

GLOBAL AI AGENT GOVERNANCE

THE REGULATORY
LANDSCAPE



COGNAONE.COM

SEPTEMBER 2026



■ REGULATORY REVIEW · SEPTEMBER 2026 · V3.2

Global AI Agent Governance: The 2026 Regulatory Landscape

Twelve instruments, five jurisdictions and one practical question: who authorised the action?

Daniel Mackle · CognaOne · cognaone.com

A cross-jurisdiction review of twelve instruments that shape how AI agents may act, from the EU AI Act and DORA to APRA CPS 230, Singapore's agentic AI framework and the UK's reformed automated-decision rules. It separates binding law from supervisory expectation and voluntary guidance, sets out what is fixed and what is expected through 2028, and maps what regulators ask for to the controls that make it real at the point of action.

Policy can be well written while the authority to act remains poorly controlled. That gap is where agent risk lives.

CONTENTS

1. The question has moved to execution
2. How to read this paper
3. Twelve instruments, with different force
4. The dates that matter
5. What supervisors are saying about agents

6. Where policy meets an agent's authority
7. The compliance gap
8. Conclusion
About the author and CognaOne
References

1. The question has moved to execution

An AI system that drafts a reply creates one kind of risk. An agent that sends it, amends a customer record, approves a payment or pushes code to production creates another.

The model is the visible part of the system. The consequence sits in the action path: the identity the agent uses, the tools it can reach, the state it relies on, and the point at which someone, or something, lets the operation proceed.

Regulators have noticed. In 2026 the EU enacted new application dates for its high-risk AI regime¹. Australia's prudential regulator wrote to every regulated entity that AI governance, risk management and assurance are not keeping pace with adoption². Singapore updated the agentic AI framework it describes as the first of its kind^{3,4}. A review commissioned by the UK's Financial Conduct Authority concluded that firms and individuals remain accountable for AI outcomes, even where model behaviour sits outside their direct control⁵. None of them created a single "agent law". All of them point at the same place.

Policy can be well written while the authority to act remains poorly controlled. That gap is where agent risk lives.

This review covers twelve instruments across the EU, the UK, Singapore, Australia and the US, sets out what is fixed and what is expected through 2028, then maps what the instruments ask for to an engineering design that can carry it. It is the Global Focus edition; an Australian Focus edition will follow.

2. How to read this paper

The twelve instruments carry very different force. Some are binding law for a defined class of organisation or use case. Some are supervisory expectations. Some are voluntary frameworks, and one is a proposal not yet finalised. We state the force of each plainly and do not treat them as equivalent.

Every factual statement about an instrument is referenced to the issuing body's own text wherever that text was accessible, and all sources were checked in September 2026. Where a primary source restricts automated access, we say so and identify the secondary source used to confirm its content.

The mapping from requirements to engineering controls is our analysis. CognaOne builds controls of the kind described in Section 6. Product descriptions are identified separately from regulatory findings, in the note on the author and CognaOne after the conclusion. No regulator reviewed here mandates a particular technical architecture, and nothing in this paper is legal advice. Whether an instrument applies depends on jurisdiction, sector, role in the AI supply chain, the data processed and the use case. A read-only internal assistant and an agent that can change a customer's financial position will not have the same control case.

