



NEWS RELEASE

Corvex Launches Security-First AI Inference Platform for Open-Weight Models

2026-09-30

Corvex Token Factory combines zero data retention by default with API-based access

ARLINGTON, Va., Sept. 30, 2026 /PRNewswire/ -- Corvex, Inc. (Nasdaq: MOVE), an engineering-led AI computing company, today announced the launch of [Corvex Token Factory](#), its serverless inference platform built to help developers and enterprises use open-weight AI models with greater control over their data.

Corvex Token Factory gives customers access to open-weight AI models through an API without requiring them to deploy or operate GPU clusters. Initial model availability includes GLM 5.3 from Z.ai and DeepSeek V4 Flash 0731. The launch extends Corvex's reach to developers and enterprise technology teams that want to use AI as a service.

"As companies put AI to work on their source code, customer records and internal documents, they need to know where that information goes and what happens to it," said Seth Demsey, co-CEO and co-founder of Corvex. "With Corvex Token Factory, by default, we do not store end-user prompts or responses nor do we train on any user data. We have built these controls into a service that developers can easily use with their existing coding harnesses and AI API calls so that they can safely enjoy the power of leading open-weights models at a fraction of the cost of closed-source alternatives."

Inference is the process that turns an end-user request into a model-generated response. Corvex runs every request on hardware it manages, and it operates the models itself. Requests are never routed to third-party inference providers or forwarded to the model developers.

At launch, Corvex Token Factory is built around four principles:

- **Assurance** – Zero data retention is the default. Prompts and responses are processed in memory, never logged or stored, and never used to train models. Corvex retains limited operational metadata for security, service operations and billing, but that metadata does not include prompt or response content.
- **Reliability** – Corvex Token Factory is engineered to hold up under sustained workloads so it runs as a dependable execution layer in production workflows.

- **Performance** – Corvex tunes its inference stack, balancing first response and throughput under load. This supports interactive applications and agent workflows.
- **Simplicity** – OpenAI- and Anthropic-compatible APIs let developers connect existing clients by changing the base URL, API key and model name. Corvex manages the underlying infrastructure; feature compatibility varies by client and model.

Corvex is SOC 2 Type II certified. The Company also signs a Business Associate Agreement with customers that are subject to HIPAA, with security controls documented in the [Corvex Trust Center](#).

Customers pay for the input and output tokens they use, without idle GPU charges. The Corvex Token Factory also provides a starting point for deeper customer relationships. Corvex's roadmap includes dedicated enterprise deployments and secure private inference environments, giving customers a path to adopt additional services as their usage and requirements grow.

The launch follows Corvex Token Factory's closed alpha period. Customers can register at no cost at <https://tokenfactory.corvex.cloud/app> and get an API key. Product documentation is available at <https://docs.corvex.cloud/>.

About Corvex, Inc.

Corvex, Inc. (Nasdaq: MOVE) is an AI cloud computing company specializing in GPU-accelerated infrastructure for AI workloads, and a publicly traded pure-play AI compute platform. The company provides secure, scalable and cost-efficient compute resources through GPU-accelerated clusters, high-throughput storage and a layered architecture engineered for reliability, performance and efficiency at scale. Corvex's product suite includes AI Factories and GPU Clusters, the Assured AI confidential-computing platform, as well as the Corvex Token Factory, a serverless inference platform. For more information, visit corvex.ai.

Cautionary Statements Regarding Forward-Looking Statements

This press release contains forward-looking statements within the meaning of the Private Securities Litigation Reform Act of 1995 that are based upon current expectations or beliefs, as well as assumptions about future events. Forward-looking statements include all statements that are not historical facts and can generally be identified by terms such as "anticipate," "believe," "could," "estimate," "expect," "intend," "may," "plan," "potentially," "seek," "should," "target," "will" or "would" or similar expressions and the negatives of those terms. These statements include, but are not limited to, statements relating to product capabilities, customer deployment, business strategy, growth plans and objectives for future operations, anticipated power and customer-readiness dates, expected megawatt capacity and expansion options, and statements regarding anticipated customer offtake discussions, contract duration, pricing and attach opportunities. Actual results could differ materially as a result of, among other factors: the Company's inability to enter into customer offtake agreements at all, or on the duration, pricing or take-or-pay terms it anticipates; changes in prevailing market pricing for GPU-as-a-Service capacity; delays in GPU and long-lead equipment delivery; delays in utility energization or interconnection; local zoning, permitting and moratorium risk affecting expansion phases at either site; the Company's dependence on third-party site partners and the execution of definitive agreements with them; the Company's ability to obtain or maintain partner certifications and co-selling support; and the Company's limited operating history. Actual results could differ materially from those expressed in or implied by the forward-looking statements due to a number of risks and uncertainties, including the risks and uncertainties described in the Company's SEC reports, and under the heading "Risk Factors" in its most recent Annual Report on Form 10-K and Quarterly Reports on Form 10-Q, which are available at www.sec.gov and in other filings the Company makes and will make with the SEC. The forward-looking statements contained herein speak only as of the date of this press release.

Except as required by law, the Company does not undertake any obligation to update or revise its forward-looking statements to reflect events or circumstances after the date of this press release.

Media Contact

Chris Donahoe, Stillpoint

corvex.media@stillpointglobaladvisors.com

View original content to download multimedia: <https://www.prnewswire.com/news-releases/corvex-launches-security-first-ai-inference-platform-for-open-weight-models-302894497.html>

SOURCE Corvex