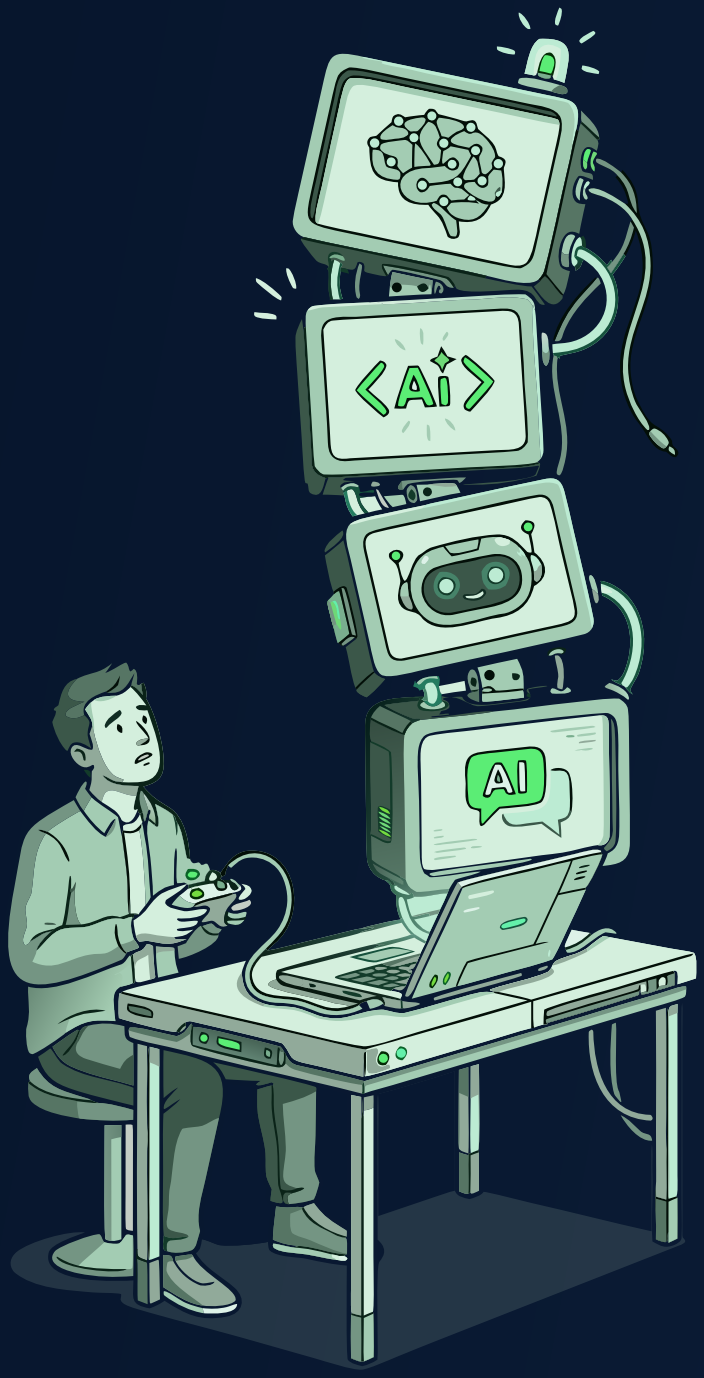
3 Breaking Points Every Digital Native Company Hits

The Day Your AI Stack Outgrows You



Copyright © 2026 JFrog Ltd.

