



25RevolAI

WHITE PAPER

Provenance Is the New Perimeter

A postmortem of last year's Illumina False Claims Act settlement, further connections to FDA's emerging questions about AI, and why the durable answer is to eliminate whole classes of attestation risk instead of managing them all.

09/15/2026

Dave Stephens | Operating Partner, 25RevolAI

dave.stephens@25revol.com | +1 801 870 1449

Executive Summary

On July 31, 2025, the Department of Justice announced that Illumina Inc. agreed to pay \$9.8 million to resolve False Claims Act allegations arising from cybersecurity vulnerabilities in genomic sequencing systems sold to Federal agencies [1]. The government asserted that between February 2016 and September 2023 the company sold systems with software vulnerabilities, without an adequate security program or sufficient quality systems to identify and address them, and that it falsely represented the software's adherence to ISO and NIST cybersecurity standards.

The most important fact is that the case did not involve any type of breach or related incident. No incident is described and no data loss is alleged. The exposure was the claim: conformance represented to a customer without demonstrable controls and evidence of compliance.

This is the third paper in a series. The first examined the first FDA warning letter to specifically "name" AI, where the missing ingredient was an accountable human and traceable, documented decision logic. The second examined what open-weight models do and do not give a regulated operator, where the missing ingredient was ownership [2]. This paper looks at cybersecurity, quality, and AI. The same issue runs through all three: provenance. Prove where the data came from, where it lives and computes, and who touched the output, or do not use it. While the issue is not novel and has existed for decades in regulated systems, it is more fraught and increasingly hard to prove or provide as AI expands rapidly.

There are really only two ways to manage the risk for your organization. Continue attesting into an environment you do not control, and carry the gap between claim and fact on your books as a perpetual liability. Or run on infrastructure where the controls are simply properties of the place, and the list of things you have to attest to is small, controlled by your organization and clearly provable across the enterprise. You can manage attestation risk that will continue to grow and increase in complexity (and cost) across multi-tenant, public cloud, SaaS and associated AI environments, or you can eliminate whole classes of it. We have now seen the price of poorly addressing the risk in one customer environment, and Illumina has not been the last.

1. What Happened: The Non-Compliance Is the Claim

The government's claims describe a program failure, not a hack: Illumina knowingly failed to incorporate product cybersecurity in software design, development, installation, and monitoring. It failed to properly support and resource the personnel, systems, and processes tasked with product security. It failed to adequately correct design features that introduced vulnerabilities into the sequencing systems. And while those failures stood, it continued to represent that the software adhered to recognized cybersecurity standards [1].

Illumina ... falsely represented that the software on the genomic sequencing systems adhered to cybersecurity standards, including standards of the

Two key details are worth examining in more depth. First, the buyers included the Department of Defense, and the investigating agencies included its criminal investigative service, which framed the matter around safeguarding research data that supports the warfighter. Genomic data raised the stakes, and the HHS Office of Inspector General said so directly. Second, the whistleblower ran the on-market portfolio. The person best positioned to see the distance between a marketing claim and the monitoring reality is exactly the person the False Claims Act encourages to come forward, and every regulated vendor employs someone who can see that distance and understands the risks involved.

The vulnerabilities were not hypothetical. In June 2022 and again in April 2023, FDA and CISA published advisories on software vulnerabilities in Illumina sequencing instruments, including systems used for clinical diagnostics. The flaws could allow an unauthorized user to reach an instrument remotely and alter its settings, software, or data, and the company issued corrective actions in response. Neither FDA nor the company reported evidence of exploitation, and no patient harm was reported [3]. What matters is that the government never had to speculate about the gap between the attestations and the software. FDA and CISA had already documented that gap, inside the very period the settlement covers.

Illumina was not alone in 2025. DOJ resolved cyber-fraud claims with MORSECORP in March, Raytheon and Nightwing in May, and Georgia Tech Research Corporation in September [4]. Illumina was the largest of the four and the only one outside the traditional defense industry, highlighting the arrival of these risks in life sciences. The settlement resolves allegations, not proven facts, and fairness requires pointing that out. The enforcement pattern does not need a finding of guilt, though. It demonstrates that a conformance claim is an instrument. It can be audited years after it was made, compared against the program that existed at the time, and priced. Attestation is not paperwork. It is exposure.

2. The New Normal: Regulators Are Focusing on Provenance and Controls

Three documents from different corners of the government point the same direction. Two carry the force of law, and neither was written with AI in mind. The third is the FDA asking about AI head on, and it wants what the other two already take for granted.

- **Cybersecurity, 2023.** Section 524B of the FD&C Act requires cyber device submissions to include a software bill of materials, a plan to monitor and disclose vulnerabilities, and evidence of secure development [5]. A software bill of materials is provenance for software: what is in the product, where each piece came from, and how it is managed, controlled and maintained.
- **Quality, 2026.** The first FDA warning letter to name AI asked, in substance, who reviewed the output. The remedy was an authorized human approving AI output, which ensures traceability of a decision to an accountable person [6].

