



EPIC-PP-2026-008 · POLICY ESSAY · AI GOVERNANCE

Sunset Clauses for Algorithms: Designing Expiry Dates into Government AI

EPIC – Emirati Perspectives and Insight Centre

Seeing what others miss...

20 October 2026

CC BY 4.0 · Creative Commons Attribution 4.0 International · creativecommons.org/licenses/by/4.0

Sunset Clauses for Algorithms: Designing Expiry Dates into Government AI

EPIC – EMIRATI PERSPECTIVES AND INSIGHT CENTRE

EPIC-PP-2026-008 · AI Governance · Published 20 October 2026

© 2026 EPIC – Emirati Perspectives and Insight Centre. Released under Creative Commons Attribution 4.0 International (CC BY 4.0): <https://creativecommons.org/licenses/by/4.0/>

Government algorithms rarely die. A scoring tool approved for a pilot becomes a permanent queue manager. A risk model tuned to last year's caseload keeps ranking this year's residents. A vendor pushes an update, accuracy reports look fine in aggregate, and nobody is forced to ask the only political question that matters: should this system still be allowed to advise the state?

Public law already knows how to ask that question of statutes. Sunset clauses—provisions that cause a law or temporary measure to expire unless affirmatively renewed—exist precisely because legislators distrust their own foresight. Sofia Ranchordás's comparative work on constitutional sunsets and experimental legislation treats such clauses as tools of legislative quality and controlled experimentation, not as gimmicks ([Ranchordás, 2014](#)). This essay argues that Gulf governments, and Dubai's policy community in particular, should import the logic of sunsets from laws to algorithms: significant government AI systems should carry expiry dates that lapse unless a named authority renews them on evidence.

Why permanence is the wrong default

AI systems are not static infrastructures like bridges. Their behaviour depends on data distributions, user adaptations, policy changes, and model updates. The NIST AI Risk Management Framework treats risk management as lifecycle work: GOVERN, MAP, MEASURE, and MANAGE are iterative, not one-off procurement checklists ([NIST, 2023](#)). The EU AI Act likewise requires post-market monitoring for high-risk systems and imposes recurring evaluation and review duties on the Commission itself (Regulation (EU) 2024/1689, Arts. 72 and 112; [Art. 112](#)). These instruments recognise drift and context change. What they do not always force, at the level of a single public deployment, is a hard stop.

Continuous monitoring without expiry can become continuous deferral. Dashboards reassure; renewal decides. A sunset clause flips the default. Instead of remaining in force until someone finds the courage to switch it off, the system falls silent unless someone finds the evidence to switch it back on. That asymmetry matters in government, where inertia protects incumbents—including incumbent software.

What an algorithmic sunset is—and is not

An algorithmic sunset is a time-bounded authorisation to use a defined AI system for a defined public purpose. When the clock runs out, lawful use for that purpose requires affirmative renewal: a written decision by a competent official or board, based on a short evidence dossier. The dossier should cover performance against the original purpose, complaint and override patterns, known distributional effects, material model or data changes, and whether a less intrusive alternative now exists.

An algorithmic sunset is not a ban on AI. It is not a claim that every spam filter needs a cabinet vote. Dubai's AI Ethics Principles already concentrate attention on significant and critical decisions ([Smart Dubai, 2019](#)). Sunsets should follow that risk logic: shorter clocks and stricter renewal for systems that affect rights, benefits, penalties, or equal access; longer clocks for low-stakes assistive tools.

Nor is a sunset a substitute for human oversight at the moment of decision. Oversight answers who finalises a case today. A sunset answers whether the machine may still sit in the advisory chair next year. Governments need both. The UAE Charter for the Development and Use of Artificial Intelligence emphasises human oversight together with governance and accountability ([UAE Government, 2024](#); [WAM, 2024](#)). Sunsets give those principles a calendar.

Design rules for expiry dates

First, inventory before poetry. No sunset regime works without a register of government AI systems that inform significant decisions: purpose, owner, vendor, data sources, go-live date, and current expiry. The UAE's national AI governance architecture—strategy Objective 8 on strong governance and effective regulation, and the AI and Blockchain Council's review remit—offers an institutional home for such inventories ([UAE National Strategy for Artificial Intelligence 2031](#); [AI Council](#)).

Second, tier the clocks. A practical starting point, offered as policy design rather than as settled law, is twelve to twenty-four months for high-impact systems and thirty-six months for medium-impact assistive systems, with emergency extensions only for documented continuity risks. The numbers can be debated; the existence of a clock should not be.

