When Trust Breaks

Open Source Security Predictions For 2026

Executive Summary

Open-source software underpins almost every modern application. In 2025, it also became one of the most consistently exploited attack surfaces.

The shift was not driven by more vulnerabilities alone, but by how attackers exploited trust across software supply chains.

Supply-chain campaigns such as **Shai-Hulud 2.0** and cyber disruptions affecting UK public-sector organisations showed that attackers now prioritise credential harvesting, maintainer compromise, and build pipeline access, often delaying visible damage until trust can be exploited at scale.

What Changed in 2025

- Vulnerability disclosures reached record levels, while exploit timelines collapsed to days or hours
- **Most impactful incidents exploited known issues, stolen credentials, or trusted update paths**
- Public sector disruptions exposed the fragility of shared platforms and third-party dependencies
- Traditional AppSec models failed to respond at machine speed

As organisations enter 2026, the most important question is no longer whether another supply-chain incident will occur.

When trust breaks, will your organisation be ready to move faster than the attacker?

The Core Shift

Security in 2026 will no longer be measured by how many vulnerabilities are detected, but by how well organisations **preserve trust, limit blast radius, and prove control in real time**.

Key Predictions for 2026

- Application security converges into full software supply-chain security
- Regulation becomes the primary driver of security investment
- Automated fixing replaces manual remediation as the default
- Provenance and integrity become mandatory, not optional
- Developer experience becomes a deciding security control
- Security data becomes legal and financial evidence

2025 Reality	2026 Expectation
Periodic scans	Continuous enforcement
Static SBOMs	Real-time, actionable SBOMs
Manual patching	Automated remediation
Developer friction	Security embedded in workflows
Audit preparation	Continuous proof of resilience

The State of Open Source Security In 2025

By the end of 2025, open-source software had become both indispensable and increasingly fragile. Across the US and UK, nearly every production application depended on open-source components, often in volumes that far exceeded what most organisations could realistically inventory or govern.

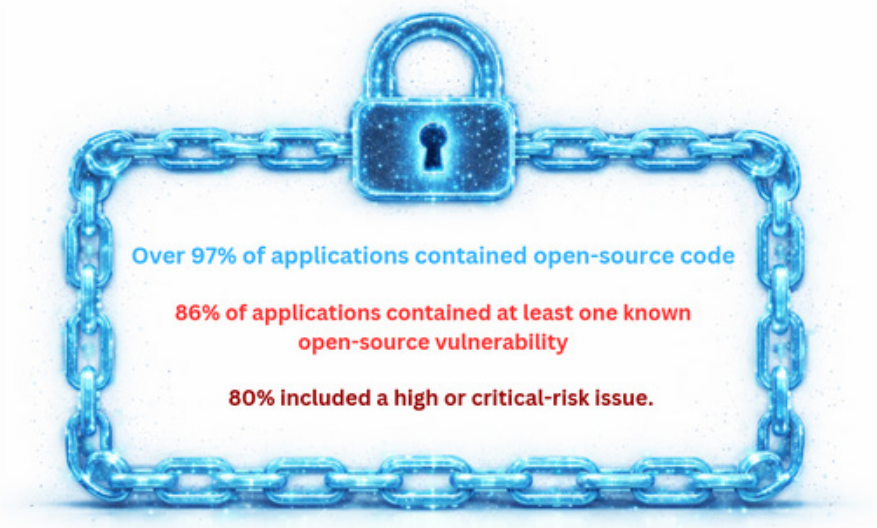

The total number of disclosed software vulnerabilities reached another record in 2025, with **global CVE counts projected to exceed 40,000 for the year**.

Open-source components accounted for a growing share of these disclosures, driven by deeper dependency trees and automated vulnerability discovery.

Key characteristics of the 2025 vulnerability landscape:

- Open-source vulnerability disclosures grew significantly faster than overall open-source adoption
- The majority of exploitable issues were not new, but previously disclosed vulnerabilities that remained unpatched
- Transitive dependencies accounted for roughly two-thirds of vulnerable components, making ownership and remediation unclear
- This created a persistent exposure problem: vulnerabilities existed, fixes often existed, but deployment lagged behind attacker timelines.

