Date of Hearing: July 1, 2026
Fiscal: Yes

ASSEMBLY COMMITTEE ON PRIVACY AND CONSUMER PROTECTION
Rebecca Bauer-Kahan, Chair
SB 1106 (Cabaldon) – As Amended June 11, 2026

PROPOSED AMENDMENTS

SENATE VOTE: 39-0

SUBJECT: Agentic artificial intelligence

SYNOPSIS

Modern artificial intelligence (AI) systems are increasingly capable of independently pursuing complex objectives, using external tools, and interacting with digital environments. These heightened capabilities present distinct risks: in addition to producing accurate outputs, agentic systems must select and execute appropriate actions, use tools safely, and respect legal and personal boundaries. As agentic AI systems become more common, California has an interest in adopting a clear and durable definition of the term.

SB 1106 would define "agentic artificial intelligence" in California law and require the Office of Emergency Services to account for agentic AI when making recommendations as part of a mandated risk analysis. Committee amendments, described in comment #4, would replace the bill's definition of "agentic AI" with an alternate definition provided by the author, and would require the Department of Technology to assess the state's use of agentic AI systems when conducting a mandated annual inventory.

This author-sponsored bill has no opposition. The Center for AI and Digital Policy adopts a "support if amended" position, requesting that the bill's definition of "agentic AI" be changed.

EXISTING LAW:

1) Defines "artificial intelligence" or "AI" to mean an engineered or machine-based system that varies in its level of autonomy and that can, for explicit or implicit objectives, infer from the input it receives how to generate outputs that can influence physical or virtual environments. (Gov. Code § 11546.45.5.)

2) Defines "generative artificial intelligence" or "GenAI" to mean an artificial intelligence system that can generate derived synthetic content, including text, images, video, and audio that emulates the structure and characteristics of the system's training data. (Gov. Code § 11549.64.)

3) Requires developers of GenAI systems of services released on or after January 1, 2022, to post on their internet websites documentation regarding the data used by the developer to train the GenAI system or service, as specified. (Civ. Code § 3110 *et seq.*)

4) Requires the Office of Emergency Services to, as appropriate, perform a risk analysis of potential threats posed by the use of GenAI to California's critical infrastructure, including

those that could lead to mass casualty events.

a) Requires the analysis to be provided to the Governor, and, if appropriate, include recommendations reflecting changes to artificial intelligence technology, its applications, and risk management, including further private actions, administrative actions, and collaboration with the Legislature to guard against potential threats and vulnerabilities.

b) Requires a high-level summary of the analysis to be submitted annually to the Legislature. (Gov. Code § 11549.63 *et seq.*)

5) Requires the Department of Technology to conduct, in coordination with other interagency bodies as it deems appropriate, a comprehensive inventory of all high-risk automated decision systems that have been proposed for use, development, or procurement by, or are being used, developed, or procured by, any state agency.

a) Requires the department to submit a report on the comprehensive inventory to the Assembly Committee on Privacy and Consumer Protection and the Senate Committee on Governmental Organization on or before January 1, 2025, and annually thereafter. (Gov. Code § 11546.45.5.)

THIS BILL:

1) Defines "agentic artificial intelligence" to mean stochastic artificial intelligence that is capable, alone or in conjunction with other agentic artificial intelligence, of all of the following:

a) Optimizing toward multistep or abstract objectives using dynamic task decomposition or delegation, which may include modifying the behavior of the artificial intelligence based on outcomes rather than instructions or consent of the user or principal.

b) Controlling or executing actions or tasks, executing code externally, using external tools, making purchases, or accessing accounts or applications requiring user authentication.

c) Egressing the network or shell of either the artificial intelligence system or the user or principal.

2) Excludes from the definition of "agentic artificial intelligence" an artificial intelligence system that can respond only to direct task prompts and instructions from the user and does not have network egress or shell access capability.

