


Yevgeniy Sverdlik



Nvidia's DGX-2 supercomputer on display at GTC 2018

DESIGN > MACHINE LEARNING

Is Your Data Center Ready for Machine Learning Hardware?

Ability to cool high-density GPU clusters for machine learning is a new battleground for colo providers.

Yevgeniy Sverdlik | Jan 31, 2019


So, you want to scale your computing muscle to train bigger deep learning models. Can your data center handle it?

According to Nvidia, which sells more of the specialized chips used in machine learning than any other company, it most likely cannot. These systems often consume so much power, a conventional data center doesn't have the capacity to remove the amount of heat they generate.

It's easy to see how customers without infrastructure that can support a piece of Nvidia hardware is a business problem for Nvidia. To widen this bottleneck for at least one of its product lines, the company now has a list of pre-approved colocation providers it will send you to if you need a place that will keep your supercomputers cool and happy.

As more companies' machine learning initiatives graduate from initial experimentation phases – during which their data scientists may have found cloud GPUs rented from the likes of Google or Microsoft sufficient – they start thinking about larger-scale models and investing in their own hardware their teams can share to train those models.

Among the go-to hardware choices for these purposes have been Nvidia's DGX-1 and DGX-2 supercomputers, which the company designed specifically with machine learning in mind. When a customer considers buying several of these systems for their data scientists, they often find that their facilities cannot support that level of power density and look to outsource the facilities part.



Nvidia's DGX-2 supercomputer on display at GTC 2018

“This program takes that challenge off their plate,” Tony Paikeday, who’s in charge of marketing for the DGX line at Nvidia, told Data Center Knowledge in an interview about the chipmaker’s new colocation referral program. “There’s definitely a lot of organizations that are starting to think about shared infrastructure” for machine learning. Deploying and managing this infrastructure falls to their IT leadership, he explained, and many of the IT leaders “are trying to proactively get ahead of their companies’ AI agendas.”

Cool Homes for Hot AI Hardware

DGX isn’t the only system companies use to train deep learning models. There are numerous choices out there, including servers by all

the major hardware vendors, powered by Nvidia's or AMD's GPUs. But because they all pack lots of GPUs in a single box – an HPE Apollo server has eight GPUs, for example, as does DGX-1, while DGX-2 has 16 GPUs – high power density is a constant across this category of hardware. This means that along with the rise of machine learning comes growing demand for high-density data centers .


The trend benefits specialist colocation providers like Colovore, Core Scientific, and ScaleMatrix, who designed their facilities for high density from the get-go. But other, more generalist data center providers are also capable of building areas within their facilities that can handle high density. Colovore, Core Scientific, and ScaleMatrix are on the list of colocation partners Nvidia will refer DGX customers to, but so are Aligned Energy, CyrusOne, Digital Realty Trust, EdgeConneX, Flexential, and Switch.

Partially owned by Digital Realty, Colovore built its facility in Santa Clara in 2014 specifically to take care of Silicon Valley's high-density data center needs . Today, it supports close to 1,000 DGX-1 and DGX-2 systems, Ben Coughlin, the company's CFO and co-founder, told us. He wouldn't say who owned the hardware, saying only that it belonged to fewer than 10 customers who were "mostly tech" companies. (Considering that the facility is only a five-minute drive from Nvidia headquarters, it's likely that the chipmaker itself is responsible for a big portion of that DGX footprint, but we haven't been able to confirm this.)

Colovore has already added one new customer because of Nvidia's referral program. A Bay Area healthcare startup using artificial intelligence is "deploying a number of DGX-1 systems to get up and running," Coughlin said.

A single DGX-1 draws 3kW in the space of three rack units, while a DGX-2 needs 10kW and takes up 10 rack units – that's 1kW per rack unit regardless of the model. Customers usually put between nine and 11 DGX-1s in a single rack, or up to three DGX-2s, Coughlin said. Pumping chilled water to the rear-door heat exchangers mounted on the cabinets, Colovore's passive cooling system (no fans on the doors) can cool up to 40kW, according to him.

In a "steady state," many of the cabinets draw 12kW to 15kW, "but when they go into some sort of workload state, when they're doing some processing, they'll spike 25 to 30 kilowatts," he said. "You can see swings on our UPSs of 400 to 500 kilowatts at that time across our infrastructure. It's pretty wild."



Cooling fans of the Nvidia DGX-2 system on display at GTC 2018

Echoing Nvidia's Paikeday, Chris Orlando, CEO and co-founder of ScaleMatrix, said typical customers that turn to his company's high-density colocation services in San Diego and Houston are well into their machine learning programs and looking at expanding and scaling the infrastructure that supports those programs.

A high-density specialist, ScaleMatrix's proprietary cooling design also brings chilled water directly to the IT cabinets. The company has "more than a handful of customers that have DGX boxes collocated today," Orlando told us.

High Density Air-Cooled

Flexential, which is part of Nvidia's referral program but doesn't have high-density colocation as its sole focus, uses traditional raised-floor air cooling for high density, adding doors at the ends of the cold aisles to isolate them from the rest of the building and "create a bathtub of cold air for the server intakes," Jason Carolan, the company's chief innovation officer, explained in an email.

According to him, this approach works fine for a 35kW rack of DGX systems. "We have next-generation cooling technologies that will take us beyond air, but to date, we haven't had a sizeable enough customer application that has required ... it on a large scale," he said. Five of Flexential's 41 data centers can cool high-density cabinets today.

As more and more companies use machine learning, it is becoming an important workload for data center providers to be able to support. Adoption of these computing techniques is only in its early phases, and they are likely to become an important growth driver for colocation companies going forward. Not many enterprises are set up to host supercomputers on-premises, and few are going to spend the money to build this infrastructure, so turning to colocation facilities that are already designed to efficiently cool tens of kilowatts per rack is their logical next step.

Source URL: <https://www.datacenterknowledge.com/machine-learning/your-data-center-ready-machine-learning-hardware>