3. Twelve instruments, with different force

INSTRUMENT	STATUS AND SCOPE	WHAT IT BRINGS INTO VIEW FOR AGENTS
1. EU Artificial Intelligence Act, Regulation (EU) 2024/1689, Articles 12 to 15 and 26 ^{6,7}	Binding. High-risk duties apply only to systems within the Act's classifications, on the timetable amended in 2026 (Section 4) ¹	Record-keeping, human oversight, accuracy, robustness and deployer duties. Article 26 requires deployers of high-risk systems to assign oversight to people with the necessary competence, training, authority and support, and to keep automatically generated logs under their control for at least six months, unless other Union or national law provides otherwise ⁷ .
2. EU General Data Protection Regulation (GDPR), Regulation (EU) 2016/679 ⁸	Binding where EU personal-data rules apply	Lawful and limited processing, security, accountability, data-subject rights and the Article 22 rules on solely automated decisions. An agent's memory, retrieval and tool access can all process personal data.
3. EU Digital Operational Resilience Act (DORA), Regulation (EU) 2022/2554 ⁹	Binding on in-scope financial entities; applies from 17 January 2025 ⁹	Information and communication technology (ICT) risk, incident reporting, resilience testing and third-party dependencies. The question is whether an agent and its providers are controlled as part of the institution's operational environment.
4. Monetary Authority of Singapore (MAS) Consultation Paper PO17-2025, proposed Guidelines on AI Risk Management ^{10,11}	Proposed. Consultation ran 13 November 2025 to 31 January 2026 ¹⁰ . In March 2026 MAS said it was still reviewing responses ¹² ; we found no final guidelines published by our check date	Board and senior management oversight, AI identification and inventories, risk materiality assessment and lifecycle controls. MAS states the guidelines are intended to cover generative AI and AI agents ¹¹ . The Project MindForge toolkit (20 March 2026) gives operational guidance across traditional, generative and agentic AI ¹² .
5. Infocomm Media Development Authority (IMDA) Model AI Governance Framework for Agentic AI ^{3,4}	Voluntary. Launched 22 January 2026 ³ ; updated version published 20 May 2026 ^{4,13}	Four dimensions: assess and bound risks up front, make humans meaningfully accountable, implement technical controls and processes, and enable end-user responsibility. Recommends defined checkpoints at which human approval is required ³ . The update encourages structural, rule-based controls over prompt-layer controls ¹³ .
6. Australian Prudential Regulation Authority (APRA) Prudential Standard CPS 230 Operational Risk Management ^{14,15,16} , read with APRA's letter to industry on AI ²	CPS 230 is binding on APRA-regulated entities. It commenced 1 July 2025; the transition for pre-existing service-provider contracts ran to the earlier of renewal or 1 July 2026, and targeted amendments took effect on 1 July 2026 ^{15,16} . The letter (30 April 2026) sets supervisory expectations, not new rules ²	Operational risk, critical operations, continuity and material service providers. The letter expects credible fallback processes where AI supports critical operations, reports that identity and access management has not yet adjusted to non-human actors such as AI agents, and finds that staff use of enterprise AI tools often relied on policy direction or detective, after-the-fact measures rather than enforceable preventative controls ² .

INSTRUMENT	STATUS AND SCOPE	WHAT IT BRINGS INTO VIEW FOR AGENTS
7. Australian Government, Guidance for AI Adoption (six essential practices) ^{17,18}	Voluntary. Published 21 October 2025; the Department describes it as updated guidance that evolves the 2024 Voluntary AI Safety Standard and its ten guardrails ¹⁸	Six practices: decide who is accountable; understand impacts and plan accordingly; measure and manage risks; share essential information; test and monitor; maintain human control ¹⁷ . Adoption supports governance; it is not statutory compliance.
8. US National Institute of Standards and Technology (NIST) AI Risk Management Framework 1.0 and Generative AI Profile (NIST AI 600-1) ¹⁹	Voluntary. Released 26 January 2023; profile released 26 July 2024. NIST states the framework is being revised under the White House AI Action Plan ¹⁹	Govern, Map, Measure and Manage offer a way to organise ownership, risk assessment, evaluation and response.
9. ISO/IEC 42001:2023, AI management systems ²⁰	Voluntary international standard; certifiable	Requirements for establishing, implementing, maintaining and improving an AI management system. Certification does not show that a particular agent action was authorised.
10. US Federal Trade Commission (FTC) Act, Section 5 (15 U.S.C. § 45) ^{21,22}	Binding consumer-protection law within the FTC's remit	Unfair or deceptive acts or practices, including deceptive AI claims, as in the FTC's 2024 Operation AI Comply actions ²² . An outcomes and representations constraint, not an agent runtime rule or a general US AI Act.
11. UK GDPR as amended by the Data (Use and Access) Act 2025 ^{23,24,25}	Binding. Section 80 (automated decision-making) commenced 5 February 2026 ²⁴ ; the ICO confirmed all data-protection provisions were in force by 19 June 2026 ²⁵	New Articles 22A to 22D set a permission-plus-safeguards model for significant decisions based solely on automated processing, defined as decisions with no meaningful human involvement, with tighter limits where special category data is used ²³ .
12. UK Financial Conduct Authority (FCA) Consumer Duty ²⁶	Binding for firms and retail activities in scope	Good outcomes for retail customers remain the obligation when AI is used. The Mills Review (6 July 2026) found the overall regulatory framework, including the Consumer Duty and the Senior Managers Regime, remains sound ^{5,27} .

