful for history but should not be automatically promoted into remediation work. An actionable record is intended for findings with a meaningful Low, Medium, High or Critical issue state.

No false precision. ASEPRF does not require every observation to become a task. A healthy/pass observation, insufficient evidence state or informational note can remain non-actionable.

Privacy and Trust Boundaries

5.1 Zero-secret configuration handling

For supported Secure Configuration workflows, the intended trust boundary is browser-local parsing. The raw configuration file is opened and processed in the user's browser. Only allow-listed, non-secret derived settings are transmitted for server-side evaluation.

Supported configuration families currently include selected WordPress `wp-config.php`, dotenv-style `.env`, Nginx, Apache/`.htaccess`/httpd and `php.ini` settings.

ASEPRF privacy principle. Passwords, API keys, database credentials, tokens, salts and other raw secret values should not be required for a configuration posture decision when a non-secret control state can answer the assessment question.

The backend re-validates the derived allow-listed payload. The security claim is therefore data minimisation by design for the configuration workflow, not a blanket claim that AssayMark never handles any raw security evidence. Log findings can contain selected evidence excerpts where required for analysis.

5.2 Authorisation boundary

Authorised Deep Assessment is intended for targets where the user has explicit permission and where target ownership/authority has been verified by the workflow. ASEPRF does not treat public reachability as authorisation.

Business View and Technical View

6.1 Why two views exist

A business owner and the person implementing the fix have different questions. ASEPRF therefore requires two contextual presentations where the product surface supports them.

6.2 Business View

Business View should emphasise:

- finding title;
- plain-English explanation;
- business impact;
- why the organisation should care;
- likely consequence if unresolved;
- priority;
- recommended timeframe;
- suggested owner;
- remediation status;
- a short “What should I do next?” section.

Business language should use cautious terms such as “may increase the likelihood of”, “could contribute to” or “may expose the organisation to”. It should not state that a breach, fine, fraud or financial loss has definitely occurred unless the evidence actually supports that conclusion.

6.3 Technical View

Technical View should emphasise:

- technical finding title and detail;
- evidence retained by the source detector;
- source/check used;
- affected asset/domain/system;
- severity;
- confidence where actually available;
- existing validated framework/control mappings;
- remediation steps supplied by the source workflow;

- verification steps;
- references where already available;
- remediation status, assignment and due date.

6.4 Session preference

The V10.3 user interface stores the Business/Technical view choice in browser session storage where practical. The preference therefore follows the user during the active browser session without being treated as permanent account policy.

Business Impact Context

7.1 Impact categories

ASEPRF may present relevant impact categories to make a technical issue easier to understand. Current categories include:

- Financial risk;
- Operational disruption;
- Account compromise;
- Data exposure;
- Fraud / impersonation;
- Regulatory / compliance relevance;
- Reputation risk;
- Business continuity.

These categories are contextual labels, not guaranteed outcomes. They should only be displayed where the finding content supports the relationship.

7.2 Example: weak email-domain protection

A weak DMARC/SPF/DKIM-related finding may be described as increasing the opportunity for impersonation or phishing. The business view may therefore surface Fraud / impersonation and Reputation risk. It should not state that fraud has occurred.

7.3 Example: authentication finding

A failed-login, privilege or credential-related finding may be described as increasing the likelihood or impact of unauthorised account access. If the affected identity is privileged, business impact may be higher, but the finding still does not prove account takeover.

Priority and Score Transparency

8.1 Priority first, score second

ASEPRF treats the question “What should I do next?” as more important than a large numerical score. Where a score exists, it remains available for transparency and trend/context, but it should not dominate the assessment summary.

8.2 Finding priority preservation

For standard finding cards, V10.3 preserves the severity assigned by the source detector/check. The “Why this priority?” panel should explain that source severity and confidence separately.

Confidence is not silently converted into a new priority unless the underlying scoring engine explicitly uses it.

8.3 Log, Deep Assessment and Secure Configuration health score

For these assessment surfaces, the existing assessment health score is separate from finding priority:

$$H = \max(0, 100 - 30C - 20H - 10M - 5L)$$

where:

- C = number of Critical findings;
- H = number of High findings;
- M = number of Medium findings;
- L = number of Low findings.

Informational findings contribute zero penalty. Confidence is displayed separately and is not part of this health-score formula.

Interpretation rule. The health score summarises assessment burden. It does not replace finding severity, business context or the remediation lifecycle.

