


Procurement as Regulation: State AI Executive Orders

Five Governors, Three Weeks, and What Vendors Should Expect

2026-10-14

 AI-assisted Rapid Research



CSAI

cloud
security
alliance®

© 2026 Cloud Security Alliance. Some rights reserved.

You may download, store, display, view, print, redistribute, and link to this document in its original, unmodified form, provided that attribution to the Cloud Security Alliance is maintained and all trademark and copyright notices remain intact.

This document may not be modified or altered. You may quote portions of the document as permitted by the Fair Use provisions of the United States Copyright Act, provided that attribution is given to the Cloud Security Alliance.

This document may be shared on professional and social media platforms in its original form with attribution.

This document was generated with AI assistance and has not undergone official CSA review and approval processes.

Key Takeaways

Between September 18 and October 6, 2026, governors in Virginia, California, Illinois, Oregon, and Utah each signed an executive order touching AI governance [1][3][5][6][8]. Only one of the five, Oregon's Executive Order 26-26, is titled as a procurement measure. The others reach procurement indirectly, through third-party verification regimes, advisory bodies tasked with considering contract requirements, or security guardrails for agency use, and one (Utah's) specifies no vendor requirements in its public announcement [8]. Describing the five as a single coordinated "procurement as regulation" movement therefore overstates what the texts say, although the direction of travel is consistent.

Two concepts recur across the orders and are likely to reach vendor contracts first: independent third-party review of AI safety, and an emergency shutdown ("kill switch") capability for frontier models. Neither is defined in the public texts we reviewed. Oregon's order directs its state Chief Information Officer to develop the criteria and to assess whether a kill-switch requirement is viable, with an implementation proposal due within 90 days [1]. California's order commissions expert recommendations by November 16, 2026 [3][4].

For enterprise security and vendor-risk teams, the practical consequence is that state contract language is one plausible early channel through which these ideas become enforceable, because governors can direct agencies without new legislation. Organizations that sell AI products or services to state governments, or that rely on suppliers who do, should begin inventorying the evidence they can already produce on independent assurance, shutdown and containment controls, and incident reporting. Most of the substantive requirements remain unwritten, and the details will be set in agency guidance over the next several months.

Background

State executives have moved ahead of Congress on AI governance in 2025 and 2026, and procurement has been one of the more accessible levers because a governor can direct state agencies without new legislation. California's earlier Executive Order N-5-26, issued in March 2026, set the template. It directed state agencies to develop vendor certification standards requiring companies to show that their AI systems include controls addressing harmful or unlawful content, algorithmic bias, and impacts on civil

rights and civil liberties, and it directed agencies to embed disclosure and risk-management expectations into contracting [7]. A law-firm analysis also noted a provision allowing California to assess federal supply-chain risk determinations independently [7].

The September orders build on that base, but they emphasize frontier-model safety and independent oversight over the civil-rights focus of the March order. They also arrived against a backdrop of new state statutes. California enacted SB 813, which creates a framework for independent verification organizations, and AB 1405, which creates a state registry of AI auditors with standards for independence, transparency, and integrity [2][3]. Press coverage describes New York's Responsible AI Safety and Education Act, passed in 2025, as requiring developer registration by November 2026 and, from January 2027, published safety protocols and critical-incident reporting within 72 hours [2]. We could not confirm these dates against the enacted text. Illinois's order also notes that the state already has an AI Safety Measures Act with third-party audits and 72-hour incident reporting [6]. Executive orders in other states therefore sit alongside a developing body of statutory requirements, and vendors will often face both.

Timing and attribution deserve a note. The Oregon signing date appears as September 23 in the Governor's newsroom release [1]. GovTech's article is dated September 25, but its text places the signing on the Wednesday of that week, which is also September 23 [2]. We use September 23, and the apparent discrepancy appears to reflect the article's publication date rather than a conflict over the signing. Utah's release gives an October 6 signing date, although its page header displays a different date, and the announcement does not state an order number [8]. Readers needing precise dates for contractual purposes should consult the signed orders directly.

Security Analysis

What each order actually does

The five orders differ substantially in scope and in how closely they bear on vendors. The table below summarizes the public texts and announcements. It reflects our reading of those sources, not the full signed instruments in every case.

