

September 14, 2026

SUBMITTED VIA CFTC PORTAL

Secretary of the Commission
Office of the Secretariat
U.S. Commodity Futures Trading Commission
Three Lafayette Centre
1155 21st Street, N.W.
Washington, D.C. 20581

Re: KalshiEX LLC – CFTC Regulation 40.2(a) Notification Regarding the Initial Listing of the "Will OpenAI announce a safety-related halt in the <development/deployment> of <model> <time period>?" Contract

Dear Sir or Madam,

Pursuant to Section 5c(c) of the Commodity Exchange Act and Section 40.2(a) of the regulations of the Commodity Futures Trading Commission, KalshiEX LLC (Kalshi), a registered DCM, hereby notifies the Commission that it is self-certifying the "Will OpenAI announce a safety-related halt in the <development/deployment> of <model> <time period>?" contract (Contract). The Contract will initially be listed after close-of-business on September 14, 2026; it is listed as the day after because of limitations of the Commission's online submission portal. The Exchange intends to list the contract on a custom basis. The Contract's terms and conditions (Appendix A) includes the following strike conditions:

- <development/deployment>
- <model>
- <time period>

Along with this letter, Kalshi submits the following documents:

- A concise explanation and analysis of the Contract;
- Certification;
- Appendix A with the Contract's Terms and Conditions;
- Confidential Appendices with further information; and
- A request for FOIA confidential treatment.

If you have any questions, please do not hesitate to contact me.

Sincerely,

Xavier Sottile
Head of Markets
KalshiEX LLC
xsottile@kalshi.com

KalshiEX LLC

Official Product Name: "Will OpenAI announce a safety-related halt in the <development/deployment> of <model> <time period>?"

Rulebook: AISAFETYHALT

Summary: AI safety-related halt announcements

Kalshi Contract Category: Science/Technology

Kalshi Internal Category: Science and Technology

September 14, 2026

CONCISE EXPLANATION AND ANALYSIS OF THE PRODUCT AND ITS COMPLIANCE WITH APPLICABLE PROVISIONS OF THE ACT, INCLUDING CORE PRINCIPLES AND THE COMMISSION'S REGULATIONS THEREUNDER

Pursuant to Commission Rule 40.2(a)(3)(v), the following is a concise explanation and analysis of the product and its compliance with the Act, including the relevant Core Principles (discussed in Appendix E), and the Commission's regulations thereunder.

I. Introduction

The "Will OpenAI announce a safety-related halt in the <development/deployment> of <model> <time period>?" Contract is a contract relating to Science and Technology.

Further information about the Contract, including an analysis of its risk mitigation and price basing utility, as well as additional considerations related to the Contract, is included in Confidential Appendices C, D, and E.

Pursuant to Section 5c(c) of the Act and CFTC Regulations 40.2(a), the Exchange hereby certifies that the listing of the Contract complies with the Act and Commission regulations under the Act.

General Contract Terms and Conditions: The Contract operates similar to other event contracts that the Exchange lists for trading. The minimum price fluctuation is \$0.01 (one cent). Price bands will apply so that Contracts may only be listed at values of at least \$0.01 and at most \$0.99. Further, the Contract is sized with a one-dollar notional value and has a minimum price fluctuation of \$0.01 to enable Members to match the size of the contracts purchased to their economic risks. As outlined in Rule 5.16 of the Rulebook, trading shall be available at all times outside of any maintenance windows, which will be announced in advance by the Exchange. Members will be charged fees in accordance with Rule 3.13 of the Rulebook. Fees, if they are charged, are charged in such amounts as may be revised from time to time to be reflected on the Exchange's Website. A new Source Agency can be added via a Part 40 amendment. All instructions on how to access

the Underlying are non-binding and are provided for convenience only and are not part of the binding Terms and Conditions of the Contract. They may be clarified at any time. Furthermore, the Contract's payout structure is characterized by the payment of an absolute amount to the holder of one side of the option and no payment to the counterparty. During the time that trading on the Contract is open, Members are able to adjust their positions and trade freely. The Expiration Value and Market Outcome are determined at or after Market Close. The market is then settled by the Exchange, and the long position holders and short position holders are paid according to the Market Outcome. In this case, "long position holders" refers to Members who purchased the "Yes" side of the Contract and "short position holders" refers to Members who purchased the "No" side of the Contract. If the Market Outcome is "Yes," meaning that an event occurs that is encompassed within the Payout Criterion, then the long position holders are paid an absolute amount proportional to the size of their position and the short position holders receive no payment. If the Market Outcome is "No," then the short position holders are paid an absolute amount proportional to the size of their position and the long position holders receive no payment. Specification of the circumstances that would trigger a Market Outcome of "Yes" are included below in the section titled "Payout Criterion" in Appendix A.

