Digital Cooperation Organization

DCO Policy Watch

Governing the Cognitive State: AI in Government

Edition 8 | March 2026

Table of Contents

Executive Summary

Welcome to the eighth edition of the DCO Policy Watch. This edition examines the emergence of AI in government — and the governance challenge it creates.

Governments around the world are moving past experimentation. Analytical AI, which analyzes large datasets to identify patterns, make predictions, and support decision-making, is already in production across tax fraud detection, social benefit triage, and infrastructure maintenance, including in Togo's Novissi emergency cash-transfer program, which used data-driven targeting to identify vulnerable households for support. Generative AI, which produces new content such as text, code, images, or translations based on patterns learned from data, is transforming citizen-facing services, drafting, translation, and internal productivity — with the UK government reporting that a public sector trial of AI coding assistants saved developers the equivalent of around 28 working days a year, or roughly an hour per day. And early agentic AI systems — capable of initiating actions and coordinating tasks across workflows with limited per-step human oversight — are now underway in leading administrations. This is the **Cognitive State:** AI embedded not merely as a tool, but as an operating layer of the state itself.

The central challenge: governing AI at scale, not just adopting it.

AI deployment in government is still constrained by capacity and institutional readiness, but governance is rapidly becoming a key challenge. As AI becomes embedded into public authority — eligibility determinations, compliance workflows, procurement decisions, and operational coordination — the defining question shifts from whether governments can adopt AI to whether they can govern it: at scale, at speed, and while preserving transparency, due process, and administrative sovereignty. Recent events — including public disputes over commercial arrangements governing AI deployed by government agencies — have underlined that these are live questions for every administration, not just frontier adopters.

A passive approach carries material risks: deployments where AI systems are configured once and then left unmonitored, allowing performance to quietly deteriorate over time as real-world conditions change ("model drift"); over-dependence on a single vendor whose systems cannot be easily audited, challenged, or replaced; failures concentrated on the most vulnerable citizens; liability gaps where accountability is unclear when things go wrong; and, at the development stage, reliance on large general-purpose AI models trained on data and for purposes that may not align with public-sector values or legal requirements. When harms occur without clear accountability, public trust erodes — and trust, once lost in the public sector, is exceptionally difficult to rebuild.

Three Strategic Pillars

This edition provides a leadership framework organized around three interconnected pillars, drawing on evidence from across the DCO's membership and the wider international policy landscape — including the EU, UK, Canada, Singapore, UAE, China, and beyond:

Pillar I

Delivering Public Value — Service Transformation & Automation.

This edition examines what it takes to move from AI pilots to scaled service delivery that genuinely improves citizen outcomes. It covers the full spectrum from assistive AI through automated decision-making to early agentic deployments — and sets out the governance conditions that determine whether AI delivers public value or simply scales existing risk: defined use-case priorities, human accountability for high-impact decisions, and minimum governance standards before (not after) deployment.

Pillar II

Procurement Governance & Contracting Policy.

AI governance is not solely a regulatory question — it is a contracting reality. This edition shows how leading governments are encoding accountability into procurement: auditability clauses, data portability rights, bias benchmarks, model-change notification, and enforceable remedies. Trust cannot be bolted on after contracts are signed; it must be engineered in from the point of purchase.

Pillar III

Operational Accountability, Monitoring & Incident Pathways.

The most consequential AI governance failures emerge not at the pilot stage but in production, often months after deployment. This edition sets out the operational shift required: from one-time assessments to continuous monitoring, from advisory oversight to institutionalized assurance functions, and from informal complaint-handling to structured incident and redress pathways that make accountability real in day-to-day administration.

Rather than treating these pillars in isolation, this edition translates complex governance debates into practical guidance for senior decision-makers. It is grounded in real-world examples from a wide range of jurisdictions — advanced economies and emerging markets alike — and concludes with a tiered roadmap to help governments sequence priorities according to their starting point and institutional maturity.

1 Introduction:

From Digital Government to the Cognitive State

Over the past decade, most public sector innovation has been defined by digitization: moving services online, migrating systems to cloud infrastructure, and building the "plumbing" of interoperable digital government. That work remains unfinished in many countries. But the global frontier is shifting.

Governments are now beginning to **automate cognition** – not just processing transactions faster, but using AI to reason over data, interpret language, triage cases, and increasingly orchestrate administrative workflows. This is the emergence of a **Cognitive State** on top of the digital infrastructure: a model of government where AI augments (and, in some cases, automates) the tasks that previously depended on human judgement, analysis, and drafting.

The shift is not AI adoption – it is a governance transition

In the private sector, the primary value proposition is efficiency and profit. In government, AI operates under a dual mandate: **deliver public value** and **uphold accountability**. This is why the central policy question is not "How fast can we deploy AI?" but:

How do governments scale AI capability while preserving transparency, due process, and administrative sovereignty?

A core risk highlighted across jurisdictions is the rise of a "black box state": where core decisions about eligibility, enforcement, or rights become dependent on opaque models procured from vendors – models that civil servants, auditors, and affected citizens cannot meaningfully interrogate. The result is not just technical dependency; it is **delegated governance.**

The tension between AI capability and administrative sovereignty is not hypothetical. In February 2026, the U.S. Department of Defense (now rebranded the Department of War) issued Anthropic — maker of the Claude AI model — an ultimatum: remove usage restrictions that prohibited the military from deploying Claude for domestic mass surveillance and fully autonomous lethal weapon systems, or face contract termination and designation as a national security supply-chain risk. Anthropic refused; the Trump administration designated it a supply-chain risk, directed all federal agencies to cease use of its products, and OpenAI promptly announced a replacement deal with the Pentagon. The episode crystallizes a governance dilemma that every government deploying commercial frontier AI will eventually face: vendors embed values in their models through training, acceptable-use policies, and contractual terms. When government requirements conflict with those values, who sets the limits? The answer has profound implications for procurement design, contract architecture, and the broader question of whether states can exercise meaningful sovereignty over AI systems they do not build themselves — all three themes this edition takes up directly.

The Cognitive State: the AI capability spectrum

The Cognitive State is being built through **two foundational classes of AI technologies**, each at different maturity levels and governance risk profiles:

1. Analytical AI: Uses structured and unstructured data to predict outcomes and allocate resources more proactively (e.g., a risk scoring, forecasting demand, identifying anomalies). Its promise is a shift from reactive services to anticipatory governance. Togo's Novissi program, for instance, used machine learning applied to mobile-phone and geospatial data to identify poorer households for emergency cash transfers, helping the government target assistance at scale beyond what traditional registries alone could support.

2. Generative AI: Synthesizes language and content across documents, cases, and policies. Internally, it can reduce technical and administrative debt (summarizing guidance, translating legacy code). Externally, it can translate bureaucracy into plain language and support inclusion through multilingual interfaces and guided life-event journeys. Singapore's AIBots platform, developed for public officers, lets agencies create retrieval-augmented generative AI chatbots using internal documents and knowledge bases.

Generative AI in government example: Singapore's AIBots Platform

By February 2025

Reached
40,000 Users

Across
115 Agencies

Created
12,000 Bots

A third mode — **Agentic AI** — is now emerging from the generative AI foundation. Unlike analytical systems that predict or generative systems that synthesize, agentic systems decide and act: they chain multi-step reasoning, invoke tools, and take bounded actions toward defined goals under human oversight.

Architecturally, agentic AI is a deployment mode built on top of generative AI — not a separate class — and in government it remains early and experimental. Singapore's Government Technology Agency (GovTech) offers the clearest public-sector benchmark: its Agentic AI Primer (April 2025) provides agencies with an introductory framework for understanding and implementing agentic systems, while the Infocomm Media Development Authority (IMDA)'s Model AI Governance Framework for Agentic AI (January 2026) — described by IMDA as the first in the world to include a comprehensive guide for responsible enterprise deployment — sets out the accountability architecture required as agentic deployments mature. As these systems spread, they will introduce qualitatively new governance challenges around delegation, auditability, and automated multi-step action: themes taken up in Pillar III of this edition.

