Openpolicy

May 28, 2025

California Privacy Protection Agency
Attn: Legal Division – Regulations Public Comment
Via: regulations@cppa.ca.gov

OpenPolicy's Public Comment on CCPA Updates, Cyber, Risk, ADMT, and Insurance Regulations

OpenPolicy appreciates the opportunity to submit these comments on the California Privacy Protection Agency's May 2025 revisions to the proposed CCPA regulations. OpenPolicy is a technology policy organization dedicated to democratizing access to policymaking for innovators and startups. We leverage data-driven insights to help companies engage with the government on critical cybersecurity and privacy issues. We are deeply committed to collaborative policymaking that strikes a balance between robust security and continued innovation. In this spirit, we commend the CPPA's diligent work on these regulations and the thoughtful incorporation of public feedback to refine the rules.

We applaud the CPPA for retaining a strong overall framework for consumer privacy and security in the revised regulations. Notably, the May 2025 draft preserves important **security measures such as multi-factor authentication (MFA) requirements and third-party identity verification services** within its provisions. By maintaining mechanisms such as phishing-resistant MFA and independent identity verification, the CPPA demonstrates a commitment to protecting consumer data through verified access controls and advanced authentication, which are critical defenses against fraud and unauthorized access.

At the same time, we note with concern that **certain key cybersecurity provisions were removed or narrowed** in the latest draft, presumably to reduce compliance burdens. In particular, the explicit reference to "zero trust architecture" was eliminated from the cybersecurity audit criteria. Zero Trust – the principle of granting the minimum necessary access and continually verifying identity and context – is widely recognized as a cornerstone of modern cybersecurity (indeed, U.S. federal agencies are required to meet specific Zero Trust objectives by FY 2024). While we understand the desire to streamline the rules, we believe the **underlying goals** of these removed provisions can still be achieved within the current regulatory text. Our comments, therefore, focus on **building upon what is still present**, recommending pragmatic enhancements and clarifications that reintroduce robust security practices in a manner compatible with the revised draft.

In the sections below, we provide feedback on **Article 5 (Verification of Requests)**, **Article 9 (Cybersecurity Audits)**, and **Article 10 (Risk Assessments)**. For each, we highlight retained elements that merit support and suggest improvements to reinforce cybersecurity safeguards without rehashing now-deleted language. Our tone is collaborative and forward-looking; we recognize the CPPA's efforts and offer constructive ideas to strengthen the regulations' security posture. We conclude by emphasizing OpenPolicy's readiness to assist in developing risk-based, future-proof solutions and by inviting continued engagement with the CPPA on these important issues.

Article 5 – Verification of Consumer Requests

We are pleased to see that Article 5 continues to prioritize robust identity verification for consumer rights requests. The revised regulations maintain requirements for businesses to verify that a person making a request to delete, correct, or know personal information is the consumer about whom the information was collected (Section 7060(a)), and explicitly permit the use of third-party identity verification services as a means to accomplish this (Section 7060(c)(1)). Allowing reputable third-party verification services, so long as they meet CCPA standards, gives businesses flexibility to employ sophisticated tools for confirming identity, which can improve accuracy and reduce fraud. We also commend the rule that verification processes must scale in stringency with the sensitivity of the data in question (Section 7060(c)(3)(A)–(D)); this risk-based approach is essential to prevent unauthorized deletions or disclosures of highly sensitive personal information.

Moreover, we appreciate the ban on consumer-paid verification fees and onerous procedures. The regulations rightly prohibit businesses from charging consumers or forcing notarization as a condition of verification (Section 7060(e)), except in cases where reimbursement is required. This protects consumers from unnecessary barriers when exercising their rights. The rules also direct businesses to implement reasonable security measures to detect fraudulent verification activity (Section 7060(f)), a critical safeguard against bad actors attempting to exploit the privacy request process.

