TECHNOLOGIES

Security Operations Series

Vulnerability Management

Jargon Buster & Quick Reference Guide

A companion reference for implementers, practitioners, and the leaders who back them.

How to Use This Guide

This is the reference companion to our Vulnerability Management series. It is not designed to be read cover to cover. It is designed to be reached for.

Use it when you encounter an unfamiliar term in a scanner report, need to explain a concept to a stakeholder who speaks a different technical language, want the formula for a metric you are presenting to leadership, or need to quickly orient yourself on a scoring system before making a prioritisation decision.

The guide is organised into eleven sections. The glossary in Part 1 is the core. The remaining sections provide structured reference material for scoring systems, metrics, tooling categories, processes, risk frameworks, maturity assessment, and compliance framework mapping. Part 9 collects every acronym referenced across the series into a single lookup table. Part 11 maps VM programme activities to the major compliance frameworks.

A note on the content: this guide draws from and supports the full series. If you need the reasoning behind a concept, the operational detail, or the worked examples that show how these pieces fit together, the relevant episode is where that depth lives. This guide gives you the what. The episodes give you the why.

Part 1: Core Terminology

Definitions are written in the context of Vulnerability Management. Where a term has broader meaning elsewhere, the VM-specific usage is given priority.

A

Agent-Based Scanning Software installed on endpoints that continuously assesses security posture from inside the system. Provides deeper visibility than network scanning but requires deployment and management overhead on every asset. Best suited to laptops, mobile devices, and systems that are not always network-connected.

Agentless Scanning Remote vulnerability assessment without installing software on targets. Uses network protocols and credentials to probe systems. Practical for servers, network devices, and environments where agent deployment is not feasible.

API (Application Programming Interface) Interface allowing software systems to communicate. In VM, APIs enable scanner integration with CMDBs, automated ticket creation, and orchestration workflows.

Asset Any component requiring protection: hardware, software, data, or services. The foundation of vulnerability management. You cannot protect what you have not inventoried. Key attributes include owner, location, criticality, function, and dependencies.

Asset Criticality Rating of how important an asset is to business operations, typically tiered (Tier 1 through 4, or Critical/High/Medium/Low). Drives remediation SLAs, scan frequency, and triage priority.

Asset Discovery The process of identifying all technology assets in your environment. Can be active (network scanning), passive (traffic analysis), or agent-based. The goal is a comprehensive, current inventory of everything in scope.

Attack Chain (Kill Chain) The sequence of steps an attacker follows from initial access to achieving their objective. Understanding chains helps prioritise vulnerabilities that enable critical steps in the sequence.

Attack Surface The sum of all points where unauthorised access or data extraction is possible, including network services, application endpoints, physical access points, people, and supply chain. A core VM goal is reducing unnecessary exposure.

Attack Vector The specific path or method used to exploit a vulnerability. Common vectors include network, email, web, physical access, supply chain, and insider threat.

Authenticated Scan (Credentialed Scan) A vulnerability scan using valid credentials to log into systems for deep inspection. Produces more accurate results with fewer false positives than unauthenticated scanning. Requires credential management and secure storage.

B

Baseline Configuration A standard, approved security configuration for a system type. Deviations from baseline frequently introduce vulnerabilities. Common sources include CIS Benchmarks, DISA STIGs, and vendor security guides.

Blast Radius The potential scope of impact if a vulnerability is exploited. Considers what systems could be compromised, what data could be accessed, and what operations could be disrupted. A key prioritisation factor.

Blue Team The defensive security team responsible for protecting systems, monitoring threats, and responding to incidents. Contrast with Red Team (offensive) and Purple Team (collaborative).

Bug Bounty A programme offering rewards to external security researchers for responsibly disclosing vulnerabilities. Requires mature vulnerability handling processes and a legal framework to operate effectively.

Business Impact The consequence to business operations if a vulnerability is exploited. Measured in financial loss, operational disruption, regulatory penalties, and reputation damage. Essential for risk-based prioritisation and executive communication.

C

Change Advisory Board (CAB) The governing body that reviews and approves significant changes to IT systems. Often involved in approving complex remediations. Can become a bottleneck without an expedited process for critical security patches.

Change Window (Maintenance Window) A scheduled time period when system changes are permitted. Types include emergency (24/7 for critical issues), standard (weekly or bi-weekly), and major (monthly or quarterly).

CI/CD (Continuous Integration/Continuous Deployment) The practice of automatically building, testing, and deploying code changes. VM integration points include security scanning in the pipeline and automated vulnerability testing (shift-left security).

CISA KEV (Known Exploited Vulnerabilities Catalog) An authoritative list maintained by the U.S. Cybersecurity and Infrastructure Security Agency of vulnerabilities confirmed to be actively exploited. Federal agencies must patch within specified timeframes. A strong prioritisation signal for all organisations. See Part 2 for detail.

CMDB (Configuration Management Database) A centralised repository tracking IT assets, configurations, and relationships. Underpins effective asset management, vulnerability correlation, and ownership tracking.

Compensating Control An alternative security measure that mitigates risk when the primary control (typically patching) cannot be implemented. Examples include network segmentation, WAF rules, IPS signatures, enhanced monitoring, and access restrictions. Must be documented, monitored, regularly reviewed, and part of a path to eventual remediation.

Compliance Framework A structured set of requirements, controls, and guidelines that organisations must satisfy to meet regulatory, contractual, or industry obligations. In the VM context, frameworks such as PCI DSS, ISO 27001, SOC 2, and NIST CSF all include vulnerability management requirements, though they express them differently. Compliance is a byproduct of effective security, not a substitute for it. See Part 11 for the cross-reference.

Configuration Drift The gradual divergence of system configurations from approved baselines. A common source of vulnerabilities, typically caused by manual changes, emergency fixes, or automation with outdated configurations.

Coverage (Scan Coverage) The percentage of in-scope assets successfully included in vulnerability scanning. Calculated as (Assets Scanned / Total Known Assets) x 100. Target: above 95% overall, above 98% for critical assets.

Critical Exposure Hours A metric that captures the total duration of exposure across all systems affected by critical vulnerabilities. Calculated as the sum of (hours each critical vulnerability remains open multiplied by the number of systems affected). A critical vulnerability open for 24 hours on 50 systems produces 1,200 exposure hours. One open for 72 hours on 3 systems produces 216. MTTR alone would rank the second as worse, but the first represents the larger risk event. This metric captures the interaction between remediation speed and breadth of exposure that no single-dimension metric can. Tracked quarterly, it provides one of the clearest trend lines for demonstrating programme improvement, and naturally resists gaming because critical vulnerabilities on high-value systems contribute disproportionately to the total. See Part 3 for formula and targets.

Critical Vulnerability The highest severity vulnerability rating, typically CVSS 9.0 to 10.0 or equivalent. Characterised by ease of exploitation and severe impact. Typical SLA: remediate within 24 to 48 hours.

CSPM (Cloud Security Posture Management) Tools that identify misconfigurations and security risks in cloud environments such as AWS, Azure, and GCP. Provides cloud-native vulnerability assessment and compliance checking.

