

# Intel Omni-Path

2020

Trusted partner for your Digital Journey

© Atos - For internal use

**Atos**

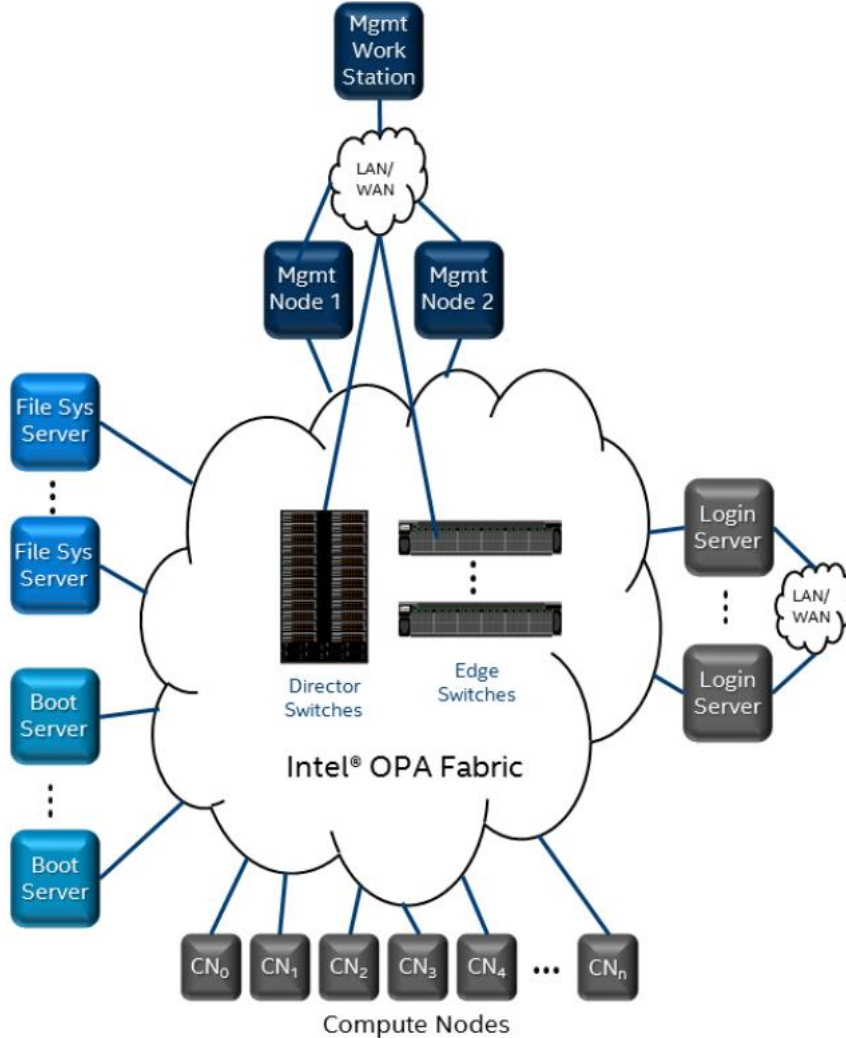
# Intel Omni-Path Architecture (OPA)

- ▶ Latency  $\sim 1\mu\text{s}$
- ▶ Bandwidth 100Gbps
- ▶ Supported MPI libraries:
  - Intel MPI, OpenMPI, MVAPICH2, ..
  - Based on open-source libpsm2
- ▶ Connected to NUMA1 (cores 24-47)
- ▶ Network also used for HOME and SCRATCH



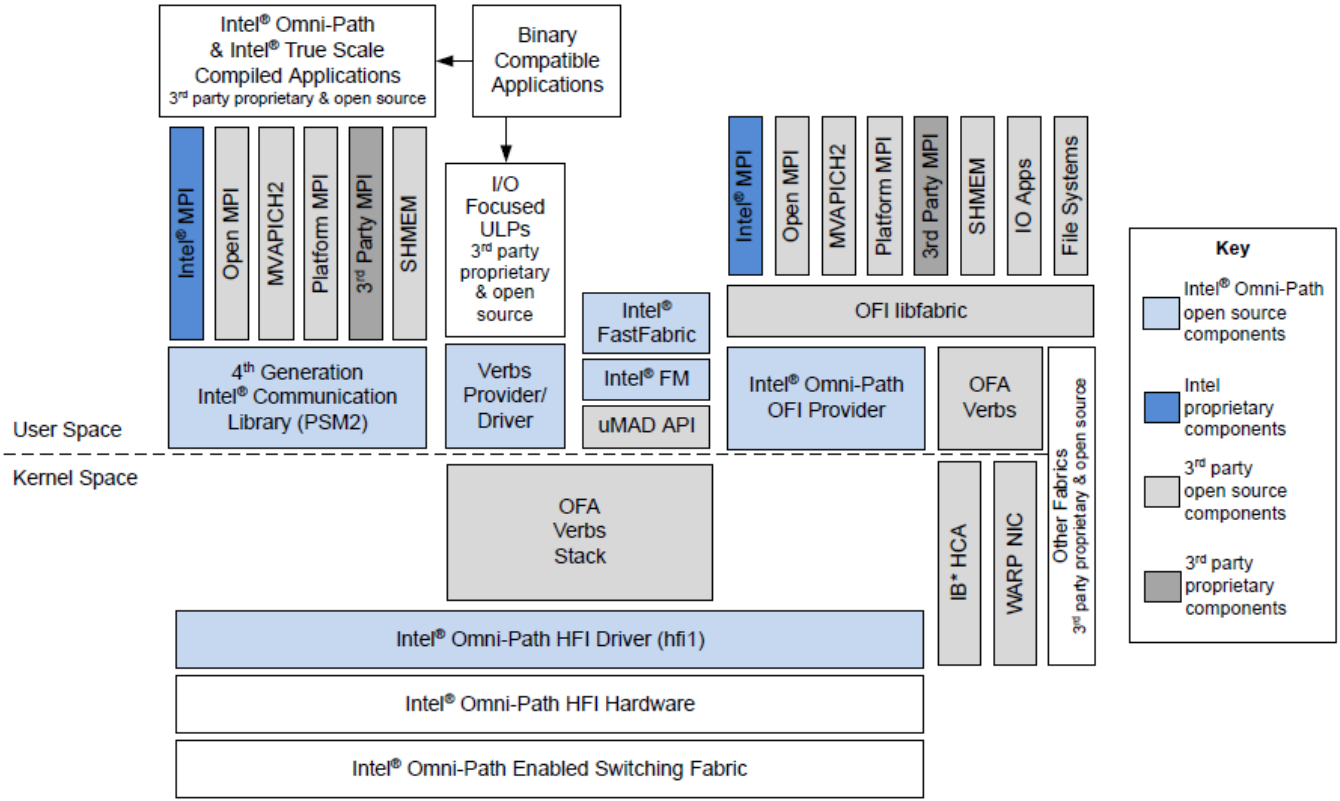
# Intel OPA Fabric

- ▶ All nodes connected via OPA
- ▶ 1 Director switch
  - 480 ports
  - 18x leaf, 9x spine




Picture from Intel Omni-Path Fabric Host software User Guide

# Intel Omni-Path Host software stack



# Node

---



# Chassis

---



# Rack

---







# System

---



# SLURM

---

- ▶ Automatically chooses nodes close together in the fabric
- ▶ Therefore, jobs run usually on either odd or even numbered nodes.

# Dispersive routing

---

- ▶ OPA Fabric Manager configures 4 routes between each set of 2 nodes
- ▶ Packets are spread over these 4 routes, to better utilize the network