To further strengthen Article 5, OpenPolicy suggests the CPPA encourage or **clarify the use of multi-factor authentication (MFA) and modern cryptographic verification methods** in the verification process. While the regulations appropriately stop short of mandating any particular method, they define "multi-factor authentication" in Section 7001 and implicitly recognize its value. We encourage the adoption of advanced identity-proofing technologies that enhance security without increasing consumer burden. The regulations already allow the use of third-party services; the CPPA might consider clarifying that such services may employ innovative techniques like **cryptographic proofs or zero-knowledge proofs to verify identity attributes**. For instance, a service could cryptographically confirm that a

consumer's government ID is valid and matches their selfie, without retaining the ID image or exposing unnecessary data, thereby preserving privacy while authenticating identity. By validating credentials or attributes in a privacy-preserving manner (e.g., confirming "age over 18" or residence in California via zero-knowledge proof), businesses can reduce the collection of sensitive data during verification, aligning with the mandate in Section 7060(c)(2) to avoid collecting sensitive personal information unless needed. We believe the CPPA could highlight these emerging solutions in commentary or future guidance, signaling that the use of privacy-enhancing verification methods is encouraged so long as they meet the regulation's standards.

Additionally, we recommend explicitly addressing identity verification challenges associated with **AI agents and Non-Human Identities (NHIs)**, which include automated bots, service accounts, and API keys increasingly used to manage or process consumer information requests. Machine identities, particularly those embedded in automated AI workflows, often maintain persistent and elevated privileges, making them prime targets for attackers. These identities, if compromised, can significantly undermine identity verification processes. To mitigate these risks, we encourage the CPPA to clarify that identity verification requirements apply equally to both machine identities and human identities. Specifically, businesses should adopt **continuous, dynamic re-authorization methods that verify each API interaction or automated request in real-time**. Such an approach aligns well with the existing principle of scaling verification stringency based on risk, ensuring that AI-driven or automated identity interactions are continuously authenticated, and reducing the window of opportunity for unauthorized or fraudulent activity.

Further, we suggest reinforcing that fraud detection measures (Section 7060(f)) should be **dynamic and adaptive**. This recommendation is especially pertinent given the increasing sophistication of AI-driven attacks, where adversaries use AI to mimic legitimate consumer behavior, manipulate verification processes, or automate reconnaissance of verification vulnerabilities. Businesses should employ **dynamic risk scoring techniques** to assess contextual factors, such as geographic anomalies, behavioral patterns, or unusual request volumes, and automatically escalate verification stringency when risks are detected. Encouraging a **proactive risk-based stance** will ensure the verification framework remains resilient to evolving threats, including sophisticated AI-driven attacks.

Article 5's verification provisions are well-crafted to balance accessibility for consumers with strong security against imposters, including automated threats. OpenPolicy supports these measures and urges the CPPA to further emphasize advanced identity verification techniques, including MFA, cryptographic methods, privacy-preserving identity proofs, and dynamic re-authorization of machine identities, as best practices under the rule. By integrating these forward-looking, adaptive security strategies, the CPPA will future-proof

the regulations, enhancing trust and robustness in consumer rights requests while proactively addressing emerging risks associated with AI and machine identity exploitation.

Article 9 – Cybersecurity Audits

OpenPolicy commends CPPA for retaining critical cybersecurity practices within Article 9, notably Multi-Factor Authentication (MFA), encryption of personal information at rest and in transit (Section 7123(b)(2)(B)), rigorous account management and access controls (Section 7123(b)(2)(D)), and mandatory security training and awareness (Section 7123(b)(2)(M)). These foundational practices set a strong baseline for organizational cybersecurity accountability. However, to further enhance resilience against increasingly sophisticated cyber threats, particularly those amplified by AI integration, we recommend embedding advanced cybersecurity principles and practices into Article 9.