Exploitation moved faster than remediation

One of the most critical shifts in 2025 was the collapse of the remediation window. Threat intelligence and government reporting showed that:

- The average time to exploit after disclosure dropped to **around five days**
- Nearly 30% of exploited vulnerabilities were **weaponised within 24 hours**
- Hundreds of vulnerabilities were added to CISA's Known Exploited Vulnerabilities (KEV) catalogue during the year, with open-source libraries repeatedly represented

By contrast, remediation timelines remained slow:

- Median patch adoption for open-source vulnerabilities ranged from several weeks to multiple months
- In npm and frontend ecosystems, full downstream remediation often **took 4 to 11 months**
- Even critical vulnerabilities frequently remained deployed long after fixes were available
- The practical outcome was predictable. Attackers consistently outran defenders, not because of superior sophistication, but because the system itself could not respond at machine speed.

Why traditional AppSec models broke down

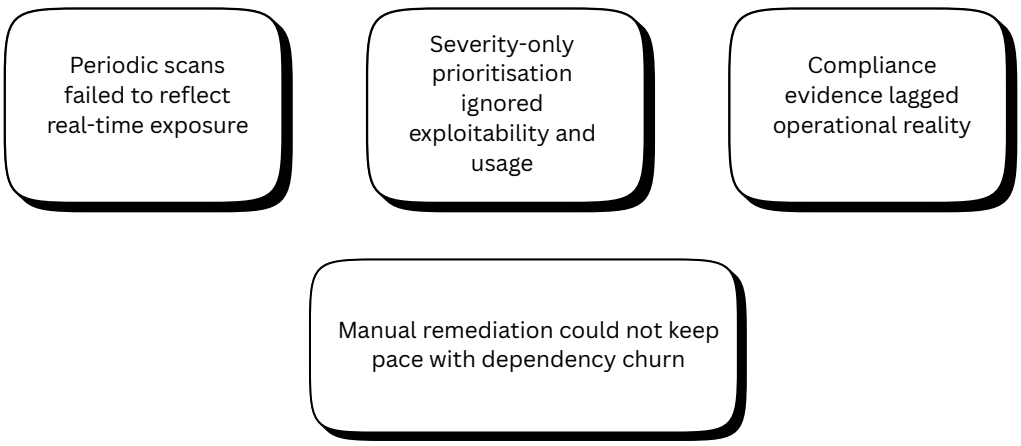

The Supply-Chain Shift: From Vulnerabilities to Trust

In 2025, the most damaging security incidents were no longer driven by unknown vulnerabilities. They were driven by **abuse of trust embedded in software supply chains**.

Attackers increasingly avoided traditional exploits and instead targeted:

- Maintainer accounts
- Package publishing permissions
- CI/CD credentials
- Automated install and build processes

Once trust was compromised, malicious code executed through normal development workflows, often without triggering conventional security controls.

From breaking software to inheriting trust

Modern software delivery relies on implicit trust:

- Dependencies update automatically
- Install scripts execute without review
- CI systems authenticate using long-lived tokens

Supply-chain attacks in 2025 exploited these mechanics. No misconfiguration or zero-day exploit was required. Attackers simply operated **inside trusted systems**.

This marked a fundamental shift in risk. Security failures moved upstream, where visibility and enforcement were weakest.

npm exposed the scale of the problem

The npm ecosystem became the clearest stress test of this shift.
Key observations from 2025:

- Malicious package publication accelerated sharply
- High-profile, well-maintained packages were targeted
- Compromised updates propagated rapidly through dependency trees
- Credential harvesting became a primary objective

These incidents demonstrated how a single trust failure could impact **thousands of downstream projects**.

