VERTICAL 01 — INSURANCE · COMPANION TO BLOG 06

NAIC AI Model Bulletin: What Insurers Must Document

The AI Systems Program, which states have adopted it, and what a market conduct examiner asks to see

Tanner Beetge Co-founder, laibrary 8 min read

SHORT ANSWER

The **NAIC AI model bulletin** requires insurers to maintain a written **AI Systems Program** covering the full AI lifecycle, with documented governance, model validation and third-party oversight. It was adopted by the NAIC on 4 December 2023, it is principles-based rather than prescriptive, and it is not a model law — it binds an insurer only in states whose insurance department has adopted it. More than twenty US jurisdictions had done so by mid-2026. Compliance is tested through market conduct examination.

ADOPTED	JURISDICTIONS	REQUIRES	ENFORCED BY
4 Dec 2023	20+	AI Systems Program	Market conduct exam

Insurance is the most regulated use of AI in the United States, and it got there without a single AI statute.

There is no new regulator here, no new filing, and no new enforcement mechanism. What there is instead is a bulletin — guidance issued by state insurance commissioners — that tells insurers what their existing obligations mean when the decision is made by a model.

That makes it easy to underestimate, and it is the reason most insurers I speak to have a policy but not a programme.

What it is, and what it is not

The NAIC Model Bulletin on the Use of Artificial Intelligence Systems by Insurers was adopted on 4 December 2023.

Three things about its legal character matter more than anything in its text.

01 **It is not a model law.** The NAIC produces both, and the difference is substantial. A model law is drafted for state legislatures to enact. A bulletin is guidance a commissioner issues under existing authority — which is why it arrived in months rather than years, and why it reaches conduct that no new statute has had to define.

02 **It binds you only where it has been adopted.** The bulletin has no effect in a state until that state's insurance department issues it. By mid-2026 more than twenty US jurisdictions had adopted it, in full or in substantially similar form. If you write in thirty states, your obligations differ across them.

03 **It is principles-based.** It expects your programme to be proportionate to your use of AI and to the potential for consumer harm. Two insurers of similar size can legitimately land in very different places — which is freedom if you have documented your reasoning, and exposure if you have not.

The AI Systems Program

The central requirement is a **written AI Systems Program** covering the AI lifecycle. In practice, an examiner is looking for four things.

01

Governance with names attached.

Who owns AI risk, who approves deployment, who can stop a model, and where that sits relative to the board. "The data science team handles it" is not an answer that survives an examination.

02

Risk management and internal controls.

How AI systems are assessed before deployment and monitored after it. Validation, performance monitoring, drift detection, and a defined route for escalating a model that has started behaving differently from the one you approved.

03 · Thinnest area

Third-party and vendor oversight.

This is the one most insurers are thinnest on. Most insurance AI is bought, not built — pricing models, fraud detection, claims triage, marketing tools. The bulletin expects you to have diligence over the vendor, over the model, and over the data the vendor supplies, with contractual rights to the information you would need in an examination.

04

Documentation that a regulator can review.

Development process, data sources, validation results, risk mitigation measures. Maintained, current, and producible on request.

How it is enforced

Through **market conduct examination** — the same examination you already undergo.

This is the part that catches people out, and it cuts both ways. There is no separate AI audit to prepare for, which sounds like relief. But it also means the evidence request arrives inside a process you cannot decline, on a timetable you already know, from examiners with decades of practice extracting documentation from regulated entities.

> **An examiner who asks how your pricing model treats a protected class is not going to accept a description of the model.**
>
> They will ask for the validation record, the date it was run, who reviewed it, and what happened to the finding.

Two differences from the banking regime

Insurers and banks are often lumped together as "financial services AI." Their AI obligations are not the same, and the differences run in the insurer's disfavour.

SCOPE	BREADTH
In April 2026 US banking supervisors issued SR 26-2, replacing SR 11-7, and expressly placed generative and agentic AI outside its formal scope. The NAIC bulletin has no such carve-out. It reaches generative and operational AI as readily as a traditional pricing model.	SR 26-2 governs models as defined in model risk management. The bulletin applies across the insurance lifecycle — underwriting, pricing, claims handling, fraud detection, marketing — wherever an AI system informs a decision affecting a consumer.

An insurer follows the 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 bulletin. Advice written for one is frequently wrong for the other.

Two state variations worth knowing

Colorado	New York
Layers SB 21-169 on top, with quantitative testing regulations for life insurance underwriting that uses external consumer data and predictive models. That is a materially heavier evidentiary burden than the bulletin alone. Insurers complying with the insurance-specific regime were exempted from the broader Colorado AI Act.	Did not adopt the bulletin at all. DFS issued Circular Letter 2024-7 instead, finalised in July 2024, which sets fairness principles and proxy assessment expectations for insurers using AI or external consumer data in underwriting and pricing. It works through supervisory oversight rather than prescriptive rules, and it applies whether or not you have an AI Systems Program in the NAIC sense.

If you write in both states plus the twenty-odd that adopted the bulletin, you are running three overlapping regimes.

The preemption question

In March 2026 the White House published a national AI legislative framework recommending that Congress preempt state AI laws deemed unduly burdensome, and channel oversight through existing sector agencies.

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 and likely to stay unresolved for some time. McCarran-Ferguson gives insurance a stronger claim to state primacy than most sectors, which makes this the industry where the preemption fight is most likely to be tested.

The practical posture is unchanged: comply with the bulletin in adopting states while tracking the argument. A programme built on the assumption that preemption arrives is a programme built on a contested proposal.

What to do first

1. **Inventory before policy.** You cannot document governance over AI systems you have not listed. Start with procurement records and vendor contracts — in insurance, most of the AI arrived that way.

2. **Map your adopting states.** Your obligations are state-specific. Know which of your writing states have adopted the bulletin, which have variations, and which have nothing.

3. **Fix vendor oversight early.** It has the longest lead time, because it usually means going back to contracts that were signed without AI diligence rights in them.

4. **Build the evidence, not the binder.** An examiner does not ask whether you have an AI Systems Program. They ask you to show that a control operated, on a date, for a named system, with an artefact attached.

Questions people ask

What is the NAIC AI model bulletin?

Guidance adopted by the NAIC on 4 December 2023 requiring insurers to maintain a written AI Systems Program covering governance, risk management, internal controls and third-party oversight across the AI lifecycle. It is principles-based, takes effect only in states that adopt it, and is enforced through market conduct examination.

Is the NAIC AI model bulletin legally binding?

Only in states whose insurance department has adopted it. It is a bulletin issued under existing commissioner authority, not a model law enacted by legislatures.

How many states have adopted the NAIC AI model bulletin?

More than twenty US jurisdictions as of mid-2026, in full or in substantially similar form. New York did not adopt it, issuing DFS Circular Letter 2024-7 instead.

Does the NAIC AI model bulletin cover generative AI?

Yes. Unlike SR 26-2 in banking, which expressly excludes generative and agentic AI from its formal scope, the bulletin has no such carve-out.

COMPANION PIECE · BLOG 06

The AI rules your regulator already enforces

This is a companion to the AI rules your regulator already enforces, which covers the same pattern in banking, healthcare and employment.