**CERTIFICATIONS PURSUANT TO SECTION 5c OF THE COMMODITY EXCHANGE
ACT, 7 U.S.C. § 7A-2 AND COMMODITY FUTURES TRADING COMMISSION RULE
40.2, 17 C.F.R. § 40.2**

Based on the above analysis, the Exchange certifies that:

- ☐ The Contract complies with the Act and Commission regulations thereunder.
- ☐ This submission (other than those appendices for which confidential treatment has been requested) has been concurrently posted on the Exchange's website at <https://kalshi.com/regulatory/filings>.

Should you have any questions concerning the above, please contact the exchange at ProductFilings@kalshi.com.



By: Xavier Sottile

Title: Head of Markets

Date: September 14, 2026

Attachments:

Appendix A - Contract Terms and Conditions

Appendix B - Trading Prohibitions

Appendix C (Confidential) - Further Considerations

Appendix D (Confidential) - Source Agency

Appendix E (Confidential) - Compliance with Core Principles

APPENDIX A – CONTRACT TERMS AND CONDITIONS

**Official Product Name: "Will OpenAI announce a safety-related halt in the
<development/deployment> of <model> <time period>?"**

Rulebook: AISAFETYHALT

AISAFETYHALT

Scope: These rules shall apply to this contract.

Underlying: The Underlying for this Contract is whether OpenAI has announced a safety-related halt in the <development/deployment> of <model> <time period>. Revisions to the Underlying made after Expiration will not be accounted for in determining the Expiration Value.

Source Agency: The Source Agencies are, in hierarchical order, the official websites, official blogs, official research and safety publications, official regulatory and securities filings, and public communications (including press releases, official corporate accounts on social media platforms, and recorded public statements by authorized officers) of OpenAI, Anthropic, Google DeepMind (including Alphabet Inc. and Google LLC), Meta Platforms, Inc., xAI, Microsoft Corporation, Amazon.com, Inc., Mistral AI, Safe Superintelligence Inc., Cohere Inc., AI21 Labs Ltd., Alibaba Group Holding Limited, DeepSeek, Moonshot AI, Zhipu AI, MiniMax, ByteDance Ltd., Baidu, Inc., and Tencent Holdings Limited, the Associated Press, Bloomberg News, Reuters, the New York Times, the Washington Post, CNBC, BBC, ABC, CBS, CNN, Fox News, MS Now, NBC, and the Wall Street Journal.

Type: The type of Contract is an Event Contract.

Issuance: After the initial Contract, Contract iterations will be listed on an as-needed basis at the discretion of the Exchange and corresponding to the risk management needs of Members.

<development/deployment>: <development/deployment> refers to the category of activity specified by the Exchange with respect to which the halt must be announced. "Development" means pre-training, training, post-training, fine-tuning, scaling, or capability-advancing research and engineering work directed at <model>. "Deployment" means the release, launch, distribution, provision, or making available of <model> to users, customers, or the public, including through an application programming interface, a product or consumer interface, a hosted or enterprise service, or the release of model weights. The Exchange may specify <development/deployment> as "development," as "deployment," as "development or deployment," in which case a halt of either the development or the deployment of <model> satisfies the Payout Criterion, or as "development and deployment," in which case the halt must extend to both. The Exchange may list iterations of the Contract corresponding to variations of <development/deployment>.

<model>: <model> refers to the artificial intelligence model, model family, or category of model-development activity specified by the Exchange, the <development/deployment> of which must be halted (e.g., "GPT-6," "any successor to Claude Opus 5," "any frontier model,"

"any model trained using more than 10^{26} floating-point operations," "all artificial intelligence models"). Where <model> names a specific model, a halt of the <development/deployment> of that model satisfies <model> whether or not the announcement uses the same internal or public name, provided the Source Agencies identify the halted model as that model, including under a different internal or public name. A halt of a successor to <model> is treated as a halt of <model> only in the circumstances described in the Additional Clarifications below. Where <model> names a category, <model> is satisfied by a halt covering any one model within that category, unless the Exchange specifies that the halt must cover all models within the category, in which case the halt must extend to every model within that category then under development by, or then deployed by, OpenAI, as applicable. <model> may also take the forms "Any," meaning any artificial intelligence model developed by OpenAI, or "None." The Exchange may list iterations of the Contract corresponding to variations of <model>.