Why Visibility Alone was Insufficient

SBOM adoption increased significantly in 2025, driven by regulatory and procurement pressure. However, visibility did not prevent execution.
SBOMs answered what was present, but not:

- Whether a package could be trusted at runtime
- Whether its provenance had changed
- Whether it posed an immediate execution risk

Without enforcement and remediation, awareness did not reduce exposure.

The strategic takeaway

By the end of 2025, supply-chain security could no longer be treated as a subset of application security.

It became a systemic trust problem, spanning code, identity, automation, and governance. This shift explains why 2026 security strategies will prioritise provenance, integrity, automated fixing, and real-time proof of resilience over traditional vulnerability counting.

Public Sector Disruption and the Shai-Hulud 2.0 Wake-Up Call

By late 2025, software supply-chain risk moved from boardroom discussion to operational disruption.

A series of cyber incidents affected multiple London councils, forcing public services offline and exposing weaknesses across shared digital infrastructure.

While investigations avoided premature attribution, one fact was clear: **software dependencies and shared platforms had become single points of failure.**

Why the Public Sector Became a Stress Test

Public sector organisations did not fail because of negligence. They were exposed because of structural realities that mirror large enterprises.

Structural Reality	Why It Amplifies Risk
Shared platforms	One compromise cascades across agencies
Mixed legacy and modern systems	Inconsistent security coverage
Heavy third-party reliance	Limited visibility into dependency risk
Slow remediation cycles	Attackers move faster than fixes

Shai-Hulud 2.0: The Defining Supply-Chain Incident of 2025

Weeks before the London council disruptions, the JavaScript ecosystem experienced a second wave of the Shai-Hulud supply-chain attack.

This attack did not exploit a vulnerability. It exploited trust.

Attackers compromised maintainer accounts and published trojanised versions of legitimate npm packages. When developers or CI systems installed these packages, malicious code executed automatically as part of standard workflows.

Once active, the malware harvested high-value credentials, including npm tokens, GitHub access keys, CI/CD secrets, and cloud credentials.

Crucially, the attack impacted well-maintained, widely trusted projects, demonstrating that maturity and best practices alone were no longer sufficient defenses.

From Developer Tooling to Enterprise Exposure

Shai-Hulud 2.0 demonstrated how quickly supply-chain compromise escapes the boundaries of developer tooling.

A single infected dependency or leaked token could:

- Move laterally across repositories
- Poison CI/CD pipelines
- Enable cloud and production access
- Remain undetected until damage surfaced

Within days, the campaign had spread into tens of thousands of repositories, showing how dependency scale magnifies even small trust failures.

The Pattern That Emerged in 2025

Taken together, the public sector incidents and Shai-Hulud 2.0 revealed a consistent attacker model.
Supply-chain attacks now unfold in phases:

1. Initial trust compromise
2. Credential harvesting and environment mapping
3. Deferred, large-scale exploitation

This phased approach explains why organisations entering 2026 can no longer ask simply whether they were affected.

Open Source Security Predictions for 2026

By the end of 2025, the direction of travel was clear. Open-source security was no longer evolving incrementally. It was undergoing a structural reset.

The events of the past year — rising vulnerability volumes, shrinking exploit windows, public-sector disruption, and supply-chain attacks such as Shai-Hulud 2.0 — exposed the limits of existing security models. In 2026, organisations will be forced to adapt, not by preference, but by necessity.

The following predictions describe the forces that will shape open-source and software supply-chain security over the next 12–18 months.

Prediction 1: AppSec Will Converge Into Full Software Supply-Chain Security

In 2026, application security will no longer be viable as a collection of point tools. Market forces will drive a clear convergence toward end-to-end software supply-chain security. This includes:

- Open-source dependencies
- Containers and images
- Infrastructure-as-code
- CI/CD pipelines
- Build and publishing credentials

This convergence is not theoretical. The attacks of 2025 showed that vulnerabilities, credentials, automation, and delivery pipelines are exploited together.

Security tools that only detect issues will be insufficient. Platforms must detect, prioritise, and fix across the full delivery lifecycle, without fragmenting developer workflows.