Our primary concern in Article 9 is the removal of the explicit **"zero trust architecture"** provision from the list of security program components. In the initial draft, Section 7123(b)(2)(C) had called for businesses to implement a zero trust architecture (described as ensuring internal connections are encrypted and authenticated). The May 2025 modified text deletes this item. We understand that this change was intended to ease prescriptive burdens; however, we believe the *principles* of Zero Trust are too important to be lost. As noted, Zero Trust has become a foundational strategy in cybersecurity, moving beyond perimeter-based defenses to **assume no implicit trust** and constantly enforce least-privilege access. The White House's federal Zero Trust strategy emphasizes that incremental improvements are not enough against modern threats and mandates a "dramatic paradigm shift" toward continuous verification of each user, device, and transaction[1].

Rather than reinsert the exact "Zero Trust" language, we recommend that the CPPA incorporate the *spirit* of zero trust into the remaining provisions on access control. For example, Section 7123(b)(2)(D) already requires granular account privilege restrictions; this could be augmented with a comment that businesses should **continuously verify user access** and network integrity, and not rely solely on network location or single authentication events. In practice, this means encouraging measures like: **dynamic risk-based authentication** (re-authenticating or challenging users when context changes or anomalies are detected), **attribute-based access control (ABAC)** policies that evaluate a user's role, device security, location, and other attributes before granting access, and network segmentation such that being "inside" the network grants no blanket trust.

[1] See Office of Management and Budget's (OMB) Memorandum M-22-09 https://www.whitehouse.gov/wp-content/uploads/2022/01/M-22-09.pdf

Additionally, businesses utilizing **hardened virtual appliances with tiered component positioning, assume-breach architecture, and internal service isolation**—where each service treats communications from other services as untrusted and communicates through cryptographically secure channels—should be explicitly recognized as meeting **zero-trust architecture standards**. Implementations that incorporate embedded security controls, sandboxed open-source libraries, and customer-owned encryption keys inherently satisfy zero-trust requirements through a comprehensive architectural design.

We suggest clarifying that **"restricting access to what is necessary"** includes ongoing monitoring of access sessions and automatic blocking of unauthorized lateral movement. By embedding these concepts, the regulation would still promote a Zero Trust mindset (continuous verification, least privilege by default) without necessarily using that exact term. This approach imposes minimal new burden—it clarifies how to implement existing listed controls rather than adding new ones—but importantly signals to businesses that simply having perimeter defenses is insufficient; they must actively reinforce internal defenses as well.

Likewise, in Section 7123(b)(2)(I) on network monitoring and defenses, we support the requirement for intrusion detection/prevention systems and would encourage the inclusion of **“real-time, continuous monitoring”** as an objective. Businesses should utilize modern security information and event management (SIEM) tools or extended detection and response (XDR) systems to audit their network and system logs for suspicious activity continuously. In today’s threat environment, real-time visibility is crucial – waiting for a periodic check could miss fast-moving breaches. NIST's various guidelines on information security continuous monitoring highlight the value of automated tools that can audit system configurations and controls on an ongoing basis. We recommend that the CPPA emphasize that **continuous security monitoring** is a best practice that complements the annual audit. For instance, an **audit report could note how the company uses automated monitoring to maintain compliance between audits**, which would demonstrate proactive risk management. By encouraging continuous audit automation, the regulations can drive organizations toward more resilient, always-on security oversight, thereby catching issues early and reducing the likelihood of large breaches.

Recognizing the growing reliance on AI agents and automated systems, we recommend that Article 9 mandate governance of non-human identities, such as AI agents, API keys, automated bots, and service accounts. Each non-human identity should undergo continuous re-validation and context-aware risk assessments.