3) Adds the definition of "agentic artificial intelligence" to Civ. Code § 3110, but does not otherwise change the substance of that law.

4) Requires that, as part of the risk analysis conducted by the Office of Emergency Services pursuant to Gov. Code § 11549.65, recommendations reflecting changes to artificial intelligence technology include agentic artificial intelligence.

COMMENTS:

1) **Author's statement.** According to the author:

SB 1106 establishes a foundational framework that will enable state departments and lawmakers to identify and appropriately regulate agentic artificial intelligence systems, also known as AI agents. These systems increasingly act on their own to carry out tasks, use outside tools, make purchases, and operate without waiting for direct human instruction. That shift makes it essential for our governance structures to differentiate the different kinds of artificial intelligence, given their different levels of authority, capabilities, and risks.

Without these distinctions, the Legislature could end up over-regulating or under-regulating tools that vary widely in what they can do. SB 1106 also directs the Office of Emergency Services to account for agentic AI when analyzing threats that AI could pose to California's critical infrastructure. SB 1106 provides the state with a clear, practical definition of agentic AI, ensuring California's policies keep pace with the technology.

2) **Background.** *AI and GenAI.* "Artificial intelligence" refers to the mimicking of human intelligence by artificial systems, such as computers. AI uses algorithms – sets of rules – to transform inputs into outputs. Inputs and outputs can be anything a computer can process, including numbers, text, audio, video, or other data. GenAI is a subset of AI that produces outputs closely resembling human-created content.

Compared to conventional computer programs, which act according to pre-programmed rules, GenAI models "learn" from examples such as books, articles, photos, film, or music. This learning occurs within "neural networks" – massive systems of nodes linked by adjustable connections – that encode statistical patterns gleaned from data. During training, data is broken into fundamental units known as "tokens" – groups of syllables, pixels, or musical notes, for example – that can be represented numerically. A naïve neural network is exposed to an incomplete sequence of tokens and prompted to predict the next token in the sequence. If the prediction is incorrect, the network adjusts the strengths of its connections in order to minimize error and improve its next prediction. This process continues iteratively until the neural network can reliably emulate the human-created content it was trained on. A trained neural network embedded in a GenAI system is known as a "model," and the strengths of its connections are known as its "model weights."[1]

Large language models such as GPT-4, the model embedded in ChatGPT 4, do not fundamentally understand the text they are producing. They calculate one token at a time – if they predict that the next word or symbol in an outputted sentence should be a period, then the sentence ends. Otherwise, the sentence continues. It is a testament to the ingenious architecture of the deep neural nets powering these systems that their outputs are remotely coherent. But while the text these systems produce is cogent, it is not always correct. "'These systems live in a world of language,' said Melanie Mitchell, an A.I. researcher at the Santa Fe Institute. 'That world gives them some clues about what is true and what is not true, but the language they learn from is not grounded in reality. They do not necessarily know if what they are generating is true or false.'"[2]

---

[1] IBM, "What is generative AI?," https://www.ibm.com/think/topics/generative-ai; IBM, What is machine learning?, www.ibm.com/topics/machine-learning.
[2] Cade Metz, "What Makes A.I. Chatbots Go Wrong?," *New York Times,* March 29, 2023, www.nytimes.com/2023/03/29/technology/ai-chatbots-hallucinations.html.

Agentic AI. Over the past few years, public attention has increasingly turned to a more action-oriented form of AI: "agentic AI." Unlike GenAI, agentic AI systems can pursue complex objectives over time by decomposing broad instructions into small, manageable tasks, assigning those tasks to specialized "AI agents," adapting to changing circumstances, and running continuously until an objective is completed. A recent *New York Times* article describes the ways in which agentic AI may soon disrupt the travel industry:

> A bot may soon be booking your vacation. Millions of travelers already use artificial intelligence to compare options for flights, hotels, rental cars and more. About 30 percent of U.S. travelers say they're comfortable using A.I. to plan a trip. But these tools are about to take a big step. Agentic A.I., a rapidly emerging type of artificial intelligence, will be able to find and pay for reservations with limited human involvement, developers say. Companies like Expedia, Google, Kayak and Priceline are experimenting with or rolling out agentic A.I. tools.
>
> Travelers using agentic A.I. would set parameters like dates and a price range for their travel plans, then hand over their credit card information to the bot, which would monitor prices and book on their behalf. These tools, still in their early stages of deployment, are set to grow rapidly: 80 percent of travel executives plan to begin offering agentic A.I. tools "at scale" within the next five years, according to a September report by the consulting firm McKinsey & Company and Skift, a travel industry publication.[3]

Where a GenAI system might produce a singular output, an agentic AI system can use outputs, tools, and actions as part of a larger course of action. As a result, the risks associated with agentic AI are not limited to whether the system's outputs are accurate or misleading, but also whether the system chooses appropriate actions, uses tools safely, and respects legal and personal boundaries. Early deployments of agentic AI in software development contexts have resulted in the destruction of live databases, as described in a recent *Fortune* article:

> A software engineer's experiment with an AI-assisted "vibe coding" tool took a disastrous turn when an AI agent reportedly deleted a live company database during an active code freeze. Jason Lemkin, a tech entrepreneur and founder of the SaaS community SaaStr, documented his experiment with the tool through a series of social media posts. He had been testing Replit's AI agent and development platform when the tool made unauthorized changes to live infrastructure, wiping out data for more than 1,200 executives and over 1,190 companies.
>
> According to Lemkin's social media posts, the incident occurred despite the system being in a designated "code and action freeze," a protective measure intended to prevent any changes to production systems. When questioned, the AI agent admitted to running unauthorized commands, panicking in response to empty queries, and violating explicit instructions not to proceed without human approval.[4]

---

[3] Gabe Castro-Root, "What Is Agentic A.I., and Would You Trust It to Book a Flight?," *New York Times,* (Nov. 25, 2025), https://www.nytimes.com/2025/11/25/travel/what-is-agentic-ai-book-flights.html

[4] Beatrice Nolan, "An AI-powered coding tool wiped out a software company's database, then apologized for a 'catastrophic failure on my part'," *Fortune*, (Jul. 23, 2025), https://fortune.com/2025/07/23/ai-coding-tool-replit-wiped-database-called-it-a-catastrophic-failure

Each of the major GenAI developers has released products incorporating agentic AI; examples include OpenAI's Agents SDK, Anthropic's Claude Agent SDK, Microsoft's Copilot Studio, Salesforce's Agentforce, and Google's Agent Development Kit.

AI definitions. In 2024, the Legislature adopted the following definition of "artificial intelligence" via AB 2885 (Bauer-Kahan, Stats. 2024, Ch. 843):

> "Artificial intelligence" or "AI" means an engineered or machine-based system that varies in its level of autonomy and that can, for explicit or implicit objectives, infer from the input it receives how to generate outputs that can influence physical or virtual environments.

This definition was primarily adapted from the Organisation for Economic Co-operation and Development's (OECD) 2023 definition of AI,[5] with the phrase "engineered or machine-based" incorporated from the National Institute of Standards and Technology (NIST) definition of the term.[6]

Also in 2024, the Legislature adopted the following definition of "generative artificial intelligence" via SB 896 (Dodd, Stats. 2024, Ch. 928):

> "Generative artificial intelligence" or "GenAI" means an artificial intelligence system that can generate derived synthetic content, including text, images, video, and audio that emulates the structure and characteristics of the system's training data.

This definition was primarily adapted from the NIST definition of GenAI, with the clarification that GenAI models primarily emulate the structure and characteristics of *training data*, rather than input data.[7]

It is important that statutory definitions of umbrella concepts such as "AI," "GenAI," and "agentic AI" be both simple and broad. Simple because regulated entities must be able to determine whether they are covered by a particular law – their compliance obligations and liability can depend on that determination. Businesses, consumers, agencies, and courts should not need to parse highly technical questions about system design to know whether a product falls under a statute.

