

The Pentest Management Maturity Model


A 2027 benchmark for turning offensive testing into measured,
provable risk reduction.

Research by

SecurityWall's award-winning offensive security researchers

A decade of penetration testing research, presented at Black Hat, OWASP, BSides and infosec conferences worldwide

For CISOs, offensive security and GRC leaders planning 2027



Contents

01	Executive summary	3
02	Why 2027 is different	4
03	The Pentest Management Maturity Model	5
04	Inside each level	6
05	The seven dimensions of maturity	8
06	Self-assessment	9
07	Where organisations sit today	10
08	A worked example	11
09	The 2027 metrics that matter	12
10	What PM3 means for each role	13
11	A 90-day plan to move up one level	14
12	The regulatory lens for 2027	15
13	Eight predictions for 2027	16
14	How SecurityWall runs the model with SLASH	17
15	Frequently asked questions	18
16	Glossary, methodology and sources	19
17	Services, compliance and free tools	20

Disclosure: SecurityWall builds SLASH, which appears in this paper as a worked example. The model, self-assessment and metrics apply to any platform or in-house workflow.

Executive summary

In 2027, the organisations that stay ahead will not be the ones that find the most vulnerabilities. They will be the ones that close, and prove the closure of, the right ones fastest.

Discovery has become abundant. Verification, remediation and evidence have not. The gap is already measurable: in Cobalt's 2026 State of Pentesting data, top performers close half their high-risk findings in 10 days while the bottom tier takes 249 days, a 25x spread built from the same tests and the same tools. AI-scale discovery is widening it. Of 26,153 findings generated under Project Glasswing, fewer than 0.8% were patched by September 2026.

What separates the leaders is not budget or tester talent. It is how pentest work is managed after the tester finds something: how findings reach engineers, how fixes are verified, and how proof reaches auditors, boards and customers.

The **Pentest Management Maturity Model (PM3)** gives security leaders a shared language for that layer. It defines five levels, from document-driven delivery to continuous assurance, across seven dimensions, with a self-assessment, a 2027 metrics set, a regulatory map and a 90-day plan to move up one level.

25x

Gap in high-risk finding half-life between leaders (10 days) and laggards (249 days)

Cobalt, State of Pentesting 2026

57% vs 15%

Executives versus practitioners who believe remediation SLAs are consistently met

Cobalt, State of Pentesting 2026

4.5x

More likely: programmatic teams closing critical findings in three days or less

Cobalt, State of Pentesting 2026

38%

Resolution rate for high-risk AI and LLM findings, the lowest of any category

Cobalt, State of Pentesting 2026

32 → 43

Days, median time to fully patch a vulnerability, year on year

Verizon DBIR 2026, via Eclipsium

< 0.8%

Share of 26,153 Project Glasswing findings patched by September 2026

VulnCheck analysis, via Dark Reading

Who this paper is for. CISOs and heads of offensive security planning 2027 budgets for [penetration testing](#), PTaaS and [vulnerability assessment](#), GRC leaders who own audit evidence, and engineering leaders who receive pentest findings. Consultancies and MSSPs can use the same model to benchmark their own delivery.

Reach Level 3 everywhere, Level 4 for every crown-jewel application, and measure it by verified fixes rather than reports delivered.

Why 2027 is different

2027 is the first planning year in which AI-scale vulnerability discovery is a given rather than a forecast. Finding flaws is getting cheaper every quarter. Validating, fixing and proving the fix is not.

Fewer than 1 in 100 AI-found vulnerabilities reached a patch



Source: VulnCheck analysis of Anthropic's Vulnerability Disclosure Ledger, April to September 2026, reported by Dark Reading, 9 September 2026. Bar lengths to scale.

Even with frontier compute and major vendors as partners, human validation and coordination became the constraint, not discovery. Four forces will define offensive security programmes in 2027.

Discovery is now an industry capability

Within weeks of the Mythos announcement, OpenAI and Microsoft launched their own AI vulnerability discovery efforts, and smaller models are reproducing flagship results. Attackers will run the same techniques.