Third, make renewal affirmative and public enough to be real. A renewal that is a silent checkbox in an IT ticket is not a sunset; it is a subscription. The renewing authority should be senior to the system's day-to-day operator. Where feasible, a short public notice—system name, purpose, renewal term, and a non-sensitive summary of performance—would align with Dubai's transparency and accountability principles ([Smart Dubai, 2019](#)).

Fourth, couple expiry to kill-switches and version locks. If authorisation lapses, the system must not continue through a back-end cron job. Material model changes during a licence period should restart or abbreviate the clock; otherwise vendors can evade sunsets by perpetual quiet updates. OECD AI Principles on accountability and lifecycle risk management support this disciplined change control ([OECD, 2024](#)).

Fifth, fund the review. Sunsets fail when renewal dossiers are unbudgeted. A fraction of every significant AI project's operating budget should be reserved for independent evaluation capacity—inside government or commissioned—before the expiry date arrives.

Sixth, distinguish system sunsets from vendor relationships. A government may renew a contract while retiring a model, or retire a contract while rebuilding an in-house tool. The expiry date attaches to the authorised decision-support function—the specific model family, purpose, and deployment context—not merely to a purchase order. Otherwise agencies will confuse procurement continuity with policy consent.

These rules are deliberately boring. Sunset design fails when it becomes metaphysical. The point is not to declare AI illegitimate after two years. The point is to force a scheduled confrontation with evidence while officials still remember why the system was bought.

Failure modes sunsets are built to catch

Purpose drift. A tool approved to prioritise building inspections begins to shape enforcement intensity in ways nobody defended at launch. Without renewal, purpose creep becomes custom.

Population drift. A model trained on one demographic mix faces another after housing, labour, or tourism patterns change. Aggregate accuracy can remain stable while error concentrates on a newly exposed group. Lifecycle frameworks warn that risk measured early can differ from risk measured later ([NIST, 2023](#)). Sunsets compel someone to look.

Automation of the exception. Overrides dry up because refusing the model is slower than accepting it. Dubai's guidelines already envision challenge rights for significant automated decisions ([Smart Dubai, 2019](#)). Renewal dossiers should report challenge and override rates; near-zero overrides are not proof of perfection. They may be proof of intimidation by interface.

Vendor opacity. Critical logic sits behind commercial confidentiality. A sunset cannot invent explainability that does not exist, but it can refuse renewal when the government cannot obtain enough intelligibility to exercise oversight. That is a buyer's power worth scheduling.

Why this fits Dubai and the wider UAE agenda

Dubai's ethics guidance already states that humanity should retain final decision power, with AI in an assisting role, and that accountability for outcomes should not lie with the system ([Smart Dubai, 2019](#)). The UAE's National AI Strategy treats governance as a strategic objective, not a brake ([UAE Cabinet, 2019](#)). Sunsets are how a fast adopter stays a deliber-

ate adopter. They protect innovation by making continuation earned, which is more credible than promising permanence at launch.

They also travel well regionally. Saudi Arabia's AI ethics framework is organised across the system lifecycle and stresses continuous monitoring within that lifecycle ([SDAIA, 2023](#)). Qatar's national AI strategy flags ethics and public policy as a pillar and calls for governance frameworks suited to sensitive domains ([MCIT, 2019](#)). Expiry dates are a concrete mechanism those frameworks can share without waiting for identical statutes.

Compare three imperfect controls. A one-time ethics self-assessment is necessary and insufficient: it freezes judgment at launch. Continuous monitoring is necessary and still insufficient if nobody must act on the red indicators. A sunset is the scheduled obligation to act—to renew, redesign, or retire. In administrative life, scheduled obligations outperform ambient virtues.

There is a regional narrative advantage as well. Gulf strategies often present AI adoption and responsible AI as joint ambitions rather than as rivals ([UAE National Strategy for Artificial Intelligence 2031](#); [MCIT, 2019](#)). Sunsets make that joint ambition legible to the public: we adopt quickly, and we re-consent openly. That is a more adult story than either technoptimism without brakes or precaution without deployment.

Policymakers sometimes fear that publishing renewals will expose weakness. The opposite is likelier. Hidden permanence teaches cynicism; visible renewal teaches that government still chooses. Where security or privacy forbids detail, publish the fact of renewal or retirement and the length of the new term. Opacity about existence is harder to justify than opacity about sensitive metrics.

