Lovable and Cerebras Partner to Power AI Software Creation on the World's Fastest Inference Platform

August 5, 2026

Partnership brings industry-leading Cerebras inference to one of the world’s fast-growing AI software creation platforms, enabling more interactive software development experiences

STOCKHOLM, Sweden and SUNNYVALE, Calif., Aug. 05, 2026 (GLOBE NEWSWIRE) -- Lovable and Cerebras Systems (NASDAQ: CBRS) today announced a partnership to power Lovable's software creation platform with the world's fastest inference platform built by Cerebras. Under the agreement, Lovable will run selected latency-sensitive workloads on dedicated Cerebras capacity, and the two companies will explore product experiences that faster inference makes possible.

Since its launch in November 2024, more than 50 million projects have been built on Lovable—internal tools, new product lines, entire companies—by the people closest to the problems that software is meant to solve. That work depends on inference speed. Every instruction, every revision, every correction is a round trip to a model, and the wait between them is where momentum is lost.

Cerebras built its Wafer-Scale Engine to remove speed as a constraint on what AI can do. Where GPU-based systems split a model's weights across many chips and pay a networking tax on every token, the Wafer-Scale Engine keeps an entire model's weights on a single wafer. This delivers far greater memory bandwidth than the fastest GPU and generating tokens fast enough to make multi-step AI workflows feel instantaneous.

That speed will be available to the millions of people who have built projects on Lovable. By dramatically reducing response times, Cerebras will enable developers and business users alike to stay in flow, transforming software creation from a sequence of waiting periods into a seamless, interactive experience.

“Lovable was built on the idea that the person closest to a problem should be able to solve it,” said Anton Osika, co-founder and CEO, Lovable. “Cerebras helps us make Lovable respond as fast as our customers think, so instead of wondering if an idea’s worth trying, they just build it. Our users deserve infrastructure that keeps up with how they work.”

Creating software is not a single request but a chain of them—planning, scaffolding, writing components, catching an error, rewriting. Each link is a decode-bound workload, where output tokens must be produced one after another, and where GPU-based inference is at its slowest. This is precisely the workload the Wafer-Scale Engine was designed to accelerate: by keeping model weights and compute on a single wafer, Cerebras delivers orders of magnitude more memory bandwidth than the fastest GPU, turning what used to be a batch process into something that feels conversational and interactive at every step of the chain.

“Fast AI is more valuable than slow AI. When AI responds in real-time, users do more with it, stay longer, and run higher value workloads. Software creation is one of the clearest examples of the importance of speed,” said Andrew Feldman, CEO and co-founder, Cerebras. “Creators don’t want to wait. Lovable has already put building software into the hands of millions of people. At Cerebras, we make the AI so fast that users are delighted by the experience.”

The collaboration marks another step in Cerebras' expansion into new AI-native markets, demonstrating how ultra-fast inference enables entirely new categories of interactive applications beyond traditional enterprise AI.

The companies are jointly exploring product experiences that faster inference makes possible. Further details, including technical results and availability, will be shared as the work progresses.

About Cerebras Systems
Cerebras Systems (NASDAQ: CBRS) builds the world’s fastest AI infrastructure. The Cerebras team of pioneering computer architects, computer scientists, AI researchers, and engineers of all types came together to make AI blisteringly fast through innovation and invention. We believe that when AI is fast, it will change the world. Leading global corporations, research institutes, and governments choose Cerebras to run their AI workloads. Cerebras solutions are available on premises and in the cloud. Visit cerebras.ai for more.

About Lovable
Lovable is the software creation platform that gives people the power to act on the problems closest to them. From solopreneurs to small business owners to teams at companies like Adidas and Zendesk, people have built over 50 million projects on Lovable

since its launch in November 2024. Headquartered in Stockholm with teams in London, Boston, New York, and San Francisco, Lovable is backed by leading investors including Menlo Ventures, CapitalG, and Accel. Learn more at https://lovable.dev.

CEREBRAS DISCLOSURE INFORMATION

Cerebras uses its investor relations page (investors.cerebras.ai), its X account (@cerebras), and its LinkedIn page (linkedin.com/company/cerebras-systems/) to disclose material non-public information and for complying with its disclosure obligations under Regulation FD. Accordingly, investors should monitor these channels, in addition to following Cerebras' press releases, Securities and Exchange Commission (SEC) filings, public conference calls and public webcasts.

Forward-Looking Statements

This press release contains "forward-looking statements" within the meaning of applicable securities laws. All statements other than statements of historical fact could be deemed to be forward-looking, including, but not limited to, statements regarding the implementation and anticipated benefits of the partnership between Cerebras and Lovable, including Lovable's plans to run latency-sensitive workloads on dedicated Cerebras capacity; the joint development of platform capabilities, model integrations, and product experiences; Cerebras' continued expansion into new AI-native markets; the enablement of entirely new categories of interactive applications beyond traditional enterprise AI; the timing of further details, including technical results and availability; and any assumptions relating to the foregoing. The words "may," "will," "shall," "should," "expects," "plans," "anticipates," "could," "intends," "target," "projects," "contemplates," "believes," "estimates," "predicts," "potential," "objective," or "continue," or the negative of these words or other similar terms or expressions that concern our expectations, strategy, plans, or intentions are intended to identify forward-looking statements, although not all forward-looking statements contain these identifying words. These forward-looking statements are subject to a number of risks and uncertainties, many of which involve factors or circumstances that are beyond Cerebras' control. These risks and uncertainties include, but are not limited to: Cerebras' ability to sustain and manage its growth, access borrowings and other sources of capital on acceptable terms, and deploy available capital to support growth; its history of net losses and ability to achieve and maintain profitability; its limited operating history at its current scale and ability to accurately forecast revenue and appropriately budget and manage expenses; its dependence on a limited number of significant customers, including OpenAI, Group 42 Holding Ltd, Mohamed bin Zayed University of Artificial Intelligence, and AWS, and the potential impact of any reduction in demand from, material adverse development in its relationships with, or failure to meet its obligations to, such customers, including under its Master Relationship Agreement with OpenAI; the timing, execution and expected benefits of its strategic customer, partner and financing arrangements; its historical reliance on sales of hardware systems and the early-stage, rapidly evolving market for its cloud-based offerings and AI infrastructure; its ability to secure sufficient data center capacity and capital to support its cloud-based offerings; its ability to launch new offerings and add new product capabilities; and its ability to compete effectively in the rapidly evolving and competitive market for AI computing solutions.

Cerebras' actual results could differ materially from those stated or implied in forward-looking statements due to a number of factors. Accordingly, undue reliance should not be placed on such statements. These forward-looking statements are made as of the date they were first issued and are based on information available to Cerebras together with Cerebras' expectations, estimates, forecasts, projections, beliefs, and assumptions as of such date. These forward-looking statements should not be relied upon as representing Cerebras' views as of any date subsequent to the date of this press release. Past performance is not necessarily indicative of future results. Cerebras undertakes no intention or obligation to update or revise any forward-looking statements, whether as a result of new information, future events, or otherwise, except as required by law.

Further information on potential risks that could affect actual results is included in Cerebras' most recent filings with the SEC, including in Cerebras' most recent Quarterly Report on Form 10-Q, copies of which may be obtained by visiting Cerebras' Investor Relations website at investors.cerebras.ai or the SEC's website at www.sec.gov.

Contacts

Cerebras
Kriselle Laran
Media Relations
pr@cerebras.ai

Sean Dorsey
Investor Relations
investors@cerebras.ai

Lovable
press@lovable.dev