www.pipelinepub.com
Volume 20, Issue 7

The Race to Build AI-optimized Data Center Networks

By: Aniket Khosla

It's hard to believe that less than two years ago most of the tech industry considered artificial intelligence (AI) a niche topic. A fascinating and important topic, but one discussed in mostly future-focused terms. Fast forward to mid-2024, and the future is now. Practically every solution area in tech now includes a strategy for incorporating generative AI (GenAI) or other AI-enabled applications. And businesses in every industry are racing to put AI to work — automating operations, lowering costs, improving quality, and more.

All this enthusiasm can make it seem like the transformative potential of AI is practically limitless. Inside the world's largest data centers, however, where AI application and training workloads run, the excitement is tempered by questions about how, exactly, operators can keep up with exploding demand. As AI workloads grow larger and more complex, how many processors and accelerators will AI clusters require, and how can data center fabrics best interconnect them? How will data center architectures meet the exacting throughput, latency, and lossless transmission requirements of real-time AI applications? What types of interfaces are best suited for front- and back-end infrastructures, and when should operators expect to need next-generation interfaces?

Many of these questions don't yet have firm answers. However, by examining the requirements for processing AI application workloads and the ways these requirements are influencing data center network designs, we can start to get a clearer picture.

AI Puts the Squeeze on Data Center Architectures

According to Dell'Oro Group, global AI application traffic is increasing by a factor of 10x every two years. Hyperscale cloud providers are already approaching terabit networking thresholds for AI workloads, and other data center operators aren't far behind. As the effects of AI adoption ripple throughout the communications industry, hyperscalers are adding bandwidth and compute infrastructure as quickly as vendors can deliver it.

There's a reason for this accelerating infrastructure investment: AI applications bring requirements on a totally different scale than traditional data center workloads. (Figure 1) These requirements are growing for both AI model training (when an application ingests vast amounts of data to train its algorithm on) and inferencing (when the AI model puts this training to work on new data).

The scale of the networking fabric needed for a given AI cluster depends on the size and complexity of the applications that infrastructure supports. But with AI models growing 1,000x more complex every three years, we should expect that in the near future, clusters will need to support models with trillions of dense parameters. In practical terms, this will require a fabric connecting thousands, even tens of thousands of central processing units (CPUs), graphics processing units (GPUs), field programmable gate arrays (FPGAs), and other "xPU" accelerators.

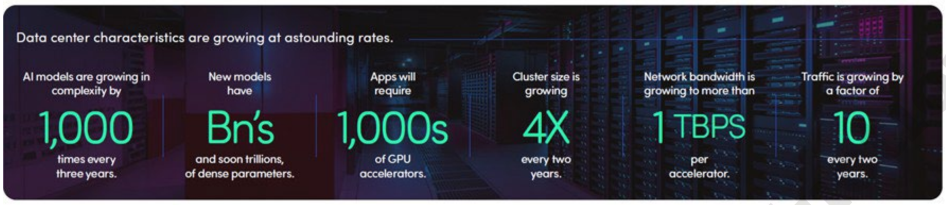

Figure 1. AI impact on data centers, by the numbers. (Source: Dell'Oro Group)

These requirements won't have much immediate effect on data center front-end access networks used to ingest data for AI algorithm training. Back-end infrastructures, however, are another matter. Here, operators can't meet exploding AI workload demands by simply throwing more hardware at the problem. They need separate, scalable, routable back-end infrastructures designed explicitly to interconnect xPUs for AI training and inferencing.

Evolving Infrastructure Requirements

New back-end infrastructures needed to support AI applications have very different, far more demanding performance requirements than traditional data center networks. To avoid bottlenecks in data transfers among compute and accelerator nodes, back-end infrastructures must support 5x more traffic and 5x higher network bandwidth per accelerator than typical front-end access networks. These networks also need to be able to handle thousands of synchronized jobs in parallel, along with far more compute- and data-intensive workloads. They must support new traffic patterns such as micro data bursts. And they must provide consistent high-throughput, low-latency connectivity between servers, storage, and xPUs. Network latency is among the most

AI application workloads running in the data center are also highly sensitive to packet loss. Missing or out-of-order packets can dramatically increase latency due to network buffering and retransmission. Many AI operations, such as model training, can't even complete until all packets are received, and packet latencies of even 0.1 percent can cause extreme delays that disrupt mission-critical AI applications. Given these factors, back-end AI fabrics must also be lossless.

Finally, model training and other complex, distributed AI applications require the ability to scale efficiently to trillions of parameters and beyond. Traditional fabric architectures using Compute Express Link (CXL), Nvidia's NVLink, or PCI Express (PCIe) can meet the connectivity needs of smaller workloads. As operators look to support larger AI applications, however, they will need to consider scale-out AI leaf and spine architectures. Moderate "rack-scale" applications will require thousands of xPUs connected via an AI leaf layer. Large AI applications employing tens of thousands of accelerators will need data center-scale architectures with a routable fabric and AI spine.