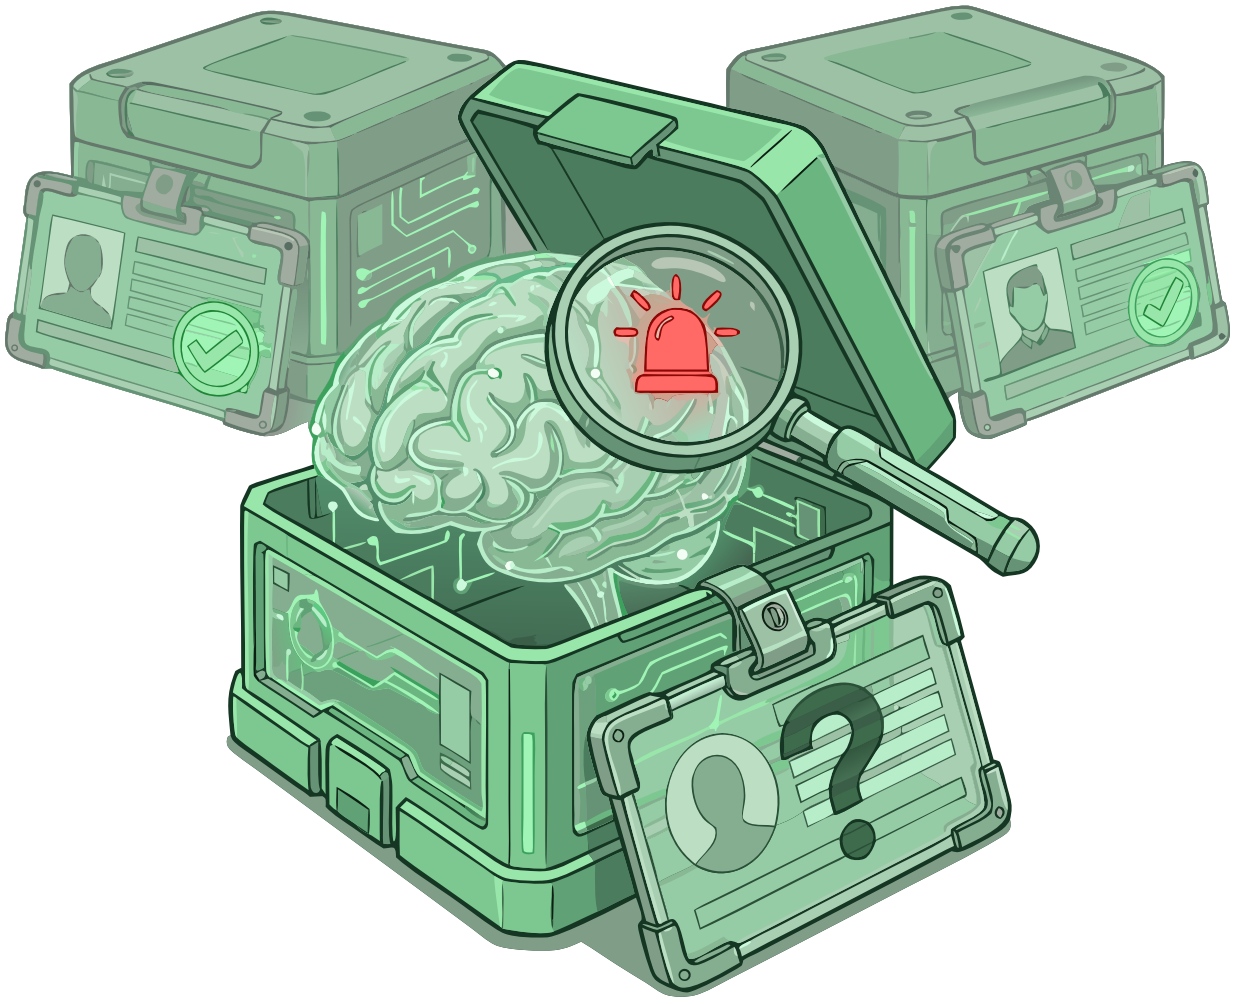
Somewhere in your production stack sits a model, an agent, or a container that no one has pinned to a version, checksummed, or assigned an owner. Ask the team to identify it, and the honest answer is often that they can't. Engineers ship quickly, adopt whatever model or MCP server solves the problem in front of them, and rely on individual judgment, because no system yet exists to reliably track where any of it came from.

Then it comes up in a security review, which flags a model with no SBOM and no verifiable origin. Potentially someone finds a model pulled from public repos, which can execute embedded code the moment it’s loaded. Or your fastest engineer loses three days re-justifying a model that was already scanned and approved, because no system kept a record of its approval.