<time period>: <time period> refers to a specific range of time as specified by the Exchange. This may be defined by exact dates (e.g., "between January 1, 2026, and December 31, 2026"), relative markers (e.g., "before July 1, 2027"), an event, or broader intervals (e.g., "Q1–Q2 2027," "January–March 2026"). "Between" is inclusive of both endpoints, while "before" and "after" exclude the specified date unless stated otherwise. <time period> may also refer to "Any" or "None," to multiple dates (even if non-consecutive), or to a singular and discrete date or time.

Payout Criterion: The Payout Criterion for the Contract encompasses the Expiration Values that OpenAI has announced a safety-related halt in the <development/deployment> of <model> <time period>

Additional Clarifications:

- A "halt" means a stop or suspension of the <development/deployment> of <model>. Where <development/deployment> is specified as development, this means the cessation of pre-training, training, post-training, fine-tuning, scaling, or capability-advancing research and engineering work directed at <model>. An announcement that OpenAI has decided, on a ground enumerated below, not to commence or not to continue the training of <model> (or, where <model> is a category, of any model within that category) is a halt of development whether or not training had begun, provided it is announced as decided and not as a change to a planned schedule. Where <development/deployment> is specified as deployment, this means the cessation of the release, launch, distribution, provision, or making available of <model> as described in <development/deployment> above, including the withdrawal of <model> from availability to users, customers, or the public to whom it has already been made available. Where <model> has not yet been made available, a postponement or rescheduling of its planned release does not constitute a halt of deployment unless OpenAI announces that the release is suspended indefinitely, or until a stated safety condition, evaluation, or safeguard is satisfied or completed, on a

ground enumerated below, and not merely as a change to a planned release date. The halt must be announced as already in effect, as effective immediately, or as effective on a defined, certain date in the future. Where the halt is announced as already in effect, the Payout Criterion is satisfied only if the halt took effect <time period>, whether or not the halt remained in effect at the time of the announcement. If no date on which the halt took effect is stated in the announcement or reported by the Source Agencies, the halt is treated as having taken effect on the date of the announcement.

- The halt need not be announced as lasting for any minimum or stated duration, and the announcement need not specify a number of days, an end date, or a date of resumption (for the avoidance of doubt, a postponement of a <model> that has not yet been made available is governed by the postponement provision above). A halt announced as indefinite, as of unspecified or undetermined duration, or as continuing until a stated condition, milestone, evaluation, or safeguard is satisfied or completed satisfies the Payout Criterion. A halt announced as lasting for a stated fixed period also satisfies the Payout Criterion, regardless of the length of that period, provided that it is announced as an actual cessation of the <development/deployment> of <model> and not as a scheduled break, maintenance or service window, staged or planned interruption, or other interruption forming part of OpenAI's ordinary course of operations.
- The halt need not extend beyond <model>. Where <model> is a single named model, or a category that may be satisfied by a single model, a halt covering that model alone satisfies the Payout Criterion even if OpenAI continues to develop or deploy its other models.
- A halt is "safety-related" if OpenAI, in the announcement itself, attributes the halt in whole or as a principal reason to safety-related risks arising from the capabilities, behavior, or controllability of artificial intelligence systems. The following stated grounds qualify:
 - (i) risk of severe or catastrophic harm to human life, health, public safety, or critical infrastructure;
 - (ii) uplift to the development, acquisition, or use of chemical, biological, radiological, or nuclear weapons;
 - (iii) offensive cyber capability, including the autonomous discovery or exploitation of vulnerabilities or the autonomous conduct of cyberattacks;
 - (iv) loss of human control, including autonomous replication, self-exfiltration, resistance to shutdown, modification, or correction, sabotage or subversion of oversight, or deceptive, manipulative, or scheming behavior toward the developer, its evaluators, or its users;
 - (v) the model having exceeded, or having been assessed as likely to exceed, a capability or risk threshold set out in OpenAI's own published safety framework (such as a responsible scaling policy, preparedness framework, frontier safety framework, or equivalent) for which adequate safeguards are not yet in place;