Patching is slowing while exploitation rises

Exploitation became the leading initial access vector in the 2026 Verizon DBIR at 31% of breaches, while the median time to fully patch rose from 32 to 43 days. Only 26% of critical KEV vulnerabilities were fully remediated, down from 38%.

Triage costs more than finding

In one test, scanning a 2-million-line codebase with AI cost about \$315 in API charges. Triaging the results cost about \$128,000. Three AI scanners run against the same codebase agreed on only 5% of their findings.

Analysts expect remediation to be wired in


Gartner predicts that by 2029, 30% of organisations will link adversarial exposure validation results to automated remediation or orchestration workflows.

Buying more discovery without upgrading how findings are managed adds to the backlog, not to security.

The Pentest Management Maturity Model

PM3 measures one thing: how reliably a validated finding becomes a verified, evidenced fix. Each level removes one human wait from that path.

Each level adds one thing: proof that a fix happened, sooner



The levels are cumulative: on-demand verification depends on findings already living inside engineering's workflow.

Level 1, Document. Testing exists to produce a report. Findings arrive as a PDF weeks after testing ends, are retyped into tickets by hand, and are not checked again until next year's test.

Level 2, Portal. Findings are visible in a vendor portal, but the work of fixing them happens somewhere else. Retests are booked on request and evidence lives in exported files.

Level 3, Integrated. Validated findings land in engineering's ticketing system as they are found. Status still flows one way, and verification waits for tester availability, typically one to three weeks.

Level 4, Closed-loop. Ticket status and comments flow back automatically. Fixes are verified on demand, each verdict is timestamped, and the board and engineering read separate views of the same record.

Level 5, Continuous assurance. Human and automated testing run as one programme. Automatable fixes are verified at deployment, complex logic stays with human testers, and proof of testing can be shared with customers and regulators on request.

Inside each level

Each level has a recognisable operating rhythm, a characteristic risk, and one move that unlocks the next. Use these profiles to sanity-check your self-assessment.



Level 1

Document

The pentest is booked to meet an audit date. A PDF arrives two to four weeks after testing ends, someone copies findings into tickets, and the report is filed.

Signs you are here


- Nobody can say which of last year's findings were fixed.
- The report goes to the board unchanged.
- Next year's test re-discovers last year's findings.

The risk

Exposure windows measured in months, and a signed record of risks you cannot show you closed.

The move to Level 2

Deliver findings through a shared platform, not by email, so each finding has one record.



Level 2

Portal

Findings live in a vendor portal. Engineers are given logins, or an analyst exports findings into tickets. Retests are booked as a separate request.

Signs you are here

- Engineers rarely log in; the real conversation happens in tickets.
- Retests need a scope call or a purchase order.
- Portal and ticket status often disagree.

The risk

Visibility improves, but fixes still depend on manual handoffs and tester availability.

The move to Level 3

Connect the platform to engineering's ticketing system and push findings as they are validated.



Level 3

Integrated

Validated findings appear as tickets during the engagement. Engineers often start fixing before the final report exists.

Signs you are here

- Many findings are fixed before the report, but nobody can prove which.
- Tickets close on an engineer's say-so.
- Audit evidence is rebuilt from ticket history.

The risk

Ticket closure is mistaken for risk closure. This is where most 2026 programmes plateau.

The move to Level 4

Make verification a workflow step, not a separate engagement, and timestamp every verdict.



Closed-loop

Status moves both ways between the testing platform and the ticketing system. An engineer marks a fix deployed, a retest runs on demand, and the verdict is recorded against the finding.

Signs you are here


- Every closed critical has a dated retest verdict.
- Auditors get evidence the day they ask.
- Board and engineers read views of one record.

The risk

Coverage, not closure: assets outside the testing cadence stay unverified.

The move to Level 5

Extend cadence to continuous, hybrid testing of the assets that matter most.



Continuous assurance

Automated testing covers deterministic checks at deployment speed, while human testers focus on business logic and chained attacks. Proof of testing can be shared externally.