- **Generative AI, 2026.** FDA's discussion paper on GenAI-enabled devices asks whether the traceability of an output to primary source materials belongs on the risk framework itself. Its proposed Foundation Model Master Files would carry training data provenance, known limitations, and update commitments. The critical open question asks how a manufacturer detects and controls changes to a model it did not build [7].

Regulators are unlikely to inspect the internal workings of your models. The workable, enforceable rule available to them is a provenance rule: prove what data the system used, prove what sources were used to inform the decision, prove where that data has been, prove who cleared the output. Organizations that fall short on any of the four carry a risk that will continue to grow and compound. The solution is a fence around AI built entirely from existing regulatory expertise; the solutions and tools have changed but the spirit and intent of regulations and statutes have not. Solving these risks does not require new science for model evaluation, and every part of it can be enforced with the same instruments that produced the Illumina settlement and Warning Letter 320-26-58. It is about controlling the environment, tools and the behavior of those tools. The compliance question of the next decade is not what your system does. It is what you can prove, trace, document and stand behind.

3. Where Boundary Controls Stop Working

Traditional cybersecurity is perimeter thinking: control the boundary of the environment, harden the walls, watch the doors. The convergence above asks a different question. Provenance and lineage are properties of the data rather than the walls, driven forward by two powerful forces.

The first is the data itself. A genome is not like a name or an account number. It is an identifier in its own right, stable for life and traceable to people who never consented to anything. Researchers demonstrated this in Science in 2013, recovering surnames from anonymous genomes using short tandem repeats on the Y chromosome and free public genealogy databases, then triangulating identity with nothing more than age and state [8]. Strip the demographics from a sequencing record and the sequence remains a fingerprint. Masking, the workhorse of traditional privacy control, does not de-identify it. For data in this class, the controls must attach to the data and its movement. That is the definition of provenance and lineage.

The second is the environment most regulated data lives in. In multi-tenant cloud and hyperscale solutions, there is no de facto compliant state. Compliance is an audited state. Organizations must attest and audit continuously: shared responsibility matrices, configuration reviews, evidence collection that never ends. The substrate is not yours and changes without your consent and often without your control. The scale of that problem is now a matter of public record. After the Schrems II judgment [9] struck down the EU-US Privacy Shield over data transfer protections, the largest hyperscaler concluded that residency inside its global multi-tenant environment could not satisfy European sovereignty requirements, and built a physically and logically separate European Sovereign Cloud instead, a planned investment of 7.8 billion euros [10] that will not be free to customers nor

will it be simple to manage as lineage questions continue to drill down to lower and lower levels across regulators. In the US Federal context, the same logic is explicit: FedRAMP authorization requires continuous monitoring, because in shared infrastructure the compliant state must be re-proven on a permanent unending cycle. Every one of those attestations is a claim. Section 1 highlights the risks and impacts of not meeting the burden of auditable proof.

4. Three Ways to Hold Regulated Data

The honest comparison is not security versus insecurity. All three of the models below can be run securely. The comparison is what compliance looks like in each one, how much it costs, and the complexity and risks of claims and attestations to regulators about things that are not always dimensions you can control.

Deployment model	What compliance looks like	What it costs	The attestation surface
Multi-tenant public cloud	A continuously audited condition. Shared responsibility, configuration review, evidence collection on a permanent cycle.	Subscription pricing plus audit labor that never ends.	Largest. Many controls are claims about elements someone else operates and changes.
Private environment on a hyperscaler	Isolation approximated through contract and architecture. Closer, never closed.	The few organizations that achieve it sustain standing programs at costs most operators cannot carry.	Smaller, still open. You attest to arrangements, not properties.
Owned sovereign infrastructure	Controls are properties of the environment you own. Isolation does not depend on someone else's configuration.	Fixed, depreciable infrastructure cost with a known run rate.	Smallest. What remains to attest is your own program, not your landlord's.

One caution, so nobody reads our paper as selling a shortcut. Infrastructure is not a substitute for a security program. The alleged failures in Section 1 were program failures: design, resourcing, correction, monitoring and auditable documentation. Sovereign deployment does not excuse any of that work. What it does is shrink the attestation surface. The solutions you use and operate will now support claims about things you control and can prove. Eliminate the risks the architecture can eliminate, and run a real program for what remains.

Three questions separate architectures faster than any demo, and they are worth asking of every vendor on your list, including us. Where was the system's knowledge built, before the model or bolted on after it? Who decides when the model changes, you or the vendor? And when an output is challenged, what does the reviewer actually hold, a citation or a probability?

This is how we build at 25RevolAI. Our platform is agnostic to the deployment model, and is run on a sovereign cloud infrastructure from our partner SoftIron. The Ground Truth Graph™

gives every statement a source an auditor can open. Source Constrained AI™ is designed to restrict output and solution behavior to verified, auditable sources and refuse actions when coverage is insufficient. Informed operators clear the output, so provenance and decision-making lineage and documentation run from data to decision. What we deliver, as a result, is validation- and audit-ready at all times. The evidence for the claim exists before the claim is made.