- (vi) the inability to conduct, complete, or rely upon adequate safety evaluations, interpretability, alignment, or control measures for <model>; or
 - (vii) severe, societal-scale risk expressly identified by OpenAI as arising from the capabilities of <model>.
- Where the announcement states more than one reason, the Payout Criterion is satisfied only if at least one of the grounds enumerated above is stated as a principal reason for the halt. An enumerated ground is not a principal reason if OpenAI characterizes it as subordinate, incidental, or pretextual, or identifies a non-enumerated ground as the primary or main reason for the halt.
- The following stated grounds, standing alone, are not safety-related for purposes of the Payout Criterion: commercial, competitive, financial, valuation, or capital-allocation considerations; compute, energy, hardware, data-center, or supply-chain constraints; export controls, sanctions, or licensing restrictions; personnel changes, labor disputes, or organizational restructuring; litigation, intellectual-property, copyright, contractual, privacy, or data-protection disputes; product quality, accuracy, latency, or performance shortfalls; content moderation, bias, or brand or reputational concerns not involving a ground enumerated above; or ordinary-course roadmap, scheduling, or release-planning changes.
- Where a governmental, regulatory, or intergovernmental authority orders, mandates, or negotiates a halt, the Payout Criterion is satisfied only if the developer or developers constituting OpenAI themselves announce or publicly confirm the halt and either (i) in that announcement or confirmation, state a ground enumerated above as a principal reason for the halt (whether as their own reason or as the authority's stated basis); or (ii) the authority's order, mandate, or agreement, as published by that authority or as reported by a Source Agency, states a ground enumerated above as a principal basis for the halt. Clause (ii) applies notwithstanding any requirement in these Additional Clarifications that the ground be stated by OpenAI in the announcement itself.
- The halt must be a halt in the activity specified by <development/deployment>. Where <development/deployment> is specified as development, a pause, delay, withdrawal, or restriction of the release, launch, deployment, availability, or distribution of <model>, where capability-advancing development of <model> continues, does not satisfy the Payout Criterion. Where <development/deployment> is specified as deployment, a stop or suspension of the development of <model>, where <model> remains available to those to whom it has been deployed, does not satisfy the Payout Criterion.
- A statement of intent, consideration, contemplation, willingness, or conditional commitment to halt the <development/deployment> of <model> — including a statement that OpenAI would halt if some future condition were met, if a competitor did likewise, or if a regulator so required — does not satisfy the Payout Criterion. The halt must be announced as decided and as already in effect, effective immediately, or effective on a date certain as described above.

- A call, petition, open letter, policy proposal, regulatory submission, testimony, or public advocacy by OpenAI or by its officers, directors, or employees urging that OpenAI or any other party halt the <development/deployment> of <model> does not, by itself, satisfy the Payout Criterion.
- A halt announced before the first day of <time period> does not satisfy the Payout Criterion, and neither does a reaffirmation, extension in duration, or continuation <time period> of a halt announced before <time period>. An announcement <time period> that broadens a stoppage announced before <time period>, which did not itself satisfy the Payout Criterion, so that it satisfies the Payout Criterion is treated as a new announcement made <time period>.
- Routine safety holds forming part of OpenAI's ordinary evaluation, review, or release process do not satisfy the Payout Criterion. A safety hold is routine if OpenAI describes it as a standard, pre-planned, or recurring step in its evaluation, review, or release process; a stoppage that OpenAI attributes to a specific safety concern identified with respect to <model> is not routine, regardless of its stated duration.
- An announcement satisfies the Payout Criterion only if it is made publicly by OpenAI through an official channel of OpenAI listed in Source Agency above, or by an officer, director, or authorized spokesperson of OpenAI speaking in that capacity and not disavowed by OpenAI prior to Expiration. A statement by an officer or director of OpenAI announcing a halt of the <development/deployment> of <model>, including through that person's personal social-media account, is treated as made in an official capacity unless disavowed by OpenAI prior to Expiration. Statements by other employees in a personal capacity, anonymous or unnamed sources, leaked internal documents, and third-party reports not attributed by name to an officer, director, or authorized spokesperson of OpenAI and not confirmed by OpenAI through an official channel do not satisfy the Payout Criterion. A statement attributed by name to an officer, director, or authorized spokesperson of OpenAI in a report published by a Source Agency is treated as made in an official capacity, unless disavowed by OpenAI prior to Expiration.
- The announcement must occur <time period>. Where the halt is announced as effective on a defined, certain date in the future, whether that date falls <time period> is not determinative.
- If OpenAI announces a qualifying halt <time period> and subsequently resumes, narrows, shortens, or reverses that halt — whether before or after the end of <time period>, and whether because safeguards were implemented, because of competitive, commercial, or regulatory pressure, or for any other reason — the Contract for OpenAI, <development/deployment>, and <model> <time period> resolves on whether the qualifying announcement was made at any point <time period>, irrespective of whether the halt remained in effect thereafter.
- If OpenAI is dissolved, wound down, acquired, merged, renamed, or reorganized <time period>, any successor entity that continues the artificial intelligence model development

of OpenAI is treated as OpenAI. If OpenAI ceases artificial intelligence model development without announcing a halt that satisfies the Payout Criterion, the Payout Criterion is not satisfied.