Signs you are here

- Crown-jewel releases are tested and verified within days.
- Customers get verifiable proof without an NDA cycle.
- Leadership reviews exposure trends, not counts.

The risk

Over-automation: automated verdicts are only as good as their validation discipline.

Staying here

Re-assess twice a year and keep a human in every loop that involves judgement.

The seven dimensions of maturity

Maturity is rarely even. Most programmes sit at Level 3 on integration but Level 1 on verification. Scoring each dimension separately shows where the next quarter's effort pays off most.

Dimension	Level 1, Document	Level 2, Portal	Level 3, Integrated	Level 4, Closed-loop	Level 5, Continuous
1. Testing cadence	Annual, compliance-timed	Annual plus ad hoc	Per major release	Risk-based, per critical app	Continuous, hybrid human and automated
2. Findings delivery	PDF weeks after testing ends	Portal at report time	Live, as findings are validated	Live, owner assigned on arrival	Live, deduplicated across sources
3. Workflow integration	Findings retyped by hand	Export or one-way push	One-way push to Jira or similar	Two-way sync, status flows back	Sync plus routing rules per owner
4. Fix verification	Next year's test	Retest booked on request	Retest scheduled within weeks	On demand, verdict in hours or less	Automated where safe, human for logic
5. Evidence and audit trail	Report files in shared drives	Reports archived in a portal	Ticket history as evidence	Timestamped verdict per finding	Evidence pack on request
6. Stakeholder communication	One technical report for all	Report plus summary slide	Separate executive summary	Board and engineering views from one source	Plus verifiable proof for customers
7. Measurement and governance	Count of findings	Findings by severity	SLA tracking	Half-life and verified-fix rate	Exposure trend tied to business risk

Your overall level cannot exceed your lowest level on findings delivery, fix verification and evidence.

Those three dimensions decide whether a finding becomes a closed, provable risk reduction. A programme with excellent integration and no verification is still a Level 2 programme in the eyes of an auditor.

Self-assessment: where does your programme sit?

Score each question from 1 to 5 using the matrix, then add the scores. It takes about ten minutes and works best when security, engineering and GRC answer together.

#	Question	Score (1–5)
1	How often are your critical applications tested, and what triggers a test?	
2	How long after a tester validates a finding does the owning engineer see it?	
3	Does a status change in your ticketing system flow back to your pentest record without manual work?	
4	When an engineer ships a fix, how long until it is independently verified?	
5	If an auditor asks for proof that last quarter's critical findings were fixed and re-verified, how long does it take to produce?	
6	Can your board and your engineers each read a version of the results written for them?	
7	Do you report half-life and verified-fix rate, or only finding counts?	
Total (7–35)		


Total	Level	What it typically means
7–12	Level 1, Document	Testing satisfies the audit calendar; risk reduction is assumed, not shown
13–18	Level 2, Portal	Findings are visible online, but fixing happens elsewhere
19–24	Level 3, Integrated	Engineering receives findings fast; verification and evidence lag
25–30	Level 4, Closed-loop	Every fix is verified and every verification is evidenced
31–35	Level 5, Continuous	Testing, verification and proof run at deployment speed

Your overall level cannot exceed your lowest score on questions 2, 4 or 5, whatever the total. A programme scoring 30 with a 2 on verification is a Level 2 programme.

Where organisations sit today


The public data suggests most organisations operate between Level 2 and Level 3: findings reach engineering faster than before, but verification and evidence still run on human scheduling.

Leaders close half their high-risk findings 25 times faster than laggards



High-risk finding half-life (days to close 50% of high and critical findings). Source: Cobalt, State of Pentesting 2026. The median organisation takes 39 days, against SLAs often set at 7.

Executives and practitioners see remediation very differently



Source: Cobalt, State of Pentesting 2026, survey of 450 security leaders and practitioners.

