Northlight

STRATEGY · INSIGHT · AI · APPS

AI GOVERNANCE READINESS

BASED ON THE CODER/WEAVE 7-STEP ACTION PLAN

01	Have you audited which AI tools are currently in use across all teams, including unsanctioned ones?	○
02	Do you have a centralized, policy-controlled development environment (not just local machines or personal SaaS tools)?	○
03	Are agent permissions explicitly separated from human developer permissions?	○
04	Do agents default to least-privilege access (no internet, no write access to repos)?	○
05	Is there observability infrastructure in place to log every model interaction, tool call, and resource access?	○
06	Can your compliance team produce a full audit trail of agent actions on demand?	○
07	Do you have a documented process for how agents escalate or pause when they hit a boundary?	○
08	Have you piloted governed agent use with at least one high-signal team before scaling?	○
09	Are your workspaces ephemeral (spun up per task, torn down after) rather than persistent?	○
10	Do you have defined metrics for governance effectiveness — not just productivity?	○

Strong governance isn't a blocker.
It's the foundation for safe, scalable AI.