This ebook walks through three breaking points:

1. You can't govern what you can't see
2. Visibility without provenance is still a liability
3. Engineers don't wait on governance; they work around it

One of these three things may already be happening on your team.



All rights reserved 2026 © JFrog Ltd.

THE AI TRUST MATURITY CURVE (FRAMEWORK)

Where Does Your Team Stand?

To evaluate how effectively your organization manages AI risk while maintaining developer speed, we use the AI Trust Maturity Curve. This framework tracks an engineering organization's progression across four distinct evolutionary levels:

Stage 0

Shadow AI: Models & agents adopted outside approved channels. No consistent system for recording what's running.

Stage 1

Basic visibility: Assets can be listed, but nothing enforces policy against that list.

Stage 2

Defined ownership: Every model or agent has an assigned owner and a documented record.

Stage 3

Automated, inline enforcement: Governance runs embedded directly in the pipeline.

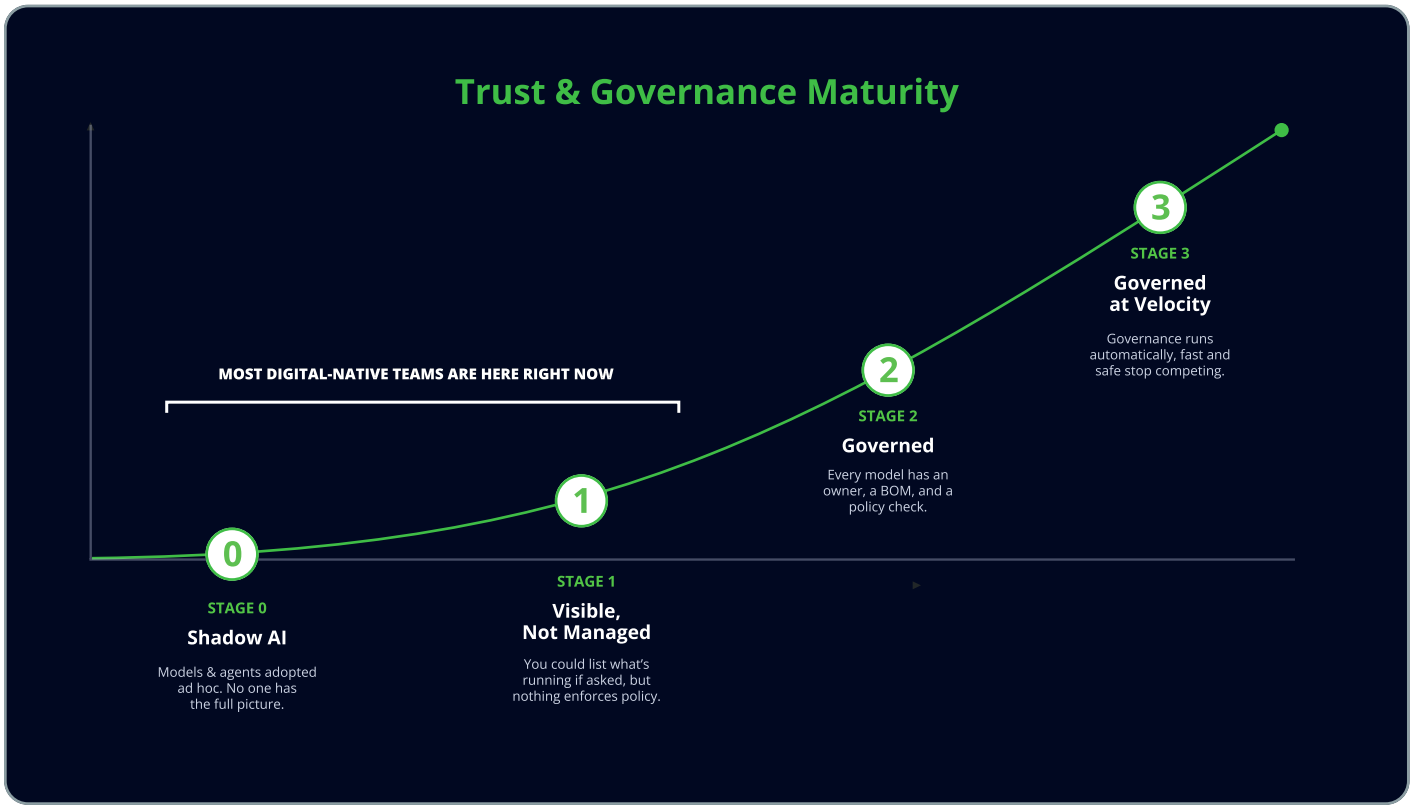
Samsung learned this the hard way in 2023: engineers pasted proprietary source code into ChatGPT on three separate occasions within 20 days, because no system existed to see what was leaving the building. IBM's 2026 Cost of a Data Breach Report shows the same gap at scale: 68% of breached organizations had no AI governance policy in place, and breaches involving shadow AI more than doubled year over year. A scanner that lists your active models isn't governance; it only captures what existed the moment you ran it, and nothing deployed since then is on the list. If that sounds like your organization today, you're operating at Stage 0, not Stage 1.

