The Unauthenticated Signal

Why clinical AI at the edge needs the signal authenticated in hardware, before the model acts on it

Dana Xiadani Founder & Principal Architect, Living Cipher LLC

The inference is solved. Clinical AI has crossed the threshold that matters: it works, and regulators have said so. AI/ML clinical inference has reached FDA marketing authorization — the first authorization for an AI diagnostic in acute care was granted through the De Novo pathway in 2024 — validated on large clinical datasets and cleared for use. The models are real, the accuracy is peer-reviewed, and the trajectory is set. What has been proven, though, was proven in a particular setting: as software, running on a trusted enterprise data pipeline, drawing its inputs from inside the hospital's electronic record.

That setting is about to change, and the change opens a gap the cloud never had to face.

The pipeline you no longer have

A model embedded in an electronic medical record inherits a trusted pipeline. The data has already been acquired, attributed to a patient, timestamped, and carried over authenticated enterprise infrastructure before the model ever sees it. Whatever else is hard about clinical inference, the authenticity of the input was somebody else's solved problem. The model could assume its signal was real because the system around it guaranteed as much.

Move that same class of inference onto hardware at the edge — an implantable, a bedside monitor, a wearable, a closed-loop stimulator, a neural interface — and the guarantee evaporates. The signal is now acquired at a physical sensor, at the outermost edge of the system, before any trusted pipeline exists to vouch for it. The model still assumes its input is authentic. Nothing at the edge entitles it to.

Connectivity is not authenticity, and neither is accuracy

A validated clinical model faithfully infers on whatever signal it is handed. It will do so whether or not the signal came from the patient it is attributed to, whether or not it is the live physiological signal rather than a recorded one, and whether or not it has been altered, replayed, or synthesized between the sensor and the model. The model does its job perfectly. The authenticity of the source is out of its scope.

Make it concrete. A physiological signal is captured and replayed at the sensor edge, or a synthetic waveform is injected in place of the live one, or a signal from one source is presented as another's. It is well-formed. It has the right shape, the right sampling, the right statistics, so it passes every signal-quality and integrity check the acquisition path applies — because those checks were built to catch a corrupted or noisy signal, not a dishonest one. The model receives a clean, plausible input and returns a confident inference. The signal-processing chain did its job perfectly. It was never the machinery that could have asked whether the signal was authentic, live, and truly from this patient, on this device, at this moment.

This is the gap. As clinical inference moves from the trusted enterprise pipeline to hardware at the point of acquisition, the authenticity of the signal becomes an unowned question. A device can run a peer-reviewed, FDA-authorized model and a fully attested boot chain, and still act on a signal that nothing in the system ever verified. When the output can change a therapy, the unauthenticated signal is the soft center of an otherwise capable device.

The question the model cannot ask itself

> *Should this model act on this signal, and on what evidence that the signal is authentic, live, and truly from this patient, on this hardware, at this moment?*

A classifier cannot answer that from the signal alone, because the very attacks that matter are the ones designed to look like a valid signal. The answer has to come from somewhere the signal itself cannot forge: from a cryptographic proof, rooted in the hardware, established before the model is allowed to act.

What trust at the source requires

The principle is simple to state and demanding to build correctly: **a signal should not be acted on until it has been proven, cryptographically and at the hardware root, to be authentic, live, and bound to this device and this moment.** Not trusted because it arrived. Trusted because it carried evidence, and the evidence was verified in hardware before any downstream model relied on it. Concretely, that means provenance (the signal bound to a specific, hardware-rooted identity that cannot be cloned or transferred), integrity (assurance the signal was not altered between acquisition and inference), liveness (evidence that it is the live signal, not a recording replayed into the input), freshness (the evidence bound to this session, so a captured exchange cannot be replayed later), and a quantum-resistant cryptographic

binding, because a device fielded now will still be in service when the threat model has moved. The trust material must never cross an exposed bus as a recoverable value; that is a property enforced in the architecture, not promised in a policy.

The approach is practical, and it is silicon

None of this is a research proposal. Living Cipher's founding work, Bio-Cryptographic Attestation, is exactly this: cryptographic proof of a signal's authenticity at the hardware root, upstream of the standard cryptographic layer, so that what a model acts on is attested before it acts. The primitives it depends on are buildable today and verifiable at the register-transfer level against published cryptographic standards, and the post-quantum binding has been carried through to timing closure on real silicon.

And it has been built across more than one kind of signal. Living Cipher has designed a family of clinical-inference chiplets on a production FPGA-SoC — each pairing an inference datapath with the trust architecture that attests the signal beneath it — spanning signal domains from bedside physiological monitoring to neural interfaces. These are logical chiplets: authored in RTL, verified, and proven on FPGA fabric, and being carried toward tapeout. The point of the range is the point of the architecture: signal-fidelity attestation is not a one-signal trick but a reusable trust layer that generalizes across the signals clinical AI at the edge will depend on. The inference in each is a validated domain in its own right; what Living Cipher adds is the layer beneath it that proves the signal was real before the model was allowed to act.

Proving these designs on FPGA before committing to an ASIC tapeout is not a limitation; it is the order the work is meant to be done in, because a tapeout is a one-way, expensive door and the design that goes through it should already have earned the trip.

One idea, stated plainly

At root this is a single principle: authenticate the signal at the hardware root, upstream, with cryptographic proof, before anything downstream trusts it. In the clinical-edge case the signal is a physiological measurement entering a model that may change a therapy, and the post-quantum cryptography belongs in the fabric because the authentication has to be bound at the root, not asserted in software an attacker may already hold. The principle is general — it is the same one that authenticates a die across a die-to-die interconnect — and this is its application where the stakes are physical. The die-to-die boundary is where the silicon is shipping now. The clinical signal is where the consequence of a spoofed input is measured in a patient.

The claim, reduced to what cannot be argued

A clinical AI model is only as trustworthy as the signal it acts on. In the cloud, a trusted pipeline made that signal's authenticity someone else's guarantee. At the edge, on hardware, that guarantee is gone, and today almost nothing fills it — the signal is trusted because it arrived, not because it was proven. It is an architectural gap that will have to be closed, deliberately, in hardware, and closing it correctly is a specialized discipline, not a feature to bolt on at the end. The sooner clinical AI at the edge treats the signal as something to authenticate rather than assume, the safer the devices built on it will be.

Dana Xiadani · Founder & Principal Architect, Living Cipher LLC · dana@livingcipher.org

Author's note: This paper describes a signal-authentication and trust architecture. It makes no diagnostic or clinical claim; the clinical inference it refers to is validated and regulated in its own right. Silicon results are substantiated by tool-generated reports, available to qualified reviewers on request.