State / Order	Date	Procurement relevance	Key mechanisms and deadlines
Virginia, EO 22	Sept 18,	Indirect. Primarily a data center accountability	Creates a Rapid-Response AI Task Force covering workforce, privacy, and

State / Order	Date	Procurement relevance	Key mechanisms and deadlines
	2026	order; its listed procurement item concerns local and sustainable procurement in workforce criteria [5]	cybersecurity risks from AI agents; bans state NDAs on data center projects; rescinds two earlier directives [5]
California, EO N-9-26	Sept 18, 2026	Indirect. Accelerates the SB 813 and AB 1405 assurance regime and references the March procurement order [3][4]	Expert recommendations due Nov 16, 2026; the kill-switch item is a feasibility evaluation, not a mandate [3][4]; later certification framework (May 1, 2027) and auditor registry (Dec 1, 2027) milestones appear in secondary summaries and were not confirmed against the primary text
Illinois, EO 2026-07	Sept 22, 2026	Evaluative. The AI Cabinet is to assess whether AI safety requirements can be built into state contracts [6]	Cabinet appointed within 30 days; sunsets Dec 31, 2027; duties include incident response procedures and risk-management recommendations to agencies [6]
Oregon, EO 26-26	Sept 23, 2026	Direct. Titled as AI procurement standards for state government [1]	State CIO to develop third-party safety review criteria and assess kill-switch viability for frontier models; proposal in 90 days; quarterly reassessment [1][2]
Utah	Oct 6, 2026	Not specified in the public announcement [8]	AI literacy training, guardrails against "rogue AI agents," a state-endorsed digital identity framework, human review of rights-affecting determinations; annual reporting from July 1, 2027 [8]

Third-party review becomes a procurement gate

A prominent pattern is the pairing of independent verification with purchasing authority. California's statutes establish who may verify AI safety claims and how auditors are registered, while Oregon's order points at the contracting stage, where a state can decline to buy frontier AI that has not undergone adequate review [1][2][3]. If these instruments mature as written, a vendor's ability to show a recognized independent assessment could become a condition of sale, much as a FedRAMP authorization functions for federal cloud purchases today. This is an inference from the direction of the orders and not a stated requirement, since Oregon's release does not specify the review criteria [1].

A second-order effect is fragmentation. Each state is defining its own verifier qualifications and timelines, and the orders do not point to a common standard. A vendor serving several states could face differing criteria for what counts as "adequate" review, which raises the value of evidence that maps to widely used control frameworks and can be reused across audits.

Kill switch requirements are undefined

Both California and Oregon raise an emergency shutdown capability for frontier models, as a matter for evaluation rather than an established mandate, but neither public description specifies its mechanism, who holds the authority to trigger it, or how it would be verified [1][4]. California's order does call for efficacy verification of any such mechanism on an ongoing basis [3]. The ambiguity matters for security teams, because a shutdown control can mean very different things: revoking API access at the provider, disabling a deployed agent within a customer environment, or halting model training or inference at the infrastructure layer. Each has different failure modes and different dependencies on the customer.

For organizations deploying agentic systems, the nearer-term question is whether existing containment controls would satisfy a reasonable reading of such a requirement. Credential revocation, tool-permission scoping, rate limits, and the ability to suspend an agent without data loss are the controls most likely to be examined. Teams that cannot currently demonstrate those capabilities should expect questions during state procurement reviews, whether or not the final language uses the term "kill switch."

Incident definitions and reporting are expanding

California's order directs an expansion of critical-incident definitions to include loss-of-control scenarios [3]. Illinois tasks its AI Cabinet with incident response procedures and with analyzing and sharing findings on emerging AI incidents [6]. Virginia's task force is charged with assessing cybersecurity risks posed by AI agents to government and critical infrastructure [5]. Taken together, these provisions suggest that vendors selling to state governments may eventually be asked to report AI-specific incidents, including

agent misbehavior that would not meet traditional breach definitions. The orders do not yet establish reporting thresholds or timelines. The only concrete benchmarks in the sources we reviewed are the 72-hour windows attributed to the New York statute [2] and to Illinois's existing AI Safety Measures Act [6].

Identity controls for agents: a signal to watch

Utah's order pairs guardrails against "rogue AI agents" with a state-endorsed digital identity framework based on cryptographically verifiable credentials [8]. We read this as an early signal that state governments may tie agent authorization to verifiable identity. The announcement does not describe technical details, so this should be treated as a direction to watch and not as a requirement vendors must meet today.

Limits of this analysis

We reviewed governor's-office releases, the Illinois order text, and law-firm and press summaries. We did not have access to the full text of every signed order, and several secondary sources could not be retrieved. Statements about deadlines come from those summaries. Where primary text and press accounts diverge, the primary text should govern.

Recommendations

Immediate Actions

Security and vendor-management teams should identify which of their AI products or suppliers fall within state government sales channels, and in which of the five states. A short register of those relationships, with contract renewal dates, lets teams align their preparation with the 90-day Oregon and November 16 California milestones. Teams should also confirm who inside the organization is monitoring state agency guidance, since procurement requirements will likely appear first in solicitations and contract amendments instead of in legislation.

