Integration Blueprint: AI Assess Technolgy × NIST AI RMF

Reference architecture for wiring the AI Assess Tech platform into an AI RMF-aligned governance program — components, data flow, and the evidence spine.

REFERENCE ARCHITECTURE — FOUR LAYERS

LAYER	FUNCTION (RMF)	COMPONENTS & DATA FLOW	EVIDENCE SPINE
LAYER 1	**Registration & Context (MAP)**	AI System Registry: system, named owner, environment, model, bound API key, deployment configuration.	Registry record enters the hash chain at creation; ownership is attributable from day one.
LAYER 2	**Assessment Pipeline (MEASURE)**	Async REST: POST assessment → 202 Accepted → poll status → fetch results. LCSH battery (120 forced-choice items) runs against the live configuration; 0–10 scores × 4 dimensions + archetype.	Question bank locked under a published bankHash and anchored to Ethereum; result record hash-chained on write.
LAYER 3	**Evidence & Verification**	Every result, attestation, and decision lands in the SHA-256 hash chain. Audit export produces the examiner package.	Public, login-free verification at aiassesstech.com/verify — third parties confirm integrity independently.
LAYER 4	**Workflow & Response (GOVERN / MANAGE)**	Drift detection compares trajectories to baseline; findings route through escalation chains (named routing, live Slack delivery, SIEM alerting). Risk Acceptance Ledger and multi-party attestation close the loop.	Escalations, acceptances, and sign-offs are themselves chained — the response is as auditable as the finding.

INTEGRATION POINTS

INTEGRATION POINT	HOW IT CONNECTS
CI/CD gate	Pipeline calls the assessment API on release candidates; deployment proceeds only on a fresh passing score — behavioral testing beside unit testing. Model-Change Reassessment Gate enforces this for model/prompt updates.
Compliance calendar	Scheduled reassessments define the MEASURE cadence; missed runs surface as governance findings, not silent gaps.
SIEM / SOC	Behavioral drift events are delivered beside security events, so the SOC sees disposition change with the same urgency as intrusion.
Board & audit surface	Board Pack PDF with public verification link; auditor role is read-only with evidence export — verify everything, alter nothing.

Alignment reference: NIST AI RMF 1.0 (NIST AI 100-1, Jan 2023) functions Govern, Map, Measure, Manage; see also NIST AI 600-1 Generative AI Profile (Jul 2024). Capability alignment only — not a NIST certification, endorsement, or conformity determination. Subcategory-level mapping is published separately.

aiassesstech.com · greg@gidanc.ai · Austin, TX

From Telemetry to the 72-Hour Clock

An incident response workflow connecting AI Assess Tech behavioral telemetry to the reporting discipline of the Illinois AI Safety Measures Act (SB 315).

THE INCIDENT CLOCK — DETECTION TO CLOSURE

PHASE	CLOCK	ACTIONS	EVIDENCE ARTIFACT
T0 — DETECT	**Clock candidate**	Drift alert fires: dimensional score crosses corridor, archetype shifts, or scheduled reassessment fails. SIEM/Slack delivery to on-call governance owner.	Assessment record + drift event, both hash-chained; alert timestamp is the provisional T0.
T0 + 4h — TRIAGE	**Clock pending**	Governance owner confirms signal vs. noise: re-run assessment for confirmation sample; pull system Registry record and recent change history.	Confirmation run result; triage note attached to the incident record.
T0 + 12h — CLASSIFY	**Clock decision**	Classify severity: (a) imminent risk of death or serious physical harm → 24-hour track; (b) critical safety incident → 72-hour track; (c) internal-only finding. Named executive signs the classification.	Signed classification attestation (sign / countersign / dissent recorded).
T0 + 24h — GATE 1	**24h track due**	Imminent-harm track: report to the appropriate law-enforcement or public-safety agency (covered developers) or per contract/customer channel. All tracks: containment decision recorded — restrict, roll back, or formally accept via Risk Acceptance Ledger (justified, time-boxed, revocable).	Filed report reference; containment or acceptance entry, chained to the triggering evidence.
T0 + 72h — GATE 2	**72h track due**	Critical-incident track: report filed — for covered developers, to IEMA and the Attorney General. Report package assembles from the chain: what was measured, when, by what instrument (bankHash), what changed, who decided what.	Audit-export incident package with public verification hashes — the regulator can verify without trusting you.
T0 + 30d — CLOSE	**Post-incident**	Root-cause review; corrective change ships through the Model-Change Reassessment Gate; baseline re-established; escalation chain and thresholds tuned from lessons learned.	Post-change assessment restores the baseline; closure attestation ends the incident chain.

