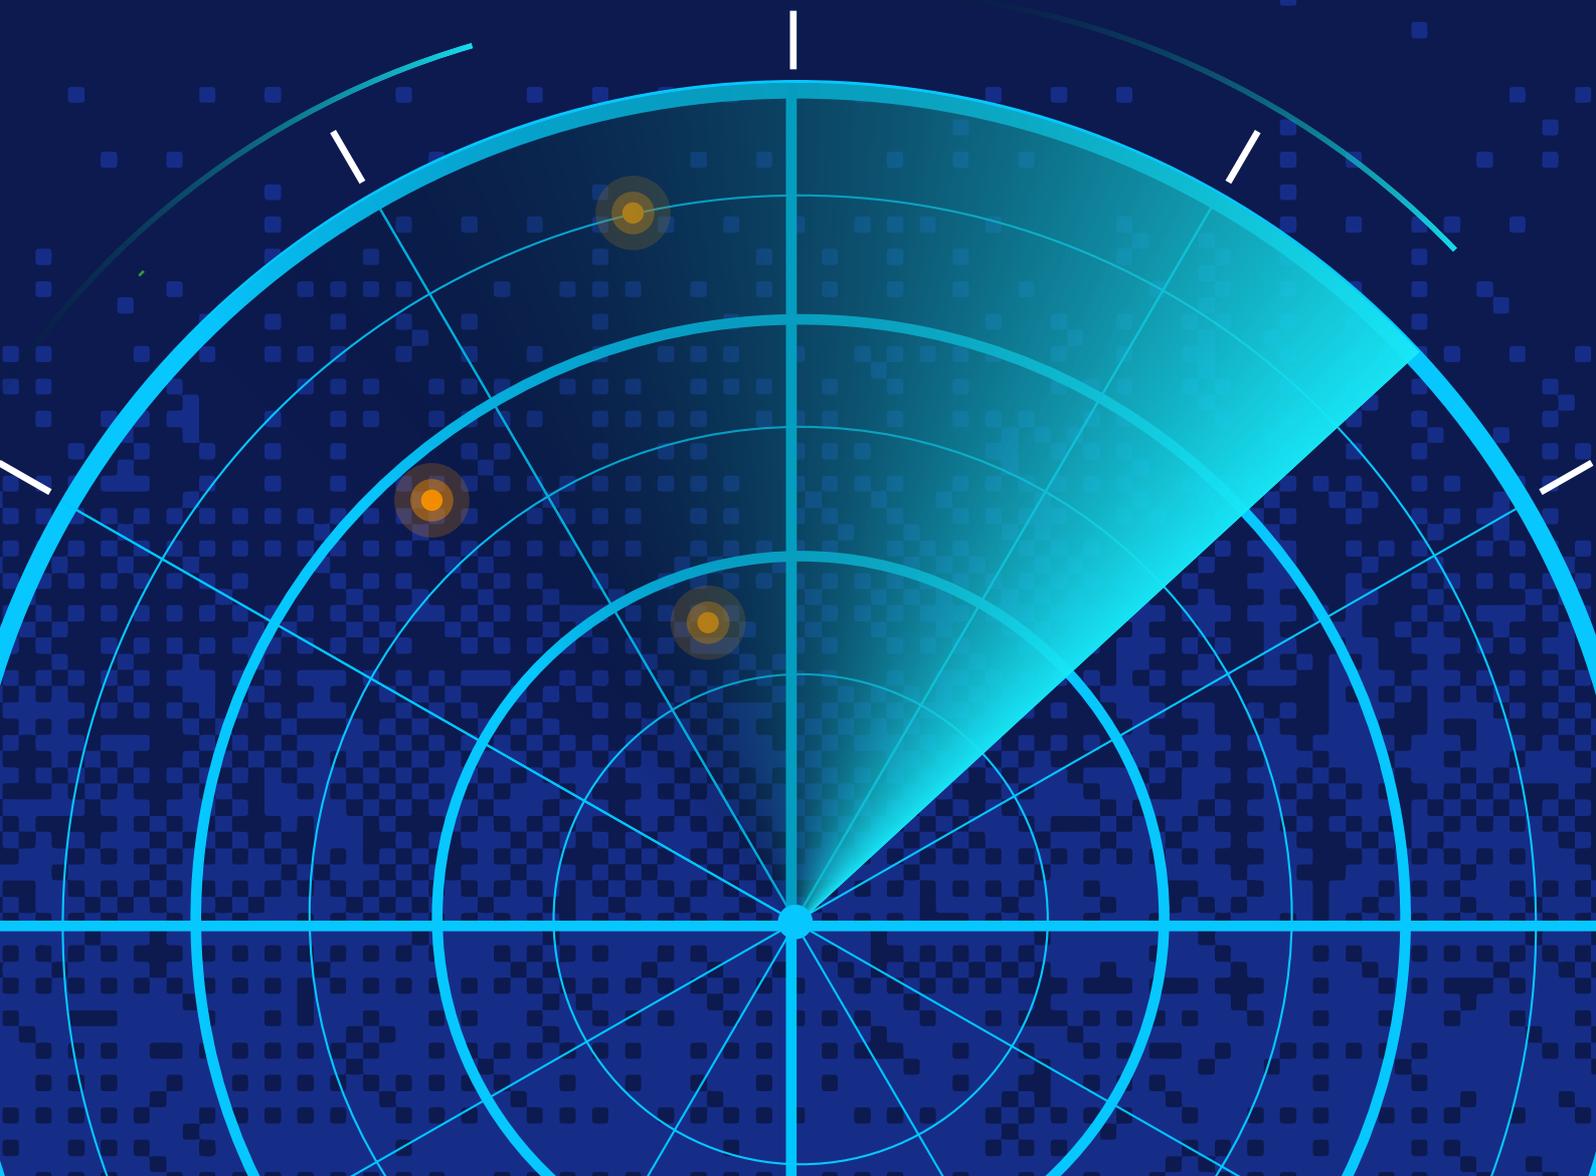


# Vulnerability management *for* platform engineers





# Introduction

Platform engineers face a stark reality. Vulnerability management has become an unsustainable burden that consumes engineering time without reducing systemic risk. You own the foundation that every application builds upon. When that foundation contains vulnerabilities, every downstream service inherits the risk, creating cascading liability across your organization.

The traditional "shift left" approach - pushing security responsibilities onto developers - has reached its limits. Developers now spend 19% of their weekly hours on security tasks, effectively losing one full workday to security toil. Vulnerability management alone costs organizations an average of \$1.9M annually, while vulnerability drag - the hidden productivity tax of constant remediation - can exceed \$31M per year.

Vulnerability management alone  
costs organizations an average of

*\$1.9M annually*

The thesis: you must shift security down into the platform itself, building secure-by-design systems that automate detection, enforcement, and remediation without developer intervention.

With 40,000+ new CVEs published in 2024 (a 38% year-over-year increase) and 84% of codebases containing at least one vulnerability, manual approaches cannot scale. The time for platform-embedded security is now.

This whitepaper provides the strategic framework and implementation roadmap for transforming vulnerability management from reactive firefighting to invisible, automated platform capability. The following sections define what security platform engineering means, why it matters for organizational velocity and risk, and how to implement it through concrete practices and cultural alignment.



# Key takeaways

Before diving into details, here are the essential insights you should take from this whitepaper:

→ **Vulnerability management is a platform capability, not a developer responsibility**

Shifting security down into the platform eliminates toil while improving protection

→ **The CVE doom cycle cannot be solved by working harder**

It requires architectural change that makes secure paths the default paths

→ **Four core capabilities define secure-by-design platforms**

Automated image hardening, policy-as-code enforcement, pre-approved service templates, and continuous secret rotation.

→ **Platform teams that embed security create a ROI loop**

Reduced friction drives developer adoption, adoption justifies investment, investment enables further automation.

→ **Cultural alignment between platform and security teams is non-negotiable**

Shared ownership and security metrics in platform KPIs transform adversarial relationships into collaborative partnerships.

These takeaways frame the strategic shift. The following sections provide the definitions, frameworks, and implementation guidance to make them actionable.



# What is vulnerability management *for* platform engineers?

Understanding vulnerability management for platform engineers requires distinguishing it from both traditional security operations and the "shift left" movement.

## Defining security platform engineering

Security platform engineering is an emerging discipline that sits at the intersection of platform engineering and security operations. It represents a fundamental shift in how organizations approach vulnerability management.