CVE (Common Vulnerabilities and Exposures) A standardised identifier for publicly known vulnerabilities, maintained by MITRE Corporation. Format: CVE-YEAR-NUMBER (e.g., CVE-2024-12345). The universal reference enabling consistent communication across tools, vendors, and organisations.

CVSS (Common Vulnerability Scoring System) A standardised framework for rating vulnerability severity on a 0 to 10 scale. Comprises Base Score (intrinsic characteristics), Temporal Score (changes over time), and Environmental Score (your context). A useful starting point for prioritisation but insufficient on its own. See Part 2 for full breakdown.

CWE (Common Weakness Enumeration) A categorisation system for types of software weaknesses. Where CVE identifies a specific instance, CWE identifies the category. Example: CVE-2024-12345 (a specific SQL injection) maps to CWE-89 (SQL Injection).

Cyber Essentials A UK Government-backed certification scheme for baseline cybersecurity. Requires organisations to keep software up to date with vendor patches applied within 14 days of release for critical and high-severity vulnerabilities, remove unsupported software, and maintain secure configuration. Cyber Essentials Plus adds independent verification through vulnerability scanning and on-site assessment. See Part 11 for cross-reference.

D

DAST (Dynamic Application Security Testing) Security testing of running applications from an external perspective. Simulates attacker behaviour by sending malicious inputs and observing responses. Finds runtime issues and tests deployed configuration, but only covers accessible code paths.

Defence in Depth A security strategy of layering multiple controls so that no single failure results in compromise. In VM terms: a vulnerability may exist, but exploitation requires bypassing firewall, IPS, endpoint protection, and monitoring. Each layer increases difficulty and detection probability.

Detection The phase in the vulnerability lifecycle where a vulnerability is identified in your environment. Methods include automated scanning, manual testing, threat intelligence, and incident investigation. Key metric: Mean Time to Detect (MTTD).

E

End-of-Life (EOL) The point when a vendor stops providing support and security updates for software or hardware. Creates permanent vulnerability exposure. Requires migration to supported versions, compensating controls, or formal risk acceptance.

EPSS (Exploit Prediction Scoring System) A data-driven probability score (0 to 100%) estimating the likelihood a vulnerability will be exploited within the next 30 days. Based on historical

exploitation patterns, exploit availability, threat actor activity, and vulnerability characteristics. Updated daily. See Part 2 for full detail.

Exception A formal, time-limited approval to deviate from standard vulnerability remediation requirements. Must be documented, justified, and regularly reviewed. Not the same as risk acceptance (which may be permanent) or ignoring vulnerabilities (which is negligence).

Exploit Code, technique, or procedure that takes advantage of a vulnerability. Forms range from proof-of-concept (demonstrates the vulnerability exists) to weaponised (ready for attack) to malware (actively used in campaigns).

Exploitability A measure of how easy it is to exploit a vulnerability. Factors include required access level, attack complexity, availability of exploit code, and skill level needed. High exploitability increases priority even when CVSS is not at maximum.

Exposure The state of being vulnerable to attack. Can refer to three distinct things: the duration a vulnerability exists (the exposure window), the accessibility of the vulnerable system (internet-facing systems carry higher exposure than segmented internal systems), or the overall attack surface. Context determines which meaning applies, but all three feed into prioritisation.

Exposure Window The total elapsed time a vulnerability exists in your environment, from introduction to confirmed remediation. Composed of two parts: the time from introduction to detection (MTTD) and the time from detection to remediation (MTTR). Reducing the exposure window requires improving both. A programme that detects quickly but remediates slowly, or remediates quickly but detects late, still carries extended exposure. Formula: Time Remediated minus Time Introduced or Discovered. Track the trend rather than targeting a fixed number, because the appropriate window varies by severity and asset criticality.

F

False Negative An actual vulnerability that scanning or testing failed to detect. Creates dangerous blind spots. Causes include scanner limitations, configuration issues, and novel vulnerability types.

False Positive A finding incorrectly identified as a vulnerability when no actual weakness exists. Wastes analyst time, erodes trust in tooling, and increases triage backlog. Target: below 5% after tuning.

Finding An individual vulnerability detection from a scan or test. Before validation, all findings are potential vulnerabilities requiring analysis to confirm or dismiss.

H

Hardening The process of securing a system by reducing its attack surface, disabling unnecessary features, and implementing security configurations. Common sources of hardening guidance include CIS Benchmarks and DISA STIGs.

Hotfix An urgent, targeted patch released outside regular update cycles to address a critical vulnerability. Typically has limited scope and less testing than standard patches.

I

IAST (Interactive Application Security Testing) A hybrid security testing approach combining SAST and DAST. Agents monitor application behaviour during testing to identify vulnerabilities with the accuracy of static analysis and the real-world validation of dynamic testing.

Incident Attribution Rate The percentage of security incidents that resulted from known but unremediated vulnerabilities. Formula: (Incidents from Known Vulns / Total Incidents) x 100. Target: below 10%. A powerful metric for demonstrating VM programme value to leadership.

Infrastructure as Code (IaC) Managing infrastructure through code files rather than manual processes. Creates opportunities for security scanning before deployment, enabling shift-left vulnerability detection and consistent configurations.

Intrusion Prevention System (IPS) A security tool that monitors network traffic and blocks suspected exploitation attempts. A common compensating control when patching is not immediately possible.

ISO 27001 The international standard for information security management systems (ISMS). The 2022 revision includes Annex A Control 8.8 (Management of Technical Vulnerabilities), which requires a risk-based, documented process for discovering, prioritising, treating, and reviewing technical vulnerabilities. Expects integration with asset management (A.5.9), change management (A.8.32), and incident response (A.5.26). Does not prescribe specific scanning frequencies but expects risk-informed scheduling. See Part 11 for cross-reference.

L

Lateral Movement An attacker's ability to move from an initially compromised system to other systems on the network. Vulnerabilities enabling lateral movement, particularly in high-value network segments, deserve elevated priority.

Lifecycle (Vulnerability Lifecycle) The complete journey of a vulnerability from introduction or discovery through remediation and closure. Stages: Introduction, Detection, Triage, Planning, Remediation, Verification, Closure. See Episode 2 for full treatment.

M

Mean Time to Acknowledge (MTTA) The average time from vulnerability detection to formal acknowledgement and assignment. Indicates triage process efficiency. Target: Critical below 4 hours, High below 24 hours.

Mean Time to Detect (MTTD) The average time from vulnerability introduction or public disclosure to detection in your environment. Target: below 4 hours for critical CVE disclosures, below 24 hours for newly deployed assets.

Mean Time to Remediate (MTTR) The average time from vulnerability detection to confirmed remediation. Must be segmented by severity. Targets: Critical below 48 hours, High below 7 days, Medium below 30 days, Low below 90 days.

Misconfiguration A security weakness resulting from incorrect system settings rather than software bugs. Cannot be fixed by patching. Common examples include default credentials, weak encryption settings, excessive permissions, and unnecessary services.