The strategic crossroads: technology takers vs. technology shapers

Governments globally could fall into either of **two camps**:

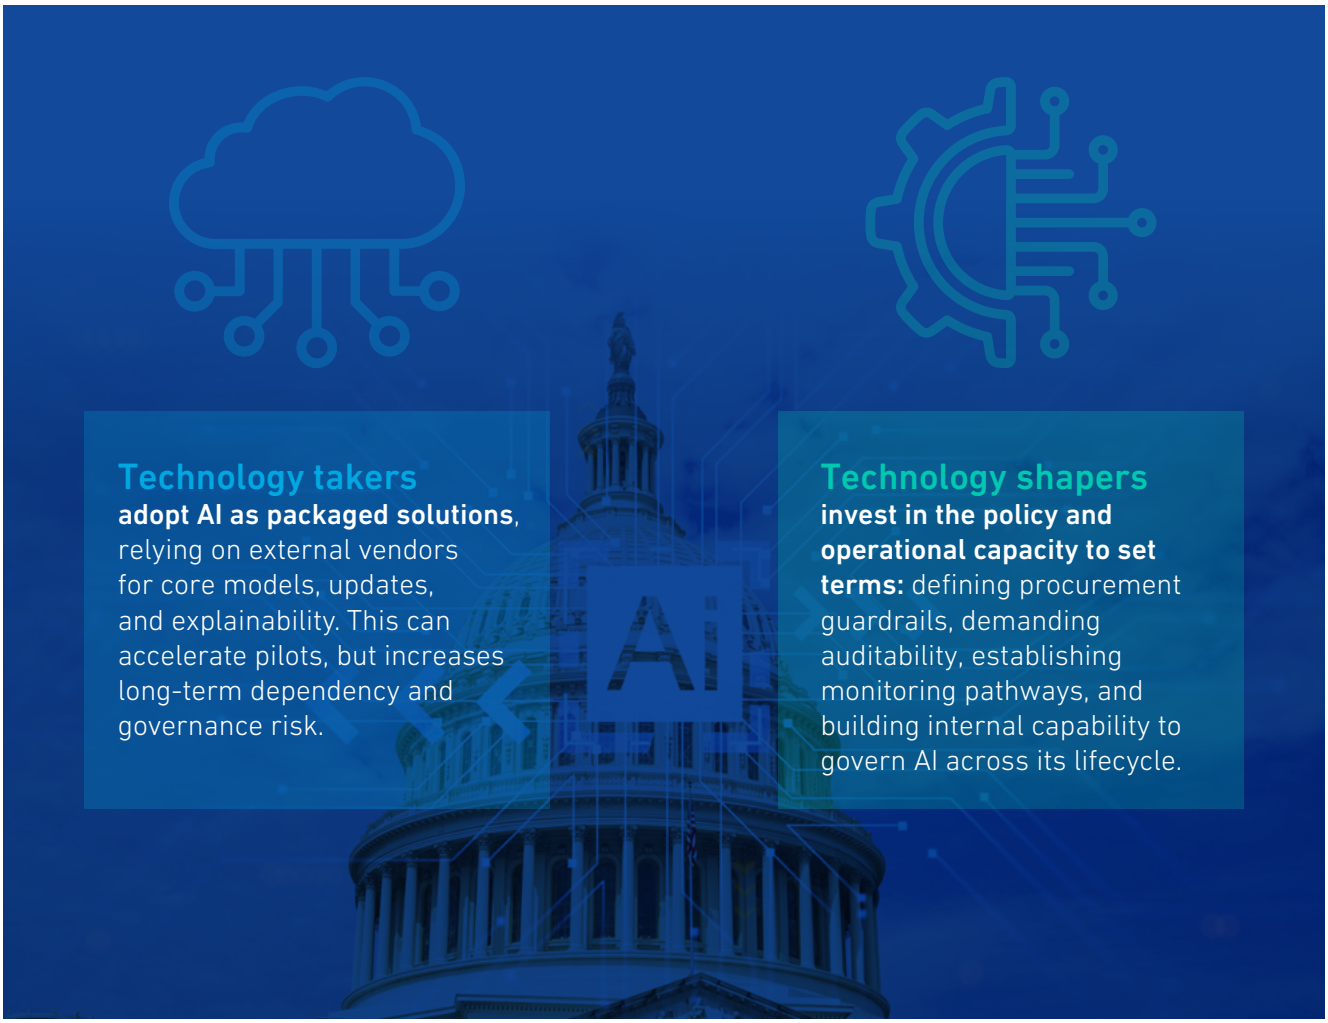

The technology taker/technology shaper spectrum plays out differently across regions, and understanding those differences matters for DCO Member States plotting their own course. The EU AI Act takes a risk-tiered regulatory approach, imposing the heaviest obligations on high-risk government applications — such as benefits eligibility systems and law enforcement tools — and banning outright uses such as real-time biometric surveillance in public spaces.

The GCC states, led by Saudi Arabia's SDAIA and the UAE's AI Office, have taken a state-capacity model: centralizing AI governance at the national level, building sovereign data infrastructure, and positioning government as the primary driver of AI adoption rather than a regulator of private-sector deployment. China offers a related but distinct model of state-led AI governance, combining strong central oversight with mandatory algorithm filing and security assessment for certain high-impact systems. Smaller economies, including many DCO Member States, face a different challenge: the risk of becoming pure technology takers by default, adopting off-the-shelf solutions without the procurement safeguards or regulatory frameworks to hold vendors accountable. The three pillars of this edition speak directly to that risk.

For leaders, the point is simple: **AI principles do not operationalize themselves**. Trust must be engineered through enforceable rules, institutional ownership, and continuous oversight.

2 Delivering Public Value:

Service Transformation & Automation

If the Cognitive State is the destination, **public value** is the only legitimate route. In government, AI cannot be justified on novelty alone; it must translate into outcomes citizens can feel: **faster decisions, fewer errors, better access, and more consistent services, while protecting fairness, due process, and trust.**

Yet the dominant pattern globally is not scale, it is experimentation. Many administrations have launched AI pilots through innovation labs and challenge funds but struggle to integrate those pilots into core service delivery. This is the **Pilot Trap**: **proof-of-concepts multiply, but production impact remains limited.**

Market Analysis: From pilots to production

Three recurring blockers explain why promising AI initiatives stall before they become real services:

1. Weak product governance (the most common failure mode)

Pilots often start as "innovation experiments" without a stable problem statement, a service owner with decision rights, or success metrics tied to public value. The result is scope creep: projects expand sideways rather than scaling into a service. The UK National Audit Office has documented this pattern across government digital programs, noting that unclear ownership and absent success metrics are persistent contributors to failed scaling.

2. Infrastructure & operating-readiness gaps

Even where models perform well, agencies may lack the data pipelines, secure environments, and change management capacity to deploy responsibly at scale, particularly across legacy systems and siloed agencies. OECD analysis confirms that most generative AI experiments remain in exploratory or pilot phases, with infrastructure readiness cited as the primary brake on production deployment.

3. Trust and transparency barriers

Government AI systems operate under a higher legitimacy threshold. Without clear public communication, explainability pathways, and redress channels, even a technically "accurate" system can fail politically and operationally. New York City's small-business chatbot — which provided guidance that misrepresented local legal requirements — is a cautionary case: technically functional, but politically untenable once errors became public.

Strategic Analysis: the public value discipline

The difference between pilots and scaled impact is not the model, it is the **discipline of implementation**. Governments that scale AI effectively do four things consistently:

A. Anchor AI to measurable public value KPIs

Scaled systems begin with an explicit KPI set that defines "value" from a public service lens, often combining efficiency and legitimacy indicators, such as:

Time-to-decision/time-to-service — how long citizens wait for a government response or service.

Error rate and appeal/overturn rate — the proportion of AI-assisted decisions that are factually or procedurally incorrect, and how often those decisions are successfully challenged; a high overturn rate signals systemic failure even where individual errors are hard to detect directly.

Equity and inclusion outcomes (who benefits; who is excluded).

User satisfaction and complaint volumes.

Caseworker workload and throughput — whether AI is freeing staff capacity or adding to it.

Crucially, these KPIs also enable stage gates: clear thresholds to **scale, pivot, or stop** a system before it becomes embedded and politically costly to unwind. The UK public sector's 2025 AI coding-assistant trial demonstrates the approach: it published average daily time-saved metrics before approving wider rollout, making the scaling decision contingent on measured public value.

B. Prioritize "bounded" (clearly scoped, limited-risk) use cases to build capability safely

Early public sector wins tend to be in bounded, low-to-medium risk applications, especially internal productivity tools and informational assistants, because they build capability without placing rights-sensitive decisions immediately on automated outputs. Cyprus deployed its gov.cy digital assistant as an informational guide — scoped to answering common citizen queries — before attempting higher-stakes automation.

C. Treat human oversight as an operating model, not a slogan

"Human-in-the-loop" is frequently assumed to be a safety net. In practice, it can fail through alert fatigue or automation bias, where humans rubber-stamp outputs at scale. Oversight only works if the government defines a Minimum Operating Model: roles, review thresholds, escalation triggers, and decision rights for pausing or withdrawing a system. The UK Government AI Playbook formalizes this requirement, making a named accountable owner and pre-defined review thresholds a prerequisite for any deployment.

D. Design for inclusion, not just efficiency

Generative AI introduces a powerful (and often underappreciated) inclusion pathway: translating administrative complexity into plain language, local dialects, and guided journeys across agencies ("starting a business," "retirement"). In many jurisdictions — including Singapore, where Virtual Intelligent Chat Assistant (VICA) serves 60+ agencies in multiple languages, and Estonia, where Bürokratt provides multilingual virtual assistance nationally — this may deliver higher public value than automating the decision itself, especially where legitimacy and trust are fragile.

