

**CORNELIS**<sup>TM</sup>  
NETWORKS

# Update on Omni-Path Express Support for DAOS

Paul Stasurak, Program Manager for Storage

Brian Wilson, Principal Engineer

DUG'21 – November 19, 2021

# Notices and Disclaimers

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH CORNELIS NETWORKS PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN CORNELIS NETWORKS'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, CORNELIS NETWORKS ASSUMES NO LIABILITY WHATSOEVER, AND CORNELIS NETWORKS DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF CORNELIS NETWORKS PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

CORNELIS NETWORKS PRODUCTS ARE NOT INTENDED FOR USE IN MEDICAL, LIFE SAVING, OR LIFE SUSTAINING APPLICATIONS.

Cornelis Networks may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Cornelis Networks reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

All products, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice. Roadmap not reflective of exact launch granularity and timing. The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Any code names featured are used internally within Cornelis Networks to identify products that are in development and not yet publicly announced for release. Customers, licensees and other third parties are not authorized by Cornelis Networks to use code names in advertising, promotion or marketing of any product or services and any such use of Cornelis Networks' internal code names is at the sole risk of the user.

All products, computer systems, dates and figures specified are preliminary based on current expectations and are subject to change without notice. Material in this presentation is intended as product positioning and not approved end user messaging.

Performance tests, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

Cornelis Networks technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration.