Mitigation Reducing the likelihood or impact of vulnerability exploitation without fully eliminating the vulnerability. Contrast with remediation, which eliminates the vulnerability entirely.

MSSP (Managed Security Service Provider) A third-party provider offering security services including vulnerability scanning, monitoring, and reporting. Valuable when internal expertise is limited. Critical point: you cannot outsource risk acceptance decisions, remediation execution, or business context.

N

Network Segmentation Dividing a network into isolated sections to limit lateral movement and reduce blast radius. A critical compensating control that can reduce the effective priority of vulnerabilities in isolated segments.

NIS2 (Network and Information Security Directive 2) The EU directive (effective October 2024) that imposes cybersecurity risk management obligations on essential and important entities across member states. Requires vulnerability handling and disclosure processes, supply chain security measures, and incident reporting. Applies broadly across sectors including energy, transport, health, digital infrastructure, and public administration. Non-compliance carries significant financial penalties. See Part 11 for cross-reference.

NIST CSF (Cybersecurity Framework) A voluntary framework published by the U.S. National Institute of Standards and Technology, now at version 2.0. Organised around six core functions: Govern, Identify, Protect, Detect, Respond, and Recover. Vulnerability management activities map across multiple functions, particularly Identify and Protect. Widely adopted internationally as a reference for structuring cybersecurity programmes. See Part 11 for cross-reference.

O

OWASP Top 10 A list of the most critical web application security risks, published by the Open Web Application Security Project and updated periodically. A standard benchmark for web application vulnerability assessment.

P

Patch A software update specifically designed to fix vulnerabilities or bugs. Types include security patches (fix vulnerabilities), cumulative patches (bundle previous patches), and service packs (major collections).

Patch Management The subset of vulnerability management focused specifically on deploying vendor-provided patches. Necessary but not sufficient. VM encompasses much more than patching, including assessment, prioritisation, compensating controls, and risk acceptance.

PCI DSS (Payment Card Industry Data Security Standard) The security standard for organisations handling payment card data. Version 4.0.1 is the current release. Requirement 6.3.1 mandates a vulnerability identification and risk-ranking process; Requirement 11.3 mandates internal and external vulnerability scanning at least quarterly; Requirement 11.4 mandates penetration testing. PCI DSS 4.0 introduced authenticated scanning requirements (11.3.1.2), risk-based remediation timelines for non-critical vulnerabilities (11.3.1.1), and targeted risk analysis. External ASV scans must resolve all vulnerabilities with CVSS 4.0 or above. Critical and high-risk internal vulnerabilities must be patched within one month. See Part 11 for cross-reference.

Penetration Test (Pen Test) An authorised simulated attack against systems to identify exploitable vulnerabilities. Categorised by knowledge (black box, white box, grey box) and by scope (external, internal, web application). Differs from Red Teaming in that pen tests aim to find all vulnerabilities in scope, while Red Teams pursue specific objectives.

Prioritisation The process of determining which vulnerabilities to remediate first based on risk, not severity alone. Factors include severity, exploitability, threat intelligence, asset criticality, exposure, compensating controls, and remediation complexity.

Privilege Escalation Exploiting vulnerabilities to gain higher access levels than originally granted. Vertical escalation moves from user to administrator. Horizontal escalation moves between accounts at the same privilege level.

R

Red Team An offensive security team conducting goal-oriented adversary simulations. Unlike penetration testing, Red Teaming focuses on achieving a specific objective using any means available, not cataloguing all vulnerabilities.

Regression Testing Verifying that changes such as patches or configuration updates have not broken existing functionality or introduced new issues. Critical for quality assurance and maintaining stakeholder trust in the remediation process.

Remediation The act of fixing or eliminating a vulnerability, typically through patching or reconfiguration. Strategies include direct patching, configuration changes, workarounds, compensating controls, and system replacement.

Recurrence A previously remediated vulnerability reappearing in the environment. Causes include configuration drift, deployment issues, code regression, and automation reverting changes. Target: below 2% recurrence rate.

Risk The potential for loss when a threat exploits a vulnerability. Commonly expressed as Risk = Likelihood x Impact. The VM goal is systematic risk reduction through prioritised vulnerability management.

Risk Acceptance A formal decision by business stakeholders to acknowledge and tolerate vulnerability risk without remediation. Requires documentation, business owner approval, clear understanding of consequences, regular review, and ideally a path to eventual remediation.

Rollback Reverting changes to a previous state when a deployment causes problems. An essential component of every remediation plan.

S

SAST (Static Application Security Testing) Analysis of source code without executing the programme. Finds vulnerabilities during development. Strengths: early detection at lowest fix cost, full code coverage. Weaknesses: higher false positive rate than DAST, cannot find runtime issues.

SBOM (Software Bill of Materials) An inventory of all software components and dependencies in an application or system. Critical for supply chain security and for rapidly assessing exposure when new vulnerabilities are disclosed in common libraries.

SCA (Software Composition Analysis) Tools that identify and analyse third-party and open-source components in applications. Critical because modern applications are typically 70 to 90% third-party code, meaning you inherit their vulnerabilities.

Scan An automated process of probing systems to identify vulnerabilities by comparing findings against vulnerability databases. Types include network, web application, cloud, container, authenticated, and unauthenticated.

Security Champion A developer or engineer who advocates for security within their team, bridging the gap between security and development. A practical ratio is one champion per 10 to 15 developers.

Shadow IT Technology deployed without IT or security knowledge or approval. Creates blind spots in VM because you cannot scan what you do not know exists. Common examples include rogue cloud accounts, unsanctioned SaaS subscriptions, and developer-deployed infrastructure.

Shift-Left Integrating security earlier in the development lifecycle to catch vulnerabilities before production deployment. Reduces production vulnerability volume and lowers the cost of fixes.

SLA (Service Level Agreement) A policy or contractual commitment for remediation timeframes, typically tied to severity. Common example: Critical within 48 hours, High within 7 days, Medium within 30 days, Low within 90 days.

SOAR (Security Orchestration, Automation, and Response) A platform for automating security workflows across multiple tools. VM use cases include automated ticket creation, threat intelligence enrichment, remediation verification, and notification workflows.

SOC 2 (Service Organisation Control 2) An audit framework developed by the AICPA for service providers, built on the Trust Services Criteria. The Security criterion (Common Criteria series) is mandatory for every SOC 2 engagement. VM-relevant criteria include CC3.2 and CC3.3 (risk assessment), CC7.1 (vulnerability detection and configuration monitoring), CC7.2 (anomaly detection and vulnerability assessment), and CC8.1 (change management including patching). Auditors expect documented scanning processes, evidence of remediation, and risk-based prioritisation. Quarterly scanning is the common baseline, with critical systems requiring more frequent assessment. See Part 11 for cross-reference.

T

Technical Debt Shortcuts or suboptimal decisions that create future security and maintenance burden. In VM, accumulated debt makes remediation harder, creates persistent exceptions, and eventually requires expensive system replacements.

Threat A potential danger that could exploit a vulnerability. Includes intentional threats (attackers) and unintentional threats (disasters, user errors).

