# Verify Your Artifacts. Protect Your Patients.

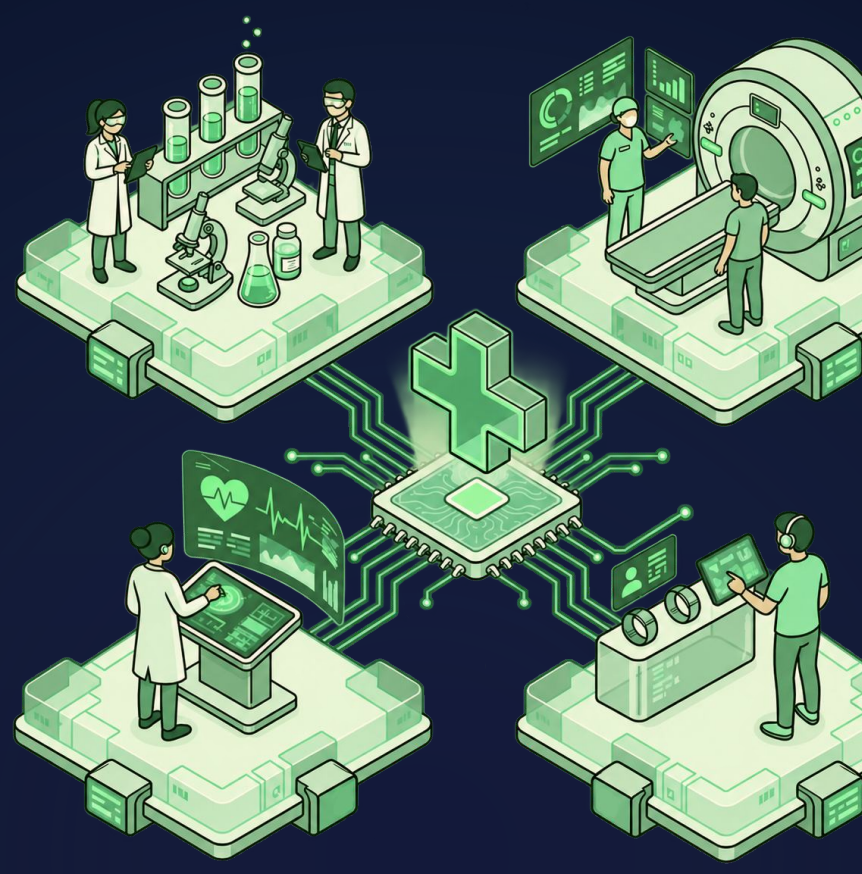
### Securing the Software and AI Behind Patient Care

## Executive Summary

Software is no longer an operational layer behind patient care. It is patient care. It's what powers the EHR a clinician relies on mid-shift, the infusion pump firmware at a bedside, the AI model reading a scan, the trial-management system a regulator will later audit. When that software fails or is compromised, the impact doesn't announce itself as an outage ticket. It shows up later, as a delay, a missed window, sometimes a death. In June 2025, King's College Hospital NHS Foundation Trust confirmed that a long wait for a blood test result, caused by the June 2024 ransomware attack on pathology provider Synnovis, was one of a number of contributing factors in a patient's death. The impact reached patients across multiple hospitals, far beyond the systems that were breached.

At the same time, regulatory pressure is widening: FDA requires cyber device manufacturers to include a software bill of materials in premarket submissions, and HHS has proposed HIPAA Security Rule changes that would require a technology asset inventory. Meanwhile, AI-assisted coding, diagnostic models, and autonomous agents are entering the pipeline faster than many governance programs can track. The threat landscape hasn't eased, either: a single unvetted open-source package or unmonitored model can still reach a clinical system unnoticed.

If software is patient care, then the integrity of that software is a patient safety issue. This brief outlines how healthcare and life sciences organizations are approaching that responsibility, from knowing what is in the software and AI they run, to verifying where it came from, to keeping that trust intact as it changes. It also covers what this means for providers, device makers, and life sciences teams shipping regulated software.



# The New Standard: Every Artifact Must Be Verified

In a hospital system, the distance between a developer's commit and a bedside device, patient portal, or diagnostic model is measured in risk. Closing that gap requires a unified chain of custody: every binary, package, container, and AI component accounted for from first commit to the point of care. Leading health systems and device makers are moving beyond point-in-time scanning toward continuous, binary-level verification.

## Lives are at Stake

The consequences of failure in healthcare are different from other industries. A stalled pump, a failed fetal heart-rate monitor, or a corrupted diagnostic feed doesn't just crash a system; it puts human lives in danger. Ungoverned components in the pipeline, AI entering the software supply chain (SSC) without oversight, and compliance treated as an afterthought, are three forces compounding this risk today:

- **Ungoverned components:** Malicious open-source packages identified in 2025 grew by 1,444% over 2024 (OpenSSF / Google Cloud Threat Intelligence, 2026).
- **Autonomous AI in the pipeline:** AI now writes or assists roughly 42% of code, and 97% of AI-related breaches hit teams with no AI access controls (Sonar 2025 / IBM 2025).
- **Governance retrofitted, not designed in:** A disconnected governance process turns an FDA submission or HHS audit into a last-minute scramble.

