

EXHIBIT

15



IMAGE CREDITS: BRYCE DURBIN/TECHCRUNCH

TRANSPORTATION

f x in e m o

Tesla Dojo: The rise and fall of Elon Musk’s AI supercomputer

Rebecca Bellan · 9:18 AM PDT · September 2, 2025



For years, Elon Musk has spoken of the promise of Dojo, the AI supercomputer that was supposed to be the cornerstone of Tesla’s AI ambitions. It was important enough to Musk that in July 2024, he said the company’s AI team would “[double down](#)” on Dojo in the lead-up to [Tesla’s robotaxi reveal](#), which happened in October.

After six years of hype, Tesla decided last month to [shut down Dojo](#) and disband the team behind the supercomputer in August 2025. Within weeks of projecting that Dojo 2, Tesla’s second supercluster that was meant to be built on the company’s in-house D2 chips, would reach scale by 2026, Musk reversed course, declaring it “an evolutionary dead end.”

Ad



Manage your stocks and ET

Fidelity Investments

[Learn more](#)

This article originally set out to explain what Dojo was and how it could help Tesla achieve full-self driving, autonomous humanoid robots, semiconductor autonomy, and more. Now, you can think of it more as an obituary of a project that convinced so many analysts and investors that Tesla wasn’t just an automaker, it was an AI company.

Dojo was Tesla’s custom-built supercomputer that was designed to train its “Full Self-Driving” neural networks.

Beefing up Dojo went hand-in-hand with Tesla’s goal to reach full self-driving and bring a robotaxi to market. FSD (Supervised) is Tesla’s advanced driver assistance system that’s on hundreds of thousands of Tesla vehicles today and can perform some automated driving tasks, but it still requires a human to be attentive behind the wheel. It’s also the basis of similar technology powering Tesla’s limited robotaxi service that the company launched in [Austin this June](#) using Model Y SUVs.

Even as Dojo’s *raison d’être* started to come to life, Tesla failed to attribute its self-driving successes — [controversial as they were](#) — to the supercomputer. In fact, Musk and Tesla had barely mentioned Dojo at all over the past year. In

PAIGE

Find Your Perfect Fit



SHOP NOW

DISRUPT

October 13 - 15 | San Francisco

Scale faster. Grow your portfolio. Gain practical expertise. No Disrupt can empower you.


Save up to \$400 today!

REGISTER NOW >

0:03

August 2024, Tesla began promoting Cortex, the company's "giant new AI training supercluster being built at Tesla HQ in Austin to solve real-world AI," which Musk has said would have "massive storage for video training of FSD and Optimus."

In Tesla's Q4 2024 shareholder deck, the company shared updates on Cortex, but nothing on Dojo. It's not clear whether Tesla's Dojo shutdown affects Cortex.



The response to Dojo's disbanding has been mixed. Some see it as another example of Musk making promises he can't deliver on that comes at a time of falling EV sales and a lackluster robotaxi rollout. Others say the shutdown wasn't a failure, but a strategic pivot from a high-risk, self-reliant hardware to a streamlined path that relies on partners for chip development.

Dojo's story reveals what was on the line, where the project fell short, and what its shutdown signals for Tesla's future.

A recap of Dojo's shutdown

Tesla disbanded its Dojo team and shut down the project in mid-August 2025. Dojo's lead, Peter Bannon, left the company as well, following the departure of around 20 workers who left to start their own AI chip and infrastructure company called DensityAI.

Analysts have pointed out that losing key talent can quickly derail a project, especially a highly specialized, internal tech project.

The shutdown came a couple of weeks after Tesla signed a \$16.5 billion deal to get its next-generation AI6 chips from Samsung. The AI6 chip is Tesla's bet on a chip design that can scale from powering FSD and Tesla's Optimus humanoid robots to high-performance AI training in data centers.

"Once it became clear that all paths converged to AI6, I had to shut down Dojo and make some tough personnel choices, as Dojo 2 was now an evolutionary dead end," Musk posted on X, the social media platform he owns. "Dojo 3 arguably lives on in the form of a large number of AI6 [systems-on-a-chip] on a single board."

Tesla's Dojo backstory



IMAGE CREDITS:SUZANNE CORDEIRO / AFP / GETTY IMAGES

Musk has insisted that Tesla isn't just an automaker, or even a purveyor of solar panels and energy storage systems. Instead, he has pitched Tesla as an

Most Popular

■ ChatGPT brings unlimited text chats to free users

■ Ford's new electric truck, 'Fathom,' starts at \$28,350

■ Bending Spoons to buy Airtable for \$1.28B

■ Influencers draw backlash for attending OpenAI's first luxury trip

■ Sequoia's Shaun Maguire leads \$1B round for nuclear startup Valar Atomics

■ Malaysia is reportedly shutting down Balaji Srinivasan's Network School

■ YouTuber Hank Green says his AI usage is 'not healthy'

Ad



Fidelity Basket Portfolios

Manage your stocks and ET

Fidelity Basket Portfolios is a faster and easier way to build a basket of securities and

Fidelity Investments

Learn more

Advertisement

AI company, one that has cracked the code to self-driving cars by mimicking human perception.