Organizations should also document their current containment capabilities for deployed AI systems, including the ability to revoke credentials, suspend agents, and disable specific tools. Capturing this evidence now, in a form a customer or auditor could review, is likely to be inexpensive relative to responding under deadline, and should shorten responses if state buyers request it.

Short-Term Mitigations

Over the next quarter, vendors should map their AI assurance evidence to a recognized control framework so that it can be reused across state-specific requests. Providers should consider engaging independent assessors on a pilot basis before criteria are finalized, which gives them early insight into what verifiers will examine. Where contracts with AI suppliers lack incident-notification terms, procurement teams should seek language covering AI-specific events such as loss of agent control, unexpected autonomous actions, and model behavior that bypasses safeguards.

Teams may also be able to offer input. Oregon's CIO proposal, California's expert process, and Illinois's AI Cabinet are each positioned to draw on outside expertise, although the sources do not establish a formal comment process for each, and industry detail on what is technically feasible for shutdown and review mechanisms is likely to shape the outcome [1][3][6].

Strategic Considerations

The longer-term question is whether state requirements converge or diverge. If California's verifier registry becomes a de facto reference, other states may adopt it by reference, which would reduce the compliance burden. If each state defines its own criteria, multi-state vendors will need a modular evidence package. Organizations should plan for both outcomes and avoid building single-state compliance processes.

Enterprises that are not state contractors should still track these orders. State purchasing requirements often become commercial baseline expectations as large buyers and insurers adopt them, and the concepts of independent AI review, shutdown capability, and AI incident reporting are likely to appear in private-sector vendor questionnaires. That outcome is an inference and not a certainty, and we have not cited a specific precedent.

CSA Resource Alignment

CSA's recent regulatory analyses provide relevant prior work. *Executive Order 14409: AI Cybersecurity Deadlines Take Effect* [9] examines the June 2, 2026 federal order, which rests on voluntary public-private collaboration while noting that a voluntary framework can still create de facto obligations through government contracting. The state orders discussed here add a second, state-level layer to that picture, and the same approach of separating binding requirements from voluntary signals applies. *AI-Era*

Compliance: Executive Order and CISA BOD 26-04 [10] addresses compliance with the federal order and CISA BOD 26-04, and its treatment of how government policy shapes enterprise obligations is a useful parallel to the way state procurement rules may reach vendors.

For concrete control mapping, the AI Controls Matrix v1.1 [11] provides a vendor-neutral set of 247 controls across 18 domains, aligned with ISO 42001 and ISO 27001, and is a reasonable foundation for the reusable evidence package recommended above. Its guidance for different roles in the AI supply chain is relevant to the question of who bears responsibility for shutdown and containment capabilities between model providers and deploying organizations.

References

- [1] Oregon Office of the Governor. "[Governor Kotek Issues Executive Order to Advance AI Safety and Oversight](#)." State of Oregon, September 23, 2026.
- [2] GovTech. "[Oregon Is Latest State to Move on Frontier AI Regulation](#)." GovTech, September 25, 2026.
- [3] Office of Governor Gavin Newsom. "[Governor Newsom Issues Executive Order to Accelerate Independent Oversight and Advance the Creation of an AI Kill Switch](#)." State of California, September 18, 2026.
- [4] PPC Land. "[Newsom Sets November 16 Deadline to Study Frontier AI Kill Switch](#)." PPC Land, September 2026.
- [5] Greenberg Traurig. "[Virginia Executive Order 22 Sets New Requirements for Data Center Development and AI Governance](#)." Greenberg Traurig, September 2026.
- [6] State of Illinois. "[Executive Order 2026-07](#)." Office of the Governor, September 22, 2026.
- [7] Morgan Lewis. "[California Executive Order Expands AI Oversight Through State Procurement](#)." Morgan Lewis, April 2026.
- [8] Office of Governor Spencer J. Cox. "[Gov. Cox Signs Executive Order Establishing Utah's Pro-Human Approach to AI in State Government](#)." State of Utah, October 2026.
- [9] Cloud Security Alliance. "[Executive Order 14409: AI Cybersecurity Deadlines Take Effect](#)." CSA, July 2026.
- [10] Cloud Security Alliance. "[AI-Era Compliance: Executive Order and CISA BOD 26-04](#)." CSA, June 2026.
- [11] Cloud Security Alliance. "[AI Controls Matrix v1.1](#)." CSA, June 2026.