8.4 Website assessment health

The website scanner retains its existing aggregate scoring behaviour. Current rating thresholds are:

- 85–100: Good;
- 65–84: Needs attention;
- below 65: High attention.

The current stored website result does not retain a trustworthy per-check numeric contribution for every finding. ASEPRF therefore does not fabricate a per-check point breakdown.

8.5 V8 attack-story legacy risk formula

The V8 decision response documents an existing story-level formula consisting of:

- technical severity: 30%;
- anomaly: 15%;
- evidence quality: 15%;
- asset criticality: 15%;
- privilege: 10%;
- internet exposure: 5%;
- cross-source correlation: 10%;
- plus a bounded analyst-feedback adjustment.

This formula belongs to the V8 intelligence/attack-story layer. It must not be implied to be the scoring formula for every AssayMark finding.

8.6 Deterministic contextual production-gate engine

The proprietary contextual engine is a separate deterministic decision layer. A signal begins with a technical score (or the mapped technical severity base) and applies business-context, Security Memory, correlation and bounded confidence adjustments:

$$R = \text{clamp}_0^{100} (0.64T + B + M + X + A_c)$$

where:

- T = technical score;
- B = business-context points;
- M = Security Memory points;
- X = correlation points;
- A_c = bounded confidence adjustment.

The engine uses a deterministic risk band:

- Critical: score ≥ 85 ;
- High: score ≥ 70 and < 85 ;
- Medium: score ≥ 45 and < 70 ;
- Low: score < 45 .

The production gate is intentionally stricter than a simple band label. Current gate logic reserves FAIL for clear high-risk conditions, including very high contextual risk and selected public/sensitive/production

recurrence conditions. CONDITIONAL_PASS begins at the documented conditional threshold; otherwise the deterministic result can remain PASS.

The V8 legacy formula and proprietary contextual engine are related R&D/decision mechanisms but are not interchangeable. Public explanations should identify which mechanism generated the displayed result.

8.7 Observe and enforce discipline

Where contextual/proprietary scoring operates in observe mode, it is benchmarked alongside the existing V8 score and does not change production risk. In enforce mode, contextual evidence may raise the final V8 score; it must not silently lower the existing V8 result.

Chapter 9

Security Memory and Longitudinal Context

9.1 Purpose

Security Memory allows a finding to be understood as part of a history rather than as a one-off snapshot. Supported memory can include recurrence count, unresolved history, prior confirmed threat information, prior false-positive information, trend/change features and lifecycle state.

9.2 Why recurrence matters

A finding that repeatedly returns after attempted remediation may deserve greater operational attention than a one-off finding of the same technical severity. ASEPRF therefore permits recurrence to break ties when ordering next actions, while preserving the underlying detector severity.

9.3 False-positive history

Historical false-positive evidence can reduce confidence in the contextual engine. It must not delete or rewrite the original evidence record.

Chapter 10

Owner, Timeframe and Next Best Action

10.1 Recommended timeframe

V10.3 maps existing priority to an operational target:

Priority	Recommended timeframe
Critical	Immediate
High	Within 7 days
Medium	Within 30 days
Low	Monitor / planned improvement
Informational	Monitor / no immediate remediation

The timeframe is guidance, not a legal SLA. An organisation may legitimately tighten or extend a target based on business risk, maintenance windows, compensating controls and specialist review.

10.2 Suggested responsible roles

The current product supports the following responsibility labels:

- Business Owner;
- Internal IT;
- MSP;
- Web Developer;
- Email Administrator;
- Security Professional.

A suggested role is derived from the finding type. It is a workflow aid and can be overridden by the organisation.

10.3 Your next three actions

At assessment level, active actionable findings are ordered using existing severity first. Recurrence and evidence confidence are tie-breakers. Resolved, Accepted Risk and Dismissed findings are excluded from the active next-action list.

No arbitrary action generation. ASEPRF does not permit the next-action list to create a new remediation that is unsupported by the underlying finding. It orders existing actionable findings.

Remediation Model

11.1 Remediation record

The V10.3 implementation extends the existing V10.1 remediation lifecycle rather than creating a parallel task system. A remediation record can retain:

- user scope;
- source kind and asset key;
- canonical finding key;
- title/category/severity;
- status;
- assignee;
- responsible role;
- due date;
- remediation note;
- verification note;
- Accepted Risk reason;
- Dismissed / False Positive reason;
- first seen and last seen;
- current/last source identifiers;
- verification result and evidence source;
- resolved/verified/accepted timestamps;
- reopened count.