5. What to Do Now

Whether or not you ever touch our platform, five actions follow from the settlement and the clear and emerging regulatory direction. A regulated operator should be able to close the list within a quarter.

- **Inventory your attestations.** Every conformance claim in a Federal contract, a premarket submission, a security questionnaire, or a customer certification letter is an instrument that can be audited later against the program that existed when it was made.
- **Locate the evidence.** For every claim, ask where the proof lives and who owns it. A control satisfied by a vendor's assertion is a claim you have inherited without the evidence.
- **Treat provenance as a design input.** Where the data came from, where it was computed, and who cleared the output should be known for every output, not reconstructed per audit.
- **Price your attestation surface.** Assume your deployment model is a liability. Count the controls that are claims about someone else's substrate, and cost the continuous audit required to keep them true.
- **Build for the rules that already apply.** Section 524B is law. The quality unit requirement behind Warning Letter 320-26-58 predates generative AI by decades. The False Claims Act needed no AI provision to reach Illumina. What is still being written is guidance on how to demonstrate compliance when AI is involved, not whether the obligation exists. Architect so the proof exists now.

Summary

The perimeter is still there, it moved into the data. Illumina's settlement shows what it costs when conformance is asserted instead of possessed, and the FDA's open questions on generative AI show the same requirements arriving for every output an AI touches. The operators who come through the shift effectively and compliantly will be the ones whose architecture makes provenance a property rather than an audited state. Choose infrastructure and systems where the proof exists before the claim is made, and entire classes of attestation risk will no longer exist. Compliance will cease to be solely an audited state.

There are two ways to start, depending on where you are. If you suspect exposure but cannot yet size it, we will run a working session on your attestation surface: which claims

you have made, where the evidence lives, and which of them depend on someone else's operations.

If you already know where the gap is and want proof it can be closed, bring us one real use case and we will deliver a working, human-reviewed prototype against it in 24 to 48 hours at no cost, with every statement carrying a source an auditor can open. We call that the Zero Imagination Gap™, where compliance is inherent to the design.

Dave Stephens | Operating Partner, 25RevolAI

dave.stephens@25revol.com | +1 801 870 1449 | www.25revol.ai

Sources

[1] U.S. Department of Justice, Illumina Inc. to Pay \$9.8M to Resolve False Claims Act Allegations Arising from Cybersecurity Vulnerabilities in Genomic Sequencing Systems, 07/31/2025, and the settlement agreement referenced therein. United States ex rel. Lenore v. Illumina Inc., No. 1:23-cv-00372 (D.R.I.).

<https://www.justice.gov/opa/pr/illumina-inc-pay-98m-resolve-false-claims-act-allegations-arising-cybersecurity>

[2] 25RevolAI, Publications, including Open Is Not an Outcome and other white papers.

<https://www.25revol.ai/publications/>

[3] U.S. FDA, safety communication on cybersecurity vulnerabilities in Illumina Local Run Manager software, 06/2022; U.S. FDA and CISA, advisories on Illumina Universal Copy Service, CVE-2023-1968 and CVE-2023-1966, 04/2023; FDA recall classification, 07/2023. <https://www.fda.gov/medical-devices/letters-health-care-providers/illumina-cybersecurity-vulnerability-affecting-universal-copy-service-software-may-present-risks>

[4] U.S. Department of Justice, civil cyber-fraud False Claims Act settlements, 2025: MORSECORP Inc., \$4.6M, 03/26/2025; Raytheon Company, RTX Corporation and Nightwing Group, \$8.4M, 05/01/2025; Georgia Tech Research Corporation, \$875,000, 09/30/2025.

[5] Section 524B, Federal Food, Drug, and Cosmetic Act, Ensuring Cybersecurity of Devices; U.S. FDA, Cybersecurity in Medical Devices: Quality System Considerations and Content of Premarket Submissions, guidance.

[6] U.S. Food and Drug Administration, Warning Letter 320-26-58, Purolea Cosmetics Lab, 04/02/2026, examined in our white paper The AI Never Told Me It Was Required, 07/2026. <https://www.fda.gov/inspections-compliance-enforcement-and-criminal-investigations/warning-letters/purolea-cosmetics-lab-722591-04022026>

[7] U.S. FDA, Center for Devices and Radiological Health, discussion paper on generative AI-enabled medical devices, 08/2026, public comments open through 10/19/2026. <https://www.fda.gov/media/194242/download>

[8] M. Gymrek, A. L. McGuire, D. Golan, E. Halperin, Y. Erlich, Identifying Personal Genomes by Surname Inference, Science 339:321-324, 2013. <https://www.science.org/doi/10.1126/science.1229566>

[9] Court of Justice of the European Union, Case C-311/18 (Schrems II), 07/16/2020.

[10] Amazon Web Services, AWS plans to invest 7.8 billion euros into the AWS European Sovereign Cloud, 05/2024, general availability announced 01/2026. <https://www.aboutamazon.eu/news/aws/aws-plans-to-invest-7-8-billion-into-the-aws-european-sovereign-cloud>