Impact at a Glance: Where AI creates public value fastest

AI application type	Typical government benefit	Primary governance risk	Suitable for early deployment?
Citizen navigation assistants (service finders, FAQs, guided journeys)	Faster access, reduced admin friction, inclusion via language support (e.g., Singapore VICA, Cyprus gov.cy)	Misinformation/harmful guidance; unclear accountability	Yes (bounded + high value)
Internal productivity tools (drafting, summarizing guidance, triage support)	Reduced backlog, less time on routine admin (e.g., UK public sector AI coding-assistant trial)	Data leakage; overreliance; undocumented model updates	Yes (with strong controls)
Analytical risk scoring (fraud detection, targeting inspections/ resources)	Better targeting, proactive services, cost reduction (e.g., Togo's Novissi emergency targeting)	Bias, opacity, due process challenges	Conditional (needs strong guardrails)
Agentic workflow automation (multi-step execution across systems)	Step-change in throughput; automation of back-office chains (e.g., Singapore's Agentic AI Primer/early public sector agentic experimentation)	Delegated authority; auditability; "who decided?"	No (emerging frontier)

DCO Member Focus 1: Saudi Arabia's "AI Agents as Government Partners" (Early thinking on guardrails for agentic AI)

Saudi Arabia provides a useful indication of the direction frontier government AI governance may take, even as a fully operational model has yet to emerge. Under Vision 2030 and the National Digital Government Strategy, the Kingdom is positioning agentic AI as a potential next step in public sector modernization, particularly for citizen support, cross-agency coordination, and data-driven administration.

The Model:
SDAIA and the Digital Government Authority (DGA) are exploring "AI Agents as Government Partners including through the launch of the "Building AI Agents" camp: concepts where autonomous agents support workflow execution, drafting, and administrative coordination.

Key Governance Innovation:
SDAIA and DGA recognize that more capable agents will require stronger oversight, testing, security, and human control as systems move from information support toward action and delegation.

Strategic Impact:
Saudi Arabia's contribution is best seen as a frontier governance signal: as public sector AI moves from answering questions to taking multi-step actions, governments will need to define much more clearly where human authority ends, where machine autonomy begins, and how accountability is preserved.

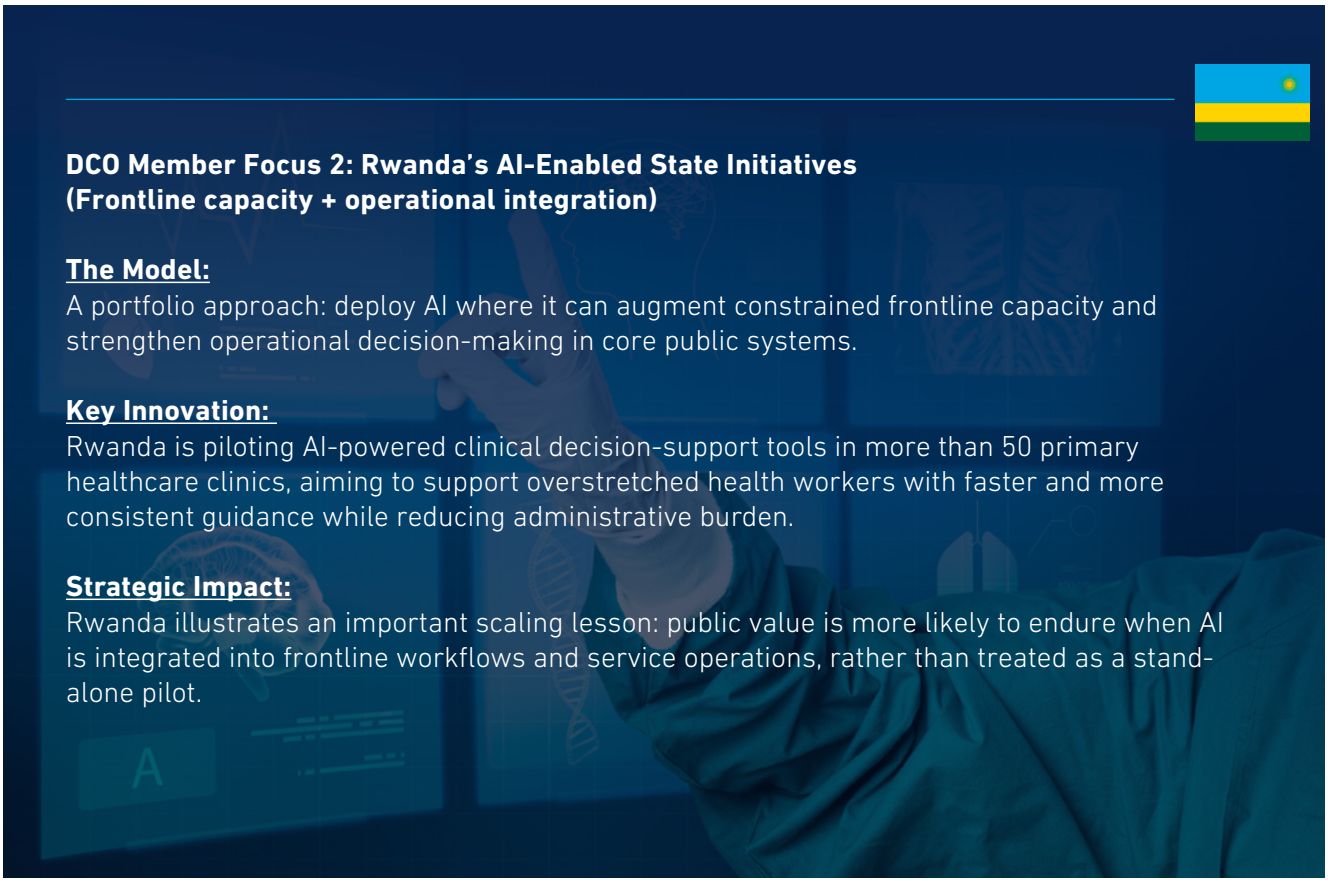

DCO Member Focus 2: Rwanda's AI-Enabled State Initiatives (Frontline capacity + operational integration)

The Model:
A portfolio approach: deploy AI where it can augment constrained frontline capacity and strengthen operational decision-making in core public systems.

Key Innovation:
Rwanda is piloting AI-powered clinical decision-support tools in more than 50 primary healthcare clinics, aiming to support overstretched health workers with faster and more consistent guidance while reducing administrative burden.

Strategic Impact:
Rwanda illustrates an important scaling lesson: public value is more likely to endure when AI is integrated into frontline workflows and service operations, rather than treated as a stand-alone pilot.

Global AI-in-government service-delivery approaches illustrate the range of paths to scale. **Singapore** emphasizes consistency through centralized infrastructure, including GovTech's shared AI/MLOps platform and common AI products for agencies. The EU pairs deployment with binding ex ante obligations under the AI Act, while the UK relies more on principle-led government guidance that encourages departments to build governance in at the design stage. **GCC states** prioritize state-led capacity: SDAIA and the UAE's AI Office build sovereign data infrastructure and skills pipelines ahead of wider deployment. China, meanwhile, illustrates a more centralized oversight model, where state authorities rely on filing, assessment, and supervisory mechanisms to keep closer visibility over high-impact algorithmic systems. For DCO Member States with more constrained resources, Rwanda offers a useful example of linking AI adoption to frontline service capacity rather than treating it as a parallel innovation exercise.

Key Takeaway

AI creates legitimacy only when it creates public value that can be measured, without turning government into a black box. The fastest path to scale is not ambitious automation; it is disciplined service design: bounded use cases, measurable KPIs, and an operating model for meaningful oversight.

