radware

Guide

CISO's Guide to Protecting First and Patching Safely

A practical guide to reducing runtime exposure while permanent remediation is validated and deployed.





Contents

Introduction: The Security Problem Has Changed 3

Key Insights 4

What This Means for CISOs 5

Part I: The Vulnerability Window Is Shrinking for Attackers, Not Defenders 6

Part II: Why Your WAF Is Essential, but Not Enough 8

Part III: A Protect-First Operating Model 10

Part IV: How AI Xploit Shield Closes the Vulnerability Window 12

Part V: CISO Value and Program Measures 13

Part VI: CISO Readiness Checklist 14

Part VII: Building a More Resilient Vulnerability-Response Program 15

Introduction: The Security Problem Has Changed

AI is accelerating vulnerability discovery, attack development and exploitation. Yet the work required to fix a vulnerability safely still takes time. Security and application teams must validate findings, identify affected assets, test dependencies, schedule change windows and avoid disrupting critical services.

The result is a vulnerability window attackers can exploit.

This guide explains how CISOs can build a protect-first operating model that reduces exposure while teams work on permanent remediation. It also outlines how Radware AI Xploit Shield evaluates vulnerability findings and turns relevant, protectable vulnerabilities into tailored, customer-specific runtime protections at cloud scale.

CISO imperative: Don't leave your business exposed or rush an unsafe patch. You can protect first and patch safely.



Key Insights

The time between vulnerability discovery and safe remediation is becoming one of the most consequential gaps in enterprise cyber defense.

AI gives attackers greater speed, scale and automation. Defenders still face testing, ownership, change-control and availability constraints that cannot simply be automated away.

A resilient vulnerability-response program should therefore operate on two parallel tracks:

- **Immediate protection:** Reduce exploitability at runtime as soon as a relevant exposure is validated.
- **Permanent remediation:** Develop, test and deploy the appropriate patch or code change through established processes.

AI Xploit Shield supports this model by analyzing penetration-testing and vulnerability-scanner findings, separating noise from security-relevant issues and determining what is relevant, exploitable and protectable at Layer 7.

It first checks whether existing Radware protections already block the exploit. When new protection is needed, it automatically generates and validates a tailored, customer-specific signature, policy or supported configuration change and deploys it through the Radware cloud. These protections can remain in place while remediation is tested and deployed.

Core principle: Runtime protection and patching are complementary. The goal is to protect first, then patch safely.



What This Means for CISOs

AI Xploit Shield can help security leaders:

- Reduce the period during which known and validated vulnerabilities remain exploitable.
- Give application owners time to test fixes without accepting unnecessary business risk.
- Scale response across large vulnerability backlogs without depending on manual rule creation.
- Distinguish findings that are already covered, require new runtime protection or still require a code-level fix.
- Prioritize legacy, business-critical and operationally constrained environments where patching is especially difficult.
- Improve governance by linking vulnerability findings, runtime protection and permanent remediation in one operating process.



Part I: The Vulnerability Window Is Shrinking for Attackers, Not Defenders

Vulnerability management was built around an assumption that defenders would have enough time to discover, prioritize and remediate vulnerabilities before widespread exploitation. That assumption is under scrutiny.

AI can help identify flaws, reason across complex code paths, generate exploit logic and scale reconnaissance. The result is a faster path from disclosure or discovery to active attack. According to Radware's H1 2026 Global Threat Analysis Report, the average time from disclosure to exploitation has dipped into negative hours.

The defender's process remains constrained by business reality. A finding may be urgent, but the fix may affect dependencies, require regression testing, depend on a vendor update or need approval for a carefully scheduled change window.

For critical services, an untested patch can create its own outage risk.



Why Patching Takes Time

Organizations may need to:

- Validate whether the finding is accurate and applicable to the deployed environment.
- Map the vulnerability to affected applications, services, versions and owners.
- Assess exploitability, exposure and potential business impact.
- Develop or obtain the permanent fix.
- Test functionality, integrations, performance and rollback procedures.
- Deploy during an approved change window and confirm successful remediation.

These steps are necessary. Rushing them can cause outages, broken integrations, degraded performance or unexpected business disruption.



Where the Risk Is Most Acute

The vulnerability window is particularly difficult to manage for:

- **Business-critical applications** where downtime or failed changes carry material consequences
- **Legacy applications** that are difficult to modify, poorly documented or dependent on unsupported components
- **Third-party software** where the organization must wait for a vendor fix.
- **On-prem and hybrid environments** with complex dependencies and slower change processes
- **Regulated environments** that require formal testing, approvals and evidence before production changes
- **Operationally constrained systems** with infrequent maintenance windows

The real risk: A vulnerability may be known and prioritized, yet still remain exploitable for days, weeks or longer while remediation proceeds.



External Guidance Reinforces Urgency

As summarized in Radware's companion guide, UK National Cyber Security Centre and its Five Eyes partners note that AI is shortening the time between vulnerability discovery and exploitation.