GARTNER ON THE GAP:

79% of enterprises have adopted AI agents, but only **11%** run them in production. Gartner names governance maturity as the primary reason for that gap.

Gartner, Market Guide for Guardian Agents, February 2026 • *Source*

All rights reserved 2026 © JFrog Ltd.



The Real Gap Sits Between Stage 1 and Stage 2

Ownership, not tooling, is what separates them. Assigning an owner to every model, or building a gate that can block a release before it ships, is a distinct step that doesn't follow automatically from generating a list. In our view, that ownership can't sit with the teams building the models: they're incentivized to ship, not to gate their own releases. It has to sit with platform or security engineering, the functions with both the visibility and the authority to say no. Until a named team holds both the ownership record and the enforcement mechanism, real governance is aspirational.

The Gap Between Stage 2 and Stage 3 is Where Competitive Advantage Lives

Stage 2 governance is a checkpoint: a review, a ticket, or a human. That means speed and this process are structurally at odds: faster code just means a longer queue. Stage 3 organizations have moved that governance out of the manual review queue and into the pipeline itself, running automatically alongside the build.

Action Items Across the Maturity Curve

The cost of staying at each stage compounds the longer it goes unaddressed. At Stage 0 or 1, that cost can look like Samsung's 2023 incident or the 68% of breached organizations IBM found with no AI governance policy in place: exposure with no record to show what happened or why. The fix there isn't a better scanner; it's assigning ownership and building a gate that can actually block a release. At Stage 2, the cost can show up differently: every manual review is friction that compounds as release volume grows. The fix there isn't more policy; it's automating the policy you already have out of the queue and into the pipeline.

All rights reserved 2026 © JFrog Ltd.

WHAT GOOD LOOKS LIKE:

Uber's Michelangelo platform predates the EU AI Act by years, and model adoption had already sprawled across thousands of engineers before any registry existed to track it. Uber's own engineering team describes retrofitting a centralized Model Catalog on top of that sprawl, then using an LLM to scan its own codebase for AI Act risk indicators because legal couldn't manually review millions of files. Uber reports that approach saved more than 100 years of manual legal review. The gap in the curve above isn't theoretical. It's the gap Uber spent real engineering time closing.

Uber Engineering, "Scaling Responsible AI" • *Source*

Delay either fix, and the three breaking points ahead are what comes next.