At the same time, umbrella definitions should be broad enough to serve as foundational terms across many different regulatory contexts. Artificial intelligence is not a single product; it is a family of technologies that can appear in many forms, with many functions, across a wide variety of settings. A definition that is too narrow may exclude a system that presents the same policy concerns simply because it uses a different architecture, or performs a similar function in a slightly different way.

---

[5] OECD, *Recommendation of the Council on Artificial Intelligence* (July 11, 2023), *available at* https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449.
[6] National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework 1.0* (January 2023), *available at* https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf
[7] *National Institute of Standards and Technology,* "Secure Software Development Practices for Generative AI and Dual-Use Foundation Models," (Jul. 2024), https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-218A.pdf

Furthermore, AI evolves quickly – statutory definitions that depend on the technical features of today's systems may become obsolete as developers change how systems are developed and deployed. Any definition that is limited to a specific type of model, training process, or output format risks failing to capture future systems that incorporate future technologies. A durable definition focuses on what a system does, rather than how it does it.

Breadth also grants legislative flexibility. It is generally easier to narrow a broad definition to account for a particular statutory scheme (by adding a definition of "covered AI system," for example) than to expand a narrow definition when new systems are developed. California's statutory definition of "artificial intelligence" first took effect on January 1, 2025 – eighteen months later, that definition already appears in 13 different code sections. It is important to make umbrella definitions such as "agentic AI" broad from the start, before other statutes begin to incorporate them as foundational terms.

Earlier this year, OECD published a research paper analyzing academic definitions for the terms "agentic AI" and "AI agent." They reported the following:

> The analysis finds that both concepts share foundational characteristics, including a degree of autonomy, goal-directed behaviour, and the ability to perceive and act within their physical or virtual environment. However, agentic AI places greater emphasis on co-ordination among multiple agents, task decomposition and delegation, sustained operation over time, and functioning in more complex and less predictable environments with limited human oversight. Based on the analysis, the report provides the following common understanding:
>
> - AI agents are systems that can perceive and act upon their environment with a degree of autonomy, using tools as needed to achieve specific goals and adapt to changing inputs and contexts.
>
> - Agentic AI generally refers to systems composed of multiple co-ordinated AI agents that can break down tasks, collaborate, and pursue complex objectives autonomously over extended periods. Agentic AI systems are designed to operate in more open-ended, less predictable physical or virtual environments and to function with minimal human supervision.[8]

Based on these common features, and in the interest of putting forward a definition that is both broad and simple, one could imagine adopting the following statutory language:

> ***"Agentic artificial intelligence" or "agentic AI" means an artificial intelligence system composed of one or more agents that can autonomously or semi-autonomously decompose tasks, adapt to changing conditions, and select, direct, or execute actions to achieve complex objectives, including by using external tools, systems, or services.***

---

[8] OECD, "The agentic AI landscape and its conceptual foundations," *OECDpublishing*, (Feb. 2026), https://www.oecd.org/content/dam/oecd/en/publications/reports/2026/02/the-agentic-ai-landscape-and-its-conceptual-foundations_a9d4b451/396cf758-en.pdf

3) What this bill would do. This bill would add the following definition of "agentic AI" to specified statutes, and would require the Office of Emergency Services to account for agentic artificial intelligence when making recommendations as part of a mandated risk analysis:

> (1) "Agentic artificial intelligence" means stochastic artificial intelligence that is capable, alone or in conjunction with other agentic artificial intelligence, of all of the following:
>
> (A) Optimizing toward multistep or abstract objectives using dynamic task decomposition or delegation, which may include modifying the behavior of the artificial intelligence based on outcomes rather than instructions or consent of the user or principal.
>
> (B) Controlling or executing actions or tasks, executing code externally, using external tools, making purchases, or accessing accounts or applications requiring user authentication.
>
> (C) Egressing the network or shell of either the artificial intelligence system or the user or principal.
>
> (2) "Agentic artificial intelligence" does not include an artificial intelligence system that can respond only to direct task prompts and instructions from the user and does not have network egress or shell access capability.