Surveying Interface Trends

There's an easy way to gauge the speed at which AI applications are coming to dominate modern data centers: follow the interfaces. As AI models grow more complex, they require more bandwidth per accelerator. We can watch the market responding to this pressure in real time, as data center operators increasingly adopt next-generation interface speeds, moving from 400 to 800 Gbps, and soon, 1.6 Tbps. (Figure 2)

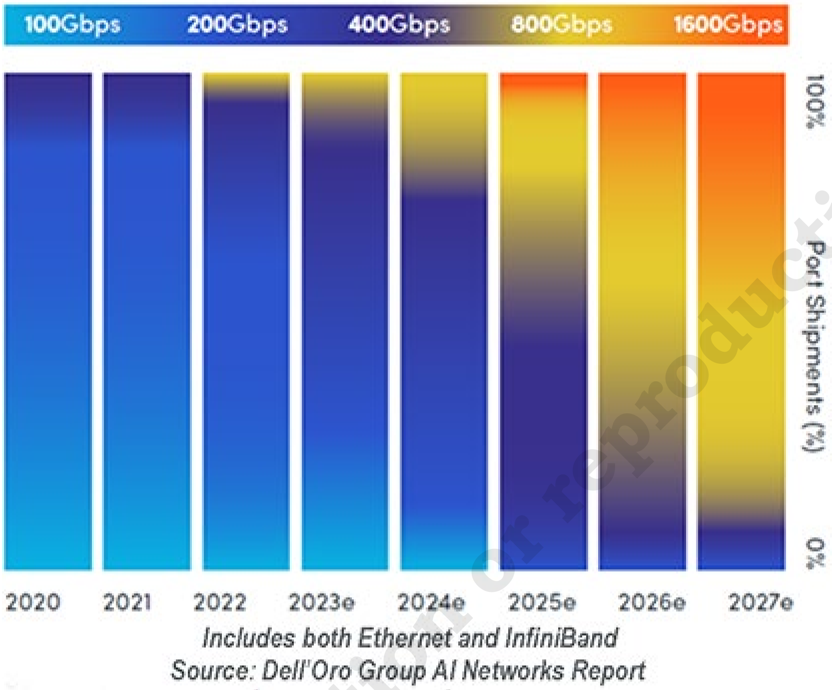

Figure 2. Accelerating adoption of high-speed interfaces.

The specific interface technologies that AI clusters will use remains an open question. Many operators prefer Ethernet when possible, as it's a widely adopted, standardized, and cost-effective option. However, Ethernet is lossy by design, using priority-based queuing and pausing to handle congested links. Operators will therefore need to deploy it in conjunction with protocols like RDMA over Converged Ethernet, version 2 (RoCEv2) — or in the future, Ultra Ethernet Transport (UET) — to deliver the performance needed for applications like AI training. Alternately, proprietary InfiniBand can provide a lossless medium with more deterministic flow control. Where will the market ultimately settle? According to a recent Dell'Oro Group forecast, the answers differ for front- and back-end data center networks. The analyst expects that network ports providing front-end connectivity to AI clusters will remain Ethernet, with data ingestion requirements initially driving the transition to next-generation speeds. By 2027, Dell'Oro expects that one third of front-end Ethernet ports will be 800G or higher.

Back-end infrastructures will evolve more quickly, with operators adopting next-generation speeds at a triple-digit compound annual growth rate (CAGR). By 2027, nearly all back-end network ports will be 800G or higher. Here though, where lossless transmission is essential, Dell'Oro forecasts that interface technologies will remain mixed, with Ethernet and InfiniBand coexisting for the foreseeable future.

Weighing Infrastructure Options

Given the pace at which customers are adopting AI applications, and the lead time needed to build new data centers, operators have little choice but to move forward with new fabric strategies now. However, while uncertainty remains regarding the best approaches to support future AI workloads, we can draw one big conclusion now: one size will *not* fit all. Different operators will follow different paths depending on a variety of factors unique to their business and technology strategies.

The size of a given deployment, the number of clusters it will support, and of course, price, will all influence infrastructure decisions. Yet operators will also need to consider many other factors. What kinds of AI applications and workloads does the operator plan to focus on? For example, will they cater to compute- and time-intensive model-training or outsource that phase of AI applications? How complex do they expect their customers' applications and workloads to be? What will those workloads' bandwidth and load-balancing requirements look like, and how important will low and deterministic latency be for their processing? Where does a given operator fall on the question of standardized versus proprietary interface technologies? How important is it to maintain a diversified multivendor supply chain? And how comfortable are they with the future roadmaps of the technologies they're considering?

Many of these questions don't have objectively correct answers. And beyond the variability of individual data center strategies, AI applications themselves, along with the technologies supporting them, will continue to evolve. So, even the best-laid plans of today will likely change considerably over the next several years and beyond.

Looking Ahead

Even with some questions still unanswered, data center operators have little choice but to push ahead with new network infrastructures. The market for AI applications has grown so huge, so quickly, that it would be foolish to do anything else. To actually deliver on the customer expectations that accompany exploding AI demand, however, rigorous testing and validation becomes absolutely essential.

Data center operators and their vendors must be able to validate next-generation Ethernet products, assure multi-vendor interoperability, and test timing and synchronization with a precision that legacy tools simply can't support. Vendors will also need the ability to emulate the unique network behavior and traffic patterns in AI clusters with lifelike accuracy. Fortunately, a new generation of testing and emulation solutions designed for the speeds and scale of AI network infrastructures has already emerged to reflect these changing requirements. With a bit of foresight and informed decision-making from data center operators and vendors, the transformative potential of AI applications really can be practically limitless.