Governance Architecture Checklist

Is your governance written as an architecture, or only as a policy?

A policy document does not stop an ungoverned call. Use this list to map every written rule onto the component that actually enforces it at runtime.

MAP POLICY TO CONTROL

- [ ] For each policy statement, name the component that enforces it.
- [ ] Mark rules enforced only by training or trust — those are gaps, not controls.
- [ ] Define the failure mode: what happens when a control is bypassed.
- [ ] Record the evidence each control emits, and where it lands.

LAYER THE CONTROLS

- [ ] Intelligence layer: model access, key custody, provider terms, egress rules.
- [ ] Operating layer: permissioning, prompt/response logging, evaluation, rate limits.
- [ ] Application layer: user-facing disclosure, human-in-the-loop, override paths.
- [ ] Confirm no control is enforced only in the UI.

PROVE IT

- [ ] Run one adversarial case per control and keep the result.
- [ ] Version the architecture and date every change.
- [ ] Assign an owner to each layer, not just to the document.
- [ ] Schedule a re-test cadence tied to model or vendor changes.

Compare your model against the published A/DVNT AI Governance Framework at advnt.ai/resources/ai-governance-framework.