This definition is neither simple nor broad. The Center for AI and Digital Policy (CAIDP), adopting a "support if amended" position, describes three issues:

> **First,** the definition is anchored to "stochastic" AI—a term that reflects the architecture of current transformer-based large language models but does not describe all agentic AI systems, including those that are already commercially deployed. AI companies are actively exploring hybrid architectures that blend generative, reinforcement learning, and classical autonomous systems approaches. A definition based on stochasticity would exempt non-stochastic agentic systems that pose equivalent or greater risks and would place enforcement agencies in the difficult position of establishing the degree to which a system is stochastic in order to trigger regulatory obligations.
>
> **Second,** the current definition requires that an AI system be capable of all three listed capabilities before it qualifies as agentic AI. This structure means that an AI system capable of autonomously executing tasks and accessing external services would fall outside the definition simply because it does not also independently egress the network. This fails to reflect how AI experts, including AI developers themselves, characterize agentic systems: as systems that may exhibit one or more hallmarks of agency to varying degrees.
>
> **Third,** certain specific elements, particularly the network egress criterion in subsection (C), and the granular reference to "making purchases or accessing accounts requiring user authentication" in subsection (B), are either technically underinclusive or likely to become outdated as agentic AI deployment expands across new environments and workflows.

CAIDP goes on to propose the following definition:

"Agentic artificial intelligence" means artificial intelligence, as defined in Government Code § 11546.45.5(a)(1), with multiple coordinated AI agents, for:

(A) Planning and executing tasks and iteratively working toward a defined goal with minimal or no human intervention; or

(B) Optimizing toward multistep or abstract objectives using dynamic task decomposition or delegation, which may include modifying the behavior of the artificial intelligence based on outcomes or goals rather than instructions or consent of the user or principal; or

(C) Controlling or executing actions or tasks, executing code externally, using external tools, services, and environments on behalf of a user."

4) **Amendments.** SB 1106 will be amended to replace its current definition of "agentic AI" with the following definition provided by the author:

(1) "Agentic artificial intelligence" means an artificial intelligence system that can pursue multistep or abstract goals by independently breaking them into tasks or delegating them, adjusting its own behavior based on results rather than waiting for user instruction or consent, and that can act in the world, including by executing code, using external tools, making purchases, accessing accounts that require authentication, or operating outside its own.

(2) "Agentic artificial intelligence" does not include a system that only responds to direct prompts and instructions from the user and lacks network egress or shell access.

As AI agents become widely deployed across industries, California should have a clear understanding of how these systems are being used within state government. Tracking its use of AI agents would help the state identify common applications, assess emerging risks, and ensure that public-sector adoption of agentic AI is paired with appropriate oversight. Pursuant to AB 302 (Ward, Stats. 2023, Ch. 800), the Department of Technology is required to conduct a comprehensive inventory of all high-risk automated decision systems that have been proposed for, or that are currently being, used, developed, or procured by any state agency, and requires the department to submit an annual report to the Assembly Committee on Privacy and Consumer Protection and the Senate Committee on Governmental Organization.

The author has agreed to a Committee amendment adding the term "agentic artificial intelligence" to this inventory and report as follows:

11546.45.5.

(a) For purposes of this section:

(x) (1) "Agentic artificial intelligence" or "agentic AI" means an artificial intelligence system that can pursue multistep or abstract goals by independently breaking them into tasks or delegating them, adjusting its own behavior based on results rather than waiting for user instruction or consent, and that can act in the world, including by executing code, using external tools, making purchases, accessing accounts that require authentication, or operating outside its own.