- If <model> is cancelled, abandoned, renamed, or superseded before any announcement, a halt announced with respect to its direct successor or replacement in OpenAI's model line is treated as a halt of <model>, provided the Source Agencies identify it as such.
- If two or more Source Agencies conflict as to whether, when, by whom, or on what stated grounds a halt in the <development/deployment> of <model> was announced, the record of the Source Agency highest in the hierarchical order set out in Source Agency above governs. If a Source Agency publishes a record and subsequently corrects, retracts, or removes that record prior to Expiration, the corrected record governs.
- For the avoidance of doubt, where the Source Agencies are available and reporting, the absence of any announcement satisfying the Payout Criterion <time period> is a determinate outcome, and the Contract will resolve No.

If a natural person who is the primary subject of a Contract's Underlying or Payout Criterion dies prior to Expiration, or if an act or attempted act of violence, assassination, or terrorism directly prevents, interrupts, or causes the event that is the primary subject of a Contract's Underlying or Payout Criterion, Kalshi may, in its sole discretion, settle the Contract at the last traded price prior to the death or act. If Kalshi determines that trading activity was materially affected by the circumstances giving rise to the death or act, Kalshi may instead use the last traded price prior to such circumstances becoming known or reasonably anticipated by market participants. For purposes of this Rule, Kalshi may rely on an approximate time where the precise time is unknown or cannot be reasonably determined. If Kalshi determines that no last traded price represents a fair settlement, the Outcome Review Committee shall determine the fair settlement price. Kalshi may halt or pause trading in any such Contract (with public notice on Kalshi's website or the Platform) if it reasonably believes that a death or act of terrorism subject to this Rule has occurred, is imminent, or that circumstances giving rise to such an occurrence may be happening. Determinations of Kalshi and the Outcome Review Committee under this Rule are final and not subject to review.

Minimum Tick: The Minimum Tick size for the Contract shall be \$0.01.

Position Accountability Level: The Position Accountability Level for the Contract shall be \$25,000 per strike, per Member.

Last Trading Date: The Last Trading Date of the Contract will be the final day of <time period>. The Last Trading Time will be one minute prior to the end of <time period>.

Settlement Date: The Settlement Date of the Contract shall be no later than the day after the Expiration Date, unless the Market Outcome is under review pursuant to Rule 7.1.

Expiration Date: The latest Expiration Date of the Contract shall be one week after the final day of <time period>. If an event described in the Payout Criterion occurs, expiration will be moved to an earlier date and time in accordance with Rule 7.2.

Expiration time: The Expiration time of the Contract shall be 10:00 AM ET.

Settlement Value: The Settlement Value for this Contract is \$1.00.

Expiration Value: The Expiration Value is the value of the Underlying as documented by the Source Agency on the Expiration Date at the Expiration time.

Contingencies: Before Settlement, Kalshi may, at its sole discretion, initiate the Market Outcome Review Process pursuant to Rule 7.1 of the Rulebook. If an Expiration Value cannot be determined on the Expiration Date, Kalshi has the right to determine payouts pursuant to Rule 7.1 in the Rulebook.

APPENDIX B – TRADING PROHIBITIONS

In addition to the general prohibition against trading on material nonpublic information, the Exchange will institute additional prohibitions for trading the contract. Persons under 18 years of age are not permitted to create Kalshi accounts. The following individuals will be prohibited from trading:

- Any employee, officer, director, contractor, or agent of OpenAI, including any research, safety, evaluation, policy, legal, or communications function responsible for assessing, authorizing, announcing, or overseeing a halt in the <development/deployment> of <model>.
- Any member of the board of directors, board committee, safety or oversight body, or trustee body of OpenAI, and any staff thereof, with access to material nonpublic information regarding the timing, scope, or status of a halt in the <development/deployment> of <model>.
- Any employee, officer, contractor, or agent of a cloud provider, compute provider, model evaluator, red-teaming organization, third-party auditor, or other vendor engaged in building, training, hosting, evaluating, deploying, or administering <model>, with access to material nonpublic information regarding the timing, scope, or status of a halt in the <development/deployment> of <model>.
- Any employee, officer, or agent of a governmental, regulatory, or intergovernmental authority with access to material nonpublic information regarding an order, mandate, negotiation, or agreement that would cause OpenAI to halt the <development/deployment> of <model>.
- Any employee, officer, contractor, or agent of a Source Agency listed in Appendix A with access to material nonpublic information relating to the Underlying.
- Any person who otherwise has the ability to influence, or has advance knowledge of, OpenAI's decision to announce, accelerate, delay, narrow, suspend, or reverse a halt in the <development/deployment> of <model>.
- The immediate family members and members of the same household of any person described above