Threat Actor An individual or group presenting a threat. Categorised by sophistication (script kiddies through to nation-states) and by motivation (financial gain, espionage, activism, destruction).

Threat Intelligence Information about current and emerging threats used to inform security decisions. In VM, it feeds prioritisation through data on active exploitation, vulnerability trending, and adversary tactics.

Triage The process of analysing and prioritising vulnerabilities based on multiple risk factors. The goal is intelligent resource allocation, not CVSS-based sorting.

U

Unauthenticated Scan A vulnerability scan without credentials, simulating the external attacker perspective. Shows what attackers see but offers limited visibility and a higher false positive rate compared to authenticated scanning.

V

Verification Confirming that remediation successfully eliminated the vulnerability and did not introduce new issues. Methods include rescanning, regression testing, and monitoring. Target: above 95% of remediations verified.

Virtual Patch A security policy or rule deployed in a WAF or IPS that blocks exploitation attempts without changing the vulnerable system. Used when no vendor patch is available or when testing and approval are still in progress. Does not fix the underlying vulnerability.

VPR (Vulnerability Priority Rating) Tenable's dynamic scoring system (0 to 10) that combines CVSS with real-world threat intelligence. Updates continuously based on exploitation activity. Only available within Tenable products. See Part 2 for detail.

Vulnerability A weakness in a system, application, or process that could be exploited to violate security policy. Sources include software bugs, misconfigurations, design flaws, and operational gaps.

Vulnerability Assessment The systematic process of identifying security weaknesses through scanning, testing, and analysis. Methods include automated scanning, manual testing, code review, configuration auditing, and architecture review.

Vulnerability Management The continuous discipline of discovering, assessing, prioritising, remediating, and verifying security weaknesses before adversaries exploit them. Broader than patching: includes assessment, prioritisation, compensating controls, risk acceptance, and process improvement.

W

WAF (Web Application Firewall) A security tool that protects web applications by filtering malicious traffic. A common compensating control for web application vulnerabilities that can implement virtual patches.

Workaround A temporary measure reducing risk when a patch is not available. Examples include disabling a vulnerable feature, restricting network access, and deploying IPS signatures. Must be temporary with a documented path to a permanent fix.

Z

Zero-Day Vulnerability A previously unknown vulnerability for which no patch exists. The vendor has had zero days of advance notice. May be actively exploited before discovery. Requires emergency response processes, workarounds, and enhanced monitoring.

Part 2: Vulnerability Scoring Systems

Four scoring systems appear most frequently in VM practice. Each serves a different purpose. None is sufficient alone.

CVSS (Common Vulnerability Scoring System)

A standardised severity rating on a 0 to 10 scale. The most widely referenced scoring system and the one most frequently misused as a sole prioritisation mechanism.

Severity Bands

Score Range	Severity	Typical SLA
9.0 - 10.0	Critical	24-48 hours
7.0 - 8.9	High	7 days
4.0 - 6.9	Medium	30 days
0.1 - 3.9	Low	90 days
0.0	None	N/A

Base Score Components

The Base Score measures intrinsic characteristics that do not change over time. It is the score most organisations use and frequently the only score they reference.

Component	Values

Attack Vector (AV)	Network (N), Adjacent (A), Local (L), Physical (P)
Attack Complexity (AC)	Low (L), High (H)
Privileges Required (PR)	None (N), Low (L), High (H)
User Interaction (UI)	None (N), Required (R)
Scope (S)	Unchanged (U), Changed (C)
Confidentiality Impact (C)	None (N), Low (L), High (H)
Integrity Impact (I)	None (N), Low (L), High (H)
Availability Impact (A)	None (N), Low (L), High (H)

Key Limitations

CVSS does not consider whether you are actually vulnerable, does not factor in asset criticality, does not reflect real-world exploitation likelihood, and assigns the same score to theoretical and actively exploited vulnerabilities. Most organisations use only the Base Score, ignoring Temporal and Environmental context entirely. Use it as a starting point, never as a sole prioritisation mechanism.

EPSS (Exploit Prediction Scoring System)

A machine-learning-driven probability score (0 to 100%) estimating the likelihood a vulnerability will be exploited within the next 30 days. Updated daily. Free and publicly available.

EPSS Range	Interpretation	Action
Below 5%	Low probability	Routine prioritisation
5% - 20%	Moderate probability	Elevated attention
20% - 50%	Elevated probability	Prioritise remediation
Above 50%	High probability	Urgent remediation

The practical value: a vulnerability with CVSS 7.0 and EPSS 65% should typically be prioritised over one with CVSS 9.5 and EPSS 0.8%. EPSS predicts exploitation probability, not impact. Combine it with CVSS and asset criticality for effective prioritisation.

CISA KEV (Known Exploited Vulnerabilities)

An authoritative catalogue of vulnerabilities confirmed to be actively exploited, maintained by the U.S. Cybersecurity and Infrastructure Security Agency. Inclusion requires an assigned CVE ID, reliable evidence of active exploitation, and a clear remediation action.

Federal Civilian Executive Branch agencies must remediate KEV vulnerabilities within specified timeframes (typically 14 to 21 days). For all organisations, any vulnerability appearing on the KEV should receive emergency or urgent priority regardless of CVSS score. Subscribe to KEV updates and automate flagging in your triage process.

VPR (Vulnerability Priority Rating)

Tenable's proprietary risk-based scoring system (0 to 10) that combines CVSS with threat intelligence, exploit availability, malware kit usage, and vulnerability age. Updates continuously as the threat landscape evolves. Severity bands mirror CVSS. Only available within Tenable products.

Combining Scores for Prioritisation

EPSS	CVSS	Asset Tier	Priority
High	High	Critical (T1)	Emergency
High	Low	Critical (T1)	Urgent
Low	High	Critical (T1)	Important
Low	Low	Non-critical	Routine

Part 3: Metrics Quick Reference

Metrics are organised by the programme capability they measure. Formulas, targets, and alert thresholds are provided for each. See Episode 5 for the reasoning behind metric selection and the operational detail on building a measurement practice that drives decisions.

A note on compliance: programmes that solve for security produce the evidence auditors need as a natural byproduct. Most of the metrics below satisfy requirements across multiple compliance frameworks simultaneously. Where a metric is particularly relevant to a specific framework, this is noted. Part 11 provides the full cross-reference.

Detection and Coverage

Metric	Formula	Target
Asset Coverage	(Assets Scanned / Total Known Assets) x 100	>95% overall, >98% critical
Scan Success Rate	(Successful Scans / Total Scheduled) x 100	>95%

Mean Time to Detect (MTTD)	Time Detected - Time Introduced or Disclosed	Critical CVEs: <4h
False Positive Rate	(False Positives / Total Findings) x 100	<5% (mature), <3% (optimised)

Compliance relevance: Asset Coverage satisfies ISO 27001 A.5.9 (asset inventory), PCI DSS 11.3.1 (scanning scope), and SOC 2 CC7.1 (monitoring scope). Scan Success Rate and MTTD provide evidence for NIST CSF ID.RA and DE.CM. False Positive Rate supports SOC 2 CC4.1 (control effectiveness).

