

PAPER TIGER TECHNOLOGIES

Vulnerability Management Policy

1. Purpose

The purpose of this Vulnerability Management Policy is to define how Paper Tiger finds, prioritizes, remediates, and verifies weaknesses in the illustrated environment. A defined cadence keeps findings from accumulating and ensures remediation is confirmed rather than assumed.

Every environment has weaknesses; the difference between a well-run one and a neglected one is whether those weaknesses are found on a schedule, ranked by risk, and closed with verification. This policy sets that cadence and makes the findings tracker the single source of truth for what is open and who owns it.

This policy is fictional and illustrative. It describes no real scan, finding, or system.

The policy deliberately separates finding, ranking, remediating, and verifying so that no single step can hide the others: a finding that is ranked but never remediated is visible, and a fix that is applied but never verified is caught.

The policy turns a firehose of scan output into a prioritized, dated, and verified queue, so effort goes to the weakness that matters most.

The policy keeps the queue honest, so effort lands on weaknesses that are real and dangerous, and the queue reflects risk rather than noise.

Vulnerability management exists to find and close weaknesses before they are exploited, and the policy makes the fix order a decision driven by severity and exploitability rather than by arrival order.

2. Definitions

Term	Definition (illustrative)
Vulnerability	A weakness in a system, application, or dependency that could be exploited.
Finding	A recorded instance of a vulnerability with a severity and an owner.
Severity	The ranked impact of a finding, used to set remediation priority.
Remediation owner	The role accountable for fixing a finding and recording the change.

Deferral	A documented decision to accept residual risk and delay remediation.
Re-scan	A follow-up scan to verify a finding is resolved.
Scanner	The tooling that detects vulnerabilities on a schedule.
Triage	The process of deduplicating and ranking new findings.
Aging	The time a finding has been open relative to its target.
Exploitability	The likelihood a weakness can actually be exploited in context.
Correlation	The deduplication and linking of findings across sources.
Exploitability context	The note on whether a weakness has a reachable path.
Disclosure posture	The stated approach to externally reported weaknesses.
Exploit	Working code or a technique that takes advantage of a vulnerability.
False positive	A finding that turns out not to be a real vulnerability on closer inspection.
Known exploited	A vulnerability for which active exploitation has been observed in the wild.
Remediation	The fix that removes a vulnerability.
Severity scoring	The rating of a vulnerability by the risk it carries.
Attack surface	The set of entry points an attacker could use.
Remediation SLA	The committed time to fix by severity.
Penetration test	An authorized simulated attack on a system.
Security assessment	The review of a system's security posture.
Dependency scan	The check of third-party components for known weaknesses.
Disclosure	The published posture on reporting and fixing weaknesses.
Taxonomy	The named weakness types used to classify findings.
Scanner coverage	The proportion of in-scope assets that a vulnerability scanning tool can reach and assess.
Business context	The operational importance of an asset, including whether it is public-facing or processes regulated data.

3. Scope and Boundaries

This policy applies to the fictional platform, its dependencies, and the supporting infrastructure inside the service boundary. It covers both application and infrastructure weaknesses.

It applies to findings from any source, whether a scheduled scan, a dependency advisory, or a report from a developer, so no finding bypasses the tracker.

Out of scope are real systems and real vulnerabilities. This policy creates no real obligation and asserts no certification.

Escalation for a vulnerability issue follows a fixed chain: a past-due critical finding goes to the vulnerability owner, then to the Information Security Lead, and finally to the governance forum when the exposure is material and unmitigated.

The program covers the demonstration's own code and configuration; a vulnerability in a customer-facing service is prioritized over one in an internal tool.

The vulnerability program is fictional; its findings and remediation are illustrative of a real program's shape.

Severity is the starting point, but exploitability can move a finding up the queue.

This policy covers vulnerabilities in operating systems, application dependencies, containers, and cloud resources, including assets discovered outside the standard inventory.

4. Roles and Responsibilities

Role	Illustrative responsibility
Vulnerability Management Lead	Owns the policy and the findings tracker.
Scanner Operator	Runs scans and maintains the tooling.
Remediation Owner	Fixes findings and records the change.
Risk Owner	Accepts residual risk for deferred findings.
Engineering Lead	Coordinates remediation across development teams.
Trust Center Coordinator	Maintains the fictional label on shared copies.

The following RACI is illustrative.

Illustrative RACI for vulnerability lifecycle:

Activity	Scanner Operator	Remediation Owner	Vuln Mgmt Lead	Risk Owner
Run the scan	R	I	C	I
Triage findings	C	I	R	I
Remediate	I	R	C	I
Verify closure	R	I	A	I

Accept deferral	I	C	C	A
-----------------	---	---	---	---

5. Policy Requirements and Controls

PTT-POL-011-C01. Scan on a defined cadence

The environment is scanned regularly and after material change so new weaknesses are found promptly rather than lingering until the next annual review. The schedule is recorded and monitored for completion. The scan schedule is published and a missed scan is flagged.

