



CONFERENCE

—NORTH AMERICA 2026—

XCCL: Scaling PyTorch Collectives to Exascale on Intel GPUs with TorchComms

Panagiotis Kourdis and Tanima Dey
Intel Corporation

XCCL: PyTorch Distributed for XPU


- **Native PyTorch Distributed support** on Intel GPUs, scaling seamlessly from a single node to exascale-class AI infrastructure
- **Broad LLM ecosystem support** spanning TorchTitan pre-training, fine-tuning workflows, and vLLM/SGLang production inference
- **NCCL-like user experience** for PyTorch distributed applications
- **Proven AI-for-Science impact** on the Aurora Supercomputer, a US DOE leadership-class exascale system powering ACM Gordon Bell Prize finalist-scale research
- **Production-ready TorchComms foundation** for current and next-generation Intel GPU architectures

XCCL: Bringing Intel XPU Support to TorchComms

- **oneCCL Integration:** Leverages Intel® oneCCL as the native communication engine through a lightweight C API layer
- **Collective Operations:** Delivers full collective communication support across scale-up and scale-out deployments
- **Upstream Enablement:** Established CI coverage and validation for the XCCL backend on Intel hardware
- **Platform Validation:** Validated on both Intel client and data center GPU platforms (Intel® Arc Pro B* Series and Intel® Data Center Max 1550 Series)
- **Training Validation:** Demonstrated end-to-end distributed training with high scaling efficiency and proven convergence

Aurora: Exascale Distributed Training on Intel GPUs

One of the world's first **exascale supercomputers**, serving as a proving ground for **TorchComms** and **XCCL** to enable **scalable PyTorch distributed training** on **Intel GPUs**.




Nodes	10,624
No. of CPUs	21,248
No. of GPUs	63,744
DDR5 Memory Capacity	10.62 PB
DDR5 Memory Bandwidth	5.31 PB/s
HBM2e Memory Capacity	9.52 PB
HBM2e Memory Bandwidth	147.46 PB/s
Injection Bandwidth	2.12 PB/s
Global Bandwidth	1.37 PB/s

Benchmark	Performance	Rank	Nodes
HPL	1.012 EF/s	#3 at SC'24	9234
HPL MxP	11.64 EF/s	#1 at SC'24	9500
HPCG	5.613 PF/s	#3 at ISC'24	4096
Graph500	24250 GTEPS	#6 at ISC'24	4096
Graph500	69373 GTEPS	#6 at ISC'24	8192
IO500	32,165.90	#1 at SC'23	300





DDP: CosmicTagger Scaling Results

- **Open-Source Model:** Image segmentation model developed by Argonne National Laboratory
- **Distributed Training Enablement:** Enabled PyTorch DDP training using the TorchComms communication backend
- **Proven Convergence:** Demonstrated convergence with both **FP32** and **BF16** training precision
- **Aurora Validation:** Evaluated end-to-end performance on the Aurora Supercomputer
- **Extreme-Scale Training:** Scaled to **2,560 nodes (30,720 ranks)**, demonstrating **large-scale distributed training**



TorchTitan: Llama3-8B Results

- **Goal:** Functionality and performance comparison with default C10D based implementation
- Llama3-8B model using the base FSDP2-configuration (with no parallelism)
- Convergence and performance evaluation use the same C4-dataset
- TorchComms shows **same functionality and performance parity** as default C10d-based distributed backend on XPU device



- Execution is configured to exercise eight GPUs of a node
- Scaling efficiency is calculated as performance on N nodes, normalized w.r.to 1 node
- The 1-node throughput for both backends have the same parity
- Higher numbers represent better scaling, variations within 5% difference is expected

XCCL at a Glance

XCCL is a first-class, feature-complete torchcomms backend that brings the entire collective API — P2P, collectives, scatter/gather, splitting — to Intel XPU hardware by wrapping Intel's oneCCL library behind the same TorchCommBackend interface NCCL and Gloo implement.

- **Roadmap:** Enabling advanced communication capabilities, including **RMA (Remote Memory Access)** and **fault tolerance** support.
- **Upstream Integration:** Migrating **TorchComms-XCCL** into upstream PyTorch as the next-generation **XCCL2 backend** within `torch.distributed`.
- Deep dive on **XCCL** and **Aurora** at **October 21, 5:30 PM**, “*Scaling PyTorch on Intel XPUs: Field Notes from AI for Science on Aurora*”.
- Learn More: Scan the QR codes to explore the **Aurora website** and read the accompanying **technical paper**.



Q/A?