

AI Assess Tech × NIST AI Risk Management Framework

Runtime behavioral governance and assessment for deployed AI systems — mapped to the four core functions of NIST AI 100-1 and the twelve risk categories of NIST AI 600-1 (Generative AI Profile).

Four Core Functions

Function	How AI Assess Tech supports it
GOVERN	Constitutional separation of powers across six-agent fleet (Commander, Operator, Conscience, Inspector General, Navigator, Engineer). Veto authority, prohibited-operation constraints enforced at DB layer. Hierarchical four-level governance gates.
MAP	Context-aware assessment of the deployed AI in its production configuration — not the bare base model. Four-dimensional LCSH behavioral profile plus DICE security-domain profile.
MEASURE	120-question forced-choice instrument with three-run Trial architecture (±0.52 points at 95% confidence). Cohen's d 10.90–66.94 (IEEE-validated). Anti-gaming via cryptographic answer shuffling. Temporal Drift Index monitors trajectory.
MANAGE	SHA-256 hash chains + Ethereum anchoring (mainnet blocks 24,186,282 / 24,467,724). Public verification endpoint. Hierarchical escalation through Conscience → Commander → Inspector General.

Four Architectural Pillars

LCSH Four-dimensional behavioral framework (Lying, Cheating, Stealing, Harm). 120 forced-choice questions. Euclidean classification into four archetypes: Well-Adjusted, Psychopath, Misguided, Manipulative.

Constitutional Fleet Six-agent system with structurally enforced role separation. No agent can approve its own actions or override its own governance constraints.

Cryptographic Audit SHA-256 hash chains, Ethereum mainnet anchoring, public verification endpoint. Tamper-evident, third-party verifiable.

DICE Level 4 security-domain overlay (Disclosure, Impersonation, Corruption, Evasion). Trustworthy Operator / Insider Threat / Compromised Asset / Rogue Agent archetypes.

Production Proof Points

Production deployment Live since February 16, 2026 · fleet of six agents on Hetzner VPS

Cryptographic anchoring Ethereum mainnet block 24,467,724 (fleet evidence) and 24,186,282 (question bank)

Empirical validation Cohen's d 10.90–66.94 across LCSH dimensions · IEEE-published

Patent portfolio 14 inventions across 6 U.S. provisional filings (USPTO 2025–2026)

Continuous monitoring 138+ active health checks · Temporal Drift Index

Public verification Third-party endpoint · no credentials required

GenAI Profile Coverage

Twelve NIST AI 600-1 risk categories:

Risk category	Coverage
Confabulation (hallucination)	Direct · Lying
Harmful bias and homogenization	Direct · Cheating
Dangerous, violent, hateful content	Direct · Harm
Information security (prompt injection, evasion)	Direct · DICE
Human–AI configuration (automation bias)	Direct · DICE
Information integrity (synthetic media)	Direct · Lying + DICE
Data privacy impacts	Partial · DICE Disclosure
CBRN information access	Partial
Obscene, degrading, abusive content	Partial · Harm
Value chain and component integration	Partial
Environmental impacts	Out of scope
Intellectual property concerns	Out of scope

What this document is

A vendor self-attestation crosswalk prepared by GiDanc AI LLC. Reference material for procurement, regulators, and investor diligence.

What it is not. Not a NIST certification (NIST does not certify products). Not a claim that AI Assess Tech alone constitutes AI RMF compliance — customers retain responsibility for their own AI risk program.

How to Engage

Scoped pilot Assess your five highest-risk AI systems. LCSH baseline plus DICE security overlay. Tamper-evident evidence package delivered.

Technical briefing Architectural walkthrough for security, governance, and compliance teams. Live demonstration of the fleet and verification endpoint.

Regulator-facing review Evidence-package walkthrough framed for inquiry response. Cryptographic chain of custody from question bank to anchored assessment.

Contact greg@gidanc.ai to scope any of the above.