Triage and Prioritisation

Metric	Formula	Target
Mean Time to Acknowledge (MTTA)	Time Acknowledged - Time Detected	Critical: <4h, High: <24h
Triage Backlog	Count of unreviewed vulnerabilities	Critical: 0, High: <50
Triage SLA Compliance	(Triaged Within SLA / Total Requiring Triage) x 100	>95%

Remediation Speed

Metric	Formula	Target
MTTR (Critical)	Total Remediation Time / Vulns Remediated	<48 hours
MTTR (High)	Total Remediation Time / Vulns Remediated	<7 days
MTTR (Medium)	Total Remediation Time / Vulns Remediated	<30 days
MTTR (Low)	Total Remediation Time / Vulns Remediated	<90 days
SLA Compliance (Critical)	(Remediated Within SLA / Total) x 100	>95%

MTTR is the metric most VM programmes track and the one most frequently misused. The misuse almost always takes the same form: reporting a single aggregate figure across all severities. An overall MTTR of 30 days tells you very little, because it blends critical vulnerabilities remediated in hours with low-severity findings that sat in a backlog for months. Always segment by severity. Where MTTR becomes most powerful is in decomposition: if a third of the total remediation time sits

between planning and approval, the constraint is the change process, not scanning speed or analyst capacity.

Compliance relevance: MTTR directly satisfies PCI DSS 6.3.3 (patch within one month for critical/high), ISO 27001 A.8.8 (timely remediation), and Cyber Essentials (14-day patching requirement for critical/high). SLA Compliance provides audit evidence for all frameworks.

Exposure Window

Metric	Formula	Target
Exposure Window	Time Remediated - Time Introduced or Discovered	Minimise; track trend

The exposure window measures the full elapsed time a vulnerability exists in your environment, from the moment it is introduced (or publicly disclosed, if already present) to confirmed remediation. It is composed of two parts: the time from introduction to detection (MTTD) and the time from detection to remediation (MTTR). Reducing the exposure window requires improving both.

This distinction matters operationally. A programme that detects a critical vulnerability within four hours of disclosure but takes three weeks to remediate has an exposure window of roughly three weeks. A programme that remediates within 48 hours of detection but takes ten days to detect has a window of twelve days. Both programmes have a different problem, but the exposure window is similar. Tracking MTTD and MTTR independently reveals where the constraint sits. Tracking the overall exposure window reveals whether the combined effect is improving.

There is no universal target for exposure window because the appropriate duration depends on severity and asset criticality. The useful practice is tracking the trend by severity tier and treating a growing exposure window as an alert threshold, regardless of the absolute number.

Critical Exposure Hours

Metric	Formula	Target
Critical Exposure Hours	Σ(Critical Vuln Duration in Hours x Affected Systems)	Minimise; decreasing trend

MTTR tells you how long a vulnerability takes to remediate on average. Critical exposure hours tell you something different and arguably more useful: the total duration of exposure across all affected systems. Consider two scenarios: a critical vulnerability open for 24 hours on 50 systems produces 1,200 exposure hours; one open for 72 hours on 3 systems produces 216. MTTR alone would rank the second vulnerability as worse (72 hours versus 24), but the first represents the significantly larger risk event.

This metric captures the interaction between remediation speed and breadth of exposure in a way that no single-dimension metric can. It is also naturally resistant to gaming. A programme cannot reduce its critical exposure hours by fixing easy things while ignoring hard ones, because critical vulnerabilities on high-value systems with broad deployment contribute disproportionately to the total.

Tracked quarterly, critical exposure hours provide one of the clearest trend lines for demonstrating programme improvement to executive stakeholders. The concept is intuitive: you are measuring how long your most important systems are exposed to your most serious vulnerabilities. When that number goes down, the programme is working. When it is flat or growing despite active remediation, the programme is either fixing the wrong things or being outpaced by new high-risk findings.

Worked example: Your programme identifies three critical vulnerabilities in a month. Vuln A affects 100 servers and is open for 36 hours (3,600 exposure hours). Vuln B affects 5 workstations and is open for 48 hours (240 exposure hours). Vuln C affects 20 servers and is open for 12 hours (240 exposure hours). Total critical exposure hours for the month: 4,080. If last quarter averaged 5,500 per month and this quarter averages 4,080, the trend demonstrates measurable improvement.

Remediation Quality

Metric	Formula	Target
First-Time Fix Rate	(Successful First Attempts / Total Attempts) x 100	>90%
Verification Rate	(Remediations Verified / Total Remediations) x 100	>95%, 100% for critical
Recurrence Rate	(Recurring Vulns / Total Vulns) x 100	<2%
Rollback Rate	(Deployments Rolled Back / Total Deployments) x 100	<5%

Compliance relevance: Verification Rate satisfies PCI DSS 11.3.1 (rescan after remediation) and ISO 27001 A.8.8 (confirm vulnerability resolution). Recurrence Rate provides evidence for ISO 27001 continual improvement (Clause 10.2) and SOC 2 CC4.1.

Risk and Governance

Metric	Formula	Target
Risk Score Reduction	Previous Period Score - Current Period Score	>10% reduction quarterly
Incident Attribution Rate	(Incidents from Known Vulns / Total Incidents) x 100	<10%
Exception Rate	(Vulns with Exceptions / Total Vulns) x 100	<10%
Risk Acceptance Rate	(Risk Acceptances / Total Vulns) x 100	<5%
Exception Age	Days since exception granted (average)	<90 days

Automation Rate	(Automated Remediations / Total) x 100	>40% for routine patches

Exposure Window and Critical Exposure Hours are covered in detail earlier in this section as they warrant fuller explanation. Critical Exposure Hours in particular is one of the strongest metrics for executive communication because it captures both speed and breadth of exposure in a single number, and it directly resists gaming.

Compliance relevance: Exception Rate and Exception Age satisfy ISO 27001 risk acceptance documentation requirements. Risk Score Reduction supports NIST CSF GV.RM (risk management strategy). Automation Rate provides evidence for SOC 2 CC8.1 (change management maturity). Incident Attribution Rate supports NIS2 Article 21 (risk management measures).

Part 4: Tool Categories

This section identifies the categories of tooling a VM programme may require rather than recommending specific products. The market moves quickly, and the right tool depends entirely on your environment, scale, and existing stack. Focus on the capabilities you need, then evaluate vendors against those requirements.

Category	Purpose and Selection Guidance
Network Vulnerability Scanners	Scan network infrastructure for known vulnerabilities. Core capability for any programme. Evaluate on asset type coverage, authenticated vs. unauthenticated scanning, false positive rate, API quality, and integration with your CMDB and ticketing systems.
Web Application Scanners	Test web applications for OWASP Top 10 and other vulnerabilities. Evaluate on authentication handling, JavaScript/SPA support, API scanning, and CI/CD integration.
CSPM (Cloud Security Posture Management)	Identify misconfigurations and security risks in cloud environments. Evaluate on multi-cloud support, agentless scanning, IaC scanning, and integration with DevOps tooling.
SAST (Static Analysis)	Analyse source code for vulnerabilities without execution. Evaluate on language support for your technology stack, IDE integration, CI/CD integration, and false positive rate.
DAST (Dynamic Analysis)	Test running applications from an external perspective. Evaluate on coverage of vulnerability types, crawling effectiveness, authentication handling, and scan speed.
SCA (Software Composition Analysis)	Identify vulnerabilities in third-party and open-source components. Evaluate on package manager support, transitive dependency tracking, remediation guidance quality, and container support.