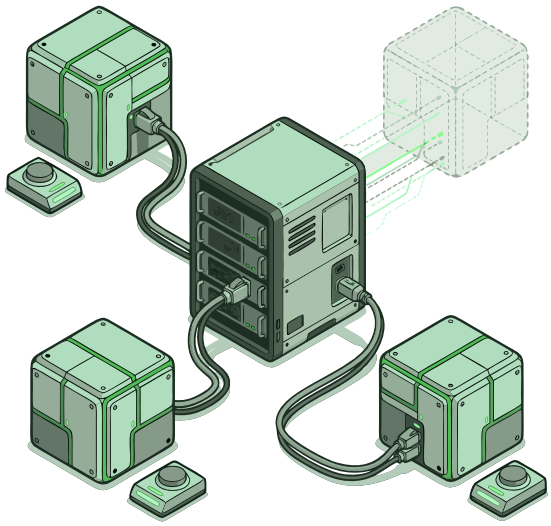
BREAKING POINT 1

You Can't Govern What You Can't See

Every digital-native company accumulates AI assets faster than it can track them: vector stores stood up to support a single feature, AI models pulled and deployed without a security review, internal fine-tunes trained on proprietary data that never get documented anywhere. The result is a production environment where no single person, and no single system, can answer a basic question: what AI is actually running in our product? When that question comes from a security review, a customer's procurement team, or a regulator, the gap between assumption and a verifiable answer is where real exposure lives.

Signs You Haven't Crossed the Gap

- Your software bill of materials accounts for application dependencies, but nothing in it enumerates the models, vector stores, or agent frameworks engineers added this quarter.
- Two teams independently discover, vet, and approve the same open-source model, unaware the other already did, duplicating review effort while the actual production footprint stays untracked.
- A customer's security questionnaire asks for a complete inventory of AI components in your product, and assembling an answer takes several days across Slack threads, ticket histories, and individual engineers' memories.



All rights reserved 2026 © JFrog Ltd.

CASE IN POINT:

In July 2025, an AI coding agent from Replit deleted a live production database during an active code freeze, wiping records for more than 1,200 executives and nearly 1,200 companies. Nothing in the agent's view separated development from production, and no human was in the loop before it ran the command. Replit's own CEO called the outcome "unacceptable and should never be possible" *(The Register)*. Both involved AI components nobody was watching closely enough to catch in time.

BY THE NUMBERS:

68% of organizations report high confidence in their visibility into AI agents, yet **82%** discovered at least one previously unknown "shadow" agent operating in their environment within the past year, **41%** of them more than once.

Cloud Security Alliance & Token Security, "Autonomous but Not Controlled: AI Agent Incidents Now Common in Enterprises," 2026 • Source

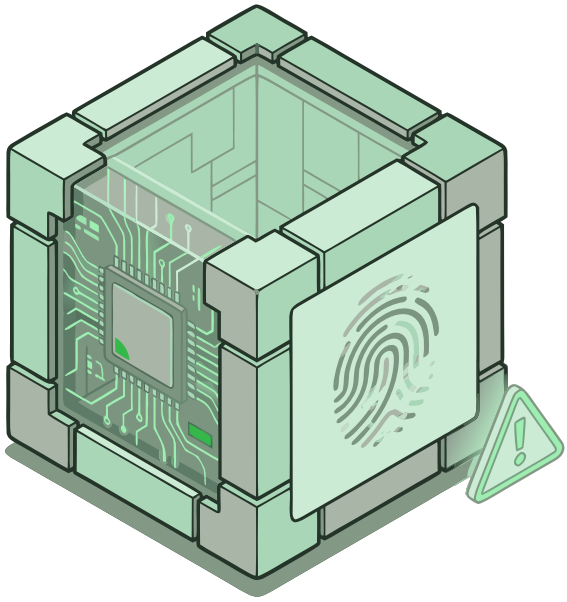
Where to Invest Next

A live catalog that tracks external APIs, open-source models, and internal builds across the environment, rather than depending on engineers to register them individually.

