[redacted]

From: Gareth Price [redacted]
Sent: Tuesday, 11 August 2026 4:06 AM
To: Policy Development
Subject: Governance Review submission — exact-action accountability for AI agents

EXTERNAL EMAIL: Do not click any links or open any attachments unless you trust the sender and know the content is safe.

Dear General Manager, Policy and Frameworks,

I am submitting a short technical contribution to APRA's current Governance Review consultation, particularly the paper's recognition that regulated entities need governance capable of overseeing the safe, effective and accountable use of AI.

I am an independent UK developer of Agent Trust Gate™, a public technical demonstrator focused on the action boundary for increasingly autonomous AI systems.

The control question is: before an AI agent performs a consequential action in banking, insurance or superannuation, can the entity prove that the exact action is within delegated authority and risk limits, supported by required evidence and, where necessary, approved by the correct currently authorised human — and later prove that execution matched the authorised action?

ATG demonstrates exact-action binding, role/scope/value checks, Verified Human Authority, short-lived one-use GatePass or refusal artefacts, expiry/replay protection and an audit-linked execution receipt.

I am not suggesting that draft CPS 510 mandates this particular architecture, nor claiming compliance or certification. The current implementation is synthetic and local-only. I offer it as an example of how broad accountability and governance principles could be translated into enforceable pre-action controls for agentic systems.

Demo:
https://gareth1953.github.io/agent-trust-gate/human-authority-demo.html

Repository:
https://github.com/Gareth1953/agent-trust-gate

I would welcome any view on whether exact-action authority and execution proof is a useful control pattern for future AI governance in APRA-regulated entities.

Kind regards,
Gareth Price
Creator, Agent Trust Gate™
[redacted]