Signal in the 2026 data	Figure	What it suggests
Organisations taking a programmatic approach	53%, first time above compliance-driven (40%)	Level 3 is becoming the norm
High-risk findings resolved, five-year dataset	52%	Half of known risk is never verifiably closed
Critical KEV vulnerabilities fully remediated	26%, down from 38%	Remediation capacity is falling behind
AI scanners agreeing on the same findings	5% across three tools	More automated findings raise triage cost

Sources: Cobalt State of Pentesting 2026; Verizon DBIR 2026 via Eclypsium; Contrast Security via Dark Reading.

Editor's note, remove before publication. Add SecurityWall's own anonymised figures from engagements delivered through SLASH here: median time from validation to engineer visibility, share of findings closed before the final report, median retest turnaround, and verified-fix rate.

A worked example: Level 2 to Level 4 in two quarters

Illustrative scenario. The organisation is a composite and the figures are hypothetical. Replace with an anonymised SecurityWall client story before publication.

A regional payments company runs annual web and API tests for PCI DSS and SOC 2. Its first self-assessment scores 16: a Level 2 programme with a 1 on fix verification.

Quarter 1: connect

The team chooses verification as its weakest dimension but starts with integration, because verification needs findings inside engineering's workflow first. Findings now push to Jira as they are validated, and ticket status flows back. Engineers begin fixing during the engagement rather than after it.

Quarter 2: prove

Retest becomes a button, not a booking. When an engineer marks a fix deployed, the finding is re-verified and the verdict is recorded with a timestamp. The CISO's board pack switches from finding counts to half-life and verified-fix rate.

Dimension	Before	After two quarters
Findings delivery	Portal at report time	Live, as validated
Workflow integration	Manual re-entry	Two-way Jira sync
Fix verification	Retest booked weeks later	On demand, verdict recorded
Evidence	Reconstructed for each audit	Generated from the record
Board reporting	Finding counts	Half-life and verified-fix rate
Self-assessment score	16, Level 2	27, Level 4 on crown-jewel apps

The biggest gain came from the least glamorous change: making verification part of the workflow. No new discovery tooling was bought.

The 2027 metrics that matter

Stop reporting how many findings were found. In 2027, report how fast verified risk reduction happens.

Seven metrics cover it. The targets are SecurityWall's recommended 2027 benchmarks, anchored where possible to published leader performance.

Metric	Definition	Level 3	Level 4	Level 5
Time to engineer visibility	Validation of a finding to its owner seeing it	< 5 days	< 24 hours	Same day, automatic
High-risk half-life	Days to close 50% of high and critical findings	< 30 days	≤ 10 days (2026 leader level)	< 7 days
Retest turnaround	Fix marked deployed to independent verdict	< 14 days	< 48 hours	Minutes for automatable checks
Verified-fix rate	Closed findings with a timestamped retest verdict	> 50%	> 90%	100% for critical and high
Evidence retrieval time	Auditor request to complete evidence pack	< 1 week	< 1 day	On demand
SLA agreement gap	Executive minus practitioner SLA confidence	< 30 points	< 10 points	One shared dashboard
AI/LLM finding resolution	High-risk AI findings resolved (2026: 38%)	> 50%	> 75%	> 90%

A fast half-life without verification measures how quickly tickets close, not how quickly risk does.

In practice with **SLASH**

In engagements delivered through SLASH, each finding moves through five tracked states (New, Triaged, Ready for Retest, Resolved, N/A) and every retest verdict is timestamped. Time to visibility, retest turnaround and verified-fix rate are read from the record, not reconstructed at quarter end.

What PM3 means for each role

The model only works when the people who own each dimension agree on the score. Each role reads it differently.

For the CISO

PM3 turns a pentest budget line into a risk-reduction line. Show the board two numbers: high-risk half-life and verified-fix rate. Both move when the programme matures; neither moves when you simply buy more testing.

First action for 2027: run the self-assessment with engineering and GRC, and set a Level 4 target for crown-jewel applications.

For the engineering leader

Maturity reduces engineering effort. Findings arrive in the tools your teams already use, written for the developer who will fix them. On-demand verification ends reopened tickets and audit-time archaeology.

First action for 2027: agree a remediation SLA by severity that engineering believes it can meet, and measure it from the record.