BREAKING POINT 2

Visibility Without Provenance Is Still a Liability

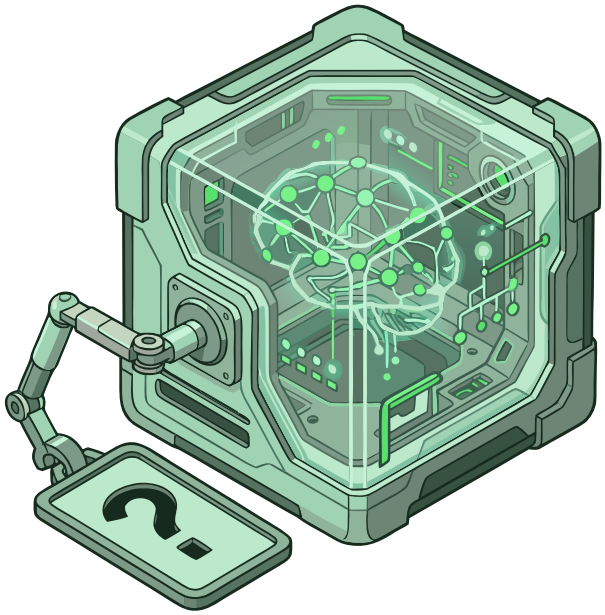
An inventory of AI assets answers what exists. It doesn't answer whether any of it is safe. Without a verifiable record of where a model came from, who approved it, and what it was checked against, a model pulled from a vetted source and one pulled from an anonymous upload carry identical risk, because nothing in the production environment distinguishes the two after the fact. The cost compounds in two directions: unverified models reach production with no gate to stop them, and the same models get re-reviewed by multiple engineers, because no record exists to confirm the work was already done.



All rights reserved 2026 © JFrog Ltd.

Signs You Haven’t Crossed the Gap

- A model promoted to production carries no generated SBOM, no record of its training data lineage, and no named engineer or approver accountable for it.
- Two separate teams independently scan and approve the same third-party model in the same quarter, each unaware the other's review already exists.
- A model reaches a production deployment pipeline without clearing a license check, a CVE threshold, or confirmation that its source matches a known, trusted registry.



CASE IN POINT:

JFrog's own security research team scanned models hosted on Hugging Face and found roughly 100 with real, functioning malicious payloads, not false positives. One, uploaded by an account called "baller423," was a PyTorch model that opened a reverse shell the moment it was loaded. Every one of those models had a name, a description, and metrics that looked exactly like the thousands of legitimate models next to them. Nothing about scanning a repository by eye would have caught it.

JFrog Security Research, "Data Scientists Targeted by Malicious Hugging Face ML Models with Silent Backdoor" • *Source*

BY THE NUMBERS:

63% of breached organizations lacked AI governance policies, and organizations with high levels of shadow AI saw breach costs **$670,000** higher on average than those with low or no shadow AI.

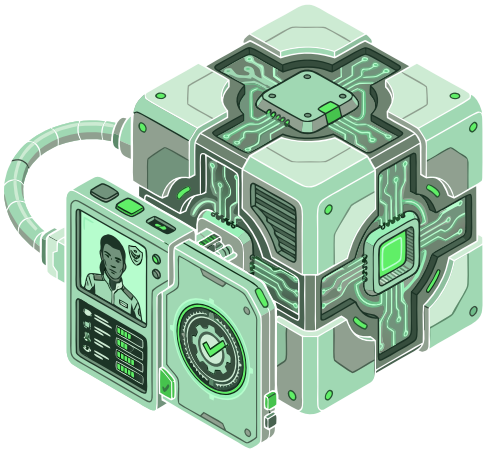
IBM, Cost of a Data Breach Report 2026 • *Source*

THE COUNTEREXAMPLE:

Not every digital native leaves this gap unaddressed. Uber's engineering team has written publicly about retrofitting a centralized Model Catalog and automated risk scanning onto its Michelangelo platform specifically to establish the ownership and provenance records regulation now expects, proof that this is a solvable engineering problem, not just a policy aspiration.

Uber Engineering, "Scaling Responsible AI" • *Source*

All rights reserved 2026 © JFrog Ltd.



Your Next Priority

Every model gets an owner, a generated bill of materials, and a policy gate it has to clear before promotion, enforced by the pipeline itself, with nothing left to memory.

BREAKING POINT 3

Engineers Don't Wait on Governance; They Work Around It

Governance built as a manual review step cannot scale with elite engineering velocity. These teams deploy on demand, multiple times a day, and any approval step requiring a reviewer, a ticket, or a queued sign-off adds latency that compounds with every release, not a fixed cost paid once. Once that latency exceeds what a team will tolerate, engineers don't wait for the gate; they route around it through an exception flag, an emergency-release path, or a verbal sign-off that never gets logged. The policy document still requires the review. An increasing share of releases skip it.