Anticipating the objections

Cost. Reviews cost money; unreviewed failures cost trust. In services that allocate benefits or burdens, trust is part of infrastructure.

Uncertainty for vendors. Good. Vendors that cannot defend renewal on evidence should not hold permanent public advisory power. Contracts can price renewal reviews in advance.

Emergency need. Sunset regimes can allow short, logged emergency extensions—thirty to ninety days—when switching off would create immediate harm, provided a full renewal or retirement decision follows. Emergency must remain the exception that proves the clock, not the loophole that abolishes it.

Overlap with existing audits. Audits ask whether rules were followed. Sunsets ask whether the authorisation should continue. A system can be compliant and obsolete at the same time.

A concrete next step

Dubai does not need a perfect federal code to begin. A pilot rule for a defined class of significant municipal or emirate AI systems—eligibility ranking, inspection targeting, or case triage, for example—could require an expiry date, a renewal dossier template, and publication of renewals. The Digital Dubai ethics self-assessment culture already habituates teams to structured questioning before and during operation ([Digital Dubai self-assessment](#)). Add a calendar and a refusal right at renewal, and the culture becomes a control.

Laws expire when democracies and constitutional systems fear their own unfinished knowledge. Government algorithms deserve the same humility. If a system is good for the public, renewal will be easy to justify. If renewal is hard to justify, expiry is not a failure of innovation. It is innovation remembering that public authority is borrowed, and borrowed power should come with a return date.

For a policy journal read by practitioners, the operational ask is small and the constitutional instinct is large. Small: put dates on significant systems. Large: refuse the fantasy that machine advice, once installed, is as permanent as geography. Dubai has spent years building a reputation for iterative government. Algorithmic sunsets are iteration with teeth.

If the pilot works, scale the rule. If it does not—if renewals become rubber stamps—the failure will at least be visible, which is more than can be said for permanent systems that never face a return date. Visibility is the beginning of governing AI rather than merely installing it.

References

- European Union. (2024). Regulation (EU) 2024/1689 (Artificial Intelligence Act). <https://artificialintelligenceact.eu/article/112/>
- Ministry of Communications and Information Technology (MCIT), State of Qatar. (2019). National artificial intelligence strategy for Qatar 2019. https://www.mcit.gov.qa/-/media/mcit/documents/strategies/national_artificial_intelligence_strategy_for_qatar_2019_en.pdf
- National Institute of Standards and Technology. (2023). Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1). <https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf>
- Organisation for Economic Co-operation and Development. (2024). AI principles. <https://www.oecd.org/en/topics/ai-principles.html>
- Ranchordás, S. (2014). Constitutional sunsets and experimental legislation: A comparative perspective. Edward Elgar. <https://www.e-elgar.com/shop/usd/constitutional-sunsets-and-experimental-legislation-9781783478941.html>
- Saudi Data and Artificial Intelligence Authority. (2023). AI ethics principles (Version 1.0). <https://dgp.sdaia.gov.sa/wps/wcm/connect/3fceaa0d-bbb6-4ead-9cab-c141c2eca42d/ai-principles-EN.pdf?MOD=AJPERES>
- Smart Dubai. (2019). AI ethics principles & guidelines. <https://www.digitaldubai.ae/docs/default-source/ai-principles-resources/ai-ethics.pdf>
- UAE Cabinet. (2019, April 21). UAE Cabinet adopts National Artificial Intelligence Strategy 2031. <https://uaecabinet.ae/en/news/uae-cabinet-adopts-national-artificial-intelligence-strategy-2031>
- UAE Government. (2024). The UAE Charter for the Development and Use of Artificial Intelligence. <https://u.ae/en/about-the-uae/strategies-initiatives-and-awards/policies/Ai/The-UAE-Charter-for-the-Development-and-Use-of-Artificial-Intelligence>
- United Arab Emirates. (n.d.). UAE National Strategy for Artificial Intelligence 2031 [PDF via OECD.AI]. https://wp.oecd.ai/app/uploads/2021/12/United_Arab_Emirates_National_Strategy_for_Artificial_Intelligence_2017-2031.pdf
- United Arab Emirates Council for Artificial Intelligence and Blockchain. (n.d.). AI Council. https://ai.gov.ae/ai_council/
- WAM (Emirates News Agency). (2024, July 30). Artificial Intelligence office launches 'UAE Charter for Development and Use of AI'. <https://www.wam.ae/en/article/b4f7n7s-artificial-intelligence-office-launches-%E2%80%98uae>