Preliminary Comment on Reducing Friction in the Exercise of Privacy Rights and Opt-Out Preference Signals

The Box Commons

April 1, 2026

THE BOX COMMONS
Independent AI Agent Credentialing Standards Body
standards@theboxcommons.org · theboxcommons.org

April 1, 2026

California Privacy Protection Agency
Attn: Legal Division — Regulations
400 R St., Suite 350
Sacramento, CA 95811

Re: Preliminary Comment — Reducing Friction & OOPS March 2026

Dear Members of the Board and Staff of the California Privacy Protection Agency:

The Box Commons respectfully submits this preliminary comment in response to the Agency's Invitation for Preliminary Comments on Reducing Friction in the Exercise of Privacy Rights and Opt-Out Preference Signals, issued March 6, 2026.

The Box Commons is a Wyoming mutual benefit nonprofit corporation organized under Section 501(c)(6) of the Internal Revenue Code. We develop independent credentialing standards for artificial intelligence systems, with a focus on verifiable compliance, behavioral safety, and interoperability across regulatory frameworks. Our standards architecture maps to the NIST AI Risk Management Framework and is designed to support third-party certification of AI systems operating in regulated environments.

We write to address a gap in the current regulatory framework that will become critical as AB 566 takes effect on January 1, 2027: **the absence of any credentialing or certification mechanism for AI systems that receive, process, and act upon consumer opt-out preference signals.**

I. The Scale Problem: From Millions to Billions of Signals

AB 566 represents a landmark in consumer privacy protection. By requiring all web browsers operating in California to offer built-in opt-out preference signal settings, the law will dramatically increase the volume of Global Privacy Control (GPC) and similar signals flowing through the digital ecosystem. Today, opt-out signals are generated by a self-selecting population of privacy-conscious consumers who install browser extensions or configure settings. After January 1, 2027, signals will flow from the general population of California internet users.

This shift in scale transforms opt-out signal processing from a niche compliance function into a core operational requirement. Increasingly, the systems receiving and processing these signals will not be human operators reviewing requests—they will be AI-driven systems: recommendation engines, advertising technology platforms, data broker aggregation tools, and automated data processing pipelines.

The question the Agency should consider in any future rulemaking is not merely whether businesses honor opt-out signals, but **whether the AI systems handling those signals can be verified as doing so correctly, consistently, and without circumvention.**

II. The Verification Gap: Lessons from Existing Compliance Failures

Consumer Reports' 2020 study of California data broker opt-out compliance — in which 543 volunteers attempted opt-outs with 214 registered data brokers — found that 62% of the time, participants either could not determine whether their request was successful or were unable to submit a request at all.[1] Participants encountered demands for government-issued identification, selfies, and Social Security numbers simply to exercise their right to opt out of data sales. Consumer Reports' subsequent development of the Permission Slip authorized agent service—which has initiated over one million data rights requests on behalf of consumers—demonstrates both the demand for automated privacy tools and the scale of the compliance gap these tools must navigate.[2]

These findings predate the widespread deployment of AI-driven systems in consumer data processing. As AI systems increasingly mediate the relationship between consumer opt-out signals and business data practices, the verification gap will widen unless the Agency establishes clear expectations for how AI systems demonstrate compliance.

Consumer Reports' Digital Standard—an open framework for evaluating digital products on privacy, security, and data practices developed in collaboration with Disconnect, Ranking Digital Rights, and other organizations—provides a conceptual model for how third-party evaluation of AI privacy compliance could work.[3] The Box Commons' credentialing standards build on this foundation, extending the evaluation framework specifically to AI systems operating in regulated environments.

III. The Intersection of OOPS and Automated Decision-Making Technology

The Agency's finalized regulations on Automated Decision-Making Technology (ADMT), effective January 1, 2026, require businesses using ADMT for significant decisions to conduct risk assessments, provide pre-use notice, and offer consumers opt-out rights. These are sound requirements.