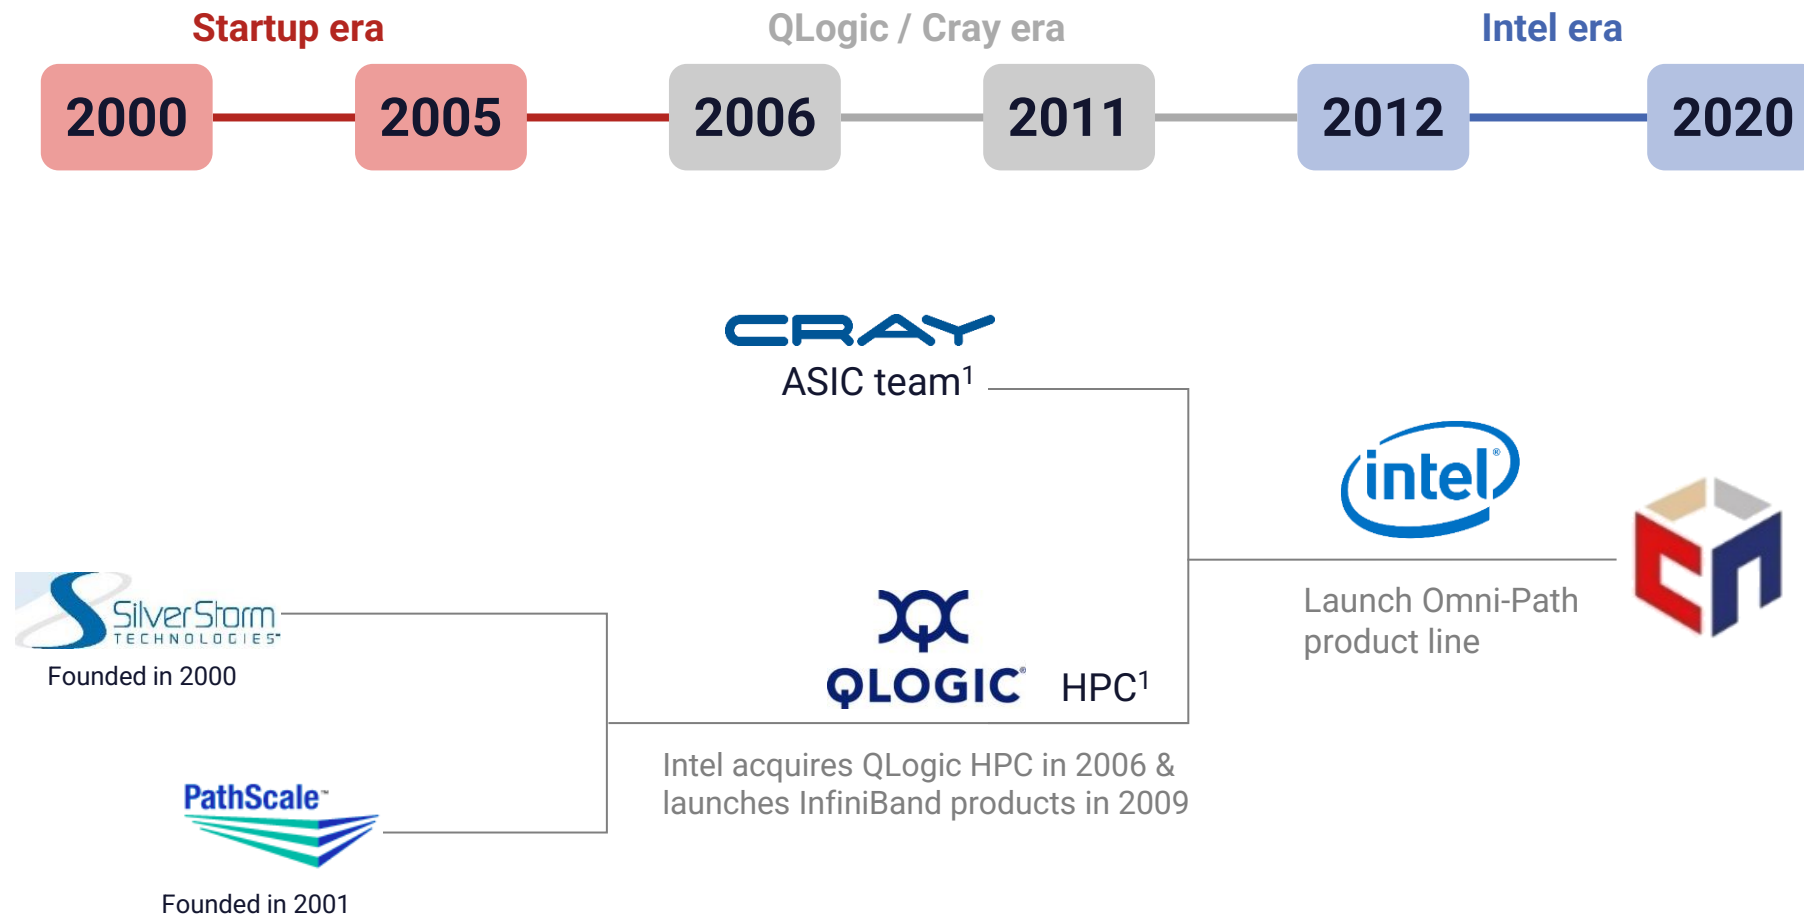
Copyright © 2021, Cornelis Networks

# Agenda

- Introduction to Cornelis Networks
- Overview of Omni-Path Express
- DAOS on Omni-Path Express
- DAOS Support Availability

# Introduction to Cornelis Networks

## Proven Technology Supplying Tier-1 Customers since 2000



*Technology built on investment over 20 years*

### Cornelis Networks era



- ✓ Acquired Intel Omni-Path business
- ✓ Complete networking fabric solutions
- ✓ Supporting Intel's 500 existing OPA customers
- ✓ Funding from tier-1 investors
- ✓ Strong ecosystem support

# Overview of Omni-Path Express



- What Are We Doing?
  - Cornelis is adding an Omni-Path Express (OPX) native provider to the OpenFabric Interfaces (OFI) libfabric software stack
- Why?
  - We recognize and want to take advantage advancements made in OFI libfabric that address the rapidly changing application environment driving greater interconnect performance requirements
  - Omni-Path Express is leveraging libfabric to access significant industry contributions across all key applications
- Benefits
  - Omni-Path Express can take advantage of continuous improvements available through OFI community on the hardware best suited to take advantage of the OFI software stack
  - Message rate, performance, and reduced latency versus current PSM2 implementation are significant
  - Broadens support for middleware applications optimized to run on OFI libfabric
  - Ideally suited to support DAOS natively for OFI communications

# Cornelis Networks OPA Roadmap

- Continued commitment to customers and HPC/HPDA/AI ecosystem
- Delivering industry leading **application performance** and unmatched **price/performance**
- Embrace OpenFabrics Interfaces industry initiative

1Q'22

2025

Today



**CN-OPA100**  
PCIeV3.0/4.0  
48p Edge Switches  
288P & 1152p Directors

## Omni-Path Express

- Dramatic performance leadership extension
- Expanded MPI support
- PGAS (SHMEM and Chapel)
- GPUs (Nvidia/AMD/Intel)
- AI Frameworks (TF, PyTorch)
- Object Storage System (DAOS)

