

EXHIBIT

14



Best Picks

CPUs

GPUs

PC Components

News

Laptops

Desktops

Software & AI

Coupons

Premium

Forums

**BEST
BUY.**Start strong
with your signature color.

MacBook Neo

TRENDING

Tom's 30th Anniversary

AMD Advancing AI

RAM Shortage

Vera Rubin

AI Data Centers

RAM Combo Deals

TH Premium

AMD Instinct MI455X

DLSS


PC Components > GPUs

ADVERTISEMENT

Tesla Brags About In-House Supercomputer, Now With 7,360 A100 GPUs

News By Mark Tyson | Published August 17, 2022

With a 28% increase in GPUs, it's now a top-7 supercomputer worldwide by GPU count



(Image credit: Tesla)

BE SUMMER TRAVEL READY

FREE PRE-TRIP SAFETY CHECK

TRAVEL

FREE PRE-TRIP SAFETY CHECK



9



Follow us




Newsletter

Tesla has boosted its in-house AI supercomputer with thousands of additional Nvidia A100 GPUs. The Tesla supercomputer had 5,760 A100 GPUs about a year ago, and that count has since risen to 7,360 A100 GPUs — that's an additional 1,600 GPUs, or about a 28% increase.

According to Tesla Engineering Manager Tim Zaman, this upgrade makes the firm's AI system a top-7 supercomputer worldwide by GPU count.


An [Nvidia A100 GPU](#) is a powerful Ampere architecture solution aimed at data centers. Yes, it uses the same GPU architecture as GeForce RTX 30 series GPUs, which are some of the [best graphics cards](#) currently available. However, there is no close consumer relation to the A100, which comes with 80GB of HBM2e memory on board, offers up to 2 TB/s bandwidth, and requires up to 400W of power. The architecture of the A100 has also been tweaked for accelerating tasks common in AI, data analytics, and high-performance computing (HPC) applications.

ADVERTISEMENT



The first system Nvidia showed wielding the A100 was the [Nvidia DGX A100](#), which packed in eight A100 GPUs linked via six NVSwitch with 4.8 TBps of bi-directional bandwidth for up to 10 PetaOPS of INT8

ADVERTISEMENT



performance, 5 PFLOPS of FP16, 2.5 TFLOPS of TF32, and 156 TFLOPS of FP64 in a single node.

That was eight A100 GPUs — Tesla's AI supercomputer now has 7,360 of these. Tesla hasn't publicly benchmarked its AI supercomputer, but the similarly-equipped GPU-based NERSC Perlmutter, which has 6,144 Nvidia A100 GPUs, achieves 70.87 Linpack petaflops. Using this and data from other A100 GPU supercomputers as performance reference points, [HPC Wire](#) estimates the Tesla AI supercomputer is capable of achieving about 100 Linpack petaflops.

YOU MAY LIKE

China bypasses US GPU bans with 1.54-exaflops 'LineShine' supercomputer



Google could build more AI accelerators than Nvidia sells in 2028, analyst claims




Nvidia's memory costs soar 485%, latest AI systems now cost \$7.8 million to build



Tesla doesn't intend to continue down the Nvidia GPU architecture path for its in-house AI supercomputers long-term. This world's top-7 machine by GPU-count is merely a precursor to the upcoming [Dojo](#) supercomputer, which was first [announced](#) by Elon Musk back in 2020. A year ago we got a look at the [Tesla D1 Dojo chip](#), which are designed to supplant Nvidia's GPUs for "maximum performance, throughput and bandwidth at every granularity."

ADVERTISEMENT



(Image credit: Tesla)

The Tesla Dojo D1 is a custom ASIC (application-specific integrated circuit) design, purposed for AI training, and it is one of the first ASICs in this field. Current D1 test chips are manufactured on TSMC N7 and pack in about 50 million transistors.

Stay On the Cutting Edge: Get the Tom's Hardware Newsletter

Get Tom's Hardware's best news and in-depth reviews, straight to your inbox.

 SIGN ME UP

By signing up, you agree to our [Terms of services](#) and acknowledge that you have read our [Privacy Notice](#). You also agree to receive marketing emails from us that may include promotions from our trusted partners and sponsors, which you can unsubscribe from at any time.

More information about the Dojo D1 chip, and the Dojo system, might be revealed at [next week's Hot Chips Symposium](#) — three Tesla presentations are schedule for next Tuesday, addressing Dojo D1 chip architecture, Dojo and ML training, and enabling AI through system integration.

TOPICS

Nvidia

Science

GeForce


 SEE ALL COMMENTS (9)



Mark Tyson News Editor



ADVERTISEMENT



REPLY ▶

cirdecus

Mandark said:

Nothing like vertical integration. ALWAYS make your OWN stuff to control your destiny. Don't believe the fools that say it's cheaper to outsource it's not. Outsource companies need to make a profit too. And when you outsource you lose the ability to innovate

I love to see Tesla vertically integrate. It's one of the main reasons for their success

Completely agree. It does, however, tend to condense power which could mean less consumer choice, but vertical integration is where it's at. I never thought I'd see Amazon buying their own fleet of transport vehicles, planes and ships lol.

REPLY ▶


jtenorj

Noticed a mistake in your article. You state that Tesla's D1 chip has 50 million transistors when Telsa's slide for the chip clearly shows 50 billion. That's a difference of 3 orders of magnitude. 50 billion is also much more in line with the chip's size in mm squared as well as its 400w power draw on a modern fairly compact process node.

REPLY ▶

SHOW MORE COMMENTS ▶

Tom's Hardware is part of Future US Inc, an international media group and leading digital publisher. [Visit our corporate site.](#)



[Terms and conditions](#)

[Contact Future's experts](#)

[Privacy policy](#)

[Cookies policy](#)

[Accessibility Statement](#)

[Advertise with us](#)

[About us](#)

[Coupons](#)

[Careers](#)

[Do not sell or share my personal information](#)

© Future US, Inc. Full 7th Floor, 130 West 42nd Street, New York, NY 10036.