Latest Developments: AI Service Transformation & Automation

- **UAE (Abu Dhabi) – AI-native government milestone:** In January 2026, Abu Dhabi's Department of Government Enablement reported that 95 per cent of public sector employees had completed AI training, and confirmed the appointment of Chief Digital and AI Officers across every government entity as part of its ambition to become the world's first fully AI-native government by 2027. (January 2026)
- **UK – CustomerFirst and AI Fellows program launched:** In January 2026, the UK government launched CustomerFirst — a new Department for Science, Innovation and Technology (DSIT) unit tasked with eliminating waiting times and repetitive form-filling in public services — alongside an AI Fellows program recruiting private-sector AI researchers directly into frontline service reform. (January 2026)
- **UK – AI service automation tools scaled to councils:** The one-year update on the UK AI Opportunities Action Plan (January 2026) confirmed that Minute, an AI tool that processes and summarizes council meeting recordings, is being rolled out as a live service to all local authorities; and that Extract, which processes planning documents into data, is expected to reach all councils by Spring 2026, freeing thousands of planning staff hours. (January 2026)
- **India – AI embedded in welfare and employment platforms:** In December 2025, India's Ministry of Labour & Employment integrated AI into e-Shram and the National Career Service — platforms covering 310 million informal workers — adding multilingual access, AI-assisted job matching, and automated resumé creation; social protection coverage has risen from 19 per cent in 2015 to 64 per cent by 2025. (December 2025)

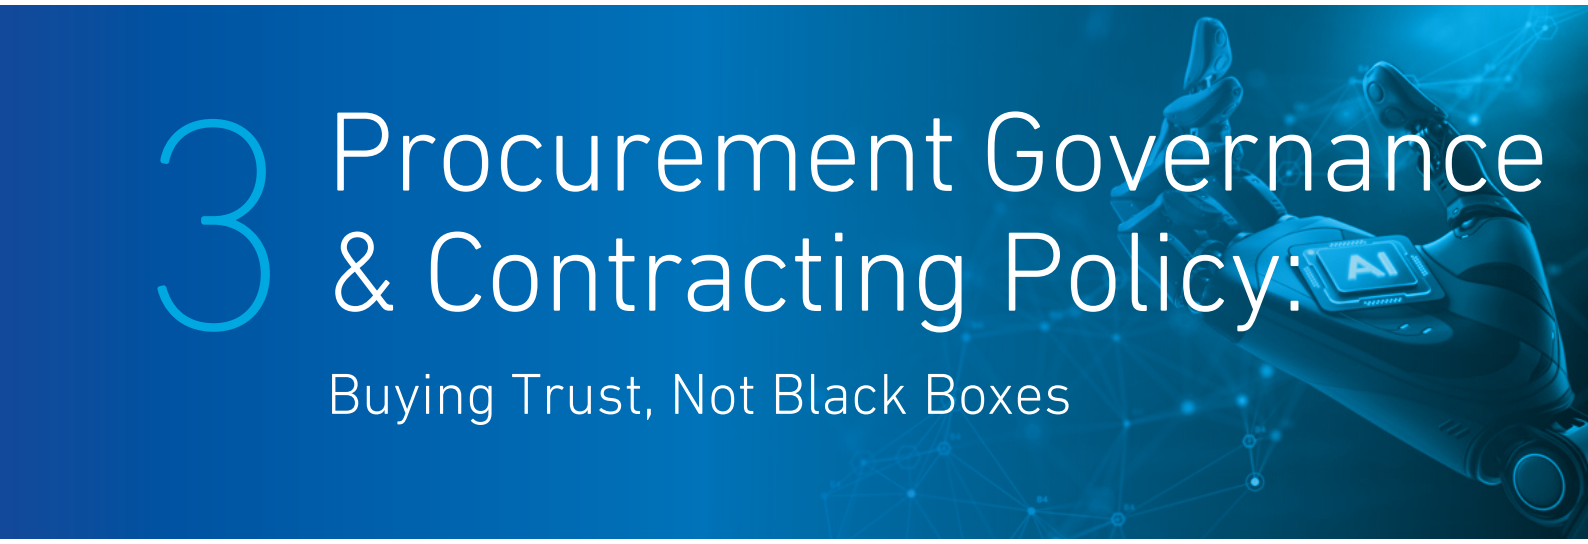

3 Procurement Governance & Contracting Policy:

Buying Trust, Not Black Boxes

Scaling AI in government is impossible if procurement remains business as usual. Traditional IT buying optimizes for price, delivery timelines, and functional requirements. AI buying must also optimize for **auditability**, **accountability**, **and sovereign control**, because procurement is where governments either prevent "black box dependency" (situations in which authorities cannot meaningfully audit a system's logic, scrutinize or contest its outputs, or disengage from a vendor without jeopardizing critical public functions), or institutionalize it.

The key shift is this: **In AI, the contract becomes a governance instrument.** Ethics principles are only real when they are translated into enforceable clauses: what must be disclosed, what must be tested, who can audit, what data can be reused, and what happens when the system fails.

Market Analysis: Procurement is becoming more agile & principle-driven

Across leading jurisdictions, procurement is evolving in two directions at once:

From Specifying the Solution → **To Specifying the Problem**

Governments are using challenge-based (problem-based) procurement: issuing outcome statements ("reduce call center wait times by 20%") and inviting vendors to propose approaches, often through proof-of-concept phases with milestone gates and outcome-based payments. This broadens supplier pools and keeps focus on public value rather than tooling. For example, Singapore's Open Innovation Platform allows government agencies to post defined challenge statements and test vendor solutions through proof-of-concept or prototype competitions, with successful solutions able to progress to pilot or larger-scale deployment.

From Buy and Deploy → **To Buy, Test, and Verify**

Procurement is increasingly incorporating staged validation and performance testing before pilot or wider deployment. Singapore's Outcome-Based Procurement model is a strong example: agencies define problem statements and KPIs, test vendor solutions through validation exercises, and then select solutions for pilot and later-stage deployment based on demonstrated performance.

Strategic Analysis:
Contracting for trust — from principles to enforceable clauses

A practical contracting toolkit is now emerging globally, with several key families of contractual clauses proving particularly important:

A. Auditability clauses

In some jurisdictions, public sector AI contracts are starting to reserve rights for procurers or independent third parties to access the documentation, datasets, logs, and performance information needed to verify compliance and explain outcomes, rather than treating trade secret claims as an absolute bar to scrutiny. **If the government cannot explain a decision that affects rights, it cannot govern that system.**

In practice, the European Commission's 2025 Model Contractual Clauses for the Public Procurement of AI codify auditability obligations — requiring vendor disclosure of system documentation, training data characteristics, and audit access rights — and are designed as a reusable template for EU public buyers.

B. Data portability and ownership

To reduce lock-in, governments are writing clauses that require vendors to return all relevant data in usable formats at contract end, and (where relevant) protect the government from training the vendor's model and losing the accumulated value when switching providers.

The EU Data Act (2023) establishes rights for public bodies to access and port (i.e., export and transfer it to another system or provider) data held by private vendors; the UK Cabinet Office Model Services Contract includes a dedicated exit-management schedule requiring data hand-back in usable formats at contract end.

C. Performance and bias benchmarks

Procurement is shifting toward explicit model performance thresholds and subgroup performance requirements (fairness benchmarks). These clauses create real remedies: fix, patch, retest, or face penalties/termination.

Canada's Directive on Automated Decision-Making is a leading example, imposing tiered bias impact assessments (Levels I–IV) as procurement prerequisites, with performance requirements that scale to the rights-sensitivity of the automated decision.

D. Alignment with ethical frameworks as a contractual duty

Rather than citing principles aspirationally, some governments are embedding **mandatory assessment mechanisms** into contracts, such as requiring scored evaluations pre-deployment and periodically thereafter.

DCO Member States can embed a scored ethics evaluation framework based on the DCO AI Ethics Evaluator as an enforceable procurement condition — assessment at award, and periodically thereafter. Canada's Algorithmic Impact Assessment — a mandatory scored evaluation required before deploying automated decision systems — is the clearest example of ethics assessment hardwired into procurement as an enforceable condition, not an aspirational principle.

E. Liability, remedies, and redress pathways

AI contracts that cap liability at low fee-linked levels can leave governments carrying much of the operational and political risk of failure. In response, public-sector procurement models are beginning to seek stronger supplier warranties, audit cooperation duties, corrective-action obligations, and enforceable remediation terms so that vendors remain accountable after deployment.

Market analysis of AI vendor contracts suggests that 88% of vendors cap liability, often at monthly subscription-fee levels, while only 17% commit to full regulatory compliance. Separately, a July 2024 ruling by a California federal court allowed discrimination claims against Workday to proceed on the theory that an AI vendor may in some circumstances be directly liable for downstream harms caused by its system.

DCO Member Focus: Bahrain's AI Procurement Guidelines (Turning ethics into procurement practice)

Bahrain's Information & eGovernment Authority has published AI Procurement Guidelines developed with the World Economic Forum through the AI Procurement in a Box initiative, aimed at helping government entities translate high-level AI ethics principles into practical procurement processes.

The Model:
A practical procurement framework that emphasizes proportionality, encouraging ministries to calibrate requirements to the potential risks and public impact of the system rather than applying a one-size-fits-all approach.

Key Innovations:

- Emphasizes ethics-by-design in procurement, including transparency, auditability, and early consideration of governance safeguards.
- Encourages buyers to consider human oversight, especially for consequential uses, and to address these issues during pre-market engagement and system design.
- Highlights the importance of interoperability and explicitly warns against vendor lock-in, helping public buyers preserve flexibility over time.

Strategic Impact:

Bahrain shows how governments can use procurement not just to buy AI systems, but to shape how they are designed, governed, and deployed — making transparency, accountability, and responsible innovation part of doing business with the state.

Global AI-in-government procurement approaches reflect a widening range of governance models. The EU is building the most systematic framework: the EU AI Act establishes common obligations for high-risk AI systems, including conformity-assessment requirements, while the European Commission's 2025 model contractual clauses give public buyers a practical template for contracting around those obligations.