Patch Management	Automate patch deployment and management. Evaluate on OS and application support, staged rollout capabilities, rollback functionality, and reporting.
VM Platforms (Aggregation)	Aggregate findings from multiple scanners, centralise prioritisation, and orchestrate workflows. Typically needed when using three or more scanning tools or managing 5,000+ assets. Evaluate on multi-scanner support, risk-based prioritisation, and ticketing integration.
SOAR	Automate security workflows across multiple tools. VM use cases include automated ticket creation, threat intelligence enrichment, and remediation verification.
Threat Intelligence Platforms	Provide exploit availability data, active exploitation intelligence, and vulnerability trending. Free sources (CISA KEV, EPSS API, NVD) may be sufficient for programmes that are not yet at enterprise scale.

Scaling Guidance

Under 1,000 assets: Network vulnerability scanner, basic patch management, ticketing system. Add a web application scanner if you have web-facing applications. Skip VM platforms and SOAR.

1,000 to 5,000 assets: Add agent-based scanning, CSPM if cloud-hosted, application security testing (SAST/DAST/SCA), and a CMDB. Consider a VM platform if using three or more scanners.

5,000+ assets: A VM platform becomes essential. Add SOAR, advanced analytics, comprehensive threat intelligence, and consider managed services for continuous coverage.

Part 5: Vulnerability Lifecycle Summary

The lifecycle is covered in depth in Episode 2. This section provides the quick reference.

Stage	Description
1. Introduction	Vulnerability enters environment through deployment (new software, misconfiguration, provisioning outside baselines) or discovery (researchers find flaw in existing software, new techniques render current configurations vulnerable).
2. Detection	Vulnerability identified through automated scanning, manual testing, threat intelligence, or incident investigation. Key metric: MTTD.
3. Triage	Finding validated (is it a false positive?), risk assessed (severity + context), priority assigned, owner assigned. Key metric: MTTA.

4. Remediation Planning	Approach determined (patch, configuration change, workaround, compensating control), testing plan documented, impact assessed, change scheduled, rollback procedure defined.
5. Remediation Execution	Change request approved, deployed through environments (dev, staging, production), monitoring active. Key metric: MTTR.
6. Verification	Rescan confirms vulnerability eliminated, regression testing confirms no new issues, monitoring confirms stability. Key metric: Verification Rate.
7. Closure	Documentation completed, ticket closed, patterns analysed for process improvement.

Exception Path

When standard remediation is not feasible, the exception path requires: technical justification for non-remediation, proposed compensating controls, business and security risk assessment, approval at the appropriate level (severity-dependent), documented monitoring plan, scheduled review cycle, and a target date for eventual remediation where possible.

Part 6: Risk Assessment Quick Reference

Risk Matrix (Likelihood x Impact)

	Low Impact	Medium Impact	High Impact	Critical Impact
High Likelihood	**Medium**	**High**	**Critical**	**Critical**
Medium Likelihood	**Low**	**Medium**	**High**	**Critical**
Low Likelihood	**Low**	**Low**	**Medium**	**High**

Low = Routine (90 days) **Medium =** Important (30 days) **High =** Urgent (7 days) **Critical =** Emergency (24-48 hours)

Asset Tier Definitions

Tier	Description	Patch SLA	Scan Frequency
Tier 1 (Critical)	Revenue-generating, compliance-critical systems	Critical: 24h, High: 3d	Daily or Continuous

Tier 2 (Important)	Key business support systems (CRM, email, HR)	Critical: 48h, High: 7d	Weekly
Tier 3 (Standard)	General business systems, standard workstations	Critical: 72h, High: 14d	Weekly
Tier 4 (Low Value)	Dev/test environments, lab systems	Critical: 7d, High: 30d	Bi-weekly or Monthly

Compensating Control Effectiveness

Control Type	Risk Reduction	Primary Use Case
Network Segmentation	40-60%	Isolating vulnerable systems
WAF	40-70%	Web application vulnerabilities
IPS	40-70%	Network-exploitable vulnerabilities
MFA	60-80%	Authentication-related vulnerabilities
Access Restrictions	30-50%	Limiting exposure
Enhanced Monitoring	20-30%	Detection (not prevention)

Part 7: Maturity Model

Programme maturity is not a score to chase. It is a diagnostic tool. Understanding where your programme sits helps you identify what to invest in next and set realistic expectations for progress. The levels below describe typical characteristics. Most organisations sit somewhere between levels rather than neatly within one.

Level	Characteristics	Key Indicators
Level 1: Initial / Reactive	No regular scanning, ad-hoc patching, no prioritisation framework, incident-driven remediation.	Patch compliance <60%, unknown asset inventory, no MTTR tracking.
Level 2: Managed / Compliance-Driven	Regular scanning, compliance-focused patching, CVSS-based prioritisation, documented processes (inconsistently followed).	Compliance 60-80%, partial asset inventory (70-85%), basic MTTR tracking.

Level 3: Defined / Risk-Based	Continuous monitoring, risk-based prioritisation (severity + exploitability + asset criticality), threat intelligence integration, metrics used for improvement.	Compliance 80-90%, asset coverage >90%, SLA compliance >85%.
Level 4: Measured / Proactive	Automated assessment, significant workflow automation, security integrated into SDLC, mature exception management, strong cross-team collaboration.	Compliance 90-95%, coverage >95%, 40%+ remediation automated, incident attribution <10%.
Level 5: Optimised / Predictive	Predictive management, extensive automation, security-by-design culture, prevention prioritised over remediation, self-healing capabilities.	Compliance >95%, coverage >98%, 60%+ automated, incident attribution <5%.

Progression Expectations

Level 1 to Level 2 typically takes 6 to 12 months and requires tool procurement, basic process documentation, and dedicated staffing. Level 2 to Level 3 takes 12 to 18 months and requires process maturation, threat intelligence integration, and a shift from compliance-focused to risk-focused thinking. Level 3 to Level 4 takes 18 to 24 months and requires significant automation investment, SDLC integration, and cultural change. Level 4 to Level 5 takes 24 months or more and requires advanced automation, organisation-wide security culture, and substantial investment. Smaller organisations may progress faster. Larger organisations typically require more time and resources at each stage.

Part 8: Scanning Method Comparison

Method	Coverage	Accuracy	False Pos.	Speed	Best For
Authenticated	High	High	Low	Moderate	Internal assets, servers
Unauthenticated	Medium	Medium	High	Fast	Perimeter testing, discovery
Agent-Based	Very High	Very High	Very Low	Continuous	Endpoints, mobile devices
Manual Testing	Targeted	Very High	Very Low	Slow	Complex applications, validation

Part 9: Acronyms and Abbreviations

A consolidated lookup for every acronym referenced across the series.

