Submission to APRA's Governance Review

Draft CPS 510 Governance: oversight and use of AI and related technologies

PROJECT ARTESIAN · PUBLIC SUBMISSION · **Aug 25, 2026**

Submitted by Project Artesian
[redacted]

Executive summary

Project Artesian supports APRA's proposed principles-based approach to CPS 510. The draft standard appropriately places responsibility for governance with boards and senior managers. The accompanying prudential practice guidance should make this responsibility more concrete where human authority is exercised through AI-enabled processes.

AI-mediated finance means financial activity in which software or AI materially recommends, initiates, mediates or executes a decision or action. Project Artesian's distinctive proposition is that prudential governance should be organised around the complete chain of authority for consequential activity, not only around the model, system or vendor that performs it.

The relevant governance question is whether an institution can demonstrate who or what acted, for whom, under what mandate and limits, what occurred, what evidence supports the action, and how intervention and recourse remain possible.

Project Artesian recommends that CPG 510:

1. Govern material AI-enabled financial activity as a complete authority-to-recourse chain.
2. Require evidence that is decision-useful for boards and fit for reliance across institutional boundaries.
3. Treat human intervention as effective only where it is operationally executable at the speed and scale of the activity.
4. Treat the capacity to understand and challenge AI-enabled operations as a collective board capability.
5. Connect governance guidance to existing prudential requirements for risk management, operational resilience, information security and accountability.

These points can be addressed within guidance without prescribing a particular technology or creating a separate AI governance regime.

Why Project Artesian is responding

Project Artesian is an Australian-origin, public-interest research and applied-systems programme examining trusted financial activity at the convergence of AI, tokenisation and programmable money. It is responding to Question 6 because its primary inquiry, On Whose Authority?, examines how authority, accountability, evidence and recourse operate when software materially influences or performs consequential financial activity.

The programme's unit of analysis is the complete activity chain through which a financial decision or action is authorised, executed, evidenced and, where necessary, stopped or remedied. This extends conventional AI governance beyond model inventories and lifecycle controls to the prudential question of whether an institution retains effective authority across people, systems, related entities and third-party providers.

Artesian is built on research and inquiry methods developed by the Humanware Institute, which studies the institutional layer above hardware and software: judgement, decision rights, operating rules, accountability, assurance, workforce capability and the practical capacity to intervene. Project Lead Kate Cooper has more than two decades of experience across government, banking, technology and institutional transformation, including ten years leading transformation and innovation within Westpac and NAB and subsequent chief executive and director roles in regulated financial services and digital assets.

Project Artesian's research and integrated demonstrations are in development. This submission therefore contributes a governance framework and testing method rather than completed empirical findings. Observations from stakeholder discussions are directional. It seeks neither funding nor endorsement from APRA.

Question 6: Artificial intelligence governance

Draft CPS 510 provides an appropriate principles-based foundation for the oversight and use of AI and related technologies. CPG 510 should make that foundation more concrete by describing the conditions that enable boards and senior managers to retain effective authority, not only formal accountability, where AI materially influences or performs an activity.

An applied governance framework: the Artesian Trust Chain

Project Artesian operationalises its analysis through a trust chain for consequential AI-mediated financial activity:

Principal → Actor → Mandate → Asset or currency → Action → Evidence → Recourse

Four applied missions test the principal trust conditions within that chain:

- **KYA+ | Trust the Actor.** Tests whether the actor, principal and execution chain can be established.
- **Mandate+ | Trust the Authority.** Tests whether authority, permissions, limits, approvals, expiry and revocation can be verified.
- **AUD+ | Trust the Local Currency.** Tests whether activity involving money and assets can operate through trusted and interoperable Australian-dollar infrastructure.
- **Recourse+ | Trust the Outcome.** Tests refusal, interruption, dispute, responsibility, loss allocation and remedy.

Evidence and institutional reliance run across all four missions. The relevant question is what happened, what proves it, which institutions may rely on that proof and where reliance ends.

The Artesian Trust Chain is developed from Humanware Institute research and applied through Project Artesian.

1. Govern the complete authority chain

Delegating a task to software does not delegate prudential responsibility. For each material AI-enabled activity, an institution should be able to identify the principal, the human, organisational and technological actors, the mandate under which the activity occurs, the accountable role, the limits and escalation conditions, and the available recourse.

Knowing which system or provider was involved is not the same as knowing what it was authorised to do at that time. CPG 510 should therefore encourage entities to map material activity across the complete authority-to-recourse chain, including where activity is designed, provided or operated by internal teams, related entities or third parties.

This does not require boards to approve individual models or routine automated decisions. It requires a governance chain that remains intelligible and capable of being tested from initiation through outcome.

2. Require evidence fit for institutional reliance

A board cannot exercise effective oversight through a list of systems or a generic statement that AI is used. Reporting should enable directors to understand where material reliance exists, whether operation remains within authority, and whether controls and interventions work.

For material AI-enabled activity, board information should address:

- the purpose, decision or process affected, and accountable owner;
- the level of autonomy or influence and the approved mandate and limits;
- material data, model, infrastructure and service-provider dependencies;
- performance, exceptions, overrides, incidents and emerging concentrations; and
- the available intervention, fallback and customer or member recourse pathways.