Canada provides a strong example of impact-assessment-led governance, with its mandatory Algorithmic Impact Assessment and Directive on Automated Decision-Making embedding transparency, quality, and recourse requirements into the lifecycle of automated decision systems. The UK's approach is more principle-driven, combining procurement guidance with mandatory algorithmic transparency requirements for central government, but it is not best described as a system of mandatory AI impact assessments across procurement.

Singapore adds a test-driven procurement model through Outcome-Based Procurement, under which agencies can validate vendor solutions before pilot and wider deployment, while Project Moonshot serves as a structured pre-deployment testing toolkit for generative AI. For DCO members, Bahrain's AI Procurement Guidelines — developed with the World Economic Forum — offer an accessible proportionality-based starting point for translating responsible AI principles into procurement practice.

Key Takeaway

Trust is not a policy statement – it is a procurement output. Governments that scale AI responsibly are becoming smart buyers: problem-led, test-driven, and contractually explicit about auditability, data control, benchmarks, and remedies.

Latest Developments: Procurement Governance & Contracting Policy

- **US – New federal LLM transparency requirements for contractors:** On December 11, 2025, the Office of Management and Budget (OMB) issued Memorandum M-26-04, requiring federal agencies to include contract clauses for procured LLMs addressing the "Unbiased AI Principles" of **truth-seeking** and **ideological neutrality**. Agencies must update procurement policies by **March 11, 2026**, and solicitations must request documentation sufficient to assess compliance, including items such as acceptable-use policies, model/system/data cards, evaluation information, and disclosures about relevant model modifications. The memo builds on OMB Memorandum **M-25-22**, whose AI acquisition guidance took effect for relevant solicitations and contract actions in **October 2025**. (December 2025)
- **US – FY2026 Defense Authorization Act combines broader acquisition reform with new AI strategy requirements:** Signed into law on December 18, 2025, the FY2026 NDAA overhauls key parts of the Department of Defense acquisition lifecycle, including a shift toward a portfolio-based acquisition model. Separately, it directs the establishment by April 1, 2026 of an AI steering committee to analyze advanced and emerging AI technologies, including systems that could enable artificial general intelligence, and to develop a strategy for the risk-informed adoption of AI. (December 2025)
- **EU – High-risk AI compliance deadlines are beginning to reshape procurement practice:** Under the EU AI Act, many obligations for high-risk AI systems begin to apply from 2 August 2026, while high-risk systems embedded in regulated products have until 2 August 2027. In practice, this is pushing public buyers to revisit contract templates and due diligence processes, even though conformity assessments remain primarily a provider obligation and registration duties depend on the role of the provider or deployer. Since 2 August 2025, the AI Act's governance regime has been in force, with the AI Office overseeing general-purpose AI model rules and national competent authorities enforcing most system-level obligations and penalties. (August 2025)
- **South Korea – Next-generation KONEPS launched:** South Korea's Next-generation KONEPS, launched in March 2025, integrates 25 procurement platforms and additional data systems into the national e-procurement platform. According to the OECD, it uses AI and Big Data to predict bidding congestion, monitor procurement operations, and recommend products, while also linking procurement data with taxation records and other government databases to strengthen efficiency and accountability across the procurement lifecycle. (June 2025)
- **UK – Government adopts a more market-first AI procurement approach: In its 29 January 2026 AI Opportunities Action Plan: One Year On update**, the UK government said it had developed an AI Commercial Strategy that prioritizes buying from the market and innovating through challenge-led procurement. The update also points to AI Accelerator Tenders and an i.AI scan function designed to speed procurement and simplify supplier engagement across government, indicating a push toward faster and more coordinated market engagement for AI procurement. (January 2026)

4 Operational Accountability:

Monitoring, Assurance, and Incident Pathways

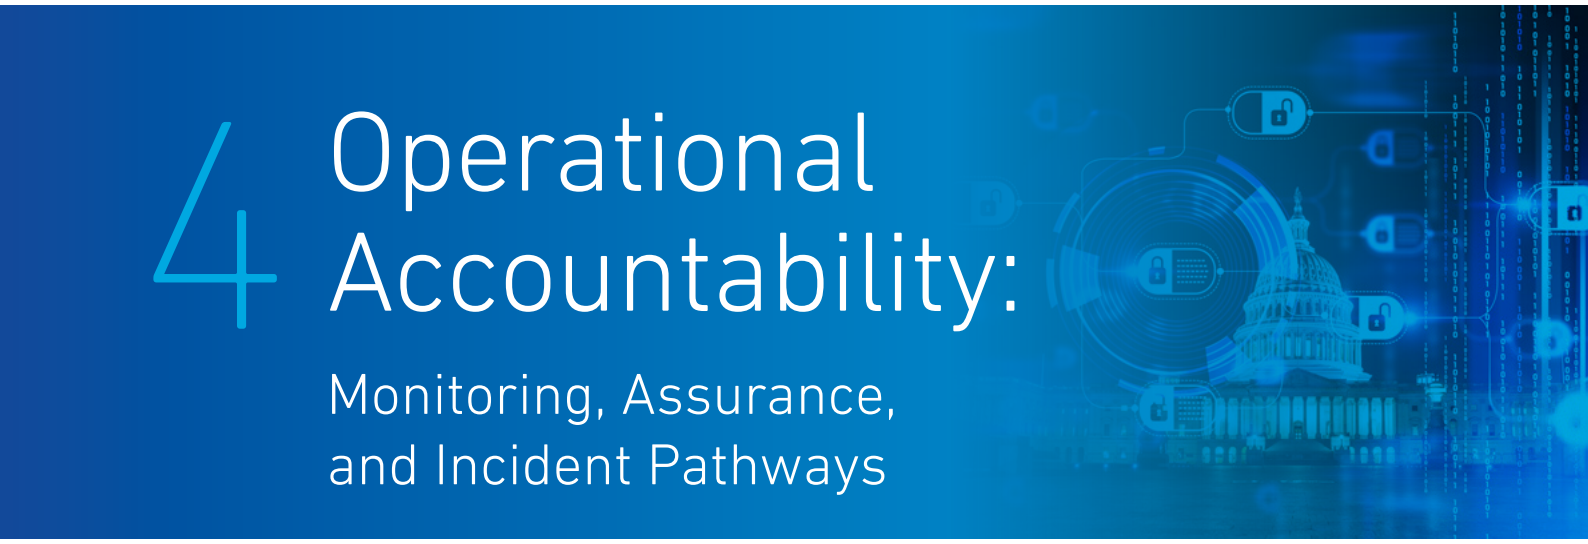

AI governance failures emerge most consequentially not at the pilot stage but after deployment, when systems drift, policies change, new data patterns emerge, and real-world edge cases accumulate. This is why the central governance challenge is not only responsible adoption, but **responsible operations**: the ability to continuously verify performance, detect harm, and intervene quickly when things go wrong.

In practice, many governments still treat AI as a deliverable rather than a living system. That creates **set-and-forget risk**: models degrade silently, accountability gets diffused across agencies and vendors, and incidents are managed ad hoc, often only once a failure becomes public.

Market Analysis: The accountability gap in production

As governments move toward the Cognitive State, accountability must shift from static checklists to **continuous assurance**, anchored in three capabilities.

Three capabilities for AI-in-government accountability

1. Ongoing monitoring (performance and harm signals)

Monitoring should extend beyond technical accuracy to include administrative and user-impact indicators that can reveal legitimacy risks early, such as complaints, appeals, override rates, delays, or disproportionate effects on particular groups. Useful monitoring dimensions include:

Model performance drift	Fairness and subgroup performance	Operational integrity	User impact signals	Security and misuse
accuracy, error types, reliability over time	disparate error rates across protected or vulnerable groups	override rates, escalation volumes, processing delays	complaints, appeal outcomes, satisfaction	prompt injection, data leakage, unauthorized access

Under Article 72 of the EU AI Act, providers of high-risk AI systems must maintain a post-market monitoring system that actively collects and analyses relevant performance data over the system's lifetime, although the Act does not prescribe these dimensions as an exact checklist. In the Netherlands, the Algorithm Register complements this by supporting public transparency about government algorithms.

China provides a more centralized version of this logic: under the Cyberspace Administration (CAC)'s algorithm governance regime, certain services with 'public opinion attributes' or 'social mobilization capacity' must complete regulatory filings and security assessments, giving authorities direct visibility into higher-impact algorithmic deployments.

2. Institutionalized assurance functions (who owns accountability)

Accountability requires named owners and clear decision rights. A recurring best practice pattern is the creation of a **central assurance capability**, either within a digital government authority, a data/AI authority, or an independent audit function, responsible for setting standards, validating deployments, and coordinating incident response across government.