Their guidance emphasizes risk-based prioritization, faster defensive action and special attention to legacy, unsupported and operationally constrained systems where updates may be delayed or unavailable.

Part II: Why Your WAF Is Essential, but Not Enough

A cloud WAF remains an essential layer of broad, continuous protection for applications and APIs. It inspects traffic and blocks known attack techniques, malicious behavior and emerging threats.

However, AI-accelerated vulnerability discovery creates a different requirement: turning specific findings about an organization's own applications into tailored runtime protection before permanent remediation is complete.



What Broad Protection Misses

A conventional WAF is designed to provide broad application protection. It is not designed to operationalize growing volumes of vulnerability findings and automatically turn each relevant finding into a validated, application-specific protection.

A WAF does not automatically:

- Ingest penetration-testing and vulnerability-scanner results
- Assess each finding in the context of the affected application
- Determine whether the finding is relevant and runtime-protectable
- Check whether existing protections already provide coverage
- Generate a tailored protection when additional coverage is required



Not Every Vulnerability Finding Should Become a WAF Rule

Some findings are irrelevant to the deployed environment. Some are already covered by existing protections. Others cannot be mitigated at Layer 7 and require a source-code, infrastructure or configuration fix.

The first challenge is therefore to filter the findings and identify which exposures can and should be addressed through runtime protection.



Manual Rule Creation Cannot Keep Pace

For a relevant and protectable finding, security specialists typically need to:

- Interpret the vulnerability
- Determine the appropriate mitigation
- Design a tailored rule or configuration change
- Validate its security efficacy
- Assess its impact on legitimate traffic
- Deploy it safely
- Monitor and maintain it until permanent remediation is complete

As AI increases the volume and speed of vulnerability discovery, this manual process becomes increasingly difficult to perform quickly and consistently at scale.



SAST, DAST and Patching Remain Essential

SAST, DAST and other testing tools improve visibility into application weaknesses. Patching, upgrading or changing the application code provides the permanent fix.

However, discovery alone not itself prevent exploitation. Safe remediation still requires validation, prioritization, development, testing and change control.

The missing capability is a scalable bridge between vulnerability discovery and permanent remediation. Organizations need a way to evaluate application-specific findings, identify those that can be mitigated at runtime and turn them into precise, tailored protection quickly and consistently.

Key distinction: AI Xploit Shield complements your WAF as an agentic force multiplier. It turns relevant, protectable findings into tailored runtime protections while permanent remediation proceeds.

Part III: A Protect-First Operating Model

A protect-first model separates the urgency of reducing exploitability from the deliberate work of permanent remediation. The two tracks begin together and remain linked through closure. Here are the steps involved in this model:

1. Ingest and Classify the Finding

Use penetration-testing results, vulnerability-scanner outputs and other supported reports. Remove noise and determine whether the finding is:

- Security-relevant
- Applicable to the protected application
- Potentially exploitable
- In scope for Layer 7 runtime protection
- Better resolved through a code-level, infrastructure or other remediation method

2. Prioritize by Business and Exploit Context

Prioritization should consider more than the scanner severity rating.

Relevant factors include:

- Internet exposure
- Application and asset criticality
- Data sensitivity
- Known exploitability
- Active exploitation
- Existing compensating controls
- Availability of a permanent fix
- Expected time to remediate

3. Check Existing Protection

Determine whether current application security controls already block the exploit before creating another rule. Reusing an existing validated control avoids unnecessary policy creation and reduces operational complexity.

4. Generate and Validate Targeted Protection

When additional protection is needed, create a tailored, customer-specific signature, policy or supported configuration change.

Confirm that it blocks the applicable exploit without unnecessarily disrupting legitimate behavior.

5. Enforce With Control

Deploy the protection through the Radware cloud using the selected enforcement mode.

Report mode can be used to assess the recommendation and potential impact before moving to blocking, depending on the organization's operating procedures and risk tolerance.

6. Patch Safely

Develop or obtain the permanent fix, complete the required testing and deploy it through established change-control processes.

The runtime protection reduces exposure during this period. It does not eliminate the need for permanent remediation.

7. Verify and Retire Deliberately

After remediation:

- Confirm that the patch or code change was deployed successfully
- Retest the vulnerability
- Preserve the required evidence
- Decide whether the runtime protection should remain in place
- Remove or adjust the protection only when it is safe to do so
- Close the related exception or remediation record

Governance requirement: Every runtime protection should have an associated finding, affected asset, owner, validation status and permanent remediation path.

Part IV: How AI Xploit Shield Closes the Vulnerability Window

AI Xploit Shield is Radware's AI-driven approach to virtual patching. It is automated, vulnerability-driven and tailored to each customer's actual exposures.

It focuses on what happens after discovery, translating supported penetration-testing and vulnerability-scanner findings into targeted runtime protection for web applications and APIs while permanent remediation continues.



