

Build the foundation for production AI

A note for CIOs. KanTec Continuum. Observable Governance.

GOVERNANCE FLOWS DOWN · OBSERVABLE FLOWS UP

Governing thought

Production AI is not a model problem. It is a governing-system problem. If policy cannot travel down the stack, and evidence cannot travel back up, the organisation has pilots — not production.

KanTec Continuum is the operating model. Observable Governance is the control standard. Together they are the foundation a CIO can put under every consequential agent, model and interface — without standing up a second plane “for AI.”

1. Why this, why now

Situation

Every CIO already funds cloud, integration, identity, security and a growing AI portfolio. Boards have approved principles. Delivery teams have shipped copilots and first agents.

Complication

- Shadow AI is already calling systems of record on paths that were never attested.
- Enterprise integration is either a golden path or a dark one. Agents find the dark path first.
- Questionnaires still assume humans click. Agents call. A deny that cannot be queried is not control.
- A second control plane for AI fragments authority and doubles exception debt.

Question

What must be true before the next wave of agents is allowed to act in production?

Answer


Governance flows down. Observable flows up. One loop. One evidence store. One named stop owner. Deployed as a beachhead path, then cloned. That is Continuum. That is Observable Governance.

2. The foundation — three claims only

Figure. Operating model — policy down the spine, evidence up the return rail. Policy version ticks when the loop closes.

Claim	CIO test
Authority is executable	Policy exists in dual form — intent and artefact, same version. A named stop authority can halt a path.
The estate is bound	Every consequential interface — including agent tool-calls — sits on one attested path. Unscoped calls fail closed.
The loop closes	A deny writes evidence with policy version and trace. The audit pack is a query. The next policy merge cites those evidence_ids.

If any test fails, you do not have production AI. You have supervised experimentation wearing a production badge.



3. KanTec Continuum as the operating stack

Continuum is not another tool category. It is the stack that already exists — made governable and observable.

- Authority plane. Appetite, stop rights, exception owner. If unnamed, stop.
- Control plane. Registry, dual-form policy, PDP/PEP, exception bus, evidence store.
- Capability plane. Six Continuum domains. Each must emit observable outcomes.
- Runtime. Applications, models, agents, pipelines — subjects of policy, not exceptions to it.
- Evidence / decision. Outcomes measured, risk evaluated, insights issued, policy evolved.

Partner products (IBM webMethods, Red Hat, security, identity) implement a plane. They do not replace it. A second plane “for AI” is a defect.

4. Observable Governance as the strategic objective

Observable Governance is the CIO’s standard, not a security workstream. It means:

- Governance flows down — policy binds PEPs on the path agents actually use.
- Observable flows up — every allow/deny is queryable with actor (including agent), asset, policy version, PEP class, exception and trace.
- Shadow AI is named traffic, not folklore. Bring it onto an attested interface or deny it.
- Enterprise integration is the fabric. Dark point-to-point is how unofficial copilots escape.

SIEM, questionnaires and model cards are inputs. They are not the system of record. The system of record is the evidence_event.

5. What the CIO funds — and what they refuse

Fund

- A named stop authority and exception owner before the first agent lands.
- One beachhead path: one process, the interfaces it uses, PEPs live, evidence query working.
- The evidence contract (append-only, policy version, trace join).
- Attested integration (webMethods or equivalent) treated as a PEP class, not a connector catalogue.

Refuse

- A parallel AI control tower.
- Unbounded agent tool-calls “while we write the policy.”
- Inventory without proof — catalogues that cannot show last month’s deny.
- Rip-and-replace of the integration estate as a prerequisite. Map dark interfaces first.

6. 90-day beachhead

Window	Done when
Day 0	Stop authority named. If not, do not proceed.
Days 1–30	Registry + dual-form policy on one path only. Shadow-AI interfaces on that path listed.
Days 31–60	PEPs live (gateway + agent tool-firewall at minimum). Unsigned artefact cannot land. Consequential + PDP timeout = deny.
Days 61–90	Weekly loop. Evidence query works. First policy merge cites production evidence_ids. Board pack is that query.

Scale by cloning the path. Do not scale by buying a second stack.

7. How a CIO should talk about this

To the board: “We are not asking you to buy observability. We are making governance true in production. Policy must be able to travel down. Evidence of whether it held must travel back up. If either path is missing, we have a toolkit, not a governing system.”

To product and data: “Agents are interfaces. They use the same golden path as every other caller. Unscoped tools fail closed.”

To security and risk: “You get a query, not a workshop. Exceptions expire. The deny is the control.”

To partners: “Implement a plane. Do not become the plane.”

8. Decision the CIO is being asked to take

Approve Continuum as the production-AI foundation: one beachhead path, Observable Governance as the standard, no second control plane. Name the stop owner this week. Fund the 90-day loop. Measure success as a working evidence query — not as a model count.

— KanTec Platform Engineering · CIO briefing —