PTT-POL-011-C02. Prioritize by severity

Findings are ranked so the highest risk is addressed first, using severity and exploitability together so that effort goes where the danger actually is. The fix order is checked against the severity rating.

PTT-POL-011-C03. Assign an owner and target date

Each finding has a named owner and a remediation deadline, because a finding with no owner and no date is effectively invisible no matter how carefully it is ranked. A finding without an owner is flagged.

PTT-POL-011-C04. Meet remediation targets

Remediation is completed within the illustrative timeline for its severity, and slippage is escalated rather than silently absorbed. The target is met, and a miss is escalated.

PTT-POL-011-C05. Verify remediation

A re-scan confirms the finding is resolved before closure, so a fix that did not actually work is caught rather than papered over. The retest result is recorded before closure.

PTT-POL-011-C06. Track deferrals

Deferred findings are recorded with a rationale and an acceptance decision, so residual risk is explicit and owned rather than implicit. A deferral records its reason and new date.

PTT-POL-011-C07. Report metrics

Aging and closure metrics are reviewed monthly in the demo, giving the program a view of whether the backlog is shrinking or growing. The metrics cover time-to-remediate and open findings, so the program is measurable.

PTT-POL-011-C08. Correlate with patches

Vulnerabilities that map to a patch are routed to the patch program, so the two processes reinforce each other. A finding is tied to its patch, so the two programs stay in step.

PTT-POL-011-C09. Cover dependencies

Third-party library advisories are triaged through the same tracker, so a vulnerable dependency is treated as a finding, not ignored. Third-party components are scanned, not just first-party code.

PTT-POL-011-C10. Escalate critical findings

Critical findings follow the emergency change or incident path rather than the routine remediation queue. A critical finding is escalated within a day.

PTT-POL-011-C11. Correlate findings across sources

Findings from scans, advisories, and manual reports are deduplicated and correlated so the tracker shows one view of the risk rather than fragmented entries. Correlation prevents double-counting and missed overlaps. Findings from different scanners are combined, so duplicates are merged.

PTT-POL-011-C12. Track exploitability context

Each finding carries an exploitability note so a high-severity weakness with no reachable path is treated differently from one with a proven path. Context keeps prioritization honest. A finding records whether it is actively exploited, so the fix order reflects danger.

PTT-POL-011-C13. Publish a vulnerability disclosure posture

The program states how it handles externally reported weaknesses, so a researcher reporting a bug knows the path and the expectation. The posture states how findings are reported and fixed.

PTT-POL-011-C14. Maintain a vulnerability taxonomy

Findings are categorized using a consistent taxonomy so the tracker can be queried and trended, and so similar findings are treated similarly. The taxonomy names the weakness types, so findings are classified consistently.

PTT-POL-011-C15. Track remediation effectiveness

The program measures whether remediations actually close findings or recur, so a fix that does not stick is investigated. The program tracks whether fixes stay fixed, so a recurrence is caught.

PTT-POL-011-C16. Publish a scanning schedule

The scanning schedule is published so owners know when their systems will be assessed and can prepare. The schedule is published, so stakeholders know when scans run.

PTT-POL-011-C17. Handle false positives

False positives are documented and suppressed with a rationale, so the tracker is not cluttered with noise that hides real findings. A false positive is closed quickly, so the queue is not polluted.

PTT-POL-011-C18. Discover vulnerabilities continuously

The environment is scanned for known weaknesses on a recurring schedule and on material change, because a one-time scan is obsolete the day it runs. Discovery is the prerequisite for every later step. Discovery is continuous, not only at the scheduled scan.

PTT-POL-011-C19. Set remediation timelines by severity

Critical findings are remediated within a short window, high within a longer one, and lower severities are scheduled, so the highest exposure is closed first and nothing drifts indefinitely. Each severity has a target, so the fix order is defined.

PTT-POL-011-C20. Prioritize known-exploited weaknesses

A finding known to be actively exploited is remediated ahead of its nominal severity, because real-world use outweighs a theoretical score. The known-exploited flag overrides the ordinary queue.

PTT-POL-011-C21. Re-scan after environment change

A material environment change triggers a fresh scan, so a new attack surface introduced by the change is found while the change is still fresh. A change triggers a re-scan, and the result is recorded.

6. Operating Procedures

Run the scan.

The scanner operator executes the scheduled or event-driven scan and confirms it completed.

Triage findings.

Findings are deduplicated, assigned a severity, and entered into the tracker.

Assign ownership.

Each finding is matched to a remediation owner and a target date.

Remediate.

The owner fixes the issue through the change process and records the change.

Verify with a re-scan.

The finding is re-checked and closed only when the re-scan confirms it is resolved.

Handle deferrals.

A deferral is documented with a rationale and an acceptance decision.

Review the tracker.

The lead reviews aging and deferrals monthly.

Report metrics.

Closure and aging metrics are reported to the program.

Correlate new findings.

New findings are deduplicated against the tracker and correlated with related entries before assignment.

Review exploitability.

The triage step records whether a finding has a reachable path, which informs its severity.

Categorize a finding.

A finding is categorized using the taxonomy at triage.

