Dell PowerEdge R7615 servers with Broadcom BCM57508 NICs can accelerate your AI fine-tuning tasks

A cluster of Dell™ PowerEdge™ R7615 servers featuring AMD EPYC processors achieved much stronger performance on multi-GPU, multi-node operations using Broadcom 100GbE NICs than the same cluster using 10GbE NICs

LLM training and inference frameworks deployed on distributed GPUs use low-level algorithms to move data between GPUs, operate on that data, and share the results with other GPUs. Our testing focused on three of these algorithms as implemented in the NVIDIA Collective Communications Library (NCCL) library: all-reduce, reduce-scatter, and send-receive. This library, which many AI frameworks use, can send data over RoCE network paths or ordinary Ethernet network paths, and can perform RDMA transfers between distributed NVIDIA GPUs.

We tested a two-node cluster of Dell PowerEdge R7615 servers with AMD EPYC™ 9374F processors and NVIDIA® L40 GPUs with two networking configurations:

- one with Broadcom® 100GbE BCM57508 NetXtreme-E network interface cards (NICs) with remote direct memory access (RDMA) over Ethernet (RoCE) support
- one with 10GbE NICs

For each configuration, we studied three multi-GPU, multi-node AI computations from the NCCL test suite at different packet sizes and measured the time to complete the task, latency, and the effective bandwidth of the network during the operation. The cluster with 100GbE networking dramatically outperformed the cluster with 10GbE networking across all packet sizes and tasks without increasing power usage.

Please note that these tests do not send enough data between servers to overwhelm the networking link. Rather, these tests comprise a sequence of computational steps on each GPU, where a given step may require data from other GPUs. In such cases, a GPU can only start the next computational step once it has the data from those other GPUs, even if that data is as small as a single byte. The operational bandwidth depends on the timely transfer of data between GPUs on different servers.

Up to 83% less time to complete multi-GPU, multi-node operations*

Send-receive performance: Time to complete task

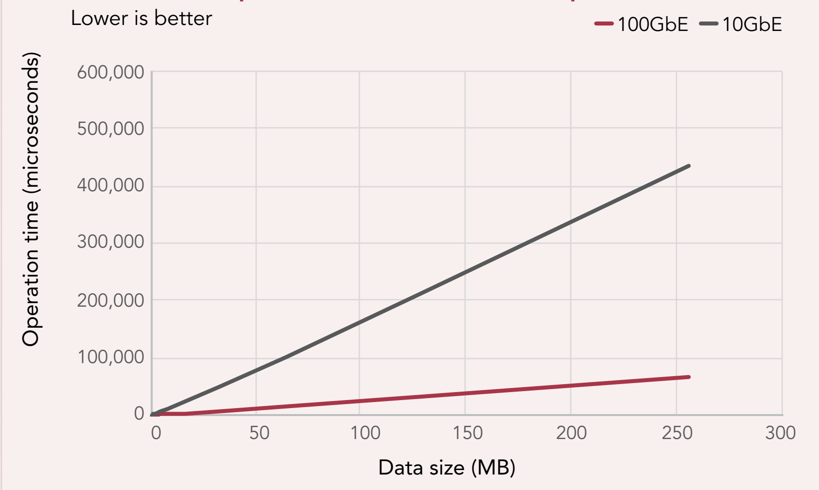

Up to 66% lower latency on multi-GPU, multi-node operations*

Multi-GPU, multi-node operation	Latency (microseconds) Lower is better		Percentage reduction Higher is better
	100GbE configuration	10GbE configuration	
all-reduce (packet size: 4 B)	40	123	67.4%
reduce-scatter (packet size: 4 B)	29	85	65.8%
send-receive (packet size: 48 B)	41	56	26.7%

Up to 6.1x the bandwidth on multi-GPU, multi-node operations*

Send-receive bandwidth

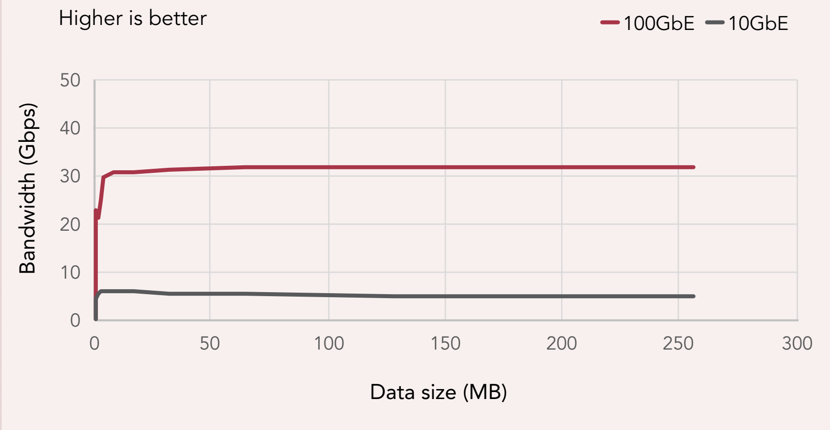

*cluster of Dell PowerEdge R7615 servers featuring AMD EPYC 9374F processors and Broadcom 100GbE BCM57508 NetXtreme-E NICs vs. the same cluster with 10GbE NICs.

The three multi-GPU, multi-node NCCL primitive operations for AI we used for testing are:

- **all-reduce**: Operate on the entire dataset, distribute across all GPUs in the cluster, and store the single result on each GPU
- **reduce-scatter**: Divide the data on every GPU into logical chunks, and operate on each chunk across the cluster to form partial results. Then send one partial result to each GPU and store it there
- **send-receive**: Send data from one GPU to another on the second server, and return a response

For full testing details and results, read our full report.

Learn more at https://facts.pt/QAauY1Y

Principled Technologies®