

FRAMEWORK ON ONE PAGE · VERSION 1.0

Agent Substrate Readiness Model

Readiness diagnostic · August 2026

TIME TO USE 30 min review / 90 min workshop	AUDIENCE Platform, security, risk, and architecture leaders	OUTPUT A substrate readiness assessment
--	--	--

?

Test whether the systems an agent touches can preserve authorization evidence and accountability.

What it is

Guidance on agentic AI addresses this gap as an observability problem. It strengthens logging requirements. It does not establish accountability where the underlying system cannot express it.

The Agent Substrate Readiness Model separates two questions that enterprise AI governance has been treating as one.

The two tiers

Tier 1: Structural Readiness	Each question has two possible answers. The system either has the property or it does not. There is no partial credit for "we're working on it."
Tier 2: Substrate Authorization Readiness	Structural readiness tells you whether a system can support agents technically.

A worked example

Situation	A generic enterprise is assessing whether its procurement system can preserve agent identity, operation-level scope, and rollback evidence.
Where it lands	Authorization readiness requires remediation
Evidence	The system records the human user but does not preserve agent identity or the full business state at the time of each write.
Gap found	Attribution and rollback evidence are incomplete for consequential writes.
Next action	Add operation-level attribution, preserve state snapshots, and define the remediation path before production deployment.
Output	A substrate readiness assessment with named engineering work.

The one action it produces

A substrate readiness assessment

What good looks like

A CISO who cannot answer "who authorized that state change, and what was the business condition that made it valid" across all five of these dimensions should not treat the authorization problem as a future-state concern.

What it cannot establish

This object can identify substrate weaknesses that affect authorization evidence. It cannot certify a platform, guarantee runtime behavior, or replace a security architecture review.

CITE AS