



**SIGNAL 1**

# The Next Healthcare AI Challenge

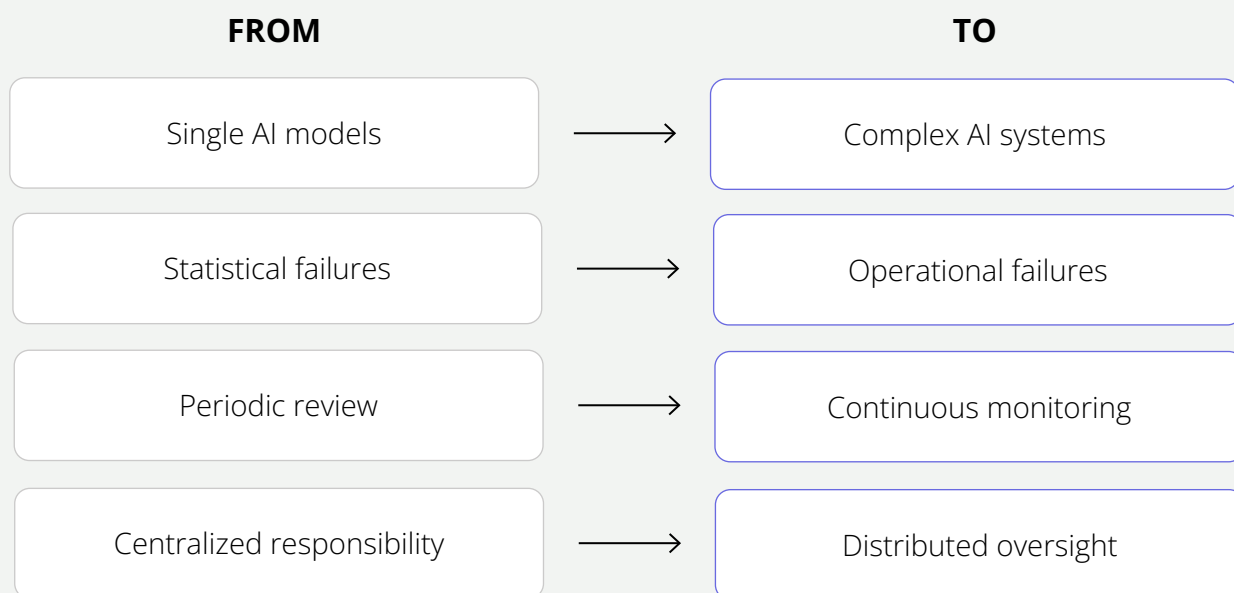
## SIGNAL 1

# The Next Healthcare AI Challenge Isn't Building AI Models. It's Managing AI Systems.

Healthcare AI is entering a new phase.

For the past few years, health systems have been deploying AI tools that support human workflows. Predictive models that generate risk scores. Documentation copilots that create drafts. In nearly every case, humans remained firmly in the loop, choosing whether and how to use the output of the AI.

Now that is changing. With the explosion of agentic AI, AI's role is evolving from **'passively assisting'** to **'actively participating'** in care delivery and health system operations — and that changes how healthcare organizations must approach AI governance and management.



### 1 From AI models to AI systems

Traditional healthcare AI governance was built around the model. A predictive algorithm or documentation copilot could be characterized by its training data, performance metrics, and intended use — captured in a model card, reviewed by a committee, and approved for deployment. The model was the unit.

Agentic AI changes the unit. Behavior now emerges from the interaction of models, data, tools, context, orchestration logic, and the environment the agent operates in. Two agents built on the same underlying model can behave very differently depending on what tools they can invoke, what data they can read, what instructions shape their scope, and what other systems they coordinate with at runtime.

This shift has direct implications for governance. Evaluating an agent means evaluating the system as a whole: its scope of work and effective job description, the data and tools it has access to, the instructions and context that shape its behavior, the workflows it participates in, and the boundaries it is expected to stay within. What was once a process of capturing model information in a model card has become substantially more complex — closer to onboarding and supervising a new role than to validating a statistical artifact.

## **2 From statistical failures to operational failures**

For traditional healthcare AI, failure is largely a statistical problem or an adoption problem. Either the model is failing to perform as expected, or users are not acting on the output as intended. Either way, the consequences of failure are bounded by the human-in-the-loop. Governance appropriately focuses on things like model accuracy, bias, data drift, and workflow integration.

Agentic systems fail in far more ways than traditional AI. An agent can reach the right conclusion and still take the wrong action — for example, by invoking the wrong tool or following an instruction it should have refused. Many failure modes only emerge from the specific combinations of context, inputs, and timing encountered once the system is operating in production. Pre-deployment testing cannot easily anticipate them all. The consequences of failure are also greater. When the agent is acting rather than suggesting, the gap between AI error and a patient-safety event narrows considerably.

Beyond increasing the consequences of failure, agentic systems introduce a new layer of organizational complexity because appropriate behavior often depends on the local clinical and operational environment. As a result, oversight cannot function as a purely centralized activity. Managing agentic AI will require a hybrid operating model that combines enterprise-level oversight with distributed operational accountability.

## **3 From periodic review to continuous oversight**

Pre-deployment review remains necessary — approval committees, validation exercises, and risk classification still matter. But for agentic systems operating in real clinical and operational environments, post-deployment monitoring becomes an absolute necessity.

Many existing governance programs are designed around periodic reviews and lightweight production monitoring. What changes with agents is the amount of oversight required after go-live, how granular that oversight must be, and the speed at which organizations need to intervene and respond. Many agentic systems require continuous evaluation of the tools and data an agent is trying to access, the tasks it is attempting to carry out, and the conformity of its behavior with policies with SOPs, escalation rules, and clinical guidelines.

The challenge is no longer simply validating a model before deployment. It is continuously ensuring that autonomous systems are behaving appropriately, safely, and within organizational policy as conditions change over time.

# Building the management layer

At Signal 1, we believe healthcare organizations need an AI management layer that is purpose-built for healthcare delivery, governance, and enterprise-scale oversight. That belief led us to build AIMS, our AI Management System.

AIMS enables organizations to centrally govern and manage AI agents with shared policies and context, auditable workflow orchestration, risk-tiered oversight, and continuous monitoring to ensure agents operate within their approved roles and boundaries.

The organizations that build this agentic management layer now will be the ones best positioned to safely scale AI, realize enterprise value, and lead the next era of healthcare delivery.

Ready to accelerate agentic deployment at your health systems? Contact us to learn how Signal 1 is helping health systems prepare for what comes next.

