BLOG 06 — THE EIGHT TIERS · PART 4

The AI Rules Your Regulator Already Enforces

Tier 4 of eight: the tier nobody writes about, and the one most likely to produce a finding against you in the next twelve months

Robin Beetge Co-founder, laibrary 10 min read

SHORT ANSWER

Sectoral regulators are already enforcing AI rules through supervisory machinery they have always had, without waiting for an AI statute to be written. The four instruments that changed most recently are the **NAIC AI model bulletin** for insurance (adopted in more than twenty US states), **SR 26-2** for US banking (April 2026, replacing SR 11-7), **revised FDA clinical decision support guidance** (March 2026), and **HTI-1** for certified health IT. If you operate in a regulated industry, at least one of them binds you today.

INSURANCE	BANKING	HEALTHCARE	HEALTH IT
NAIC AI model bulletin	**SR 26-2**	**FDA CDS guidance**	**HTI-1**

Almost every AI governance conversation skips this tier, which is remarkable given it is the one carrying enforcement machinery that already exists, staffed by examiners who already visit you.

Insurance is the sharpest illustration, so start there. It is the sector where the obligation is most explicit, where adoption is widest, and where the enforcement mechanism is something you already sit through every few years.

Insurance — the NAIC AI model bulletin

The NAIC Model Bulletin on the Use of Artificial Intelligence Systems by Insurers was adopted on 4 December 2023. It is principles-based rather than prescriptive, and it is not a model law — which means it takes effect in a state only when that state's insurance department adopts it. **By mid-2026 more than twenty US jurisdictions had adopted it** in full or in substantially similar form.

What it requires is a written **AI Systems Program** covering the full AI lifecycle. In practice: documented governance with named accountability, model validation before and during deployment, and oversight of third-party AI vendors and the data they supply. The bulletin expects that programme to be proportionate to

the insurer's use of AI and to the potential for consumer harm — the usual principles-based formulation, and the reason two insurers of similar size can legitimately land in very different places.

The enforcement mechanism is the part that catches people out. There is no new AI regulator here, and no new filing. Compliance is tested through **market conduct examination** — the same examination you already undergo. The evidence request arrives on a schedule you already know, from examiners with decades of practice extracting documentation from regulated entities.

Two features distinguish it from the banking regime, and both matter. It reaches generative and operational AI, which SR 26-2 expressly does not. And it applies across the insurance lifecycle rather than to models narrowly defined — underwriting, pricing, claims handling, fraud detection, marketing.

Two state variations are worth knowing. Colorado layers SB 21-169 quantitative testing regulations on top, for life insurance underwriting using external consumer data. And New York did not adopt the NAIC bulletin at all, issuing DFS Circular Letter 2024-7 instead, which sets fairness principles and proxy assessment expectations through supervisory oversight rather than prescriptive rules.

One more thing, because it connects to the preemption fight running through US AI policy. When the White House published its national AI legislative framework in March 2026 recommending preemption of state AI laws, the NAIC responded directly — publishing an issue brief reaffirming state authority over insurance AI oversight under McCarran-Ferguson, and urging Congress to protect the state-based system. That tension is unresolved. Insurance is the sector where it is most likely to be litigated.

> *If you are an insurer, the NAIC AI model bulletin has its own piece — what the AI Systems Program has to contain, which states have adopted it, and what a market conduct examiner asks for.*

The same pattern, four more regulators

Insurance is the clearest case, not the only one. In banking, healthcare, employment and half a dozen other sectors, the same thing is happening: a regulator that already has jurisdiction over you is applying existing supervisory machinery to AI, without waiting for an AI statute to be written.

Here is where each of them stands.

Banking — SR 26-2 replaced fifteen years of doctrine

On 17 April 2026 the Federal Reserve, OCC and FDIC issued SR 26-2 and OCC Bulletin 2026-13, **rescinding SR 11-7 and OCC Bulletin 2011-12**.

WHAT SURVIVED

The four pillars of model risk management — an enterprise model inventory, independent validation, ongoing monitoring, and documentation — and the central concept of **effective challenge**, meaning critical analysis by objective, informed parties with the standing to act on their findings.

WHAT CHANGED

The posture is now explicitly risk-based and materiality-sensitive, scaled to institution size and complexity rather than applied uniformly.

THE PART THAT MATTERS MOST

Generative and agentic AI are expressly placed outside the new framework's formal scope, described as novel and rapidly evolving.

The agencies signalled a forthcoming request for information on model risk management and bank use of AI.