11.2 Append-only remediation events

Lifecycle changes are also recorded as append-only remediation events. This creates a basic audit trail for status changes, verification events and recurrence.

11.3 Status model

The current user-facing status set is:

- Open;
- In Progress;

- Verification Pending;
- Resolved;
- Accepted Risk;
- Dismissed / False Positive.




Figure 11.1: Simplified remediation lifecycle.

11.4 Resolved vs verified

Choosing Resolved manually records implementation/closure but does not automatically mean independently verified. A separate manual-verification action requires an evidence note and records a verified timestamp/result.

Chapter 12

Verification Model

12.1 Automatic re-test eligibility

Automatic absence-based verification is only appropriate for assessment surfaces where AssayMark can repeat the same bounded check against the same asset:

Source	Automatic verification	Rule
Website	Yes, where repeat scan is comparable	Re-run the same website assessment and confirm the canonical finding is absent.
Authorised Deep	Yes, within the same verified scope	Repeat the authorised assessment against the same verified target.
Secure Config	Yes, where the same labelled asset/profile is reassessed	Re-run the configuration assessment and confirm the canonical finding is absent from derived non-secret settings.
Log Analysis	No automatic durable closure	A later log sample can support the conclusion, but absence from one sample is not proof that the control was fixed. Manual verification remains required.

12.2 Verification pending behaviour

If a Verification Pending finding is still detected on repeat assessment, the record returns to Open and the verification result records that it is still detected.

If an eligible repeat assessment no longer observes the canonical finding, the record can move to Resolved with a verified timestamp and repeat-assessment evidence reference.

12.3 Manual verification

Where automatic verification is unavailable, the user must record an evidence note. A non-empty verification note distinguishes evidence-backed manual verification from simply selecting a closed status.

Accepted Risk and Dismissed findings must be reopened before they can be verified as resolved. This prevents a business decision or false-positive decision from being silently converted into technical remediation evidence.

12.4 Recurrence after resolution

If a resolved canonical finding reappears in a later assessment, the remediation record is reopened, resolved/verified timestamps are cleared for the new active state, and the reopened count is incremented.

Chapter 13

Framework and Control Mapping Discipline

13.1 General rule

ASEPRF separates a finding from any external control mapping. A finding can be technically valid even when no external framework mapping exists.

Mapping is not certification. A mapping means the finding is relevant to a control area. It does not prove that the organisation passes or fails the external standard.

13.2 Cyber Essentials alignment

The current raw-finding mapper uses Cyber Essentials v3.3 control areas conservatively:

- Firewalls;
- Secure Configuration;
- Security Update Management;
- User Access Control;
- Malware Protection.

Relationship labels are:

- **Direct:** the finding closely corresponds to the control theme;
- **Supporting:** the finding provides useful evidence but is not a complete scheme decision;
- **Contextual:** the signal may inform the control area but requires additional validation.

Current product wording should use “Cyber Essentials Alignment” or “Relevant Cyber Essentials Control”, not “Cyber Essentials certified”, “compliant” or “failed Cyber Essentials” unless a qualified certification process independently supports such a statement.

13.3 Other mappings

Existing attack-story records can expose Cyber Essentials, ISO 27001:2022 and NIST CSF 2.0 mappings where those mappings are already stored by the intelligence workflow. ASEPRF v1.1 does not infer new ISO/NIST relationships for raw findings solely for presentation purposes.

13.4 ASEPRF control catalogue status

The current AssayMark V10.3 source package does not contain an authoritative machine-readable ASEPRF control-ID catalogue. This document therefore defines the methodology and lifecycle but does not pretend that every product finding is already mapped to a numbered ASEPRF control identifier.

A future machine-readable catalogue can be introduced as a separately versioned artefact after independent review and validation.

Advisory Machine Learning and Human Judgement

14.1 Isolation Forest role

AssayMark uses Isolation Forest as an advisory anomaly signal when a compatible historical baseline exists. The model does not autonomously control the production gate.

A baseline can be unavailable for early pilot activity. This is an expected state, not an error. The deterministic engine remains authoritative.

