The Governance Institute

BOARD BRIEFING

Who checks AI before it acts on a patient?

Human oversight means different things before and after a prescription is sent.

5–6 MINUTE READ

In This Briefing

A clinician can check a proposed prescription before it is issued or review it afterward. Those arrangements offer different opportunities to catch an error. Utah's limited pilots make the distinction concrete as software takes on more responsibility in care.

About Board Briefings

Board Briefings offer deeper explanations of healthcare issues and their implications for governance. Each includes a practical example, authoritative sources, and key questions to support informed board discussion.

What Happened

On October 5, Utah announced an expansion of its healthcare artificial intelligence pilots, including prescription refills from August AI and selected acne treatments from Nolla Health. Both begin with clinician review before prescriptions proceed. These are limited tests, not evidence that independent AI prescribing has proved safe. (1), (2)

August AI was authorized October 2, although its demonstration had not begun in the state register. Nolla's pilot became active October 5. Later stages can reduce the amount of review only with further written state approval. (2)

The Issue, Explained

A patient asks for another supply of an existing medicine. Software gathers information and prepares a proposed refill. A clinician checks it, changes it if necessary, and decides whether it should proceed. The software has helped with the work, but a person has made the decision.

In a different arrangement, software might be permitted to authorize the refill within a defined process, with a clinician reviewing it afterward. The service would still have human oversight. It would not have the same opportunity for that person to stop an error before the prescription is issued.

Utah is testing narrowly defined versions of this change. August AI's pilot concerns selected maintenance medicines an adult already takes under an established plan, not unrestricted new prescribing or dose changes. Nolla's pilot concerns selected adults with mild to moderate acne and a limited set of medicines applied to the skin. (2)

The launch safeguards are specific. August AI starts with a licensed provider checking refills before they proceed. Nolla starts with two licensed physicians reviewing every prescription before it is sent. Later phases may use review afterward or sampling, subject to separate state permission. (2)

Those checks are part of the authorized arrangement, not an incidental extra. Permission to test that arrangement does not establish that it has succeeded, or authorize the same use elsewhere. The development is consequential because it explores clinical authority, not because clinicians have simply been removed.

The Issue, Explained (continued)

A review also depends on the circumstances in which it takes place. The clinician needs relevant information, time, and authority to intervene. Voluntary federal AI guidance emphasizes clear responsibilities and evaluation suited to the intended use. It does not prescribe one review arrangement for every clinical task. (3)

An overall accuracy score can hide the error that matters most. A tool may answer many harmless questions correctly while making mistakes in the smaller set that leads to a prescription. The proportion of correct answers does not explain the seriousness of the remaining errors.

Benefits also extend beyond the speed of the first step. Faster preparation may save time while checking exceptions or answering patients' questions adds work elsewhere. The gain depends on whether the whole service delivers appropriate care more reliably or sooner.

A limited rollout can produce evidence before the software's role expands. Its pace alone does not establish whether adoption is appropriate. The more relevant question is whether results support the specific authority being granted. New software versions, different medicines, or a broader patient group can change what earlier testing establishes.

A Practical Example

Two hypothetical services use tools that make the same mistake in a refill recommendation. At the first, a clinician checks every proposal before it is sent and catches the error. At the second, the refill proceeds and the error is found during later review.

The tools could have identical accuracy in a test set. Both services could say a clinician oversees the software. Only the first arrangement offers that particular opportunity to stop the error before it reaches the patient.

This is not an incident reported in either Utah pilot. It isolates the importance of timing without deciding that one method fits every task. The seriousness of the possible error, the reliability of the check, and the benefit of quicker care all remain relevant.

Governance Implications

Using software to prepare information is different from permitting it to authorize a treatment step. The latter changes the responsibilities attached to the care process. Expected time savings need to be understood alongside the clinical action, the people affected, and the safeguards around it.

Clinical and operating leaders evaluate and manage individual uses. The board oversees significant safety risks, accountability, and major investments. That does not require trustees to approve individual prescriptions. It requires a clear understanding of the responsibilities attached to the organization's commitments.

Evidence from a narrow use may support a next step without establishing suitability for every clinical decision. Results on safety and access, together with authority to change or pause a use, make expansion more understandable. The relevant assurance is not simply that a person is involved, but what that person can review and when.

Questions for the Board

1. **When reviewing clinical AI strategy:**
   Which significant clinical AI uses prepare information, and which can authorize or initiate an action affecting care?

2. **When assessing clinical AI safeguards:**
   How are clinical accountability and opportunities to detect or stop errors maintained as the software's role changes?

3. **Before considering broader clinical AI use:**
   What results and approvals would support broader authority, and what findings would lead to a pause or narrower use?

Sources & Further Reading

1. Utah Department of Commerce. Healthcare AI pilot expansion. October 5, 2026. Announcement of the limited healthcare pilots.

2. Utah Office of Artificial Intelligence Policy. Authorized Pilots. Pilot boundaries and review stages; later phases require written approval.

3. National Institute of Standards and Technology. Generative Artificial Intelligence Profile. July 26, 2024. Voluntary guidance, not approval of the pilots.

Board Briefings provide governance education and do not constitute legal, clinical, accounting, or investment advice.