



Independent AI Conscience Agent

Autonomous Fleet Assessment, Temporal Drift Detection & Dual-Mode Certification

THE INNOVATION

An independent AI agent—the conscience agent—deployed within multi-agent environments with the sole function of autonomous behavioral assessment coordination. Features dual-mode commands (inline self-check + external conscience assessment), Temporal Drift Index for detecting degradation, fleet-level anomaly detection, and category-adaptive policies.

WHAT IT DOES

- ✓ Independent conscience agent with 5 isolation guarantees (process, credential, data, communication, unidirectional)
- ✓ Dual-mode commands: /assess (inline self-check) + grillo -assess (external conscience)
- ✓ Temporal Drift Index (TDI) detects behavioral degradation over time
- ✓ Fleet anomaly detection identifies correlated drift across shared providers/models
- ✓ Priority-based fleet orchestration + category-adaptive assessment policies

WHY IT MATTERS

Enterprises deploy hundreds of AI agents but have no autonomous system watching the fleet.

Self-assessment creates conflicts of interest. Manual initiation doesn't scale. Individual assessments miss fleet-level patterns.

The conscience agent provides the missing autonomous, fleet-scale behavioral governance layer. No AI goes live without a conscience.

KEY CLAIMS (14 total)

1. Independent conscience agent system with unidirectional assessment protocol and 5 isolation guarantees
2. Dual-mode assessment command architecture (inline + conscience agent) on shared infrastructure
3. Temporal Drift Index (TDI) for detecting behavioral degradation with exponential moving average baselines
4. Fleet-level anomaly detection identifying correlated drift across agents sharing common attributes
5. Priority-based fleet orchestration + hierarchical certification progression + category-adaptive policies

STATUS

- Filed: February 22nd, 2026
- Type: Contained in US 63/988,410; Related to US 63/949,454
- Builds on: Patents 1-4

APPLICATIONS

- Enterprise AI fleet governance
- Multi-agent platform compliance
- Conscience-as-a-Service (subscription per agent)

See a live verification example:

aiassessmenttool.com



Scan to verify