This is also where governments are beginning to operationalize structured evaluation mechanisms (including ethics and impact assessment tools) as part of an internal assurance pipeline, rather than treating them as one-off documentation exercises. Canada's Treasury Board Secretariat provides a strong example of this model through its Directive on Automated Decision-Making and mandatory Algorithmic Impact Assessment, which establish central standards and risk-based requirements for departments deploying automated decision systems; the UK Algorithmic Transparency Recording Standard plays a related but narrower role by mandating disclosure of in-scope algorithmic tools across central government, strengthening transparency and accountability rather than serving as a full assurance or validation function. China also illustrates how centralized institutional capacity can support oversight: alongside CAC filing requirements, the AI Subcommittee of the National Ethics Committee has helped shape guidance and coordination on AI governance.

3. Clear incident pathways and redress mechanisms (what happens when AI causes harm)

When AI systems influence eligibility, enforcement, or resource allocation, governments need predictable pathways for:

Incident triage
severity levels; who is paged; required response times

Containment
pause/rollback authority; kill switches; model version control

Investigation
logs, audit trails, vendor cooperation requirements

Communication
internal escalation; external transparency where appropriate

Redress
how citizens challenge outcomes; how cases are corrected

A key leadership question here is straightforward:
If the system produces a harmful outcome today, can we identify it quickly – and can we stop it?

Strategic Analysis: the minimum operating model for accountable AI

Across jurisdictions, successful scaling converges on a minimum operating model that makes accountability real in day-to-day administration:

A. Named accountable owner (business + technical)

Every AI-enabled service requires a business owner (service accountability) and a technical owner (model/system accountability), with clear sign-off authority and escalation responsibilities.

B. Documentation that supports oversight, not bureaucracy

Governments are increasingly prioritizing documentation that enables audit and redress (what the system does, where it is used, known limitations, evaluation results, and change logs), rather than lengthy ethics statements that do not translate into operational control.

C. Human oversight designed to work at scale

Oversight must be engineered to avoid automation bias. That typically means: defined review thresholds, sampling regimes, override tracking, and periodic quality audits, especially when AI influences high-volume decisions.

D. Lifecycle governance (re-testing, re-approval, retirement)

Accountable AI requires a lifecycle view: periodic re-testing, trigger-based reassessments when data/policy changes, and explicit retirement pathways when systems no longer meet thresholds.

Accountable AI in Practice

In practice, the UK Government AI Playbook strongly supports these elements through its emphasis on accountability, transparency, contestability, testing, and responsible governance, though it does not formalize all four as a single mandatory prerequisite checklist for every deployment;

Canada's Algorithmic Impact Assessment process more clearly embeds lifecycle governance, requiring departments to review and update assessments on a scheduled basis and whenever system functionality or scope changes.

DCO Member Spotlight: Oman's National AI Programme (Operational Focus on Accountability Capacity)

Oman's approach to AI in government reflects a foundational insight: accountable deployment depends not only on technology adoption, but on whether the state has the internal capability to supervise, govern, and scale these systems effectively. Rather than treating AI as a stand-alone technology rollout, Oman links adoption to civil service capability, institutional readiness, and human-centric governance.

The Model

Through the Ministry of Transport, Communications and Information Technology's National Program for Artificial Intelligence and Advanced Digital Technologies, Oman is combining AI adoption with capacity-building, governance development, and risk-based oversight. This helps ensure that accountability does not rest solely on vendor claims or ad hoc supervision.

What's distinctive

A flagship initiative is the Ministry's Ertiqaa capacity-building program, designed to strengthen the ability of public officials and digital transformation teams to lead and supervise emerging technology adoption across government. Official materials indicate that it trains cohorts of civil servants across:

- **Digital governance and data management**, helping ministries build the institutional foundations for oversight and informed decision-making.
- **Emerging technologies, including AI**, building both conceptual literacy and practical readiness for adoption.
- I**nnovation methods, digital risk management**, **and organizational change,** so implementation is supported by governance frameworks, coordination mechanisms, and delivery discipline rather than left to isolated technical teams.

Why it matters for accountability

Ertiqaa is not simply a technical skills program. In practice, it supports the creation of the internal capability that ministries need to establish governance structures, oversee deployment, and make responsible use of AI systems. This is reinforced by Oman's broader AI governance agenda, which emphasizes human-centric and risk-based governance.

Strategic Impact:

Oman's approach demonstrates a practical scaling lesson: Operational accountability is built, not assumed. Before deploying higher-impact AI systems, governments need trained owners, clear roles, and institutional muscle memory for oversight. Otherwise, oversight becomes symbolic rather than effective

Global accountability approaches illustrate the spectrum of operational models. The EU AI Act establishes the most prescriptive post-deployment regime, requiring mandatory post-market monitoring plans and serious incident reporting obligations for high-risk systems. The UK's Algorithmic Transparency Recording Standard and the Netherlands Algorithm Register represent disclosure-led accountability: public registers that document active systems, create accountability to citizens, and enable external scrutiny. For governments with leaner institutions, Oman's model — building civil service supervision capacity before widening deployment — may be the most practical route to meaningful oversight.

Key Takeaway

In government, accountability is an operating capability – not a policy statement.
A very important question for leaders is whether AI systems can be continuously monitored, independently assured, and quickly paused or corrected when harms emerge.

Latest Developments: Operational Accountability, Monitoring & Incident Pathways

- **Canada – Strengthened Automated Decision-Making Directive enters force:** On June 24, 2025, Canada's Treasury Board Secretariat amended its Directive on Automated Decision-Making, adding stronger senior-level accountability, enhanced testing and monitoring requirements, and a more detailed Algorithmic Impact Assessment tool. Federal institutions using pre-existing in-scope automated decision systems have until June 24, 2026 to comply with the new measures, reinforcing Canada's position as one of the most operationally developed public sector AI governance frameworks.(June 2025)
- **G7/OECD – Hiroshima Process AI reporting framework goes live:** In February 2025, the OECD launched a standardized voluntary reporting framework for organizations to demonstrate alignment with the G7 Hiroshima AI Code of Conduct, covering risk management practices, incident reporting, and information-sharing mechanisms. Companies including Amazon, Anthropic, Google, Microsoft, and OpenAI have participated in the framework, helping establish a new international reference point for transparent post-deployment accountability practices. (February 2025)
- **EU – Digital Omnibus introduces conditional delays to high-risk AI obligations:** In November 2025, the European Commission's Digital Omnibus package proposed linking the application of rules for Annex III high-risk AI systems to the availability of support tools, including harmonized standards, with a long-stop date of 2 December 2027 rather than the original 2 August 2026 deadline. For public authorities using AI in areas such as benefits administration or law enforcement, this may create additional implementation time, but it does not remove the need to prepare governance, documentation, and monitoring arrangements ahead of compliance. (November 2025)
- **US (California) – New law bars "AI autonomy" as a civil defense:** California's AB 316, effective January 1, 2026, provides that in a civil action against a defendant who developed, modified, or used AI alleged to have caused harm, it is not a defense to argue that the AI system autonomously caused the plaintiff's harm. The law preserves other defenses, but it reinforces the principle that human or organizational accountability does not disappear simply because an AI system acted with some degree of autonomy. (October 2025)
- **Canada (Ontario) – Human rights and privacy regulators issue joint AI principles:** In January 2026, Ontario's Information and Privacy Commissioner and the Ontario Human Rights Commission jointly released six principles for the responsible use of AI, covering reliability, safety, privacy, human rights, transparency, and accountability, including human oversight across the AI lifecycle. The joint release reflects growing recognition that AI oversight cannot be siloed within a single regulatory function. (January 2026)

5 Conclusion:

A Roadmap for AI Leadership in Government

AI in government is no longer a future capability. It is becoming a core layer of how the state delivers services, manages resources, and exercises authority. The opportunity is significant: governments can reduce administrative friction, improve responsiveness, and unlock better outcomes for citizens.

But the decisive factor will not be how quickly AI is adopted. It will be whether governments can **govern AI at scale**. That means building the foundations for trustworthy deployment (data readiness, capable institutions, and clear boundaries), using procurement as a lever for accountability, and operationalizing continuous assurance through monitoring, incident pathways, and accessible redress.

The governments that succeed will be **technology shapers**, not technology takers: they will embed accountability into contracts, strengthen internal supervision capacity, and define what must remain non-delegable as more agentic systems emerge. The result is not only better services, but a more resilient state: one that can harness AI while preserving legitimacy, transparency, and public trust. The following roadmap provides a practical framework for sequencing that work.

A Roadmap for AI Leadership in Government

The shift from digital government to the Cognitive State is a strategic governance challenge, not only a technology adoption challenge. AI can improve public services, productivity, and state capability, but only where deployment, procurement, and oversight are built as one system.