14.2 Why ML is bounded

ASEPRF requires the user to be able to explain a priority. “The model said so” is not sufficient. ML may support prioritisation where the baseline is compatible, but deterministic evidence and explainable rules remain the basis of operational decisions.

14.3 Human validation

Human review remains appropriate for:

- context that cannot be safely inferred;
- unusual business impact;
- accepted-risk decisions;
- false-positive/dismissal decisions;
- manual verification;
- high-impact remediation that could affect production availability.

Assessment Summary Requirements

15.1 Overall posture

The assessment summary should answer “What should I do next?” before presenting a large score. A useful summary includes:

- counts of Critical, High, Medium and Low findings;
- Top Risks;
- Your Next 3 Actions;
- Business Impact Areas;
- Remediation Progress;
- Framework Mapping Coverage;
- numerical/health score as a secondary element where applicable.

15.2 Remediation progress semantics

The current progress calculation treats Resolved, Accepted Risk and Dismissed as lifecycle-completed decisions. This does not mean all three states are equivalent security outcomes. “Completed” means the finding has reached a recorded decision state rather than remaining active.

15.3 Framework coverage semantics

Framework mapping coverage counts actionable findings with at least one validated mapping. It is not a compliance percentage.

Chapter 16

Edge States and Honest Failure Modes

ASEPRF requires explicit handling of cases where the ideal workflow is unavailable.

State	Required treatment
No remediation available	State that no remediation text is currently recorded; do not generate arbitrary technical steps.
Manual remediation required	Make the manual nature visible and allow notes/assignment.
No framework mapping	Show no validated mapping rather than inventing one.
Insufficient evidence	Preserve uncertainty and avoid strong business-impact claims.
Scan cannot verify fix	Use Manual verification required.
Accepted Risk	Record the reason as a business decision; do not treat it as technically resolved.
False positive / Dismissed	Require a reason; do not count it as verified remediation.
Resolved but not verified	Distinguish manual closure from independent verification.
Reopened finding	Reopen the lifecycle and retain recurrence count/history.
Baseline building	Keep ML advisory/inactive until a compatible baseline exists.

Roles, Permissions and Security

17.1 User scoping

Remediation records are user-scoped in the current implementation. Retrieval helpers filter by the authenticated user so one user cannot simply request another user's remediation record by ID.

17.2 Existing platform protections

ASEPRF implementation must not weaken the surrounding application security. The current platform design includes authentication/authorisation, CSRF protection, rate limiting, secure cookies, trusted-host validation, CSP and other defensive HTTP headers, upload validation and audit logging.

17.3 Sensitive evidence

Business-facing summaries should not expose unnecessary raw evidence. Technical evidence is only useful when it is available to an authorised user and required to understand or implement the fix.

Reporting Standard

A report aligned with ASEPRF v1.1 should contain, where the assessment supports the fields:

1. assessment identity, scope and source;
2. executive posture summary;
3. Critical/High/Medium/Low counts;
4. Top Risks;
5. Next 3 Actions;
6. business-impact areas;
7. remediation progress;
8. each actionable finding in Business View terms;
9. technical evidence/details for implementation;
10. Why this priority? explanation;
11. owner and recommended timeframe;
12. remediation steps;
13. verification instructions/result;
14. framework/control relationships where validated;
15. evidence limitations and assessment limitations.

A report should avoid implying that the absence of detected findings proves the absence of risk.

Chapter 19

Worked Example – Email Domain Protection

This example illustrates ASEPRF structure. It is a methodology example, not evidence that a specific customer has this issue.

19.1 Business View example

Issue: DMARC policy not enforced / monitoring-only policy.

Business impact: Weak email-domain enforcement may make impersonation or phishing attempts easier and could contribute to fraud, credential theft or loss of trust.

Priority: High, if that is the severity assigned by the relevant email/domain check.

Suggested owner: Email Administrator or Internal IT.

Target: Within 7 days.

Next step: Review SPF/DKIM alignment and prepare a controlled path from monitoring to stronger enforcement where appropriate.

19.2 Technical View example

A technical view can show the DNS evidence returned by the assessment, such as an observed DMARC record, and the exact remediation guidance associated with the check. A verification step should re-run the same domain/email assessment and confirm the policy state and alignment after the change.

The framework does not require AssayMark to claim that a specific policy value is always correct for every organisation. Deployment should consider legitimate senders, alignment reports and phased enforcement.