Suppress a false positive.

A false positive is documented with its rationale and suppression date.

Triage a new finding.

The finding is scored, assigned an owner, and slotted into the remediation timeline based on its severity.

Retest a remediated finding.

The fix is verified by a fresh scan, and the finding is closed only on a clean result.

Resolve a false positive.

The finding is verified as not exploitable, documented with the evidence, and closed.

Close a finding.

The analyst confirms the fix by re-testing, records the result, and the finding is closed only on that evidence.

Run a penetration test.

The tester runs the authorized exercise, records the findings, and the owner assigns remediation.

Re-test a fix.

The analyst re-runs the check to confirm the weakness no longer reproduces.

Escalate a critical finding.

A critical finding is raised to the Information Security Lead within a day.

Scan dependencies.

The analyst runs the dependency scan and records any known-vulnerable component.

Correlate findings.

The analyst merges findings from different scanners and removes duplicates.

Publish a disclosure.

The provider publishes its vulnerability disclosure posture and updates it.

Handle a false positive.

The analyst closes the false finding quickly and records why it was not real.

Maintain a taxonomy.

The analyst classifies each finding by weakness type and keeps the taxonomy current.

Track exploitability.

The analyst records whether a finding is actively exploited, so the fix order reflects danger.

Validate scanner coverage.

The vulnerability team periodically confirms that every in-scope asset is reachable by the scanning tool and that no asset class is silently excluded from assessment.

Prioritize findings by business context.

The team ranks vulnerabilities using the business importance of the affected asset alongside technical severity, so a critical issue on a public-facing system is addressed before a lower-impact one.

7. Frequency and SLAs

Illustrative vulnerability cadence:

Activity	Frequency (illustrative)
Environment scan	Weekly and after material change
Dependency scan	On every build
Tracker review	Monthly
Deferral review	Monthly

8. Exceptions and Lifecycle

An exception permits deferring a finding past its target when the Risk Owner documents the compensating controls and accepts the residual risk. The deferral names the finding, the reason, and the new date.

Deferrals are reviewed at least monthly in the demo, and a finding deferred across two cycles is escalated for a decision rather than deferred indefinitely.

A finding deferred across two cycles is escalated for a decision rather than deferred indefinitely, because a permanent deferral is really an unstated acceptance that should be made explicit.

9. Records and Evidence

Record type	What it demonstrates (illustrative)
Scan reports	The scheduled scan output.
Findings tracker	The ranked, owned findings.
Remediation change records	The changes that closed findings.
Re-scan verification	The confirmation of closure.
Deferral approvals	The accepted residual risk.
Metrics reports	The monthly aging and closure view.
Correlation record	The deduplicated and linked findings.
Exploitability notes	The path analysis per finding.
Scan results	The raw and triaged findings.
Remediation records	The fixes applied and their retests.
Retest records	The confirmation that a fix closed the finding.
Scan output	The raw results of a scheduled scan.
SLA compliance report	The measured fix times against target.
Penetration test report	The findings of an authorized test.
Coverage validation report	Shows each in-scope asset and confirms it is reachable by the scanner.
Prioritized finding list	Ranks findings by severity and business context with an owner and target date.
False positive log	Lists suppressed findings with the evidence that justified suppression.

10. Oversight and Metrics

- Mean time to remediate: within the illustrative severity target.

- Critical findings: none open beyond the emergency window.
- Backlog: total open findings trending down.
- Verification rate: 100 percent of closures re-scanned.
- Correlation accuracy: no duplicate findings in the tracker.
- Mean time to remediate: the median time to close findings by severity trending down.
- Open critical findings: the count of critical findings past due trending toward zero.
- Mean time to remediate by severity.
- Number of findings open beyond their target date.
- Findings open beyond their target date.
- Attack surface change per period.
- SLA compliance by severity.
- Penetration test findings closed.
- Remediation SLA compliance.

11. Appendix

Illustrative remediation targets:

Severity	Target (illustrative)
Critical	Immediate or emergency path
High	Within a short defined window
Medium	Within a scheduled cycle
Low	Next maintenance window or documented deferral

12. Related Policy Dependencies

- Patch Management Policy (applies the patches that close many findings)
- Secure SDLC Policy (prevents findings at build time)
- Risk Management Policy (owns the deferral decisions)
- Change Management Policy (governs the remediation changes)
- Patch Management Policy (closes the findings this policy ranks)

13. Review Cadence

The Vulnerability Management Lead reviews this policy at least annually, after a change to the illustrated scanning tooling, and after a significant simulated vulnerability event.

Revision History

Version	Date	Change summary	Author
1.0	September 28, 2026	Original enterprise-depth fictional demo	Paper Tiger Technologies

		policy for controlled trust-center testing.	
--	--	--	--

Document Control

Field	Value
Document ID	PTT-POL-011
Document Owner	Vulnerability Management Lead
Status	Fictional demo artifact
Effective Date (illustrative)	September 28, 2026
Review Cadence	At least annually and after material fictional changes
Classification	Internal Demo
Approval	Illustrative governance approval only; no real adoption represented