



Verne Global

# Building a home for AI

Design for the future today, lest you get left behind, *Max Smolaks* warns



**Max Smolaks**  
Contributor

**D**ata center designers and builders have to stay on top of the latest developments in server hardware: the environments they create require a massive upfront investment and are expected to last at least 20 years, so they have to be ready for housing the IT equipment of the future.

The latest trend in IT workloads that's set to impact the way data centers are constructed is machine learning. The ideas fueling the boom in artificial intelligence are not new - many of them were proposed in the 1950s - and the power of AI is undoubtedly being over-hyped, but there are plenty of use cases where AI tech in its current state is already bringing tangible benefits.

For example, algorithms are much better than people at securing corporate networks, able to pick up on anomalies that humans and their rules-based tools might miss. Algorithms are also great at analyzing large chunks of

boring text: lawyers use AI-based software to scan through case files and contracts, while universities use something similar to establish whether a paper was written by the student who submitted it, or a freelancer hired online. And finally, there's plenty of research to suggest that optical image recognition will match and even surpass the best human doctors at spotting signs of disease on radiology scans.

The variety of machine learning applications is only going to increase, introducing radically different demands on storage performance, network bandwidth and compute, more akin to something seen in the world of supercomputers.

According to a recent research report by Tractica, an analyst firm focused on emerging technologies, AI hardware sales to cloud and enterprise data centers will see compound annual growth rate (CAGR) of almost 35 percent for the next six years, increasing from around \$6.2 billion in 2018 to \$50 billion by

2025. All this gear will have higher power densities than traditional IT, and as a result, higher cooling requirements.

Up until now, most machine learning tasks have been performed on banks of GPUs, plugged into as few motherboards as possible. Every GPU has thousands of tiny cores that need to be supplied with power, and on average, a GPU will need more power than a CPU. Nvidia, the world's largest supplier of graphics chips, has defined the current state of the art in ML hardware with DGX-2, a 10U box for algorithm training that contains 16 Volta V100 GPUs along with two Intel Xeon Platinum CPUs and 30TB of flash storage. DGX-2 delivers up to two petaflops of compute, and consumes a whopping 10kW of power - more than an entire 42U rack of traditional servers running at an average load.

---

*"If you don't support water-cooled processors, you're excluding yourself from the top end of the market"*

---

There are two parts to almost any machine learning project: training and inference. Inference is easy, you just take a fully developed machine learning model and apply it to whatever data you want to manipulate. This process can run facial recognition on a smartphone, for example. It's the training

that's the intensive part - getting the model to look at thousands of faces to learn what a nose should look like.

"There are large amounts of compute used and the training can take days, weeks, sometimes even months, depending on the size of the [neural] network, and the diversity of the data feeding into it and the complexity of the task that you're trying to train the network for," Bob Fletcher, VP of strategy at Verne Global, told *DCD*. Verne runs a data center campus in Iceland which was originally dedicated to industrial-scale HPC, but has recently embraced AI workloads. The company has also experimented with blockchain, but who hasn't done that?

According to Fletcher, AI training is proving to be much more compute-intensive than traditional HPC workloads that the company is used to, like computational fluid dynamics. "We have people like DeepL [a machine translation service] in our data center that are running tens of racks full of GPU-powered servers; they are running them at 95 percent, the whole time. It sounds like a jet engine and they have been running like that for a couple of years now," he said. What many people don't realize is the fact that machine learning models used in production have to be re-trained and updated all the time to maintain accuracy - creating a sustained need for compute, rather than a one-off spike.

When housing machine learning equipment, it is absolutely necessary to be religious about aisle containment and things like blanking plates. "You go from 3-5kW to around 10kW per rack for regular HPC, and if you are going into AI workloads, which are generally going to be around 15-35kW air-cooled, then you have to be more careful about air handling," Fletcher explained.

"If you look at something like DGX-2, it blows out 11-12kW of warm air. And it has a very low temperature spread between input and output, so the airflow is quite fast. If you are not thoughtful about positioning, and you have two of these at the same height, pointing back to back and maybe 30 inches between them, they are going to blow hot air at each other like crazy, and you lose the airflow.

"So you have to place them at different heights, or you have to use air deflectors, or you can spread the aisles further apart, but whatever you do, you've got to make sure that all of the hot air that's coming out of one of these devices isn't going to be banging into hot air coming out of another."

The advent of AI hardware in the data center could at last mark the moment when water cooling finally becomes necessary. "One of the real changes from a data center operator's perspective is direct water-cooled chips to support some of these applications. These products were once on roadmaps, and they are now mainstream - the reason

they are available is because of some of these workloads around AI," said Paul Finch, CEO at Kao Data.

Kao has just opened a wholesale colocation campus near London, inspired by hyperscale designs, with 8.8MW of power capacity available in Phase 1, and 35MW in total. The project was developed with an eye to the future - and according to Finch, that future includes lots of AI. He says that data halls at Kao have been designed to support up to 70kW per rack.

"Many of these new processors, the real Ferraris of chip technology, are now moving

---

*"It sounds like a jet engine and they have been running like that for a couple of years now"*

---

towards being water-cooled. If your data center is not able to support water-cooled processors, you are actually going to be excluding yourself from the top end of the market," Finch told *DCD*.

"In terms of the data center architecture, we have higher floor-to-ceiling heights - obviously water is far heavier than air, so it's not just about the weight of IT, it's going to be about the weight of the water used to cool the IT systems. All of that has to be factored into the building structure and floor loading. We see immersion cooling as a viable alternative, it just comes with some different challenges."

Floors are not the most glamorous data center component, but even floors have to be

considered when housing AI gear, since high density racks are also heavier. Like many of the recent data center projects, Kao houses servers on a concrete slab. This makes it suitable for hyperscale-style pre-populated racks like those built by the fans of the Open Compute Project, and enables the facility to support heights of up to 58U.

And in terms of networking, Fletcher said AI requires more expensive InfiniBand connectivity between the servers - classic Ethernet simply doesn't have enough bandwidth to support clusters with dozens of GPUs. Cables between servers in a cluster need to be kept as short as possible, too. "Not only are you looking at cooling constraints, you are looking at networking and connectivity constraints to keep the performance as high as possible," he said.

In the next few years, hardware for machine learning workloads will get a lot more diverse, with all kinds of novel AI accelerators battling it out in the marketplace, all invariably introducing their own acronyms to signify a new class of devices: these include Graphcore's Intelligence Processing Units (IPUs), Intel's Nervana Neural Network Processors (NNPs) and Xilinx's Adaptive Compute Acceleration Platforms (ACAPs). Google has its own Tensor Processing Units (TPUs), but these are only available in the cloud. American startup Cerebras Systems recently unveiled the Wafer Scale Engine (WSE) - a single chip that measures 8.5 by 8.5 inches and features 400,000 cores, all optimized for deep learning.

It's not entirely clear just how this monster will fit into a standard rack arrangement, but it serves as a great example of just how weird AI hardware of the future might become.

Sean Lie, Cerebras