> The recent supply-chain attacks were structured primarily to collect intelligence: credentials, secrets, pipeline metadata, and trust relationships. Even organisations that did not see compromised packages could already be exposed. Stolen tokens and credentials allow attackers to return later with legitimate access. The next phase is expected to impact thousands of packages simultaneously, using trust already harvested.

Bruno Bossola, CTO at Meterian

Prediction 2: Automated Fix-First Security Becomes the Default

By 2026, automated remediation will move from “advanced capability” to baseline expectation, with adoption accelerating toward 75% among mature organisations.
This shift is driven by simple economics. Manual remediation cannot keep pace with:

- Dependency churn
- Exploit speed
- Regulatory remediation timelines

Automated fixing operates as an invisible security layer, reducing risk without slowing development. Tools that create friction, noise, or workflow disruption will increasingly be rejected.

Market Reality: Security that requires constant human intervention does not scale. Security that fixes silently does.

Prediction 3: Compliance Evolves Into Financial and Legal Assurance

In 2026, demand for open-source security will increasingly come from CFOs and General Counsel, not just security teams.

Regulatory exposure, supplier accountability, and audit scrutiny will force organisations to demonstrate continuous, auditable proof of resilience. Fix-first security will no longer be viewed purely as risk reduction, but as a financial assurance mechanism.

Security data becomes evidence:

- Evidence of due diligence
- Evidence of supplier accountability
- Evidence of operational control

Organisations that can prove resilience in real time will gain economic and competitive advantage over those relying on manual, compliance-only processes.

Prediction 4: Regulation Will Drive Buying Decisions More Than Security Maturity

In 2026, regulatory pressure will outweigh voluntary security maturity as the primary driver of investment.

Key forces include:

- UK Cyber Security & Resilience Bill expanding mandatory responsibility not just of owned and operated systems but also that of their own suppliers
- EU Cyber Resilience Act (CRA) and NIS2 remediation and reporting expectations
- Continued SBOM and secure-by-design mandates across the EU, UK, US, and other jurisdictions

Organisations will be required to demonstrate:

- Timely remediation
- Secure development practices
- Provenance and integrity of software components

> Security platforms that cannot produce machine-verifiable compliance evidence will increasingly fail procurement and audit requirements. Best-effort security will not satisfy fixed remediation deadlines. Automation becomes mandatory.

Vivian Dufour, CEO at Meterian

Prediction 5: SBOMs Will Be Everywhere, and Still Not Enough

By 2026, SBOM generation will be nearly universal. However, SBOMs alone will not prevent supply-chain compromise.

SBOMs answer what exists. They do not answer:

- Whether a component can be trusted right now
- Whether its provenance has changed
- Whether it introduces execution risk

Buyers will demand real-time **SBOMs linked directly to enforcement and remediation**, not static inventories created for audits.

Prediction 6: Provenance and Integrity Become Non-Negotiable

The most damaging attacks of 2025 demonstrated that trust is dynamic. In 2026, organisations will be expected to continuously verify:

- Who published a package
- How it was built
- Whether it has been altered
- Whether it should be trusted at runtime

This requirement will extend beyond libraries to include containers, IaC templates, build artifacts, and CI/CD workflows.

Prediction 7: Developer Experience Becomes a Security Control

Developers will no longer tolerate fragmented security tooling.
Effective platforms in 2026 will:

- Operate directly inside the IDE
- Surface only issues relevant to the current task
- Eliminate context switching
- Minimise noise

Security that feels invisible will be adopted. Security that disrupts workflows will be bypassed.

Prediction 8: Risk Prioritisation Becomes Context-Aware by Default

Severity scores alone will continue to fail in real environments. By 2026, effective prioritisation will require:

- Exploit-aware intelligence
- Usage and reachability context
- Business impact tagging

Security teams will act on likelihood and impact, not abstract severity.

Security teams will act on likelihood and impact, not abstract severity. Not all vulnerabilities deserve equal attention.