These are twelve sources of obligation and design signal, not twelve "agent regulations". Reading them as equivalent is the first mistake.

4. The dates that matter

The calendar splits three ways: what has already taken effect or been published, what is fixed in law or official instruments for the next two years, and what regulators and governments have signalled but not yet fixed. Only the first two are dates an organisation can plan against with confidence.

Already in effect or published

DATE	INSTRUMENT	MILESTONE
17 Jan 2025	DORA	Applies to in-scope financial entities ⁹
2 Feb 2025	EU AI Act	Prohibitions, definitions and AI literacy provisions apply ²⁸
1 Jul 2025	APRA CPS 230	Commences ¹⁶
2 Aug 2025	EU AI Act	Governance rules and general-purpose AI obligations apply ²⁸
21 Oct 2025	Australian Guidance for AI Adoption	Published ¹⁸
13 Nov 2025	MAS PO17-2025	Consultation opens (closed 31 Jan 2026) ¹⁰
11 Dec 2025	US Executive Order 14365	Directs federal challenges to state AI laws ²⁹
22 Jan 2026	IMDA agentic framework	Launched ³
5 Feb 2026	UK Data (Use and Access) Act	Automated decision-making reforms commence ²⁴
20 Mar 2026	MAS Project MindForge	AI risk management toolkit published ¹²
30 Apr 2026	APRA	Letter to industry on AI; CPS 230 amendments finalised ^{2,15}
12 May 2026	UK SI 2026/425	ICO duty to prepare a statutory code on AI and automated decision-making applies ³⁰
14 May 2026	US, Colorado SB 26-189	Replacement automated decision-making law signed ^{31,32}
19 May 2026	EU AI Act	Draft Commission guidelines on classifying high-risk systems published ³³
20 May 2026	IMDA agentic framework	Updated version published ⁴
19 Jun 2026	UK Data (Use and Access) Act	All data-protection provisions in force ²⁵
1 Jul 2026	APRA CPS 230	Amendments commence; pre-existing contract transition ends ^{15,16}
6 Jul 2026	FCA	Mills Review published ²⁷
15 Jul 2026	Australian Government	Australian Standards for AI announced; Office of AI established ³⁴
27 Jul 2026	EU AI Act	Regulation (EU) 2026/1744 enters into force ¹
2 Aug 2026	EU AI Act	General date of application ^{1,28}

Scheduled: fixed in law or official instruments

DATE	INSTRUMENT	MILESTONE
2 Dec 2026	EU AI Act	Four-month transitional period for Article 50(2) marking ends for generative systems already on the market before 2 August 2026 ¹
10 Dec 2026	Australian Privacy Act	Privacy-policy transparency for automated decisions (APP 1.7 to 1.9) commences ³⁵
1 Jan 2027	US, Colorado SB 26-189	Replacement automated decision-making technology law applies to consequential decisions made from this date ^{31,32}
1 Aug 2027	EU AI Act	Commission guidelines for high-risk systems covered by Annex I product legislation to be published by this date ¹
2 Aug 2027	EU AI Act	National AI regulatory sandboxes to be operational in each Member State; Commission delegated acts on Annex I limitations due ¹
2 Sep 2027	EU AI Act	Commission guidance and a voluntary template for post-market monitoring plans due ¹
2 Dec 2027	EU AI Act	Chapter III Sections 1 to 3 apply to Annex III high-risk systems ¹
28 Jan 2028	EU AI Act	Notified bodies under Annex I Section A legislation to apply for designation under the AI Act ¹
2 Aug 2028	EU AI Act	Chapter III Sections 1 to 3 apply to Annex I high-risk systems ¹