The three pillars are interdependent: weakness in any one undermines progress in the others:

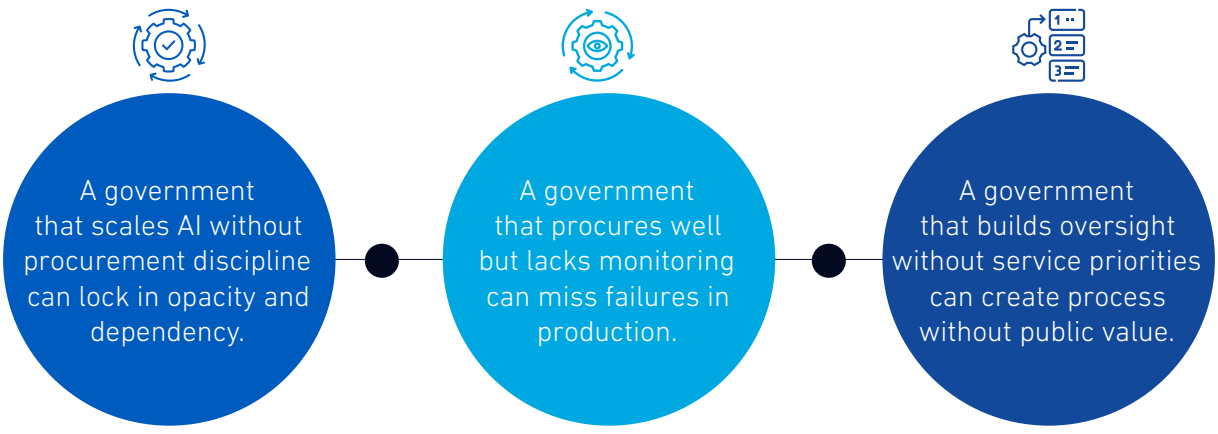

A Strategic Roadmap for DCO Member States

The roadmap below is structured in three stages — **Foundational** (readiness), **Accelerating** (scale with oversight), and **Leadership** (standard-setting and strategic capability) — across all three pillars. Governments may sit at different stages across different pillars; the purpose is to identify and close the most important gaps first. The recommendations are cumulative: earlier-stage actions remain relevant as governments advance.

Pillar I: Delivering Public Value — Service Transformation & Automation	Pillar II: Procurement & Contracting Policy	Pillar III: Operational Accountability, Monitoring & Incident Pathways
1. Foundational: Building core strategy, procurement discipline, and minimum operational safeguards		
Establish a government AI service priority plan. Identify a small number of bounded, lower-risk service domains with clear public-value outcomes – such as internal productivity tools, informational chatbots, and service-navigation assistants. AI systems should not be deployed as a patchwork of disjointed initiatives, but aligned with strategic objectives and interoperable where relevant. Define a minimum governance operating model (named owner, basic documentation, risk screen) before any pilot goes live. **In practice:** Bahrain — policy-led early adoption, with national AI strategy and ethics materials linked to practical exploration of government AI-enabled services, illustrating an early strategy-to-service pathway.	**Publish government AI procurement guidance.** Provide practical instructions for officials on scoping, risk questions, minimum documentation, and vendor disclosures. Apply proportionate requirements by risk tier rather than a uniform standard. **In practice:** Bahrain AI Procurement Guidelines (based on WEF's AI Procurement in a Box) — a practical example of adapting responsible AI procurement guidance for government use, with an emphasis on proportionality, transparency, auditability, and avoiding vendor lock-in.	**Define a minimum AI accountability operating model.** Assign named business and technical owners before deployment; require basic documentation, a risk screen, and a defined incident escalation path. Apply human review to any rights-affecting outputs. **In practice:** UK Government AI Playbook (2025) — practical guidance covering governance, procurement, human oversight, quality assurance, and lifecycle management that departments can apply before and after a system goes live.
2. Accelerating: Scaling deployments with transparency, monitoring, and internal assurance capacity		
Scale AI to priority services under explicit guardrails. Expand from pilots to production only where governance conditions are met. Introduce public-value KPIs — time-to-decision, error rates, override volumes, and complaints — to track outcomes, not only technical performance. **In practice:** Canada's immigration processing tools — advanced analytics used to streamline routine applications and triage cases by complexity, while keeping officers central to consequential decisions.	**Move from one-off buying to portfolio governance.** Maintain central visibility over major AI systems, contracts, renewals, and risks. Introduce lifecycle performance controls, model-change notifications, and portability clauses as standard contract terms. **In practice:** Singapore's Public Sector AI Playbook and central procurement pathways — including AI product catalogues, accredited supplier routes, and earlier bulk tender arrangements — provide a practical model for reducing procurement fragmentation and giving agencies more consistent access to AI services.	**Establish a public AI use register and implement continuous monitoring.** Publish material AI systems, their purposes, and accountable owners, while separately implementing standing internal monitoring for performance, overrides, subgroup impacts, and user signals. **In practice:** The UK Algorithmic Transparency Recording Standard and the Netherlands Algorithm Register are workable public disclosure and accountability models that strengthen external scrutiny, though they are better understood as complements to — rather than substitutes for — internal continuous monitoring. China's Provisions on the Management of Algorithmic Recommendations in Internet Information Services offer a more centralized variant of registry-led oversight, tying visibility into certain high-impact systems to regulatory filing and security review.

3. Leadership: Shaping standards, building reusable tools, and reducing strategic dependency		
Develop reusable government AI components and pioneer high-control frontier deployments. Build shared evaluation tools, secure retrieval layers, and reusable service modules. Pilot agentic systems in constrained workflows with defined delegation limits, audit trails, and human override authority. **In practice:** Singapore AI Verify / Project Moonshot — an open-source LLM evaluation and stress-testing toolkit that offers a reusable foundation for benchmarking and red teaming across organizations, supporting more consistent deployment and assurance practices.	**Publish shareable model procurement clauses and shape vendor markets.** Move from internal templates to reusable standards for other public buyers. Use the combined weight of government procurement to drive vendor behavior toward transparency, interoperability, and portability — making accountability a condition of doing business with the state. **In practice:** EU Model Contractual Clauses for AI Procurement and Bahrain's WEF-adapted guidelines — two examples of public procurement toolkits that other governments can adapt to their own contexts, with the EU model clauses providing the stronger standard-setting benchmark.	**Establish mature assurance regimes with formal re-evaluation triggers.** Reassess systems when models, data, or use contexts materially change. Create pathways to capture lessons from incidents and near-misses, and define formal red lines for non-delegable decisions as more autonomous systems emerge. **In practice:** Canada's Algorithmic Impact Assessment process — departments must review and update the assessment on a schedule and whenever system functionality or scope changes, providing a strong model for formal re-evaluation triggers and lifecycle governance.

Use this roadmap as a sequencing tool, not a maturity label. Progress may differ across pillars; the priority is to identify the binding constraint in each pillar and address it first. Across all stages, the core principle is unchanged: **trust must be engineered.** AI in government should operate as an accountable public system — intelligible, monitored, and aligned with the public interest.

DCO Member State Snapshot 1: Strategy & Service Readiness

- Gov AI Policy/Guidance
- GovAI Initiative Example

Bahrain

- General Policy for Use of AI+ AI Procurement Guidelines for government sector
- **Nationality, Passports and Residence Affairs AI-powered chatbot:** Answers passport, residency, visa, and ID queries

Bangladesh

- Draft National AI Policy 2026-2030
- **Aspire to Innovate:** Leads digital government innovation and citizen service transformation

Cyprus

- EU AI Act Digital Assistant Usage Policy
- **Gov.cy AI Digital Assistant:** Guides users to government information, forms, and procedures

Djibouti

- Early-stage policy work; AI governance under development
- **AI for School Mapping & Infrastructure:** Uses AI/satellite data to map schools and plan connectivitye

The Gambia

- None
- **Government Chatbot:** Planned AI-powered virtual assistant to guide citizens to e-services

Ghana

- Governance components evolving; strategy provides policy baseline (ethics/skills/data pillars)
- **GEPbot (Ghana Entrepreneurship Policy Chatbot):** Explains entrepreneurship policies and answers SME regulatory questions

Greece

- EU AI Act trajectory + national transparency/accountability measures (registry/obligations referenced in legal guidance)
- **AI digital assistant:** Helps citizens navigate gov.gr services and requirements

Jordan

- National AI Code of Ethics
- **Smart chatbot:** Directs citizens to digital services via WhatsApp/phone support

Kuwait

- National AI Strategyrk 2025-2028
- **Public Authority for Civil Information mobile applications:** Digital ID app for authentication, verification, and e-signature

Morocco

- Maroc IA 2030
- **Tax authority virtual assistant:** Answers tax questions and helps users navigate filings/services

Nigeria

- Draft AI Code of Practice; National AI Strategy
- **Youth Help Desk WhatsApp AI chatbot:** Provides youth support info, guidance, and referrals via WhatsApp

Oman

- National AI Policy 2024
- **Ask Raayid:** Provides entrepreneurs/SMEs quick guidance on services and support via WhatsApp

Pakistan

- AI Policy 2025
- **Sindh Population Welfare Dept AI Chatbot + AI call assistant:** Provides family planning guidance & service referrals via chat/voice

