FOR BOARD MEMBERS · AI AGENTS AT WORK

The Seat first. Then the Signal.

Two tests, in order. First: can the board hear? Then: is each agent fit to work, with a named human who answers for it?

TEST 1 · THE BOARD

The SEAT™ test

Can your board hear the people closest to the risk?

S — Signal
Can the person who spots an agent's mistake say so without penalty?

E — Escalation
Does a serious agent failure have a defined route upward, or does it stop in a Slack thread?

A — Authority
Who can switch an agent off, and has anyone practiced it?

T — Traceability
After the last agent error, can anyone show what was found and what changed?

Closest to the risk: the people who use agents.

THEN →

TEST 2 · EACH AGENT

The SIGNAL test™

Is each agent fit to work, and are the people who use it heard?

S — Scope
Does the agent have a stated purpose and a set autonomy tier?

I — Identity
Does it have its own identity and least-privilege access?

G — Guardrails
Does it run inside a harness and guardrails, under change control?

N — Named owner
Is a named human in the seat for it, able to pause it?

A — Assessment
Was it tested before release and on a cadence, against a pass bar set by someone outside the team that built it?

L — Listening
Do the people who use it know how to report its mistakes, and are they heard?

Named owner: fourth in the letters, first in practice.

THE PRINCIPLE

An agent is a tool, wherever it appears on the org chart. It can't sit in the seat. Every agent has a named human who answers for it.

START WITH T: ONE QUESTION FOR THE NEXT MEETING

"Show me the last five agent errors and what happened to each."

A framework by Bevin McArthur.* *A proposed framework, for discussion purposes only. Not legal, business, or tax advice. SEAT™ and the SIGNAL test™ are trademarks of Seat & Signal LLC. © 2026 Seat & Signal LLC. All rights reserved. seatandsignal.com ·