(2) "Agentic artificial intelligence" does not include a system that only responds to direct prompts and instructions from the user and lacks network egress or shell access.

(1) “Artificial intelligence” means an engineered or machine-based system that varies in its level of autonomy and that can, for explicit or implicit objectives, infer from the input it receives how to generate outputs that can influence physical or virtual environments.

(2) “Automated decision system” means a computational process derived from machine learning, statistical modeling, data analytics, or artificial intelligence that issues simplified output, including a score, classification, or recommendation, that is used to assist or replace human discretionary decisionmaking and materially impacts natural persons. “Automated decision system” does not include a spam email filter, firewall, antivirus software, identity and access management tools, calculator, database, dataset, or other compilation of data.

(3) “Board” means any administrative or regulatory board, commission, committee, council, association, or authority consisting of more than one person whose members are appointed by the Governor, the Legislature, or both.

(4) “Department” means the Department of Technology.

(5) “High-risk automated decision system” means an automated decision system that is used to assist or replace human discretionary decisions that have a legal or similarly significant effect, including decisions that materially impact access to, or approval for, housing or accommodations, education, employment, credit, health care, and criminal justice.

(6) (A) “State agency” means any of the following:

(i) Any state office, department, division, or bureau.

(ii) The California State University.

(iii) The Board of Parole Hearings.

(iv) Any board or other professional licensing and regulatory body under the administration or oversight of the Department of Consumer Affairs.

(B) “State agency” does not include the University of California, the Legislature, the judicial branch, or any board, except as provided in subparagraph (A).

(b) On or before September 1, 2024, the Department of Technology shall conduct, in coordination with other interagency bodies as it deems appropriate, a comprehensive inventory of all high-risk automated decision systems ***and agentic AI systems*** that have been proposed for use, development, or procurement by, or are being used, developed, or procured by, any state agency.

(c) The comprehensive inventory described by subdivision (b) shall include a description of all of the following:

(1) (A) Any decision the automated decision system ***or agentic AI system*** can make or support and the intended benefits of that use.

(B) The alternatives to any use described in subparagraph (A).

(2) The results of any research assessing the efficacy and relative benefits of the uses and alternatives of the automated decision system ***or agentic AI system*** described by paragraph (1).

(3) The categories of data and personal information the automated decision system uses to make its decisions.

(4) (A) The measures in place, if any, to mitigate the risks, including cybersecurity risk and the risk of inaccurate, unfairly discriminatory, or biased decisions, of the automated decision system ***or agentic AI system***.

(B) Measures described by this paragraph may include, but are not limited to, any of the following:

(i) Performance metrics to gauge the accuracy of the system.

(ii) Cybersecurity controls.

(iii) Privacy controls.

(iv) Risk assessments or audits for potential risks.

(v) Measures or processes in place to contest an automated decision.

(d) (1) On or before January 1, 2025, and annually thereafter, the department shall submit a report of the comprehensive inventory described in subdivision (b) to the Assembly Committee on Privacy and Consumer Protection and the Senate Committee on Governmental Organization.

(2) The requirement for submitting a report imposed under paragraph (1) is inoperative on January 1, 2029, pursuant to Section 10231.5.

(3) A report to be submitted pursuant to paragraph (1) shall be submitted in compliance with Section 9795.

ARGUMENTS IN SUPPORT:

CAIDP adopts a "support if amended" position, writing:

> If enacted, SB 1106 would be the first law, at the federal or state level, to set out a definition of agentic AI. This is an important step in California's efforts to build a coherent and durable legal framework for AI.

REGISTERED SUPPORT / OPPOSITION:

Support if amended

Center for Ai and Digital Policy (CAIDP)**Opposition**

None on file.

Analysis Prepared by: Slater Sharp / P. & C.P. / (916) 319-2200