

Evidence Over Attestation: The Twenty-One-Point Executive Standard for Agentic AI

Lindsay Hiebert, Founder, PIGENAI LLC · AI governance, security, AEO and networking.

Published at pigenai.com/blog/evidence-over-attestation-standard, September 4, 2026.

The canonical version is always the one on pigenai.com.

Twelve disciplines, five operating rules, four board decisions. The condensed form of the four-part essay, for CEOs, boards and founders who deploy autonomous agents and need to know what to ask for. Everything here is derived from the four parts; nothing is claimed that they do not claim.

AI already has the speed. What your customers need from you is judgment: the competence to ask the right question, the discipline to prove the answer, and the responsibility to put the person on the other side of the screen first. A green light is a claim made by an entity incapable of feeling doubt. Evidence is an artifact. The twenty-one points below are how a board tells the difference.

I. The twelve disciplines, as a scorecard

Each discipline is reduced to the question a CEO asks and the evidence to ask for. The full principle, the trap and the recorded example for each are in Part Two.

Judgment

#	DISCIPLINE	THE QUESTION A CEO ASKS	THE EVIDENCE TO ASK FOR
1	Asking the right question	When the dashboard spiked, or went green, what did you ask before you acted, and what would have had to be true for the number to be misleading?	The count that preceded the conclusion, in the incident record: the byte rate, the request log, the timestamps.
2	Deep domain competency	Can you name the pathology from the raw data, and its effect on a customer, before	An incident where the diagnosis came from the team, not the tool: the raw

#	DISCIPLINE	THE QUESTION A CEO ASKS	THE EVIDENCE TO ASK FOR
		the vendor or the dashboard explains it?	reading, and the reading of it.
3	Focus on the one thing that matters	What is the single priority this week, whose customer is it, and what valuable work did you stop to protect it?	The list of work deliberately not started, and the products held in maintenance by rule.

Character

#	DISCIPLINE	THE QUESTION A CEO ASKS	THE EVIDENCE TO ASK FOR
4	Pre-committed standards	Where is the written definition of done, in the customer's terms, dated before the work started?	The gate with its date, and a case where it held against an artifact that looked finished.
5	The customer first, before the fix	In the last incident, who was affected, and when were they told, relative to the fix?	The disclosure record, timestamped against the patch.
6	The honest number	Which of our published numbers would we rather not publish, and do we publish it anyway?	A metric published at zero or below target, with the method and the transcripts behind it.

#	DISCIPLINE	THE QUESTION A CEO ASKS	THE EVIDENCE TO ASK FOR
7	Effectiveness over efficiency	What did you decide not to do this month, and why was waiting the right decision for the customer?	Deferred and cancelled items with the reasoning, in the record, not in memory.

Systems craft

#	DISCIPLINE	THE QUESTION A CEO ASKS	THE EVIDENCE TO ASK FOR
8	Conceptual integrity	Where is the single ratified statement of what the product is, what it promises, and what it will never claim?	The ratified statement, the check that rejects drift from it, and the last change it rejected.
9	Lessons hardened into gates	Show me the last lesson learned and the check that fails if it returns.	The test or lint tied to the retrospective, and a retrospective log in which no root cause appears twice.
10	Wearing the right hat	Who inspected the last change that touched customer data, how long did approval take, and how many roles did one person cover to get there?	The review record with the time from request to decision.

People

#	DISCIPLINE	THE QUESTION A CEO ASKS	THE EVIDENCE TO ASK FOR
11	Innovation for trust, not novelty	Which of the features we celebrate made the product safer, quieter or more dependable for the person using it?	The constraint that shipped (what the product now refuses to do), and the customer-facing statement of it.
12	The fingertip standard	Walk me through our worst customer-facing failure from her side of the screen. Who fixed it, and did the fix protect her or the server?	The customer-side reproduction, and the fix described in terms of what the customer sees.

II. The five operating rules

The minimum a CEO should require of anyone, employee or partner, who builds with agents on the company's behalf. Under each rule: what to ask for at the next review.

13. Count before you conclude. The obvious narrative is a hypothesis, never a finding. *Ask for:* the counts, the records and the timestamps that preceded the remediation.

14. Define “done” before the first instruction. In the customer's terms, and before the work begins. *Ask for:* the written definition of done, dated earlier than the first line of work, readable by a non-technical executive.

15. Close only on outside-in proof. Internal tests, telemetry and an agent's attestation inform; none of them closes the case. *Ask for:* the artifact from the customer's vantage point against the version actually serving: the fetch, the timing, the diff.

16. Make courteous failures loud. A silent fallback converts your problem into your customer's discovery. *Ask for:* the inventory of fallbacks, the alert attached to each, and who was paged the last time one fired.

17. No artifact, no closure. Attestation is a claim; evidence is an artifact. *Ask for:* the identifier of the thing that is serving and the record of the check made against it, attached to every item marked complete.

III. The four board decisions

18. Audit your attestation gap. Inspect the internal AI projects. Are teams measuring agent success by generated volume, such as lines of code, automated summaries and closed tickets, or by verified outcomes for customers?

19. Ban self-certifying workflows. No autonomous agent, AI pipeline or team evaluates its own output without an external, independent verification gate stated in the customer's terms.

20. Empower the domain experts. Move generative AI budgets away from generic, centralized sandboxes and into the hands of the people who know the customer's problem from the inside.

21. Reward the honest zero. Promote the executives and teams who halt a deployment to fix a measurement flaw, rather than those who ship a flattering estimate.

The standard in one line

The machines have brought the speed. It is time for human leadership to bring the standard, and the people who can hold it. The durable moat is not access to intelligence. It is knowing what matters, knowing what good looks like, and refusing to close without evidence.

*Lindsay Hiebert, Founder, PIGENAI LLC · AI governance, security, AEO and networking.
PIGENAI LLC, in Kansas City, builds and operates products for AI governance, AI visibility*

and AI trust. The full argument is the four-part essay beginning at pigenai.com/blog/evidence-over-attestation.