Acronym	Full Term
APT	Advanced Persistent Threat
AICPA	American Institute of Certified Public Accountants
AST	Application Security Testing
ASV	Approved Scanning Vendor (PCI DSS)
CAB	Change Advisory Board
CASB	Cloud Access Security Broker
CE	Cyber Essentials (UK)
CI/CD	Continuous Integration / Continuous Deployment
CIS	Center for Internet Security
CISO	Chief Information Security Officer
CMDB	Configuration Management Database
CSIRT	Computer Security Incident Response Team
CSPM	Cloud Security Posture Management
CVE	Common Vulnerabilities and Exposures
CVSS	Common Vulnerability Scoring System
CWE	Common Weakness Enumeration
DAST	Dynamic Application Security Testing
DDoS	Distributed Denial of Service
DEP	Data Execution Prevention
DoS	Denial of Service
EDR	Endpoint Detection and Response

EOL	End of Life
EPSS	Exploit Prediction Scoring System
GDPR	General Data Protection Regulation
HIPAA	Health Insurance Portability and Accountability Act
IaC	Infrastructure as Code
IAM	Identity and Access Management
IAST	Interactive Application Security Testing
IDS	Intrusion Detection System
IPS	Intrusion Prevention System
IR	Incident Response
ISMS	Information Security Management System (ISO 27001)
KEV	Known Exploited Vulnerabilities (CISA catalogue)
MFA	Multi-Factor Authentication
MITM	Man-in-the-Middle (attack)
MSSP	Managed Security Service Provider
MTTA	Mean Time to Acknowledge
MTTD	Mean Time to Detect
MTTR	Mean Time to Remediate
NIS2	Network and Information Security Directive 2 (EU)
NIST	National Institute of Standards and Technology (US)
NIST CSF	NIST Cybersecurity Framework
NVD	National Vulnerability Database
OWASP	Open Web Application Security Project
PCI-DSS	Payment Card Industry Data Security Standard

PII	Personally Identifiable Information
PoC	Proof of Concept
RCE	Remote Code Execution
REST	Representational State Transfer
ROI	Return on Investment
SANS	SysAdmin, Audit, Network, Security (institute)
SAST	Static Application Security Testing
SBOM	Software Bill of Materials
SCA	Software Composition Analysis
SDLC	Software Development Lifecycle
SIEM	Security Information and Event Management
SLA	Service Level Agreement
SOAR	Security Orchestration, Automation, and Response
SOC	Security Operations Centre
SOC 2	Service Organisation Control 2 (AICPA audit framework)
STIG	Security Technical Implementation Guide
TLS	Transport Layer Security
TSC	Trust Services Criteria (SOC 2)
VM	Vulnerability Management (or Virtual Machine, depending on context)
VPN	Virtual Private Network
VPR	Vulnerability Priority Rating (Tenable)
WAF	Web Application Firewall
XSS	Cross-Site Scripting

Part 10: Key Resources

Vulnerability Databases

National Vulnerability Database (NVD): https://nvd.nist.gov/

CISA KEV Catalog: https://www.cisa.gov/known-exploited-vulnerabilities-catalog

MITRE CVE: https://cve.mitre.org/

Exploit Database: https://www.exploit-db.com/

Scoring and Prioritisation

CVSS Calculator: https://www.first.org/cvss/calculator/

EPSS: https://www.first.org/epss/

Standards and Benchmarks

CIS Benchmarks: https://www.cisecurity.org/cis-benchmarks/

OWASP Top 10: https://owasp.org/www-project-top-ten/

NIST Cybersecurity Framework: https://www.nist.gov/cyberframework

Compliance Frameworks

ISO 27001:2022 (overview): https://www.iso.org/standard/27001

PCI DSS v4.0.1: https://www.pcisecuritystandards.org/document_library/

AICPA Trust Services Criteria (SOC 2): https://www.aicpa.org/resources/landing/system-and-organization-controls-soc-suite-of-services

NIST CSF 2.0 Reference Tool: https://csf.tools/framework/csf-v2-0/

NIS2 Directive (EUR-Lex): https://eur-lex.europa.eu/eli/dir/2022/2555

Cyber Essentials (NCSC): https://www.ncsc.gov.uk/cyberessentials/overview

Part 11: Compliance Framework Cross-Reference

This section maps VM programme activities to the requirements of six widely adopted compliance frameworks. The position taken throughout this series is worth restating: programmes that solve for security produce compliance as a natural byproduct. The evidence auditors need is a side effect of a programme that actually measures its performance and acts on what it finds.

This cross-reference is therefore intended as a lookup, not a to-do list. If your programme addresses the capabilities described in Episodes 1 through 5, the framework requirements below will largely be met. Where specific frameworks impose prescriptive requirements beyond good practice, those are called out explicitly.

Framework Overview

Framework	Scope	Obligation	VM Approach	Key Requirements
ISO 27001:2022	International	Voluntary (often contractual)	Risk-based; no fixed frequencies	Annex A Control 8.8
PCI DSS 4.0.1	Global (payment card)	Mandatory for card data handlers	Prescriptive; quarterly minimum, 30-day patch for critical/high	Requirements 6.3, 11.3, 11.4
SOC 2	Primarily US/Global SaaS	Voluntary (often contractual)	Risk-based; quarterly minimum common	CC3.2, CC7.1, CC7.2, CC8.1
NIST CSF 2.0	US/International	Voluntary (mandatory for US federal)	Risk-based; outcome-focused	ID.RA, PR.IP, DE.CM
NIS2	EU member states	Mandatory for essential/important entities	Risk-based; proportionate measures	Article 21(2)(d), (e)
Cyber Essentials	UK	Voluntary (mandatory for UK govt contracts)	Prescriptive; 14-day patch for critical/high	Patch Management control

VM Activity to Framework Mapping

This table maps core VM activities to their corresponding framework requirements. A tick indicates the framework explicitly requires or expects the activity.

