



KubeCon



CloudNativeCon

North America 2026

#KubeCon #CloudNativeCon

Topograph - Your AI Cluster's Missing Live Transit Map

[_Dmitry Shmulevich @NVIDIA](#)



Why Topology Matters



KubeCon



CloudNativeCon

North America 2026




Multi-node jobs can span accelerated domain, rack, and switch boundaries.



Placement dictates efficiency and cost in distributed AI workloads.



Modern workload managers support topology aware scheduling, but rely on topology data



Source of Network Topology



KubeCon



CloudNativeCon

North America 2026



Node Labels and Annotations

`cloud.google.com/gce-topology-block`
`cloud.google.com/gce-topology-subblock`
`cloud.google.com/gce-topology-host`



APIs

- [DescribeInstanceTopology](#) API



Protocols and Tools

- LLDP
- ibnetdiscover



Problem: Data Heterogeneity

- Schedulers cannot directly consume all data sources

Consuming Topology



KubeCon



CloudNativeCon

North America 2026



Kubernetes Schedulers: node labels, device attributes

- [KEP-4671](#): Gang Scheduling using Workload Object
- [KEP-5729](#): DRA: ResourceClaim Support for Workloads
- [KEP-5732](#): Topology-aware workload scheduling
- [KEP-2724](#): Topology Aware Scheduling (Kueue)
- [KEP-6012](#): Composite PodGroup API
- [KEP-2549](#): NFD network topology updater



Slinky: node annotations, ConfigMaps

- Slinky Operator [Topology Guide](#)
- Slinky Bridge [Topology Config](#)



Solution: Topograph

- Discover cluster network topology
- Publish in the format scheduler understand

What Topograph does



KubeCon




CloudNativeCon

North America 2026

- **Discovers cluster topology**
Supports both cloud and on-premises environments.
- **Normalizes graph data**
Creates a robust, provider-independent canonical graph.
- **Translates output structure**
Compiles the topology into scheduler-specific outputs.
- **Dynamic regeneration**
Triggers instant pipeline updates upon cluster changes.

Key Principle

Providers discover. Engines translate.



Architectural Overview




KubeCon



CloudNativeCon

North America 2026



Topology sources

Cloud Operators

AWS, GCP, OCI, Crusoe, Nebius, Akamai,

Cloud Operators (cont.)

Nscale, Lambda, Mistral, Reflection AI, Firebird

On-Prem / Fabric

InfiniBand, SpectrumX

Accelerator domains

GPU operator, NVML

Testing & Simulation

KinD, Tilt, KWOK

Roadmap: from Nodes to Devices



KubeCon



CloudNativeCon

North America 2026

- Expand support for cloud providers and switch vendors
- Extend device-level topology discovery
 - Map GPU-to-NIC paths across network rails
 - Discover accelerator cliques and track membership
- Work with the community to improve topology-aware scheduling
- Standardize how complex topologies are represented



Thank you!



 github.com/dsx-ai-factory/topograph

 [#topograph in DSX AI Factory Slack](#)