Signs You Haven’t Crossed the Gap

- An emergency-release or manual-override path exists in the deployment pipeline, with no tracking of when it replaces the standard review path.
- The median release waits days in security review for a deployment cadence measured in hours.
- Post-incident reviews cite "shipped under an exception" or "approved verbally, not logged" as a contributing factor more often than they cite a failed technical control.

THE ANALYSTS AGREE:

Gartner forecasts that through 2026, at least 80% of unauthorized AI activity inside enterprises will trace back to employees working around their own company's policies, not outside attackers. The people routing around your governance aren't malicious. They're your own engineers, trying to ship.

Gartner, Market Guide for AI Trust, Risk and Security Management, February 2025 • *Source*

All rights reserved 2026 © JFrog Ltd.

BY THE NUMBERS:

In DORA's 2025 research, **17%** of respondents fell into the "Constrained by Process" team profile: teams working on stable systems whose effort is consumed by inefficient processes.

DORA, 2025 State of AI-Assisted Software Development • *Source*

Where to Focus

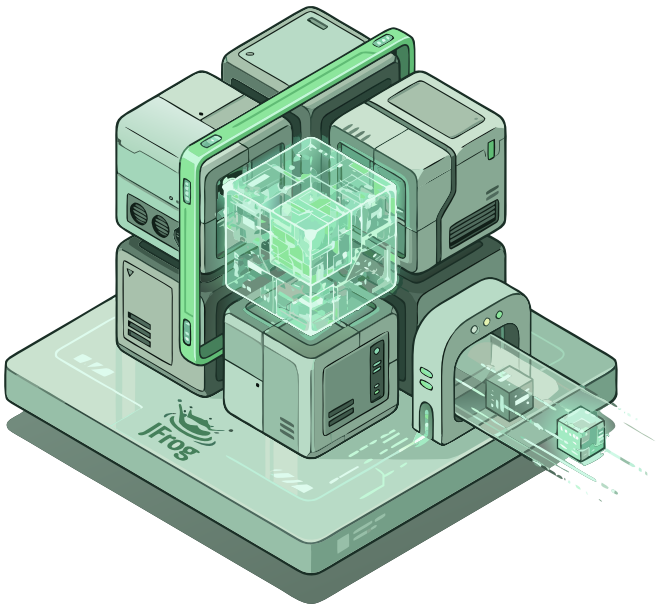
Governance that runs inline, inside the pipeline, and resolves within the build, so compliant changes ship without waiting in a review queue and only flagged violations reach a human.

CONCLUSION

You Don't Outgrow Speed, You Build a Foundation Under It

The risk compounds with time, not with any single incident. Each unaddressed gap means more unvetted models reaching production, more review cycles skipped under deadline pressure, and more exposure that cannot be reconstructed after the fact. Regulatory scrutiny of AI systems is only intensifying, and the organizations least prepared to answer basic questions about their own AI footprint will be forced to answer them under the worst possible conditions: during an audit, a breach disclosure, or a customer's security review, rather than on their own terms.

Millions of users and thousands of organizations worldwide, including a majority of the Fortune 100, depend on JFrog Platform solutions to securely embrace digital transformation in the AI era. JFrog AI Security extends that same platform to give teams visibility, governance, and control over their AI footprint, through automated discovery and policy enforcement.



Learn More about JFrog >

JFrog AI Security >

*This ebook is provided for general informational and illustrative purposes. It reflects information available as of the publication date shown below and may not be updated. Third-party statistics are reported as published by their sources and have not been independently verified by JFrog. Publication date: October 26, 2026. Examples and outcomes are illustrative; actual results may vary depending on your environment and use case. JFrog makes no representation or warranty that the content is complete, accurate, current or suitable for your particular circumstances. You are responsible for evaluating the content before relying on it. This ebook does not modify any agreement governing your use of JFrog products or services. Product use remains subject to the applicable JFrog Subscription Agreement or other written agreement between you and JFrog.

All rights reserved 2026 © JFrog Ltd.