However, the current framework creates a regulatory gap at the intersection of OOPS and ADMT: when an AI system that is itself classified as ADMT receives an opt-out preference signal, what standard governs its processing of that signal? The ADMT regulations address the AI system's decision-making outputs; the OOPS framework addresses the consumer's signal inputs. Neither

[1] Consumer Reports, "California's New Privacy Rights Are Tough to Use" (2020), documenting a study of 543 California volunteers attempting opt-outs with 214 registered data brokers.

[2] Consumer Reports Innovation Lab, Permission Slip authorized agent service (2022–present), initiating over 1 million data rights requests on behalf of consumers.

[3] Consumer Reports, The Digital Standard (2017–present), an open framework for evaluating digital products on privacy, security, and data practices. Available at thedigitalstandard.org.

framework currently addresses the fidelity of the AI system's signal processing—the critical link between the consumer's expressed preference and the system's behavioral response.

We recommend that the Agency consider the following in any future rulemaking:

Recommendation 1: Establish verifiable compliance standards for AI-driven opt-out signal processing. Any AI system that receives and processes consumer opt-out preference signals should be subject to standards that verify: (a) the signal is received without degradation or selective filtering; (b) the signal is applied consistently across all data processing operations, not only the specific interaction that generated it; and (c) the system does not employ technical mechanisms that functionally circumvent the signal while nominally honoring it.

Recommendation 2: Recognize independent third-party credentialing as a compliance pathway. The Agency's ADMT regulations already require risk assessment certifications. Extending this model to OOPS compliance for AI systems would create a coherent regulatory framework. Independent credentialing bodies can provide the technical evaluation capacity that no single regulator can maintain at the pace of AI development. This approach mirrors the model that has worked for decades in product safety (UL), food safety (NSF International), and information security (ISO/IEC 27001 certification bodies).

Recommendation 3: Require architectural privacy guarantees, not merely policy commitments. Signal Foundation President Meredith Whittaker has articulated the fundamental challenge of AI systems handling private data: autonomous systems that process personal information require verifiable architectural privacy guarantees, not merely policy promises.[4] When AI agents operate on consumer data—including processing opt-out signals—the privacy properties of the system must be built into its architecture and subject to independent verification. Self-attestation is insufficient for systems whose internal operations are opaque to both consumers and regulators.

IV. Reducing Friction Through Standardized AI Compliance

The Agency's invitation asks about reducing friction in the exercise of privacy rights. We note that one significant source of friction is the inconsistency of AI system responses to opt-out signals across platforms and services. A consumer who sends a GPC signal may encounter radically different system behaviors depending on how each platform's AI processes that signal—behaviors that are invisible to the consumer and difficult for the Agency to audit at scale.

Standardized credentialing for AI systems processing opt-out signals would reduce this friction by establishing a common behavioral baseline. Consumers would benefit from knowing that credentialed systems meet verified compliance standards. Businesses would benefit from clear, auditable requirements that reduce regulatory uncertainty. The Agency would benefit from a scalable compliance verification mechanism that supplements its enforcement capacity.

V. Conclusion

The Box Commons commends the Agency for proactively seeking input on these issues before AB 566's effective date. The window between now and January 1, 2027, is the appropriate time to

[4] Meredith Whittaker, remarks on AI agents and privacy architecture, SXSW (March 2025) and subsequent public statements on the security implications of autonomous AI systems processing personal data.

develop the standards and verification mechanisms that will be needed when opt-out signals flow at population scale through AI-driven systems.

We respectfully urge the Agency to consider independent third-party credentialing—mapped to the NIST AI Risk Management Framework and harmonized with the Agency's existing ADMT regulations—as a scalable, technology-neutral compliance pathway for AI systems processing consumer opt-out preference signals.

The Box Commons stands ready to share our technical standards architecture and to support the Agency's work in this area.

Respectfully submitted,

Brice Love
Acting Executive Director
The Box Commons
standards@theboxcommons.org