How It Works

AI Xploit Shield can:

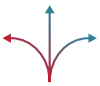
- Analyze penetration-testing, vulnerability-scanner and other supported assessment findings
- Classify findings as relevant and in scope, out of scope or not security-relevant
- Identify findings that require a code-level or other form of remediation
- Check whether existing Radware protections already block the exploit, avoiding unnecessary rule creation
- Generate a tailored, customer-specific signature, WAF rule, API-related protection or supported configuration change when additional protection is required
- Validate the protection
- Deploy it through the Radware cloud using the selected enforcement mode
- Maintain protection while the organization develops, tests and deploys the permanent fix



Designed for Scale

Large organizations manage extensive application portfolios, continuous testing results and persistent remediation backlogs.

AI-driven protection generation and cloud deployment help organizations respond at scale without increasing manual effort at the same rate.



Flexible Deployment

Offered as a standalone service, AI Xploit Shield can work with Radware Application Protection or another vendor's WAF. An existing Radware Cloud Application Protection subscription is not required.

Applications are onboarded to the Radware cloud for enforcement, either inline or through SecurePath.

In third-party WAF environments, AI Xploit Shield operates alongside that service and enforces its protections through the Radware cloud.

CISO outcome: A repeatable way to turn validated vulnerability intelligence into immediate, tailored and customer-specific protection while preserving a controlled remediation process.

Part V: CISO Value and Program Measures



Reduce Business Risk Without Increasing Change Risk

Runtime protection reduces pressure to choose between leaving a critical vulnerability exposed and deploying an insufficiently tested fix.

Application and security teams gain time to remediate deliberately while the organization maintains a stronger defensive posture.



Scale Scarce Expertise

Automated policy generation reduces dependence on specialists manually translating every finding into a protection rule.

This helps security teams respond consistently across larger application portfolios and higher finding volumes.



Improve Accountability

A protect-first workflow creates a clearer chain from finding to protection to remediation.

CISOs can ask not only how quickly teams patch, but also how quickly exposure is reduced and whether temporary protections remain governed until closure.

CISOs need visibility into actual exploitability, runtime protection and permanent remediation.

Part VI: CISO Readiness Checklist

Use these questions to assess whether your organization can close the vulnerability window at the speed the threat environment now demands.



Visibility and Prioritization

- ☑ Can we map high-risk findings to the affected application, service, environment and owner?
- ☑ Do we prioritize using business criticality and exploit context, not severity scores alone?
- ☑ Can we distinguish findings that are irrelevant, already covered, runtime-protectable or dependent on a code-level fix?
- ☑ Can we identify vulnerabilities that cannot be patched within the required risk window?



Immediate Protection

- ☑ Can we create targeted runtime protection without waiting for the permanent fix?
- ☑ How much specialist time is required to translate a finding into an enforceable policy?
- ☑ Can we validate protection before enforcement and monitor its legitimate-traffic impact?



Scale and Coverage

- ☑ Can the process support large application portfolios and continuous vulnerability findings?
- ☑ Does it work across business-critical, legacy, third-party and hybrid environments?
- ☑ Can it coexist with our existing application-protection investments?



Governance and Closure

- ☑ Is each temporary protection tied to an accountable owner and remediation plan?
- ☑ Can we demonstrate when exposure was reduced, not only when the patch was completed?
- ☑ Do we verify permanent remediation before retiring or adjusting runtime protection?

Decision point: If your organization cannot consistently protect critical exposures before permanent remediation is complete, the vulnerability window remains an unmanaged business risk.

Part VII: Building a More Resilient Vulnerability-Response Program

The acceleration of AI-driven vulnerability discovery and exploitation requires more than faster patching targets.

CISOs need a response model that acknowledges operational reality while reducing exposure as early as possible.

AI Xploit Shield adds a tailored, customer-specific runtime protection layer between discovery and permanent remediation. By evaluating vulnerability findings, reusing existing protections where appropriate and automatically creating targeted policies or configuration changes when needed, it helps organizations protect first and patch safely.

The result is not a replacement for secure development, vulnerability testing, application protection or disciplined patch management. It is the connective capability that helps these programs work together at the speed and scale demanded by the new threat environment.

Final takeaway: The strategic question is no longer only, “How fast can we patch?” It is also, “How quickly can we make a validated vulnerability non-exploitable while we patch safely?”

Selected External Guidance

- Five Eyes, *The AI Shift in Cyber Risk: Why Leaders Must Act Now*
- UK National Cyber Security Centre, *Preparing for a Vulnerability Patch Wave*

Learn how you can protect first and patch safely.

Contact Radware



This document is provided for information purposes only. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law. Radware specifically disclaims any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. The technologies, functionalities, services, or processes described herein are subject to change without notice.

© 2026 Radware Ltd. All rights reserved. The Radware products and solutions mentioned in this document are protected by trademarks, patents and pending patent applications of Radware in the U.S. and other countries. For more details, please see: https://www.radware.com/LegalNotice/. All other trademarks and names are property of their respective owners.

RWI6452 | 09.22.2026