Expected: signalled, dates may move

SIGNALLED TIMING	JURISDICTION	WHAT IS EXPECTED
During 2026	Australia, OAIC	Detailed guidance on the new APP 1 automated-decision obligations ³⁵
Winter 2026	UK, ICO	Final guidance on automated decision-making and profiling, following consultation ³⁶
Oct 2026 to Jan 2027	UK, FCA	Review of general-purpose AI tools outside the regulatory perimeter, which the Mills Review recommends within three to six months of July 2026; adoption is a matter for the FCA Board ⁵
Early 2027	Australia	Legislation for Australian Standards for AI, focused on large data centres and AI training ³⁴
Not announced	EU, Commission	Final guidelines on classifying high-risk AI systems; a draft was published on 19 May 2026 ³³
Not announced	EU, CEN and CENELEC	Harmonised standards for high-risk requirements, which give legal certainty once referenced in the Official Journal ³⁷
Not announced	UK, ICO	Statutory code of practice on AI and automated decision-making; the duty to prepare it applies from 12 May 2026 ³⁰

SIGNALLED TIMING	JURISDICTION	WHAT IS EXPECTED
Not announced	Singapore, MAS	Final Guidelines on AI Risk Management; MAS said in March 2026 it was reviewing consultation responses ¹²
Not announced	Australia, APRA	Supervisory follow-through on the AI letter, including prudential reviews, thematic activities and AI supplier engagement ²
Not announced	US, federal	Challenges to state AI laws under Executive Order 14365, through a Department of Justice AI Litigation Task Force and a Commerce evaluation of state laws ²⁹
Not announced	US, NIST	Revision of the AI Risk Management Framework under the White House AI Action Plan ¹⁹

The EU AI Act needs particular care. Regulation (EU) 2026/1744 was adopted on 8 July 2026, published on 24 July 2026 and entered into force on the third day after publication¹. It sets the application date of Chapter III Sections 1, 2 and 3 to **2 December 2027** for systems classified as high-risk under Article 6(2) and Annex III, and to **2 August 2028** for systems classified under Article 6(1) and Annex I¹. Those are the dates behind the Article 26 deployer duties discussed here. Presenting them as enforceable from August 2026 is now wrong. The Regulation also gives providers of generative systems already on the market before 2 August 2026 a four-month transitional period for the Article 50(2) marking obligation¹.

The voluntary frameworks, NIST AI RMF, ISO/IEC 42001, the Australian Guidance for AI Adoption and IMDA's framework, have no enforcement clock. An enterprise may adopt them or commit to them in contracts; publication does not make them law. MAS's proposal is not final guidance¹².

The United States is moving in two directions at once. Colorado replaced its 2024 AI Act with a narrower automated-decision law that applies from 1 January 2027^{31,32}, while a December 2025 executive order directs federal challenges to state AI laws considered inconsistent with national policy²⁹. Organisations operating in the US should plan for state obligations while expecting them to be contested.

5. What supervisors are saying about agents

Across four jurisdictions, supervisors used different words in 2026 for a similar observation.

APRA observed boards relying on vendor presentations without sufficient examination of risks such as unpredictable model behaviour, identity and access management that has not adjusted to AI agents, and controls over staff AI use that detect problems after the fact rather than prevent them². IMDA recommends bounding an agent's autonomy and access to tools and data before deployment, and defining the checkpoints at which a person must approve³. The Mills Review states that firms and individuals remain accountable for outcomes as AI systems move towards more autonomous modes⁵. The EU assigns deployer oversight to people with the competence, training, authority and support to exercise it⁷.