For the GRC lead

Auditors increasingly ask for proof of closure, not proof of testing. A Level 4 programme produces that proof as a side effect of normal work. Map your frameworks to the regulatory table in this paper.

First action for 2027: ask for last quarter's critical findings with retest proof, and time the answer.

For consultancies and MSSPs

Your clients are moving up the model whether or not you do. Delivery maturity is becoming a buying criterion, and every hour saved on reporting and retest scheduling returns to testing.

First action for 2027: score your own delivery against the seven dimensions and share the result with clients.


In practice with SLASH

SLASH serves all four roles from one record. Client roles separate owners, admins, members and auditors, with per-engagement scope for auditor access, so each stakeholder sees what they need without a separate report.

A 90-day plan to move up one level

One level per quarter is a realistic pace. Moving two levels at once usually means new tooling arrives before the workflow it is meant to support.

Ninety days, three phases, three gates



Each gate needs agreement from security, engineering and GRC, because each dimension belongs to more than one of them. Start with the dimension that scored lowest, not the one that is easiest to fix.

Common traps to avoid

- **Buying discovery before fixing delivery.** More findings into a Level 2 workflow lengthens the backlog.
- **Treating ticket closure as remediation.** A closed ticket without a retest verdict is a claim, not evidence.
- **One report for every audience.** Boards and engineers need different views of the same record, not two documents that drift apart.
- **Automating judgement.** Automated verification suits deterministic fixes. Business-logic and chained findings still need a human tester.

The regulatory lens for 2027

Regulators and auditors are converging on one question: can you prove that what you found was fixed? Level 4 is the comfortable floor for most regulated organisations in 2027. The table maps compliance penetration testing requirements to the maturity level that satisfies them most comfortably.

Framework	What it asks of offensive testing	Evidence that satisfies it	Level
SOC 2 CC4.1, CC7.1	Ongoing evaluation that vulnerabilities are identified and addressed	Test reports plus remediation and retest evidence across the audit period	3–4
ISO/IEC 27001:2022 A.8.8, A.8.29	Management of technical vulnerabilities; security testing in development	Test evidence tied to the Statement of Applicability and treatment records	3
PCI DSS v4.0.1 Req. 11.4	Annual and post-change testing; exploitable findings corrected and testing repeated to verify	Retest results for every exploitable finding	4
DORA Art. 24–26	Regular testing of ICT systems; threat-led testing for significant entities	Basic testing evidence; TLPT stays human-led under TIBER-EU	4
NIS2 Art. 21	Policies to assess the effectiveness of risk-management measures	Trend evidence that measures work over time	4
EU AI Act , ISO/IEC 42001 Art. 9, 15	Testing of high-risk AI systems across the lifecycle; robustness and cybersecurity; an AI management system	AI and LLM penetration test and remediation records; confirm current application dates	4
GDPR Art. 32	A process for regularly testing and evaluating the effectiveness of security measures	Periodic test results with remediation records for systems processing personal data	3–4
HIPAA §164.308(a)(8)	Periodic technical evaluation of safeguards protecting ePHI	Test reports and remediation evidence for in-scope systems	3–4
NCA ECC , SAMA CSF , UAE IA (NESA)	Periodic penetration testing with remediation tracking for in-scope entities	Scoped reports, remediation status and re-verification records	3–4

The artefact most often missing from evidence packs is the retest.

A report listing open critical findings with no verification attached reads as documented, unremediated risk, which is a weaker position than a clean record of closure.

In practice with SLASH

SLASH reports carry CVSS v3.1 scoring, timestamped retest verdicts and a publishing archive for every engagement, so the evidence an auditor asks for in month eleven is the same record engineering used in week one.

Eight predictions for offensive security in 2027

SecurityWall's forecasts, grounded in the 2026 data. They are judgements, not certainties, and we will score them publicly in the 2028 edition.

- 1** **Verified-fix rate becomes a board metric.** As finding volumes grow, boards will stop asking how many vulnerabilities were found and start asking how many were proven fixed.