Qatar

- Principles and Guidelines for Ethical Use of AI; National AI Strategy 2019
- **Ministry of Commerce and Industry's Smart Assistant:** Guides business setup and handles inquiries/complaints on Single Window

Rwanda

- National AI Policy 2023
- **Rwanda Biomedical Centre government chatbot:** Provides official health guidance, statistics, and service information

Saudi Arabia

- SDAIA AI Ethics Principles
- **Absher "Personal Assistant Service":** AI-enabled assistance to help users complete Absher services

DCO Member State Snapshot 2: Governance & Procurement Guardrails

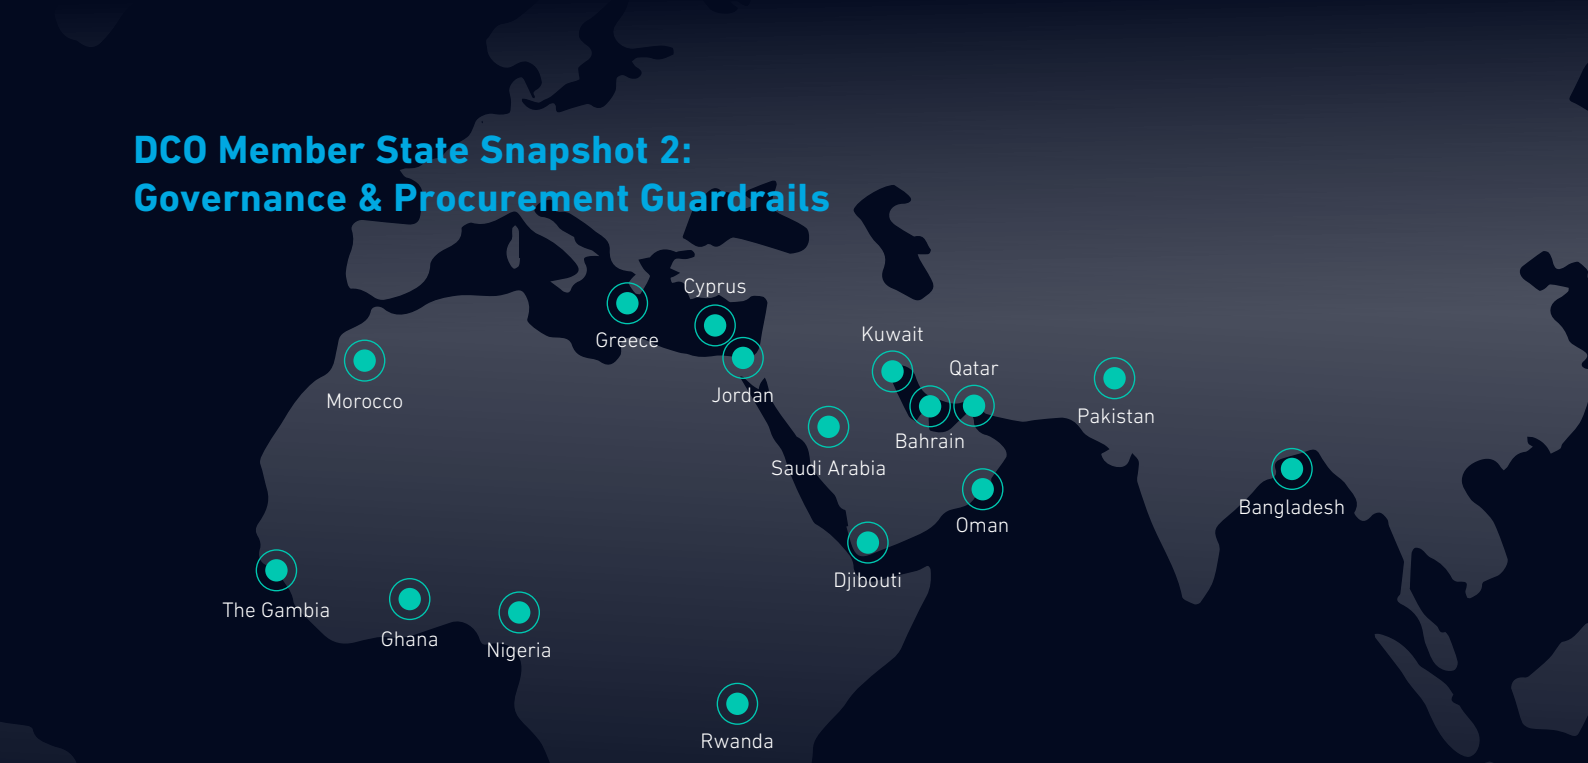

Bahrain

AI Procurement Guidelines for government sector

Bangladesh

Introduced in the draft National AI Policy

Cyprus

Planned under the EU AI Act

Djibouti

None

The Gambia

None

Ghana

AI Practitioners' Guide

Greece

Planned under the EU AI Act

Jordan

Ministry of Digital Economy and Entrepreneurship mandated to review legislative requirements related to government procurement (alongside the Procurement Department)

Kuwait

None

Morocco

None

Nigeria

None

Oman

None

Pakistan

None

Qatar

NCSA Guidelines for Secure Adoption and Usage of AI can be reflected in procurement (security-by-design expectations), but they are not framed as a procurement standard

Rwanda

The National AI Policy states that "the Government will engage local AI solution providers through innovation-friendly procurement processes." In 2026, Rwanda is expected to develop procurement guidelines for public projects including AI components

Saudi Arabia

SDAIA's GenAI Government Guidelines for Government are directly usable as procurement guardrails, but positioned as guidance rather than a procurement regulation

DCO Member State Snapshot 3: Operational Accountability Maturity

- Human-in-the-Loop Mandate Status
- Civil Service Literacy Program

Bahrain

- Recommended Bahrain Government - Bahrain's Approach to AI Governance
- Yes (AI Tools for Employee Productivity)

Bangladesh

- Recommended
- Yes (Effective e-governance: Accelerating e-government and digital public services in Bangladesh)

Cyprus

- Mandated for high-risk AI systems
- Yes (Cyprus Academy of Public Administration - Digital Skills Action Plan)

Djibouti

- None
- Yes (digital skills program for government officials)

The Gambia

- None
- Yes (e-gov strategy highlights training civil servants)

Ghana

- Not specified (strategy emphasizes responsible AI)
- Yes (TBI Digital Academy and DigSMART)

Greece

- Mandated for high- risk AI systems
- Yes ("AI for All" initiative)

Jordan

- Recommended (AI Strategy emphasizes oversight of AI adoption procedures)
- Yes (AI Raising Awareness for Public Employees)

Kuwait

- Emerging (human oversight in the upcoming AI Strategy)
- Yes (Kuwait National Skilling Initiative)

Morocco

- None
- Yes (through BMZ project)

Nigeria

- Recommended
- Yes (Capacity-building program for civil servants)

Oman

- Recommended
- Yes (Ertiqaa program offers comprehensive digital training, including AI skills)

Pakistan

- Recommended
- Yes (AI training is mandatory for all new civil servants)

Qatar

- Recommended
- Yes (National AI Strategy & Digital Agenda includes comprehensive workforce upskilling)

Rwanda

- Not specified
- Yes (Mandatory training for all civil servants)

Saudi Arabia

- Guidance-based
- Yes (Workflow Automation Using AI program)

Document Disclaimer

The following legal disclaimer ("Disclaimer") applies to this document ("Document") and by accessing or using the Document, you ("User" or "Reader") acknowledge and agree to be bound by this Disclaimer. If you do not agree to this Disclaimer, please refrain from using the Document.

This Document, prepared by the Digital Cooperation Organization (DCO). While reasonable efforts have been made to ensure accuracy and relevance of the information provided, the DCO makes no representation or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability or availability of the information contained in this Document.

The information provided in this Document is intended for general informational purposes only and should not be considered as professional advice. The DCO disclaims any liability for any actions taken or not taken based on the information provided in this Document.

The DCO reserves the right to update, modify or remove content from this Document without prior notice. The publication of this Document does not create a consultant-client relationship between the DCO and the User.

The designations employed in this Document of the material on any map do not imply the expression of any opinion whatsoever on the part of the DCO concerning the legal status of any country, territory, city, or area or of its authorities, or concerning the delimitation of its frontiers or boundaries.

The use of this Document is solely at the User's own risk. Under no circumstances shall the DCO be liable for any loss, damage, including but not limited to, direct or indirect or consequential loss or damage, or any loss whatsoever arising from the use of this Document.

Unless expressly stated otherwise, the findings, interpretations and conclusions expressed in this Document do not necessarily represent the views of the DCO. The User shall not reproduce any content of this Document without obtaining the DCO's consent or shall provide a reference to the DCO's information in all cases.

By accessing and using this Document, the Reader acknowledges and agrees to the terms of this Disclaimer, which is subject to change without notice, and any updates will be effective upon posting.

Follow us on

@dcorg

www.dco.org