
Identity is not enough.

The third layer of agent trust: behavioral conformance.

Cyril Gorlla, CTGT · ITU-T Workshop, Session 5 · Geneva, 2 June 2026

An agent can be **verified, authorized,**
and still **wrong.**

Three layers of trust. The third is missing.

01

Identity

Who the agent is. Credentials and verifiable identity, X.1285 and OpenID Connect for agents.

MATURING

02

Authorization

What the agent is permitted to do. The scopes and delegation it acts under.

EMERGING

03

Behavioral conformance

Whether the agent's outputs stayed inside those permissions, verifiable in real time and after the fact.

NO SHARED STANDARD

Identity and authorization are checked **at the door**.

Conformance must be checked **at every output**, in real time and after.

A LLOYD'S OF LONDON PORTFOLIO INSURER

A single AI-generated false positive carrying **over \$1M in exposure**, caught at the behavioral layer within weeks of deployment. The identity was valid. The authorization was valid. The output was wrong.

Three primitives worth standardizing.

- | | | |
|----|--------------------------------|--|
| 01 | Conformance attestation | A signed, portable record that an output was evaluated against a named policy set, with the result: permit, remediate, block, or route to a human, including which policies fired. |
| 02 | Policy reference | A standard way to name and version the policy set an agent acts under, so a relying party can ask "conformant against which policy?" An isolated "approved" is meaningless without its referent. |
| 03 | Observability interface | A minimal, model-agnostic way for a relying party to request and verify conformance evidence, at runtime and in audit, without depending on any single vendor's stack. |
-

The market is already reaching for this layer, without a standard.

Mastercard, Verifiable Intent

Tamper-resistant audit trails for agent transactions (2026). A proprietary attestation scheme built in the absence of a shared one.

EU AI Act

Traceability and logging duties that require a record of what an automated system did, and why.

Regulated deployment

Agentic AI is accelerating into financial services, insurance, and public institutions, each inventing its own conformance check.

Trust is the precondition for deployment.

Standardize the third layer, and every relying party inherits a verifiable trust posture from day one.

Cyril Gorlla · hello@ctgt.ai · ctgt.ai