- 2** **"Retest included" stops being a differentiator.** Buyers will expect on-demand verification as standard, and vendors that schedule retests by the week will lose renewals.

- 3** **Auditors ask for retest evidence by default.** Evidence requests across SOC 2, ISO 27001 and PCI DSS will include verification records, not just test reports.

- 4** **Hybrid testing becomes the default model.** After the 2026 retreat from fully automated testing, most programmes will pair automated checks with human testers for logic and chained findings.

- 5** **AI and LLM findings become the slowest to close.** With a 38% resolution rate in 2026, AI findings are where maturity gaps will show first.

- 6** **Pentest results go public in summary form.** Verifiable attestations will replace NDA-gated report sharing in sales cycles.


- 7** **Exposure validation and pentest management converge.** Following Gartner's AEV framing, testing results will feed remediation workflows directly.

- 8** **Triage capacity, not testing capacity, sets budgets.** Spending will shift from buying more discovery to faster validation, routing and verification.

How SecurityWall runs the model with SLASH

SecurityWall built SLASH around one design bet: most teams do not fail at documenting findings, they fail at closing them and proving it.

Every fix ends in a timestamped verdict, not a closed ticket



The SLASH closed loop, from SecurityWall product documentation. A verdict is one of three: fixed, still vulnerable, or could not verify.

PM3 dimension	How SLASH supports it
2. Findings delivery	Findings are visible as they are validated, not when the report is written
3. Workflow integration	Bidirectional Jira sync: issues push out, comments and status flow back by webhook, with per-finding sync state
4. Fix verification	On-demand retest with a verdict of fixed, still vulnerable or could not verify; human testers handle complex logic
5. Evidence and audit trail	CVSS v3.1 scoring, timestamped retest verdicts and a report archive with publishing timestamps
6. Stakeholder communication	Executive and technical views from one set of findings; PDF, Excel and Markdown exports; a publicly shareable, independently verifiable certificate
7. Measurement	Five tracked finding states give retest and remediation metrics from the record

Where SLASH is not the right fit

SLASH is cloud-only, so organisations that require self-hosted or air-gapped deployment should look elsewhere. Teams whose main need is purple-team runbooks or consolidating many scanner feeds may be better served by broader platforms. SecurityWall's published comparison with PlexTrac sets out these trade-offs openly.

Frequently asked questions

Is PM3 a replacement for CTEM or Gartner's AEV framing?

No. CTEM describes the full exposure management cycle and AEV describes automated attack validation. PM3 covers one part of both: what happens to a validated finding until its fix is proven. It fits inside a CTEM programme's validation and mobilisation stages.

Do we need to reach Level 5?

Not everywhere. Level 3 is a sound baseline for most assets. Level 4 is the 2027 target for crown-jewel and regulated applications. Level 5 is worth the investment where release frequency is high and breach impact is severe.

Can automated retesting replace a human tester?

For deterministic fixes, such as a missing header, an exposed endpoint or an injection point that is now parameterised, automated verification is reliable and fast. For business-logic flaws, authorisation chains and multi-step attacks, a human should confirm the fix. Mature programmes use both.

How does PM3 relate to bug bounty programmes?

Bug bounty findings enter the same management layer. The question is unchanged: how quickly does a validated report become a verified fix?

Is PM3 tied to a specific product?

No. Each level describes operating practice that different platforms or in-house tooling can achieve. SLASH is used as a worked example because SecurityWall builds it.

How often should we re-assess?

Twice a year, and after any major change to tooling, team structure or regulatory scope.

What is the quickest win?

Measuring evidence retrieval time. Asking for last quarter's critical findings with retest proof usually exposes the weakest dimension within a day.