Security platform engineering is the practice of embedding security controls, policies, and automated remediation directly into the internal developer platform (IDP), making security an invisible property of the infrastructure rather than an

additional task for developers. This approach encompasses automated scanning integration in CI/CD pipelines, policy-as-code enforcement at deployment boundaries, hardened base images and golden paths, and continuous monitoring with automated response.

The platform team owns the security posture of the foundation. Application teams inherit that security by default when they build on platform-provided primitives.



# Shift left vs. shift down

The "shift left" movement has dominated security conversations for years, but we are increasingly discovering its limitations.

Understanding the distinction between shifting left and shifting down is essential for effective vulnerability management.

Shift left asks developers to catch vulnerabilities earlier in the development cycle - running scans, reviewing findings, applying

patches - which adds cognitive load and context-switching to already-burdened development teams. Shift down removes that burden entirely by making the platform responsible for security; developers use secure primitives (hardened images, pre-approved templates, automated secret rotation) without needing to understand the underlying security mechanisms.



**The practical difference: shift left means developers run security tools; shift down means the platform runs security tools and developers never see the findings unless action is required.**

This is not merely a semantic distinction. It represents a change in domain ownership where security becomes part of how the platform works, not something teams must remember to do.





# Scope *and* assumptions

Security platform engineering is not a universal solution - it requires specific organizational conditions and infrastructure maturity. This section establishes the boundaries and assumptions that shape implementation. This whitepaper assumes you have:

- **Established CI/CD pipelines with automated build and deployment processes**
- **Container-based workloads (Kubernetes or similar orchestration)**
- **An internal developer platform or the intent to build one**
- **At least one dedicated platform engineer or team**

Organizations at earlier maturity stages should focus first on pipeline automation and container adoption before attempting comprehensive security platform engineering.

The scope covers vulnerability management for container images, dependencies, and infrastructure configurations. It does not cover

application-level security testing (penetration testing, threat modeling) or security operations center (SOC) functions. Platform teams own the security of shared infrastructure and golden paths; application teams retain responsibility for business logic vulnerabilities and data handling practices.



This approach works best when platform adoption is voluntary or incentive-driven. Mandated platforms with poor developer experience will not benefit from embedded security as an adoption lever.

**Mandy Hubbard**

CHAINGUARD

 **Chainguard**

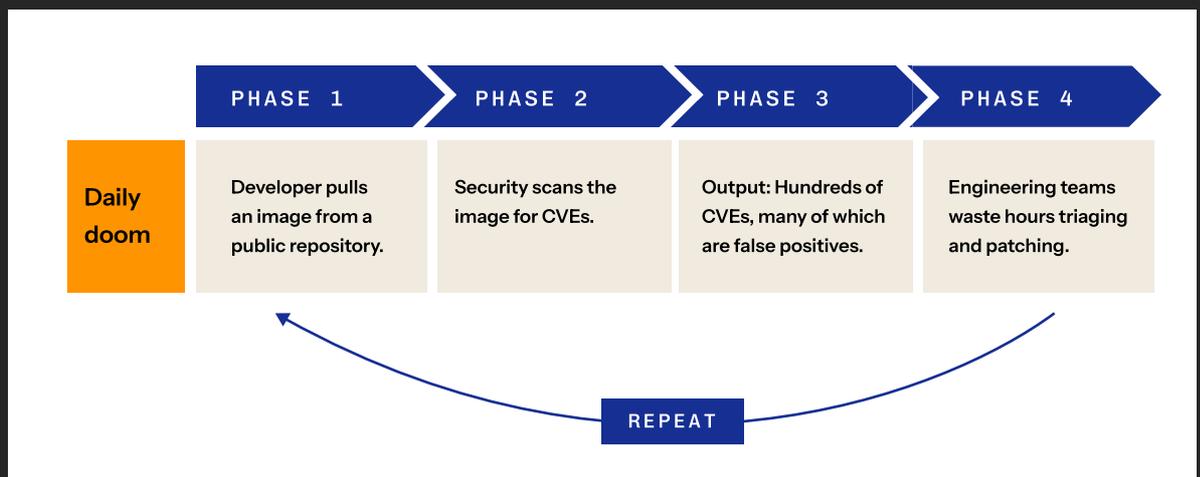


# Why vulnerability management *for* platform engineers matters

The strategic importance of platform-embedded vulnerability management becomes clear when examining both the costs of the current state and the organizational friction it creates. This section makes the business case through concrete scenarios and quantified impact.