Banks are therefore in an awkward position. Supervisors will still expect model risk discipline applied to consequential AI. There is simply no instrument specifying what good looks like. That gap is where examination findings will land.

Adjacent obligations do not wait: ECOA and Regulation B require specific reasons in adverse action notices, which bites hard on model-driven credit decisions regardless of any AI-specific rule. The EU AI Act treats creditworthiness assessment as Annex III high-risk. DORA captures AI vendors within ICT third-party risk.

Keep the two financial regimes distinct, by the way. An insurer follows the NAIC bulletin in adopting states and does not answer to SR 26-2. A bank answers to SR 26-2 and its prudential framework, not the NAIC bulletin.

Healthcare — two separate gates

GATE 1

The FDA gate

In January 2026 the FDA announced revised guidance on clinical decision support software, finalised on 11 March 2026 and replacing the September 2022 version. The material change: **single-recommendation CDS can now qualify for the non-device exemption** where clinicians can independently review the basis for each recommendation. AI or machine learning CDS that processes medical images or in-vitro diagnostic signals remains regulated as software as a medical device.

The Predetermined Change Control Plan final guidance of August 2025 remains the most operationally useful FDA instrument for AI, allowing planned and validated post-market model updates without a new submission.

GATE 2

The certification gate

ONC's HTI-1 rule has applied since 2024, with the Decision Support Interventions criterion part of the Base EHR definition from 1 January 2025. Developers of certified health IT must publish **31 structured source attributes** for predictive decision support interventions, implement intervention risk management practices, and meet the FAVES criteria — fair, appropriate, valid, effective and safe.

> **The critical point: HTI-1 applies whether or not the FDA exempts the software, if it sits inside a certified EHR.**
>
> A CDS tool can be simultaneously outside FDA device regulation and inside HTI-1 transparency obligations.

HIPAA adds no AI-specific rule, but the existing Privacy and Security Rules apply with full force. In practice that means cloud LLM processing of protected health information requires a Business Associate Agreement, which is why serious deployments run on platforms offering one or on self-hosted models.

Employment — the most consistently regulated AI use case

NYC Local Law 144 bias audits. Illinois HB 3773 and the AI Video Interview Act. Colorado SB 26-189. California's FEHA automated decision system regulations. EU AI Act Annex III category 4. The EU Platform Work Directive on algorithmic management. Korea's high-impact employment category. And underneath all of it, Title VII and ADA adverse impact analysis, which applied long before anyone wrote an AI statute.

If you screen, rank, assess or monitor workers with software, you are in scope somewhere.

Other sectors worth knowing

Automotive UNECE WP.29 regulations R155, R156 and R157, plus ISO 21448 and ISO 26262.	**Aviation** The EASA AI Roadmap and concept papers with their Level 1 to 3 classification.
Defence DoD Directive 3000.09 on autonomy in weapon systems and the Political Declaration on Responsible Military Use of AI.	**Elections** Synthetic political advertising laws in more than thirty-eight US states.

Why this tier is underweighted

Tier 4 lacks novelty. There is no launch moment, no press cycle, no countdown clock. A market conduct examination is not a news event.

It is also the tier where the enforcement infrastructure is mature. Banking examiners, insurance market conduct examiners and FDA inspectors have decades of practice extracting evidence from regulated entities. They are not waiting for an AI-specific statute to start asking.

Questions people ask

What is the NAIC AI model bulletin?

Principles-based guidance adopted by the NAIC in December 2023, requiring insurers to maintain a written AI Systems Program covering governance, model validation and third-party oversight. It takes effect in a state only on that state's adoption, and is enforced through market conduct examination.

How many states have adopted the NAIC AI model bulletin?

More than twenty jurisdictions as of mid-2026, in full or substantially similar form.

Does SR 26-2 cover generative AI?

No. Generative and agentic AI are expressly outside its formal scope, though supervisors still expect model risk principles to be applied to consequential AI.

Does HTI-1 apply if the FDA does not regulate my software?

Yes, if the software sits within certified health IT. The two gates operate independently.

NEXT IN THE SERIES

ISO 42001 certification

The first AI standard you can be audited against, and the piece of it that took two more years to arrive.

THE EIGHT TIERS SERIES

Part 0: The eight kinds · Part 1: Comprehensive statutes · Part 2: The treaty nobody is watching · Part 3: Where the obligations actually are · **Part 4: The regulators who were already there** · Part 5: ISO/IEC 42001 · Part 6: NIST AI RMF · Part 7: Security taxonomies · Part 8: Principles · Part 9: Putting it together