Glossary, methodology and sources

Term	Meaning in this paper
AEV	Adversarial exposure validation: Gartner's 2026 category for technologies that give continuous, automated evidence that an attack is feasible
CTEM	Continuous threat exposure management: scoping, discovery, prioritisation, validation and mobilisation
CVSS	Common Vulnerability Scoring System, the standard severity score attached to findings
Evidence retrieval time	Time from an auditor's request to a complete evidence pack
High-risk half-life	Days taken to close 50% of high and critical findings from an engagement
Hybrid testing	Human and automated testing in one engagement, with humans on logic and chained attacks
KEV	CISA's Known Exploited Vulnerabilities catalogue
PTaaS	Penetration testing as a service, delivered through an ongoing platform relationship
Retest verdict	The recorded result of re-verifying a fix: fixed, still vulnerable, or could not verify
TLPT	Threat-led penetration testing, required for significant financial entities under DORA
Verified-fix rate	Share of closed findings that carry a timestamped retest verdict

How the model was built

PM3 draws on SecurityWall's delivery experience across web application and API penetration testing, [mobile app penetration testing](#), [cloud security](#) assessments, [red team operations](#), [assumed breach testing](#), AI and OT engagements, and on public 2026 research on remediation performance. The five levels describe observable operating practice, not tooling. Targets in the metrics section are SecurityWall recommendations, not industry standards.

Sources

- Cobalt, [Top 15 statistics from the 2026 State of Pentesting Report](#), May 2026
- Dark Reading, [Mythos vulnerability firehose hits a human bottleneck](#), September 2026
- Eclipsium, [Mythos finds vulnerabilities. But can anyone patch fast enough?](#), June 2026, citing the Verizon 2026 DBIR
- Cymulate, [2026 Gartner Market Guide for Adversarial Exposure Validation](#), March 2026
- Matproof, [Pentest as a Service: the complete guide for 2026](#)
- SecurityWall, [SLASH product documentation](#) and [SLASH vs PlexTrac](#), October 2026
- SecurityWall, [ISO 27001 and SOC 2 penetration testing: what auditors want](#), August 2026

Services, compliance and free tools

Move up the model with help from SecurityWall: offensive security services delivered through SLASH, compliance testing mapped to your framework, and free tools to check your exposure today.

Offensive security services

Penetration testing

Web application, API and network penetration testing

Vulnerability assessment

Validated, prioritised vulnerability scanning

Mobile app penetration testing

iOS and Android applications

Cloud security

AWS, Azure and GCP penetration testing and configuration review

Red team operations

Adversary simulation against people, process and technology

Assumed breach testing

What an attacker can reach from the inside

Threat hunting

Proactive search for active compromise

Digital forensics

Incident investigation and evidence handling

Compliance penetration testing

SOC 2 compliance

ISO 27001

PCI DSS compliance

DORA compliance

GDPR compliance

HIPAA

NCA compliance (KSA)

SAMA compliance

NESA compliance

AI compliance, ISO 42001

Free security tools

SOC 2 readiness assessment

API key checker

JWT analyzer

IPA analyzer

SSL scanner

All free tools

Platforms

SLASH, the pentest management platform behind every engagement in this paper: automated retesting, bidirectional Jira sync and boardroom-ready reports. **VIGIX**, external attack surface and dark-web exposure intelligence for leaked credentials and breach data.

See a Level 4 workflow on your own scope.

Thirty minutes on a live sample engagement, then a trial run against your environment: scoping, on-demand retest, Jira sync, board and engineering reports, and a verifiable certificate.

[Book a SLASH walkthrough](#)

securitywall.co/product/slash

Run the PM3 self-assessment with your team before the call

Bring one recent pentest report; we will map it to the seven dimensions

Leave with your current level and a 90-day plan to the next one



About SecurityWall

An offensive security firm delivering penetration testing, red teaming and compliance assessments through SLASH, its hybrid offensive security platform.

Locations

Riyadh
Dubai
Abu Dhabi
London
Amsterdam

Team certifications

CISSP CISM CISA CRT CREST OSCP OSCP+ OSWE CySA+ PenTest+ + more

Standards and frameworks we test and audit against

ISO 27001 ISO/IEC 27034 ISO/IEC 27400 ISO/IEC 42001 SOC 2 PCI DSS NIST DORA
TX-RAMP SAMA NCA ECC NESA