## The CVE doom cycle

Most platform engineers recognize this scenario immediately – it plays out daily in organizations without platform-embedded security.



A developer pulls a Python base image from a public registry to build a new service. The security scanner flags 287 vulnerabilities, including 52 rated high or critical severity. The developer spends four hours triaging findings, determining which vulnerabilities are exploitable in their context, documenting exceptions for false positives, and applying patches where possible.



The next day, a new CVE is published affecting the same base image. The cycle repeats, and the developer loses another half-day to remediation instead of feature development.

Multiply this across dozens of services and hundreds of developers, and your

organization hemorrhages engineering capacity to repetitive, low-value security toil. This is not a failure of individual diligence, it's a systemic problem that requires system wide solutions.

## Quantified organizational impact

Beyond individual developer frustration, poor vulnerability management creates measurable organizational costs that compound over time.

Vulnerability management alone costs organizations an average of \$1.9M annually in direct remediation effort, tooling, and compliance documentation.

Vulnerability drag, the hidden productivity tax of constant context-switching, delayed releases, and accumulated security debt, costs organizations up to

*\$31M  
per year*

The average data breach costs [\\$4.88M](#). High-profile vulnerabilities like Log4Shell resulted in [7,386 unique incidents in Q4 2023 alone](#), averaging \$90,000 per exposure.

At the same time, fewer than 30% of platform teams report successful voluntary adoption, and security friction is a major driver of developer resistance. When security is embedded effectively, platform teams can create an ROI flywheel: reduced friction boosts adoption, adoption justifies continued investment, and that investment enables further automation.



# Decision criteria *and* trade-offs

Implementing security platform engineering requires deliberate choices about tooling, automation scope, and organizational change. This section provides the decision framework for evaluating options and understanding trade-offs.

## Evaluation framework

Not all security automation delivers equal value, so teams need clear criteria to decide where to invest limited resources and engineering capacity. Key evaluation dimensions include integration depth, meaning whether the capability embeds into existing CI/CD pipelines and IDP workflows or requires separate tools and context switching; automation completeness, assessing whether it goes beyond detection to include automated remediation and policy enforcement; developer experience

impact, evaluating whether it reduces cognitive load or introduces new gates and approvals; feedback loop quality, considering whether insights appear where developers already work, such as pull requests or dashboards, rather than in isolated reporting silos; and compliance evidence generation, measuring whether audit artifacts are produced automatically or require manual documentation. Evaluating investments against these criteria helps identify capabilities that deliver the highest platform ROI.

## Trade-offs analysis

Every implementation choice involves trade-offs. Understanding them helps teams make informed decisions instead of reacting after problems occur.

Automation involves a balance between scope and implementation complexity. Comprehensive automation such as image rebuilding, policy enforcement, and secret



rotation can deliver strong long-term value, but it requires a higher upfront investment. A practical approach is to start with high-impact, low-complexity capabilities and expand incrementally.

At the same time, security strictness must also be balanced with developer friction.



**Blocking all images that contain any CVE maximizes security, but it can create massive deployment bottlenecks. Risk-based thresholds that block critical and high issues while warning on medium ones often provide a better balance between protection and delivery speed.**

Another consideration is centralized control versus team autonomy. Platform-enforced policies improve consistency, but they may not cover legitimate edge cases. Instead of absolute blocks, provide controlled escape hatches with clear audit trails.

Finally, the build-versus-buy decision matters. Open-source tools such as Trivy, Grype, and Syft offer flexibility

and cost savings, but they take effort to integrate and maintain. Commercial platforms can be quicker to adopt, but they increase vendor dependency.

The right position on each trade-off depends on organizational maturity, risk tolerance, and platform team capacity. There is no universal right answer.



# How to apply vulnerability management *for* platform engineers

It is crucial to understand a practical maturity journey for security platform engineering. This is a battle tested seven-step roadmap that sequences capabilities from baseline assessment to continuous trust and compliance automation. Alongside a playbook of highest-impact practices, focusing on hardened base images, policy-as-code, automated SBOMs, developer-native insights, and measurable security outcomes.

## Seven-step implementation roadmap

Implementing security platform engineering is a maturity journey, not a single initiative. This roadmap sequences capabilities for incremental adoption with clear ownership boundaries.