Prediction 9: Data Governance Becomes a Core Security Requirement

As security data becomes legal and financial evidence, data governance moves to the foreground. Organisations will need to control:

- Where dependency and trust data is stored
- Who can access it
- How long it is retained
- How it is audited and shared

Poor governance turns stolen credentials into long-term systemic risk.

These predictions point to a single conclusion: awareness is no longer enough. Organisations entering 2026 must be able to prove control, rotate trust, and respond at machine speed.

Next question executives care about most: "What to do now?"

What Leaders Must Do in the Next 90 Days

The events of 2025 showed that supply-chain attacks no longer wait for organisational readiness. They move faster than manual response and exploit trust rather than vulnerabilities.

The next 90 days matter because they determine whether an organisation enters 2026 resilient or exposed.

1. Reduce Credential Risk Immediately

Assume some level of credential exposure already exists. Rotate long-lived tokens, reduce secret lifetimes in CI/CD systems, and restrict who can publish, build, or deploy software. Treat credentials and trust relationships as critical assets, not background configuration.

2. Connect Visibility to Action

If your organisation produces SBOMs or dependency inventories, ensure they drive remediation. Dependency data must be tied directly to fixing workflows, not stored as static documentation. Visibility without execution does not reduce risk. Review your current CI/CD pipeline to ensure it includes automated infrastructure misconfiguration checks before deployment.

3. Accelerate Remediation

Manual patching cannot keep pace with modern exploit timelines. Prioritise automated fixing for high-confidence issues and enforce clear ownership for dependency risk. The goal is to shorten response time from weeks to hours.

4. Integrate Security Into Developer Workflows

Security controls should operate where developers already work. Reduce alert noise, eliminate unnecessary tool switching, and focus on contextual guidance that supports fast fixes without disrupting delivery.

5. Prepare for Continuous Regulatory Scrutiny

Regulators and auditors will increasingly expect real-time evidence of control. Ensure security data, provenance information, and remediation records are auditable, retained, and accessible without manual preparation.

meterian

National Cyber Security Centre *For Startups* Alumni

The Pulse of Software Resilience

At Meterian, we believe that in 2026, software is more than code—it forms part of the bedrock of national and economic sovereignty. Founded in London in 2018, our platform was engineered to solve the most critical challenge in modern development: the invisible risk of the open-source software supply chain.

Trusted by the Most Discreet Sectors

We are the silent partner to some of the world's most demanding organisations. Meterian provides mission-critical security for enterprise-grade software giants in the financial services (including a Tier 1 UK Financial Institution and fintech providers), health, utilities, and telecommunications and technology sectors. UK and European public sector services and global critical infrastructure providers operating in high-stakes environments where detection is not enough and failure is not an option count on Meterian. Our platform delivers the "Machine-Speed Remediation" required by these leaders to preserve trust, prove continuous compliance, and limit the blast radius of evolving threats in real time.

NorthStandard

GAMING COMPANY NDA

HEALTHTECH FIRM NDA

GLOBAL TRANSPORT FIRM NDA

FINANCIAL INSTITUTION NDA

FINANCIAL SERVICES FIRM NDA

GLOBAL TELECOM FIRM NDA

CRITICAL NATIONAL INFRASTRUCTURE NDA

A Catalyst for International Growth

Beyond our work with national giants, Meterian is a dedicated driver of international digital transformation for SMBs. We empower growing businesses to compete on a global stage by providing elite-grade security tools that are automated, affordable, and seamlessly integrated into their existing workflows. By removing the security friction that often stalls innovation, we help small-to-medium businesses scale with confidence.

Our ethos: Build Fast. Fix Faster. Stay Secure.

Meterian's AI-powered platform ensures unparalleled coverage and precision, providing actionable Software Bills of Materials (SBOMs) and automated fixes directly where developers work. We don't just find vulnerabilities; we eliminate them, ensuring that your software remains a resilient asset rather than a liability.

Join the future of secure development.

meterian.io | hello@meterian.io | +44 (0)20 7112 4879 | LinkedIn | Book a demo