Most other companies building autonomous vehicle technology rely on a combination of sensors to perceive the world — like lidar, radar and cameras — as well as high-definition maps to localize the vehicle. Tesla believes it can achieve fully autonomous driving by relying on cameras alone to capture visual data and then use advanced neural networks to process that data and make quick decisions about how the car should behave.

The pitch has been that Dojo-trained AI software will eventually be pushed out to Tesla customers via over-the-air updates. The scale of FSD also means Tesla has been able to rake in millions of miles worth of video footage that it uses to train FSD. The idea there is that the more data Tesla can collect, the closer the automaker can get to actually achieving full self-driving.

However, some industry experts say there might be a limit to the brute force approach of throwing more data at a model and expecting it to get smarter.

“First of all, there’s an economic constraint, and soon it will just get too expensive to do that,” Anand Raghunathan, Purdue University’s Silicon Valley professor of electrical and computer engineering, told TechCrunch. Further, he said, “Some people claim that we might actually run out of meaningful data to train the models on. More data doesn’t necessarily mean more information, so it depends on whether that data has information that is useful to create a better model, and if the training process is able to actually distill that information into a better model.”

Raghunathan said despite these doubts, the trend of more data appears to be here for the short-term at least. And more data means more compute power needed to store and process it all to train Tesla’s AI models. That was where Dojo, the supercomputer, came in.

What is a supercomputer?

Dojo was Tesla’s supercomputer system that was designed to function as a training ground for AI, specifically FSD. The name is a nod to the space where martial arts are practiced.

A supercomputer is made up of thousands of smaller computers called nodes. Each of those nodes has its own CPU (central processing unit) and GPU (graphics processing unit). The former handles overall management of the node, and the latter does the complex stuff, like splitting tasks into multiple parts and working on them simultaneously.

GPUs are essential for machine learning operations like those that power FSD training in simulation. They also power large language models, which is why the rise of generative AI has made Nvidia the most valuable company on the planet.


Even Tesla buys Nvidia GPUs to train its AI (more on that later).

Why did Tesla need a supercomputer?

Tesla’s vision-only approach was the main reason Tesla needed a supercomputer. The neural networks behind FSD are trained on vast amounts of driving data to recognize and classify objects around the vehicle and then make driving decisions. That means that when FSD is engaged, the neural nets have to collect and process visual data continuously at speeds that match the depth and velocity recognition capabilities of a human.

In other words, Tesla means to create a digital duplicate of the human visual cortex and brain function.

To get there, Tesla needs to store and process all the video data collected from its cars around the world and run millions of simulations to train its model on the data.





Tesla relied mainly on Nvidia to power its current Dojo training computer, but it didn't want to have all its eggs in one basket — not least because Nvidia chips are expensive. Tesla had hoped to make something better that increased bandwidth and decreased latencies. That's why the automaker's AI division decided to come up with its own custom hardware program that aimed to train AI models more efficiently than traditional systems.

At that program's core was Tesla's proprietary D1 chips, which the company said were optimized for AI workloads.

Tell me more about these chips



GANESH VENKATARAMANAN, FORMER SENIOR DIRECTOR OF AUTOPILOT HARDWARE, PRESENTING THE D1 TRAINING TILE AT TESLA'S 2021 AI DAY.
IMAGE CREDITS: TESLA / SCREENSHOT

Tesla, like Apple, thinks hardware and software should be designed to work together. That's why Tesla was working to move away from the standard GPU hardware and [design its own chips](#) to power Dojo.


Tesla unveiled its D1 chip, a silicon square the size of a palm, on AI Day in 2021. The D1 chip entered into production around [July 2023](#).

The Taiwan Semiconductor Manufacturing Company (TSMC) manufactured the chips using 7 nanometer semiconductor nodes. The D1 has 50 billion transistors and a large die size of 645 millimeters squared, according to Tesla. This is all to say that the D1 promises to be extremely powerful and efficient and to handle complex tasks quickly.

The D1 wasn't as powerful as Nvidia's A100 chip, though.

Tesla had been working on a next-gen D2 chip that aimed to solve information flow bottlenecks. Instead of connecting the individual chips, the D2 would have put the entire Dojo tile onto a single wafer of silicon.

Tesla never confirmed how many D1 chips it ordered or received. The company also never provided a timeline for how long it would have taken to get Dojo supercomputers running on D1 chips.



What did Dojo mean for Tesla?



VISITORS ARE VIEWING TESLA'S HUMANOID ROBOT OPTIMUS PRIME II AT WAIC IN SHANGHAI, CHINA, ON JULY 7, 2024.

IMAGE CREDITS: COSTFOTO / NURPHOTO / GETTY IMAGES

Tesla's hope was that by taking control of its own chip production, it might one day be able to quickly add large amounts of compute power to AI training programs at a low cost.