> *1 in 4 healthcare organizations report higher patient mortality following a ransomware attack. The average healthcare data breach costs $7.42M, the highest of any industry for 14 years running, and takes 279 days on average to detect and contain.*
>
> **(PONEMON INSTITUTE; IBM COST OF A DATA BREACH, 2025)**

Protecting every patient means embedding security into every phase of the software supply chain, generating a real-time software bill of materials (SBOM), and giving security and compliance teams continuous, audit-ready evidence instead of a once-a-year forensic exercise.

# The ROI of Trust

Governed correctly, software delivery becomes a competitive advantage rather than a cost center. Health systems and device makers that treat every artifact (any software package, container, or AI model) as a verified, traceable asset turn their delivery pipeline into a source of both safety and speed.

## Eliminating the Compliance Tax on Innovation

Traditional “block-and-fail” security policies act as a tax on clinical and engineering teams alike. When a release stalls because of a hidden dependency or an unreviewed model, developer hours and clinical timelines both suffer. A curated catalog of vetted packages and AI models removes that drain: teams get safe, pre-approved alternatives automatically, so security no longer stalls the pipeline.

## Precision Risk Mitigation, Not Alert Fatigue

Legacy, source-only scanners produce an unmanageable volume of noise, burying the alerts that matter. Analysis that checks reachability at the binary level identifies which vulnerabilities are actually applicable in a given clinical system, letting security and platform teams focus on the risks that could genuinely reach a patient.

*A composite organization modeled in the Forrester Total Economic Impact™ of the JFrog Platform in January of 2026 saw a 282% ROI, payback in under 6 months, and 38 hours saved per developer, across 40+ package types.*

**COMPOSITE CUSTOMER; INDIVIDUAL RESULTS VARY.**



# Regulatory Compliance & Audit Readiness

Meeting HIPAA, FDA 21 CFR Part 11, EU MDR, and FDA §524B is now the baseline expectation, not a differentiator. Since October 2023, the FDA can refuse to accept a cyber-device premarket submission that lacks an adequate, machine-readable SBOM. Historically, assembling that evidence has been a manual, months-long exercise, repeated for every submission and every audit.

### Automating the Audit Trail

*Imagine an HHS auditor asking for a submission's full history and getting it in minutes, not months.*

Automated governance turns compliance from a point-in-time scramble into a continuous state, generating the SBOMs, provenance, and approvals an HHS audit needs at the moment software ships. That same evidence trail, an immutable record of what ran, who approved it, and its full security history, doubles as Part 11 traceability for life sciences teams.

# Future-Proofing the Agentic Frontier in Clinical Software

There are three ways AI reaches a patient's care, and today, almost none of them are governed:

01 **The coding assistant:** AI-generated code reaching a clinical system without the same provenance (a verified origin and change history) and review as human-written code.

02 **The diagnostic model:** a model influencing patient care without version history, vetting, or a way to recall it if it drifts.

03 **The autonomous agent:** agents pulling, testing, and shipping code or model updates at machine speed, faster than manual review can follow.

The answer isn't to slow AI adoption; it's to stop trusting the AI layer blindly. A curated AI catalog and model registry gives clinical and engineering teams self-service access to approved models, with lineage and version history, so shadow AI is surfaced and stopped before an unvetted model reaches clinical use.

# Pain Points by Segment

| Segment | Critical Pain Point | Impact |
|---|---|---|
| **Providers** | EHR / connected clinical system uptime and HIPAA alignment | A stalled patient portal or frozen EHR is a life at stake, not a ticket. Unmanaged AI in clinical decision tools compounds patient-safety risk. |
| **Device Makers** | FDA §524B submission readiness and SBOM integrity | Since Oct. 2023, the FDA can refuse a cyber-device submission lacking an adequate machine-readable SBOM, risking 510(k) delay or a Refuse-To-Accept (RTA) decision. |
| **Life Sciences** | Part 11 traceability across GxP (Good Practice regulations) and clinical trial systems | Manually assembled audit trails under deadline create submission risk and slow validated-environment releases. |

**Trust in clinical software has to be earned, and proven.**

Every clinician relies on software they never see, and every patient relies on that clinician. JFrog works with healthcare providers, device makers, and life sciences teams to make the trust in that software verifiable, from the first commit to the bedside.

*The Synnovis attack didn't begin in a hospital room. It began with an unverified component, several steps removed from patient care, that nobody was watching. Closing that distance is what separates a contained incident from a clinical one.*

## Verified Artifacts Can Mean Lives Saved

JFrog is built to be the single source of truth for every software and AI artifact a healthcare organization ships, from a developer's first commit to the device, app, or clinical system at the bedside. The organizations leading in 2026 recognize that security in clinical software cannot be retrofitted after the fact; it has to be built into the platform itself.

Learn More