In modern IT environments, not only human users but also machines, applications, and AI agents often have credentials and access to data. These **non-human identities** (such as AI

agents, API keys, service accounts, robotic process automation bots, and AI algorithms operating autonomously) can be targets for attackers if not properly managed. We suggest that the CPPA clarify, under Section 7123(b)(2)(D) (Account Management), that businesses should include governance of **machine and service accounts** in their access controls. Each such account should have an identified owner, minimal privileges, and be rotated or revoked when no longer needed, similar to employee accounts. Additionally, as AI agents (such as autonomous software using personal data to make decisions) become more prevalent, companies should ensure that these agents are subject to the same access restrictions and monitoring as human users. For example, if an AI process is retrieving consumer data from a database, its access should be strictly scoped to the necessary fields and logged for audit purposes. We believe explicitly acknowledging **NHI security** in the audit criteria will future-proof the regulations, encouraging businesses to extend their identity and access management practices to all entities that can interact with personal information, whether human or not.

Further, Section 7123(b)(2)(B) appropriately requires encryption of personal information. We recommend that the CPPA encourage businesses to assess their cryptographic algorithms and **prepare for emerging threats such as quantum computing**. Quantum computers in the near future could potentially break widely used encryption (like RSA/ECDSA). Leading experts at NSA, NIST, and CISA have warned that actors might harvest encrypted data now to decrypt later when quantum capabilities arise, and they urge organizations to start planning for **post-quantum cryptography** transitions today[2]. In the context of CCPA audits, this means companies should inventory where they use long-lived encryption and follow NIST's work on quantum-resistant cryptographic standards.[3] We suggest adding to the encryption requirement that businesses **use strong, modern encryption algorithms and have a migration plan for quantum-resistant encryption**, aligning with already existing best practices in the field.

Additionally, recognizing businesses that employ file and disk **double encryption**—encrypting personal information at the application level with separate file-level keys and again at the operating system level for disk storage—should satisfy enhanced encryption standards. Similarly, businesses using TLS 1.3, AES-256 encryption, and FIPS 140-3 validated encryption standards should **meet advanced encryption-in-transit requirements.** These recommendations would modify CCPA regulations to acknowledge that comprehensive security platforms with **integrated encryption approaches provide superior consumer protection**, as opposed to requiring

[2] See Post-Quantum Cryptography: CISA, NIST, and NSA Recommend How to Prepare Now https://shorturl.at/YFa9C

[3]See NIST Post-Quantum Encryption Standards https://shorturl.at/Y9CNt

multiple separate and potentially less secure encryption implementations. Adopting these forward-looking measures ensures the long-term security of Californians' personal data, aligning with global best practices and reducing the risk and expense associated with future urgent encryption updates.

Regarding vendor and AI supply chain risk management, we are pleased to see the regulations (in the prior draft's Section 7123(b)(2)(O)) emphasize oversight of service providers, contractors, and third parties for CCPA compliance. We recommend extending this concept to explicitly include **security assessments of vendors**, particularly those handling personal information or providing critical technologies, such as AI systems. Many businesses rely on third-party software or AI models (for example, a fraud detection algorithm or a cloud analytics tool) that process consumer data. These supply chain elements can introduce vulnerabilities, as seen in incidents where compromised software updates or AI biases caused harm. The audit rule should encourage companies to **inventory their critical vendors and evaluate each vendor's security practices and reliability**. This could involve requiring vendors to complete security questionnaires, adhere to the business's cybersecurity standards, or obtain certifications. In particular, if a company uses third-party **AI or automated decision systems**, it should ensure that those systems are secure (free of malware, properly handling data) and that the vendor has controls in place to prevent unauthorized access to the shared data. By incorporating vendor risk management into the audit scope, the CPPA will help close a potential blind spot and ensure that outsourcing does not become a weak link in the protection of privacy.

Also, continuous monitoring should become a requirement to complement annual audits, particularly emphasized in Section 7123(b)(2)(I). Businesses must employ modern Security Information and Event Management (SIEM) systems or Extended Detection and Response (XDR) platforms to audit and monitor network and system logs continuously. Automated control testing, centralized log management, and proactive risk management strategies significantly enhance an organization's ability to detect, respond to, and mitigate fast-moving cyber threats.