It also meant not having to rely on Nvidia's chips in the future, which are increasingly expensive and hard to secure. Now, Tesla is going all-in on partnerships — with Nvidia, AMD, and Samsung, which will build its next-gen AI6 chip.

During Tesla's second-quarter 2024 earnings call, Musk said demand for Nvidia hardware was "so high that it's often difficult to get the GPUs." He said he was "quite concerned about actually being able to get steady GPUs when we want them, and I think this therefore requires that we put a lot more effort on Dojo in order to ensure that we've got the training capability that we need."

Dojo was a risky bet, one that Musk hedged several times by saying that Tesla might not succeed.


In the long run, Tesla toyed with the idea of creating a new business model based on its AI division, with Musk even saying during a Q2 2024 earnings call that he saw "a path to being competitive with Nvidia with Dojo." While D1 was more tailored for Tesla computer vision labeling and training — useful for FSD and [Optimus](#) training — it wouldn't have been useful for much else. Future versions would have to be [more tailored to general-purpose AI training](#), Musk said.

The problem that Tesla might have come up against is that almost all AI software out there has been written to work with GPUs. Using Dojo chips to train general-purpose AI models would have required rewriting the software.

That is, unless Tesla rented out its compute, similar to how AWS and Azure rent out cloud computing capabilities — an idea that excited analysts. A September 2023 report from Morgan Stanley predicted that Dojo could [add \\$500 billion](#) to Tesla's market value by unlocking new revenue streams in the form of robotaxis and software services.

In short, Dojo chips were an insurance policy for the automaker, but one that might have paid dividends.

How far did Tesla Dojo get?





NVIDIA CEO JENSEN HUANG AND TESLA CEO ELON MUSK
IMAGE CREDITS: KIM KULISH / CORBIS / GETTY IMAGES

Musk often provided progress reports, but many of his goals for Dojo were never reached.

For instance, Musk suggested in June 2023 that Dojo had been online and running useful tasks for a few months.” Around the same time, Tesla said it expected Dojo to be one of the top five most powerful supercomputers by February 2024 and had planned for total compute to reach 100 exaflops in October 2024, which would have required roughly 276,000 D1s, or around 320,500 Nvidia A100 GPUs.

Tesla never provided an update or any information that would suggest it ever reached these goals.

Tesla and Musk made numerous other pledges for Dojo, including financial ones. For instance, Tesla committed in January 2024 to [spend \\$500 million](#) to build a Dojo supercomputer at its gigafactory in Buffalo, New York, and has already spent \$314 million of that, per a 2024 report.

Just after Tesla’s second-quarter 2024 earnings call, Musk [posted photos of Dojo 1 on X](#), saying that it would have “roughly 8k H100-equivalent of training online by end of year. Not massive, but not trivial either.”

Despite all of this activity — particularly by Musk on X and in earnings calls — mention of Dojo abruptly ended August 2024. And talk switched to Cortex.

During the company’s fourth-quarter 2024 earnings call, Tesla said it completed the deployment of Cortex, “a ~50k H100 training cluster at Gigafactory Texas” and that Cortex helped enable V13 of supervised FSD.

In Q2 2025, Tesla noted it “expanded AI training compute with an additional 16k H200 GPUs at Gigafactory Texas, bringing Cortex to a total of 67k H100 equivalents.” During that same earnings call, Musk [said](#) he expected to have a second Dojo cluster operating “at scale” in 2026. He also hinted at potential redundancies.

“Thinking about Dojo 3 and the AI6 inference chip, it seems like intuitively, we want to try to find convergence there, where it’s basically the same chip,” Musk said.

A few weeks later, he reversed course and disbanded the Dojo team.


TechCrunch confirmed in late August 2025 that Tesla still plans to commit \$500 million to a supercomputer in Buffalo — it just won’t be Dojo.

This story originally published August 3, 2024. The article was updated for a final time September 2, 2025, with new information about Tesla’s decision to shut down Dojo.

Topics: [AI](#) [Dojo](#) [Elon Musk](#) [evergreens](#) [Tesla](#) [Tesla FSD](#) [Transportation](#)

When you purchase through links in our articles, [we may earn a small commission](#). This doesn’t affect our editorial independence.

Advertisement





Rebecca Bellan

Senior Reporter |  

Rebecca Bellan is a senior reporter at TechCrunch where she covers the business, policy, and emerging trends shaping artificial intelligence. Her work has also appeared in Forbes, Bloomberg, The Atlantic, The Daily Beast, and...

View Bio >

Loading the next article



Advertisement



TechCrunch

Staff

Contact Us

Advertise

Crunchboard Jobs

Site Map

Terms of Service

Privacy Policy

RSS Terms of Use

OpenAI vs Apple

Nous Research

Space Data Centers

Staya Nadella

SpaceX Starship

Tech Layoffs

ChatGPT