The message is consistent: an organisation must be able to show control over what its agents do, not just describe it.

That is an engineering problem as much as a policy one.

6. Where policy meets an agent's authority

The instruments do not say "install an action gate". Their requirements and expectations concern oversight, risk, records, security, continuity and accountability. The engineering question is how to make those outcomes credible when an agent can initiate side effects.

Consider an accounts-payable agent. It reads an invoice, checks a supplier record and proposes a payment. A useful system may reason flexibly about a discrepancy. It should not be able to change the supplier's bank account and release the payment simply because its own plan says the step is necessary. The organisation needs a point at which the proposed operation is checked against authority and current evidence, and is then allowed, held for approval or refused.




Figure 2. An illustrative control point at the action boundary. This is our design interpretation, not a requirement of any instrument reviewed here.

CONTROL QUESTION	ENGINEERING RESPONSE	EVIDENCE TO RETAIN	RELEVANT INSTRUMENTS
Which agent, acting for whom, has authority?	A distinct workload identity; a named business owner; narrowly scoped credentials and delegated permissions	Owner, identity, granted scope, credential issue and revocation history	DORA, CPS 230 and APRA's AI letter, IMDA, NIST ^{2,3,9,14,19}
Is this specific action permitted now?	A pre-execution decision at the tool or transaction boundary, checking purpose, target, parameters, limits and current policy	Proposed action, policy version, decision, reason and execution outcome	IMDA, EU AI Act oversight duties, NIST ^{3,7,13,19}

CONTROL QUESTION	ENGINEERING RESPONSE	EVIDENCE TO RETAIN	RELEVANT INSTRUMENTS
Does a person need to decide?	Approval for defined consequential or irreversible operations, tied to the exact proposed action, with the ability to reject, pause or intervene	Approver identity, information shown, decision, timestamp and any changed action	EU AI Act Article 26 where applicable; IMDA; UK Articles 22A to 22D where applicable ^{3,7,23}
What information is the action based on?	Versioned, access-controlled memory and workflow state, with source provenance and conflict handling	State version, relevant sources, changes and unresolved conflicts	GDPR, UK GDPR, IMDA, NIST ^{3,8,19,23}
Can the event be reconstructed?	Attributable event records across instruction, policy decision, approval, tool call and result, with integrity protection proportionate to risk	Correlated timestamps, actors, versions, results and retention rule	EU AI Act Article 26 for applicable logs; DORA; CPS 230 ^{7,9,14}
What happens when a provider or agent fails?	Monitoring, incident triage, credential revocation, isolation, recovery, fallback and dependency mapping	Alerts, incident timeline, provider records, recovery tests	DORA, CPS 230 and APRA's AI letter, MAS proposal, NIST ^{2,9,11,14,19}

Article 26 requires deployers of applicable high-risk systems to assign competent human oversight and retain certain automatically generated logs. It does **not** require a human approver for every tool call, or any particular storage technology for logs. For deployers that are financial institutions, Article 26 also lets certain monitoring and log-keeping duties be met through the governance and documentation requirements of Union financial services law⁷. The right design follows the consequence of the action and the legal context.

Determinism at the boundary

The model's reasoning need not be deterministic. The permission decision for a particular operation should be. Given the same action, authority, policy and state snapshot, the gate should return the same answer every time. If an agent proposes a payment above its limit, the gateway declines it, however persuasive the surrounding text. If supplier details conflict, the action is held until the conflict is resolved.

The critical distinction is between instructions to an agent and controls on what its tools will execute. A prompt can tell an agent to seek approval. The payment service can refuse to release funds without a valid, action-specific approval. The updated IMDA framework makes the same distinction, encouraging structural, rule-based controls that operate at the system level over model-based or prompt-layer controls¹³.

