

LEAN METHOD

AI readiness and governance posture

A practical questionnaire. One room, one afternoon, evidence on the table. Mapped to Observational Governance on KanTec Continuum.

1. How to use this

This is not a maturity model with five poetic stages. It is a lean method: twenty-eight questions, three scores, one heat map, one beachhead decision.

- Room: CISO or CDO, platform lead, one AI / product owner, one risk lead. Optional: CRO delegate.
- Time: 90–150 minutes. If it needs a week, you are writing essays instead of scoring.
- Rule: score only what you can show in the room — a screen, a ticket, a repo, a dashboard. Intent scores zero.
- Scale: 0 = cannot show / not true. 1 = partial, one team or one path. 2 = true on the paths that matter, with evidence.

Maximum 56. Record owner and artefact in the last column or the score is decoration.

2. Scoring bands

0–18 Exposed	Policy and agents are disconnected. Do not scale AI. Commission a beachhead and name stop authority this month.
19–36 Partial	Fragments exist (a register, a cluster policy, a model card). The loop is not closed. Pick one path and finish it.
37–48 Governable	A path can be observed. Clone it. Do not invent a second control plane for the next use case.
49–56 Observational	Evidence can change policy. Hold the standard. Audit pack is a query. Measure bypass rate.

3. Questionnaire

A. Authority (items 1–4)

If these four are weak, stop the workshop and name people before discussing tools.

#	Question — answer 0 / 1 / 2	Score	Evidence / owner
1	Is there a single written objective for AI and automated action (what must never happen unsupervised)?		
2	Is a named human the stop authority for consequential AI / agent action?		
3	Is a named human the exception owner, with a register of live exceptions and expiry dates?		
4	Can the board see last month's AI / agent incidents and exceptions on one page?		

B. Inventory (items 5–8)

#	Question — answer 0 / 1 / 2	Score	Evidence / owner
5	Does a register exist of systems, models, agents and tools that can take action?		
6	Does every registered actor have an owner and a risk tier?		
7	Can you list agents in production in under one hour without a discovery project?		
8	Are shadow / unofficial AI tools (browser copilots, unmanaged SaaS) known and either banned or registered?		

C. Policy that executes (items 9–12)

#	Question — answer 0 / 1 / 2	Score	Evidence / owner
9	Does written policy exist in dual form — intent in language and an executable artefact with the same version?		

10	Is policy evaluated at the point of action (PDP), not only in a quarterly review?		
11	Does a PEP force the decision on at least one live path (pipeline, cluster admission, API gateway, or agent tool-firewall)?		
12	On consequential actions, does PDP timeout fail closed?		

D. Actors and supply path (items 13–16)

#	Question — answer 0 / 1 / 2	Score	Evidence / owner
13	Do models and agents have first-class identities (not a shared service account)?		
14	Are agent tool calls scoped, and is out-of-scope use blocked rather than logged only?		
15	Is there a human gate on payments, identity grants, external send, and model promotion?		
16	Can an unsigned image or unregistered agent artefact reach production today?		

Item 16 is reverse-scored in practice: if an unsigned artefact can reach production, mark 0.

E. Evidence and loop (items 17–20)

#	Question — answer 0 / 1 / 2	Score	Evidence / owner
17	Do allow / deny / exception events write to an append-only evidence store with policy version?		
18	Can a deny event be joined to an application or agent trace?		
19	Can you produce a 30-day audit pack for one system as a query, with no workshop?		
20	Has any policy change in the last quarter cited a production evidence_id?		

F. Platform and sovereignty (items 21–24)

#	Question — answer 0 / 1 / 2	Score	Evidence / owner
21	Is there one golden path for build → sign → deploy of AI / agent workloads?		
22	Is bypass of that path measured weekly?		
23	Are integration interfaces attested (no dark APIs into models or agents)?		
24	Are residency and operator-access rules enforced as controls, not only written as policy?		

G. Operating rhythm (items 25–28)

#	Question — answer 0 / 1 / 2	Score	Evidence / owner
25	Do risk / compliance run continuous tests against the evidence store, or only interviews?		
26	Is there a weekly or fortnightly govern review that looks at denials, exceptions and drift?		
27	Are SLOs defined for policy-deny handling (detect, decide, expire)?		
28	Would a new agent be refused production if it could not emit the evidence contract?		

4. How to read the heat map

Sum each block (max 8). A 7 on inventory and a 1 on evidence means you know what you have and cannot prove what it did. That is the usual pattern. Do not average it into ‘medium maturity’.

- Lowest block is the beachhead constraint. Start there.
- Authority below 4: no further AI go-lives until names exist.
- Item 14 or 16 at 0: agents and supply chain are the first PEP, not a centre of excellence.
- Item 19 at 0: you have monitoring, not Observational Governance.

5. Lean output — one page

Leave the room with five lines only.

- Score and band.
- Lowest block and the two questions that hurt.
- Beachhead named (one system, one process, or one platform path).
- Stop authority and exception owner, written.
- Day-90 tests you will hold: replay a decision, join a deny to a trace, expire an exception.

That page is the brief for Continuum delivery. The full architecture sits in the Observational Governance reference architecture. This method decides whether you need it this quarter — and where to put it.

6. What this method refuses

- Self-score by the AI team alone.
- Vendor questionnaires substituted for evidence.
- A five-level maturity narrative with no owner.
- Scoring 'in progress' as a 2.

— KanTec Continuum · Lean AI readiness and governance posture —