Article 10 – Risk Assessments

Conducting risk assessments is a proactive approach that compels businesses to **consider the potential harms to consumers** before and during high-risk processing. We commend the CPPA for retaining this requirement in the revised regulations, albeit with a narrowed scope and reduced procedural burden. By focusing risk assessments on truly significant decisions or sensitive profiling (as refined in the latest draft), the Agency ensures that effort is directed where it matters most – on processing that could seriously impact consumers' rights and freedoms.

Within the current structure, we recommend several steps to enhance the risk assessment process. Privacy and security risks are often intertwined. We suggest that businesses, when conducting a CCPA risk assessment, be explicitly encouraged to consider **cybersecurity risks posed by the processing activity** in addition to privacy impacts. For example, if a company is assessing a new system that uses sensitive personal information (like biometric data) to make automated decisions, the assessment should cover not only potential bias or fairness issues (privacy/ethics concerns) but also the security dimension e.g., could a breach of this system expose biometric identifiers, and what safeguards are in place to prevent that? The CPPA might clarify in Article 10 that a "risk assessment" should evaluate **the likelihood and severity of potential security incidents** associated with the processing, as well as misuse or unauthorized access. **Many modern privacy harms actually originate from security failures** (data breaches, malware, identity theft), so folding **security risk into the assessment will give a more holistic view of consumer risk**. We believe that this is entirely in line with the intent of the law, and it complements the cybersecurity audits (which look broadly at enterprise security) by focusing on the security of specific high-risk processing operations.

Building on our Article 9 comments, if a high-risk processing activity relies on third-party technology or data (for instance, using a third-party AI platform to analyze consumer data), the risk assessment should contemplate risks arising from that dependency. We recommend highlighting that **businesses should evaluate risks from their supply chain** in each relevant assessment – e.g., could the third-party fail to protect the data, or might the third-party's model have hidden biases or security vulnerabilities? Including a section on vendor risk in the risk assessment template will prompt companies to ensure that their partners and service providers do not undermine consumer protection. The CPPA could even reference known standards or frameworks for such evaluations (like requiring that AI systems be subjected to a vendor security review, or checking if vendors adhere to industry security certifications). Thus, we suggest that **if a service provider plays a role in the high-risk processing, its controls must be factored into the risk analysis**.

Furthermore, the CPPA may consider encouraging businesses to utilize **automated or continuous risk assessment tools** as part of their compliance toolkit. Just as continuous monitoring helps in audits, continuous risk scanning can help flag issues between formal assessment cycles. Some organizations are adopting dynamic risk scoring systems that automatically update risk levels when conditions change (for example, if a dataset grows significantly or new threat intelligence emerges about a vulnerability in an AI algorithm). While not every business will have such tools, the CPPA could promote them in guidance, encouraging **real-time detection of changes in processing or in the threat landscape, which can trigger ad-hoc risk assessments** in addition to the required periodic ones. This ensures that risk management is **not a one-time checkbox but an ongoing practice**. It is

heartening that NIST and others have been developing catalogues of AI risks and suggesting continuous risk mitigation processes. Aligning California's approach with these evolving best practices will keep the regulations forward-compatible.

In closing, OpenPolicy **applauds the CPPA** for its leadership in crafting these regulations. This revision demonstrates a responsive approach – maintaining robust consumer protections and cybersecurity expectations, while streamlining areas that posed undue burden or legal uncertainty. **Thank you for your consideration of our comments.** We appreciate the CPPA's hard work and openness to feedback in this rulemaking process. Please do not hesitate to contact us for any clarification or further information. OpenPolicy looks forward to the successful finalization and implementation of the CCPA regulations, and to working with the Agency on promoting a secure, innovative, and privacy-respectful digital ecosystem.

Sincerely,

/s/ Dr. Amit Elazari
Dr. Amit Elazari
CEO and Co-Founder of
OpenPolicy