Evidence should be sufficient to reconstruct what the system was authorised to do, what information materially informed it, what it did and what human review or override occurred. Where activity crosses institutional boundaries, the evidence should also enable an institution that did not build or operate the AI to decide whether it can accept, refuse, continue or remediate the action.

3. Make human intervention operationally executable

Human involvement is not an effective control merely because a person appears in the process. The responsible person must have sufficient information, authority, time and technical means to intervene. This becomes particularly important where AI-enabled activity operates at a speed and scale that compresses the time available for review, or where several institutions and providers are involved.

CPG 510 should encourage institutions to test each material activity under three conditions:

- an authorised action succeeds as intended;
- an action is refused, expired, altered or interrupted; and
- a completed action is subsequently found to be erroneous, compromised or disputed.

Testing should establish what evidence is created, where intervention applies, whether the action is final or reversible, who bears responsibility and what remedy is available.

4. Treat AI governance as a collective board capability

The proposed emphasis on collective board skills and experience is well suited to AI-enabled operations. Effective oversight does not require every director to be a technical specialist. It does require the board, collectively, to recognise material reliance, ask informed questions, understand uncertainty and limits, and obtain independent challenge when needed.

CPG 510 could describe capability in terms of the board's demonstrated ability to:

- identify material AI-enabled activity and associated dependencies;
- test whether management's assurance and evidence are decision-useful; and
- understand how the entity would intervene, continue critical operations and provide recourse if the activity failed.

This framing supports accountable judgement while avoiding a prescriptive qualification or a technology-specific board seat.

5. Connect governance with existing prudential requirements

Guidance should show how board oversight of material AI-enabled activity connects to the existing prudential framework.

CPS 220 and SPS 220. AI-related risks should be identified, assessed, monitored and reported through the entity's risk management framework.

CPS 230 and CPG 230. Material service providers, operational dependencies, tolerance levels, continuity arrangements and scenario analysis may all be relevant where AI-enabled processes support critical operations.

CPS 234. Information assets, security capability and control effectiveness remain relevant where AI changes data flows, access, system behaviour or dependencies.

The Financial Accountability Regime. Responsibility for material AI-enabled activities should be capable of being located within accountable-person responsibilities and accountability statements where applicable.

Cross-references in CPG 510 would help institutions avoid parallel governance structures and make clear that AI changes the application of existing obligations rather than displacing them.

Question 4: Areas where further detail in guidance would assist

APRA asks what information would help boards and institutions implement the revised governance standard. CPG 510 would be most useful if it translated principles into practical governance questions while leaving institutions flexibility in how they respond.

Materiality. Explain that assessment should consider potential prudential impact, customer or member harm, autonomy or influence, reversibility, criticality, external dependencies, speed and the practical ability to intervene.

Authority-chain mapping. Provide a worked example showing how a material AI-enabled activity can be mapped from its principal and actor through mandate, financial action, evidence and recourse, including accountable senior managers, operational owners and third-party responsibilities.

Board information. Illustrate a concise board reporting set covering authority, limits, dependencies, assurance, exceptions, incidents, interventions and recourse.

Intervention and testing. Explain that human involvement is effective only where the responsible people have the authority, information, time and technical means to act. Illustrate testing of an authorised action, a refused or interrupted action, and a completed action that is later found to be erroneous, compromised or disputed.

Capability and assurance. Explain how boards may demonstrate collective capability and obtain independent challenge without requiring every director to hold technical expertise.

Cross-standard and cross-industry application. Use short examples to connect CPS 510 with CPS 220 or SPS 220, CPS 230 and CPG 230, CPS 234 and the Financial Accountability Regime across banking, insurance and APRA-regulated superannuation. SMSFs should remain clearly distinguished as ATO-regulated entities.

Question 1: A limited proportionality point

Any proportionality guidance should allow oversight to scale with the consequence and governance characteristics of the activity, not merely the size or sector of the institution. A small institution may rely heavily on a concentrated provider, while a large institution may use low-risk automation extensively.

Relevant factors include potential prudential impact and harm, degree of autonomy or influence, reversibility, operational criticality, dependency concentration, speed and the institution’s ability to detect, intervene and provide recourse.

For RSE licensees, member interests, long time horizons and dependence on administrators, custodians and other service providers may affect materiality. SMSFs should not be treated as APRA-regulated entities; they may be relevant only as part of the wider ecosystem or where an APRA-regulated entity has a material dependency or exposure.

Conclusion

A short, targeted addition to CPG 510 can strengthen implementation of the proposed standard by making clear what accountable governance looks like when authority is exercised through AI-enabled processes. The Artesian Trust Chain offers a technology-neutral and proportionate method for connecting formal accountability to operational authority, evidence, intervention and recourse. It can be integrated with existing prudential standards without creating a separate AI governance regime.

Project Artesian is open to follow-up discussions in whatever way is useful to APRA.

[redacted], **Project Artesian**

Sources

APRA, Governance Review Consultation Paper · APRA, Draft Prudential Standard CPS 510 Governance · APRA, Letter to Industry on Artificial Intelligence · Financial Accountability Regime Act 2023 · APRA and ASIC, Financial Accountability Regime: Information for accountable entities (RG 279) · CPS 220 · SPS 220 · CPS 230 · CPG 230 · CPS 234