Implementation Alignment – AssayMark V10.3.1

20.1 What V10.3 already implements

The current product implementation aligned to this framework includes:

- Business View and Technical View per main finding surface;
- session-persistent finding-view toggle;
- cautious business-impact categories and language;
- priority-to-timeframe guidance;
- suggested owner roles;
- “Why this priority?” panels;
- explicit source-specific score-method explanations;
- next-three-action ordering;
- first-class remediation status, assignee, responsible role and due date;
- manual verification notes;
- automatic repeat-assessment verification where legitimate;
- Dismissed / False Positive and Accepted Risk states;
- recurrence reopening;
- Cyber Essentials v3.3 raw-finding alignment;
- existing attack-story framework mappings exposed without new inference;
- action-first assessment summary components;
- public V10.3.1 SEO positioning aligned to contextual remediation and verification.

20.2 Known limitations

The current framework/product boundary includes:

- website results do not retain a reliable per-check numeric contribution for every check, so exact per-finding point breakdown is not shown;
- log absence cannot automatically prove durable remediation;
- ML requires a compatible historical baseline and will often be inactive in early pilots;

- raw finding-level external mapping is currently conservative Cyber Essentials alignment; broader raw-finding ISO/NIST mapping is not automatically inferred;
- no authoritative machine-readable ASEPRF control-ID catalogue is currently integrated into the product;
- automated verification is not possible for every security control or business process;
- AssayMark is not a substitute for full penetration testing, SOC operations or specialist incident response.

Acceptance Criteria

A finding is ASEPRF-complete when the available evidence allows the user to answer the following without needing to reverse-engineer the product:

1. What is wrong?
2. Why does it matter?
3. Why was it given this priority?
4. What evidence supports the finding?
5. What should be done?
6. Who should do it?
7. When should it be done?
8. How do they fix it?
9. How do they know the fix worked?
10. Which relevant security control does it relate to, where a validated mapping exists?

Framework success condition. The product should help a business owner understand the significance and next decision, while also giving the IT/MSP/security implementer enough technical evidence and remediation context to act and verify.

Governance and Versioning

22.1 Versioning rule

Framework changes that alter terminology, scoring interpretation, remediation semantics, verification rules or external mapping discipline should increment the ASEPRF version and include a change log.

Product version and framework version are related but separate. AssayMark V10.3.1 is aligned to ASEPRF v1.1; this does not imply that every future product patch requires a framework-version increment.

22.2 Change log

Framework version	Date	Change
1.0	August 2026	Initial ASEPRF public methodology released with the AssayMark framework page.
1.1	16 September 2026	Actionability and verification update aligned to AssayMark V10.3.1: formal Business/Technical views, priority transparency, owner/timeframe, next actions, remediation lifecycle, manual/automatic verification boundaries, edge states, mapping discipline and explicit score-method separation.

Chapter A

Reference Tables

A.1 Priority and timeframe

Priority	Operational target	Typical workflow intent
Critical	Immediate	Escalate and contain/remediate as a priority; specialist review may be appropriate.
High	Within 7 days	Assign quickly and plan a verified fix.
Medium	Within 30 days	Schedule remediation within normal change planning.
Low	Monitor / planned improvement	Address as defence-in-depth or planned hardening.
Informational	Monitor / no immediate remediation	Retain useful evidence/history without forcing a remediation task.

A.2 Current external control areas exposed for raw findings

Control area	ASEPRF interpretation
Cyber Essentials – Firewalls	Relevant to supported network exposure/filtering findings.
Cyber Essentials – Secure Configuration	Relevant to supported hardening/default/configuration findings.
Cyber Essentials – Security Update Management	Relevant to supported software support/update/vulnerability indicators.
Cyber Essentials – User Access Control	Relevant to supported authentication/account/privilege findings.
Cyber Essentials – Malware Protection	Relevant to supported malware/detection/execution indicators.

Methodology Summary

AssayMark ASEPRF v1.1

Evidence should lead to a decision, a responsible action, a tracked remediation and a credible verification path.

CONTEXT → PRIORITY → ACTION → REMEDIATION → VERIFICATION

ASEPRF is intentionally conservative. It prefers a transparent limitation over a fabricated score component, a manual verification state over a false automatic closure, and “no validated mapping” over an impressive-looking but unsupported standards crosswalk.

That restraint is part of the framework: the purpose is not to make every assessment look certain. The purpose is to make the security decision understandable, actionable and auditable.