Governability Decision Scorecard

Score the controls that determine whether an agent workflow is approvable

Feature lists do not approve production deployments. Governance does. This scorecard gives security teams a compact way to compare approaches based on the controls that matter most: what can be shown, what can be sent back, how runtime is contained, how evidence is produced, and how much operational friction the model introduces.

Use it to compare vendors, architectures, or internal approaches on the factors that most often decide whether a deployment stays stuck in review or moves forward.

Use this scorecard when:

- Comparing two or more platforms or approaches
- Deciding whether a pilot is ready to move forward
- Identifying where compensating controls are still required

① ② ③ ④ ⑤	**Downstream control strength** (pre-render redaction, granularity, determinism)
① ② ③ ④ ⑤	**Upstream control strength** (input governance, tool/action gating, exfil resistance)
① ② ③ ④ ⑤	**Containment** (centralized runtime, network segmentation, “agents not in the wild”)
① ② ③ ④ ⑤	**Evidence-grade audit** (integrity + time correctness + chain-of-custody + exportable evidence packages)
① ② ③ ④ ⑤	**Operational friction** (deploy and maintain)
① ② ③ ④ ⑤	**Policy UX** (testability, change control, reviewability)

How to interpret the score

- **24–30:** Strong governability posture. The approach appears production-ready if the use case and evidence hold up in validation.
- **18–23:** Promising but incomplete. Expect targeted gaps, compensating controls, or tighter workflow scoping before approval.
- **12–17:** Material governance risk. The approach may work technically, but approval will likely stall without architectural change.
- **Below 12:** Weak governability posture. Treat as exploratory until core controls are redesigned.