VM Activity	ISO 27001	PCI DSS	SOC 2	NIST CSF	NIS2	Cyber Ess.
Asset Inventory	A.5.9	2.4, 9.4	CC6.1	ID.AM	Art 21(2)(a)	Asset mgmt control
Vulnerability Scanning (Internal)	A.8.8	11.3.1	CC7.1	ID.RA, DE.CM	Art 21(2)(d)	Patch mgmt control

Vulnerability Scanning (External)	A.8.8	11.3.2 (ASV)	CC7.1	ID.RA	Art 21(2)(d)	CE Plus only
Authenticated Scanning	A.8.8	11.3.1.2	CC7.1	ID.RA	Art 21(2)(d)	Not specified
Risk-Based Prioritisation	A.8.8, A.8.2	6.3.1	CC3.2, CC3.3	ID.RA	Art 21(2)(a)	Not specified
Patch Management	A.8.8, A.8.32	6.3.3	CC8.1	PR.IP	Art 21(2)(d)	14-day SLA (crit/high)
Compensating Controls	A.8.8	11.3.3	CC5.2	PR.IP	Art 21(2)(d)	Not specified
Penetration Testing	A.8.8, A.8.28	11.4.1-11.4.4	CC4.1	ID.RA	Art 21(2)(e)	CE Plus (external)
Exception / Risk Accept.	A.8.2	6.3.1	CC3.2	GV.RM	Art 21(2)(a)	Not specified
Remediation Verification	A.8.8	11.3.1, 11.3.2	CC7.1	PR.IP	Art 21(2)(d)	CE Plus (external)
Change Management	A.8.32	6.5.1	CC8.1	PR.IP	Art 21(2)(d)	Not specified
Incident Reporting	A.5.26	12.10	CC7.3, CC7.4	RS.CO	Art 23 (24h/72h)	Not specified
SBOM / Supply Chain	A.5.19-A.5.22	6.3.2	CC9.2	GV.SC	Art 21(2)(d)	Not specified
Documentation / Audit Trail	A.8.8	12.3.1	CC4.1	GV.OC	Art 21(2)(a)	Required for CE Plus
Metrics and Reporting	A.8.8, 9.1	6.3.1	CC4.1, CC4.2	GV.OC, ID.IM	Art 21(2)(f)	Not specified

VM Metrics to Compliance Evidence

The table below maps the metrics from Part 3 to the compliance evidence they satisfy. When preparing for audits, these metrics provide the quantitative evidence auditors expect.

Metric	ISO 27001	PCI DSS	SOC 2	NIST CSF

Asset Coverage	A.5.9 (inventory completeness)	11.3.1 (scope), 2.4	CC7.1 (monitoring scope)	ID.AM (asset management)
Scan Success Rate	A.8.8 (scanning effectiveness)	11.3.1 (scan completion)	CC7.1 (monitoring)	DE.CM (continuous monitoring)
MTTD	A.8.8 (timely identification)	6.3.1 (identification)	CC7.2 (anomaly detection)	DE.CM (detection timeliness)
MTTR (by severity)	A.8.8 (timely remediation)	6.3.3 (30-day patch SLA)	CC8.1 (change management)	PR.IP (protection processes)
SLA Compliance	A.8.8 (documented timelines)	6.3.1 (risk-ranked response)	CC7.1, CC8.1	PR.IP (protection processes)
Verification Rate	A.8.8 (confirm resolution)	11.3.1 (rescan requirement)	CC4.1 (control monitoring)	PR.IP (verification)
Exception Rate / Age	A.8.2 (risk acceptance)	6.3.1 (risk ranking)	CC3.2 (risk assessment)	GV.RM (risk management)
Recurrence Rate	10.2 (continual improvement)	6.3.1 (process effectiveness)	CC4.1 (ongoing evaluation)	ID.IM (improvement)
Risk Score Reduction	Clause 6.1 (risk treatment)	Not directly required	CC3.2 (risk evaluation)	GV.RM (risk strategy)
Critical Exposure Hours	A.8.8 (exposure management)	Not directly required	Not directly required	ID.RA (risk assessment)
Incident Attribution	A.5.26 (incident connection)	12.10 (incident response)	CC7.4 (incident response)	RS.AN (analysis)

Framework-Specific Prescriptive Requirements

Most frameworks are risk-based and allow organisations to determine appropriate frequencies, timelines, and methods. The exceptions are worth knowing because they impose fixed requirements that your programme must accommodate regardless of your own risk assessment.

PCI DSS 4.0.1

PCI DSS is the most prescriptive of the major frameworks for VM. External vulnerability scans must be performed quarterly by an Approved Scanning Vendor (ASV) and must resolve all vulnerabilities scored CVSS 4.0 or above to achieve a passing result. Internal scans must also run at least quarterly

and after any significant change. Since version 4.0, internal scans must use authenticated scanning (Requirement 11.3.1.2), with documented exceptions for systems that cannot accept credentials. Critical and high-risk vulnerabilities identified internally must be patched within one month of release (Requirement 6.3.3), with the risk ranking determined by the process defined in Requirement 6.3.1. Penetration testing is required at least annually (or semi-annually for service providers), and all exploitable vulnerabilities and security weaknesses found must be remediated and retested (Requirement 11.4.4).

Cyber Essentials

Cyber Essentials requires all high-risk and critical-severity patches to be applied within 14 days of release. Unsupported software (past end-of-life) must be removed from the environment or isolated and documented. Cyber Essentials Plus adds independent vulnerability scanning and on-site verification. While simpler than PCI DSS, the 14-day patching requirement is more aggressive than most organisations' default SLAs for high-severity findings and should be factored into remediation planning.

NIS2

NIS2 requires essential and important entities to implement risk management measures that are proportionate to the risk, including vulnerability handling and disclosure (Article 21(2)(d)) and policies on the effectiveness of cybersecurity measures (Article 21(2)(f)). Significant incidents must be reported to the national CSIRT or competent authority with an early warning within 24 hours and a full notification within 72 hours. Member state transposition varies, so organisations should verify local implementation requirements.

ISO 27001:2022 and SOC 2

Both frameworks are risk-based and do not prescribe specific scanning frequencies or remediation timelines. ISO 27001 expects organisations to determine appropriate frequencies through their own risk assessment, document their rationale, and demonstrate that the approach is effective. SOC 2 auditors expect evidence of regular scanning (quarterly is the common baseline), documented prioritisation, and demonstrated remediation follow-through. Both frameworks place significant weight on documented processes and evidence of continual improvement, which means your metrics programme (Part 3) directly serves your audit evidence needs.

NIST CSF 2.0

NIST CSF 2.0 is outcome-focused and does not prescribe specific controls. Its value for VM programmes is as an organising framework: the six functions (Govern, Identify, Protect, Detect, Respond, Recover) provide a structure for ensuring your programme covers the full scope of vulnerability management activities. The Govern function, new in version 2.0, emphasises that cybersecurity risk management must be integrated with enterprise risk management, which aligns directly with the series' position that VM decisions are business decisions.

Using This Guide

This guide is a reference, not a roadmap. If you are building or improving a VM programme, start with the episodes in the series for the operational reasoning and worked examples. Use this guide to look things up when you need them.

If you are just getting started, the five things that matter most are: an asset inventory you trust, regular scanning with authenticated credentials, prioritisation that accounts for more than CVSS, MTTR tracking segmented by severity, and documented processes that people actually follow.

Build from there. Add sophistication as your programme matures and as the organisation's appetite for risk-informed decision-making grows. Adapt everything to your context. These are guidelines, not rigid prescriptions. Your risk tolerance, available resources, regulatory requirements, and technology stack will shape what good looks like for your organisation.

On compliance: if your programme does the five things above consistently and well, most framework requirements take care of themselves. Part 11 maps programme activities to specific compliance controls for those who need to demonstrate alignment, but the principle from Episode 5 bears repeating. Solve for security, and compliance follows. Solve for compliance alone, and you will always be measuring activity rather than outcomes.

A basic VM programme that runs consistently will always outperform a sophisticated programme that exists only on paper.