Scope: SB 315's reporting duties bind covered frontier developers — 72 hours to the Illinois Emergency Management Agency and Attorney General; 24 hours to appropriate law-enforcement or public-safety agencies where there is imminent risk of death or serious physical injury. Effective January 1, 2027, with transparency and audit obligations beginning January 1, 2028; most enterprises deploying AI are not covered entities. This workflow adopts the statutory clock as an internal SLA and contract-readiness discipline — the standard regulators, insurers, and enterprise customers will increasingly expect. Definitions subject to pending IEMA/Attorney General guidance. Not legal advice.

aiassesstech.com · greg@gidanc.ai · Austin, TX

Internal Engineering AI Governance Questionnaire

Self-assessment for engineering teams that own deployed AI systems. Every 'Yes' requires an evidence reference — a verify URL, record hash, or export ID. Unevidenced answers count as 'No.'

#	QUESTION	RESPONSE	EVIDENCE REF (hash / URL / export ID)
A. INVENTORY & OWNERSHIP			
A1	Is every AI system your team operates registered in the AI System Registry with a named owner?	☐ Y ☐ N ☐ N/A	
A2	Is each production API key bound to a registered system (no shared or orphan keys)?	☐ Y ☐ N ☐ N/A	
A3	Has ownership been re-attested within the current quarter?	☐ Y ☐ N ☐ N/A	
B. CONFIGURATION & CHANGE CONTROL			
B1	Is the production system prompt and tool configuration under version control?	☐ Y ☐ N ☐ N/A	
B2	Does every model, prompt, or tool change pass the Model-Change Reassessment Gate before shipping?	☐ Y ☐ N ☐ N/A	
B3	Can you produce the assessment run that cleared your most recent production change?	☐ Y ☐ N ☐ N/A	
C. ASSESSMENT CADENCE & DRIFT			
C1	Is a scheduled reassessment cadence configured on the compliance calendar for each system?	☐ Y ☐ N ☐ N/A	
C2	Have any scheduled runs been missed in the last 90 days? If yes, is each miss documented?	☐ Y ☐ N ☐ N/A	
C3	Do you know each system's current dimensional scores and archetype — and its drift since baseline?	☐ Y ☐ N ☐ N/A	
D. EVIDENCE & VERIFICATION			
D1	Can you produce a public verification link (aiassesstech.com/verify) for your latest assessment?	☐ Y ☐ N ☐ N/A	
D2	Has an audit export been generated and test-opened within the last quarter?	☐ Y ☐ N ☐ N/A	
D3	Does the question-bank hash on your results match the published anchored bankHash?	☐ Y ☐ N ☐ N/A	
E. ACCESS & SEGREGATION OF DUTIES			
E1	Are the people who can change system configuration different from those who approve risk acceptances?	☐ Y ☐ N ☐ N/A	
E2	Is auditor access read-only, and has it been tested from an auditor-role account?	☐ Y ☐ N ☐ N/A	
E3	Are all active risk acceptances for your systems justified, time-boxed, and unexpired?	☐ Y ☐ N ☐ N/A	
F. INCIDENT READINESS			
F1	Is an escalation chain configured for your systems with named, current on-call routing?	☐ Y ☐ N ☐ N/A	
F2	Has the escalation path been exercised (test alert) in the last two quarters?	☐ Y ☐ N ☐ N/A	
F3	Do you know who owns the incident clock and the 24h/72h classification decision for your systems?	☐ Y ☐ N ☐ N/A	

Completed by (engineer): _____	Reviewed by (system owner): _____	Date / cycle: _____

Governance you don't just claim.
Governance you can prove.

aiassesstech.com · greg@gidanc.ai · Austin, TX