STEP  
01

### Baseline

**Map current risk and workflow** - audit delivery flow from commit to runtime, identify friction points, generate SBOMs with Syft or Grype, classify risk ownership. Owner: Platform team. Cadence: One-time assessment, quarterly refresh.

STEP  
02

### Guardrails

**Build policy-as-code enforcement** - adopt OPA, Kyverno, or Conftest; define baselines (no unscanned images, block CVEs  $\geq 7$ ); integrate into CI/CD. Owner: Platform team with security input. Cadence: Policies reviewed monthly.

STEP  
03

### Supply chain

**Secure the foundation** - centralize image management in registries with automatic rescanning; harden and sign base images; automate rebuilds on upstream CVE disclosure. Owner: Platform team. Cadence: Continuous automation.



STEP  
04

### IDP integration

**Embed in developer workflows** - add secure templates (Terraform, Helm) enforcing best practices; bake in scanning hooks; surface insights in dashboards and PR comments. Owner: Platform team. Cadence: Template updates quarterly.

STEP  
05

### Secrets

**Automate credential management** - integrate Vault, AWS Secrets Manager, or Doppler; enforce short-lived tokens; add secret scanning to pipelines. Owner: Platform team with security review. Cadence: Rotation policies continuous.

STEP  
06

### Culture

**Embed shared ownership** - position platform as enabler not enforcer; partner with security to codify ownership; add security metrics to platform KPIs. Owner: Platform and security teams jointly. Cadence: Quarterly alignment reviews.

STEP  
07

### Continuous trust

**Iterate toward compliance invisibility** - track reduction in manual approvals and CVE resolution time; automate compliance evidence generation; evolve toward audits generated from pipeline data. Owner: Platform team. Cadence: Continuous



# Success metrics and start now plan

Measuring progress and taking immediate action transforms strategy into results. These metrics and steps provide the accountability framework.

## SUCCESS METRICS (TRACK DIRECTION OVER TIME)



**Mean time to patch:** Target decreasing trend as automation handles routine remediation



**Percentage of vulnerabilities auto-remediated:** Target increasing toward 80%+ for routine CVEs



**CVE backlog trend:** Target decreasing as platform capabilities mature



**Developer time on security tasks:** Target decreasing from 19% baseline



**Platform adoption rate:** Target increasing as security friction decreases



**Compliance evidence generation time:** Target decreasing toward near-zero manual effort

**START NOW: SEVEN STEPS TO BEGIN**

- 01** — **First step (this week): Generate SBOMs for your three highest-traffic services using Syft or Gype - establish baseline visibility into your current dependency landscape**
- 02** — **Audit your current delivery flow and identify the three highest-friction security touchpoints.**
- 03** — **Deploy one policy-as-code rule in warning mode (e.g., flag images with critical CVEs without blocking).**
- 04** — **Configure your CI/CD pipeline to surface scan findings in PR comments for one team.**
- 05** — **Schedule a joint session with your security team to define shared ownership boundaries.**
- 06** — **Establish baseline measurements for the six success metrics above.**
- 07** — **Identify one base image used across multiple services and create a hardened, signed alternative.**

With these practical starting points, you can begin practising the incredible value generating discipline of security platform engineering.



# Conclusion

Vulnerability management is no longer a problem that can be solved with more alerts, more dashboards, or more developer effort. The scale and velocity of modern software delivery demand a structural response, and that response sits squarely with platform teams. By shifting security down into the platform, organizations replace reactive firefighting with systems that make secure behavior the default, not an exception.

This whitepaper has shown that security platform engineering is not about perfection or zero vulnerabilities. It is about leverage. Hardened base images eliminate entire classes of risk at once. Policy-as-code replaces inconsistent human judgment with predictable

enforcement. Automated SBOMs turn uncertainty into fast answers. Developer-native feedback loops close issues without friction. Measurable security outcomes transform security from a cost center into a visible platform capability.

Most importantly, embedded security changes organizational dynamics. Platform teams move from gatekeepers to enablers. Security teams move from blockers to partners. Developers regain focus on delivering customer value while inheriting stronger protection by default.

The journey outlined here is incremental and pragmatic. You do not need to solve everything at once.



**Start where the leverage is highest, automate relentlessly, and measure what matters.**

In a world of accelerating CVEs and constrained engineering capacity, the winning organizations will not work harder at vulnerability management. They will design it into the platform and make it disappear.