A prompt is a request. A control is a constraint. Only one of them holds when the agent is wrong.

Oversight that can change the outcome

An operator watching a dashboard after execution has visibility. Oversight, where required, also needs authority: to understand the proposed action, intervene in time, and prevent or suspend it. The EU AI Act assigns high-risk deployer oversight to people with competence, training, authority

and support⁷. In the UK, a decision is "based solely on automated processing" if there is no meaningful human involvement in taking it²³, so the quality of human involvement matters, not its mere presence. Neither regime means a person must click "approve" on every low-risk search.

Records that answer an investigation

For a consequential action, the record should connect the request, agent identity, applicable policy, relevant state, decision, any approval, tool invocation and actual result. That is far more useful than a transcript holding only the final response. Append-only or tamper-evident storage can strengthen assurance where the risk justifies it; it is a design recommendation, not a universal legal requirement. Retention must also respect data-protection and sector rules. Article 26's six-month minimum applies to automatically generated logs under the deployer's control, unless other applicable law provides otherwise^{7,8}.

State is part of the risk

An agent may rely on a memory entry saying a supplier is approved, a retrieved document containing altered instructions, or a stale permission inherited from an earlier workflow. No model update is needed for the action to go wrong. State therefore needs ownership, access control, source tracking, versioning, expiry and a way to surface conflict. None of the twelve instruments names "state integrity" as a requirement. The connection between governed state and action authority is our engineering conclusion from the risks they address.

7. The compliance gap

Many enterprises have model assessments, procurement reviews and responsible-AI principles. The hard part begins when an agent crosses from advice into execution. An inventory may name the model provider but omit the agent's tool credentials. A policy may require human oversight without saying which operations pause. A trace may show that a tool was called but not which policy authorised it or which account it changed.

| An inventory that lists the model but not the agent's credentials is a list, not a control.

A workable review starts with six concrete questions:

1. Which agents can write to production systems, send external messages or move value?
2. What business purpose, identity, permissions and limits apply to each one?
3. Where is an action intercepted before a side effect occurs?
4. What state and source material did the action rely on, and was any of it contested or stale?
5. Can an authorised person prevent, interrupt or reverse the action where the risk calls for it?
6. Can the organisation reconstruct what actually happened, including provider dependencies and failures?

An organisation that cannot answer these should narrow the agent's action space before expanding its autonomy. That is a risk-based engineering recommendation, not a claim that every deployment falls under the EU's high-risk provisions.

8. Conclusion

The 2026 landscape is uneven. The EU has an AI-specific law with phased application and newly enacted high-risk dates. Financial regulators in Europe and Australia already impose operational-resilience requirements that can reach agent deployments, and APRA has now told regulated entities that AI governance is not keeping pace with adoption. Singapore has published guidance written specifically for agents, while its financial-sector AI guidelines remain proposed. The UK continues to apply existing data and conduct rules, with statutory AI guidance still to come. In the US, state laws are arriving while the federal government moves to contest them.

What joins these instruments is the need to show control over consequences. The test is specific. For a material action, can the organisation establish who gave the agent authority, what it was permitted to do, which information and policy it used, whether a person could intervene, what was executed, and how the result can be investigated?

If the answer exists only in a policy document, the execution system is carrying more authority than the governance system can account for.

About the author and CognaOne

Daniel Mackle is the founder of CognaOne. CognaOne is building the active governance platform for AI in regulated financial services: an inventory of AI systems and the agents connected to them, controls held with their regulatory provenance, an action gate that allows, holds or declines proposed actions and returns the same decision on the same facts and policy, and decision records created at the moment of authorisation. These are the kinds of controls described in Section 6. This review was written to stand on its sources, not on the product, and every factual statement in it is referenced to the issuing body wherever that body's text was accessible.

AI use: research and drafting for this review were supported by AI tools. Multiple models were used for research and to cross-check facts, and every factual statement was verified by the author against the cited sources. The author is responsible for the content.

References