2023



**CN-OPA400**  
PCIeV5.0  
High-density Edge Switches  
High-density Directors



**CN-OPA800**  
PCIeV6.0  
High-density Edge Switches  
High-density Directors

200 Gbps

100 Gbps (native)

400 Gbps (native)

200 Gbps

800 Gbps (native)

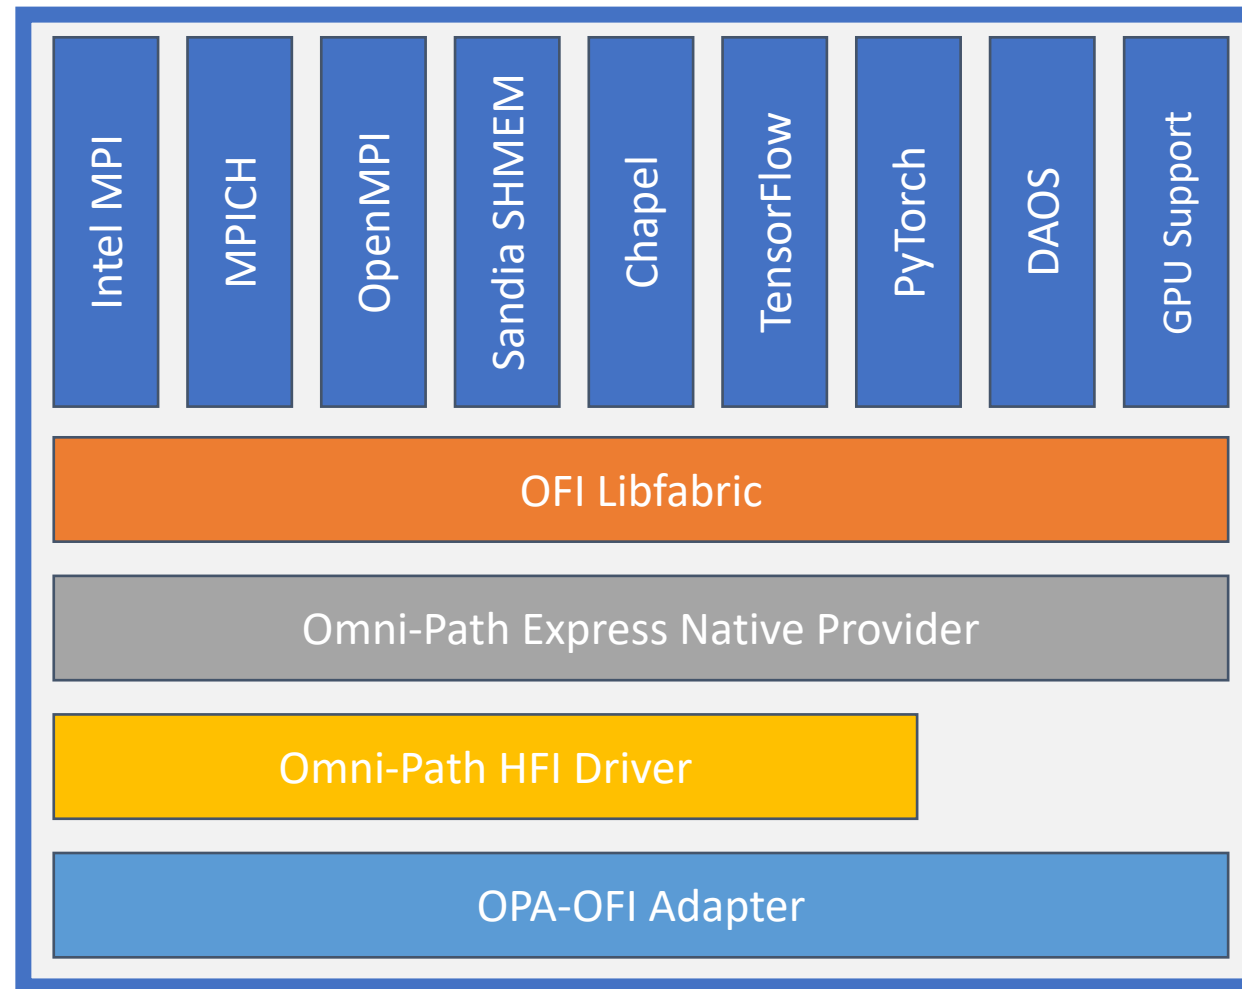
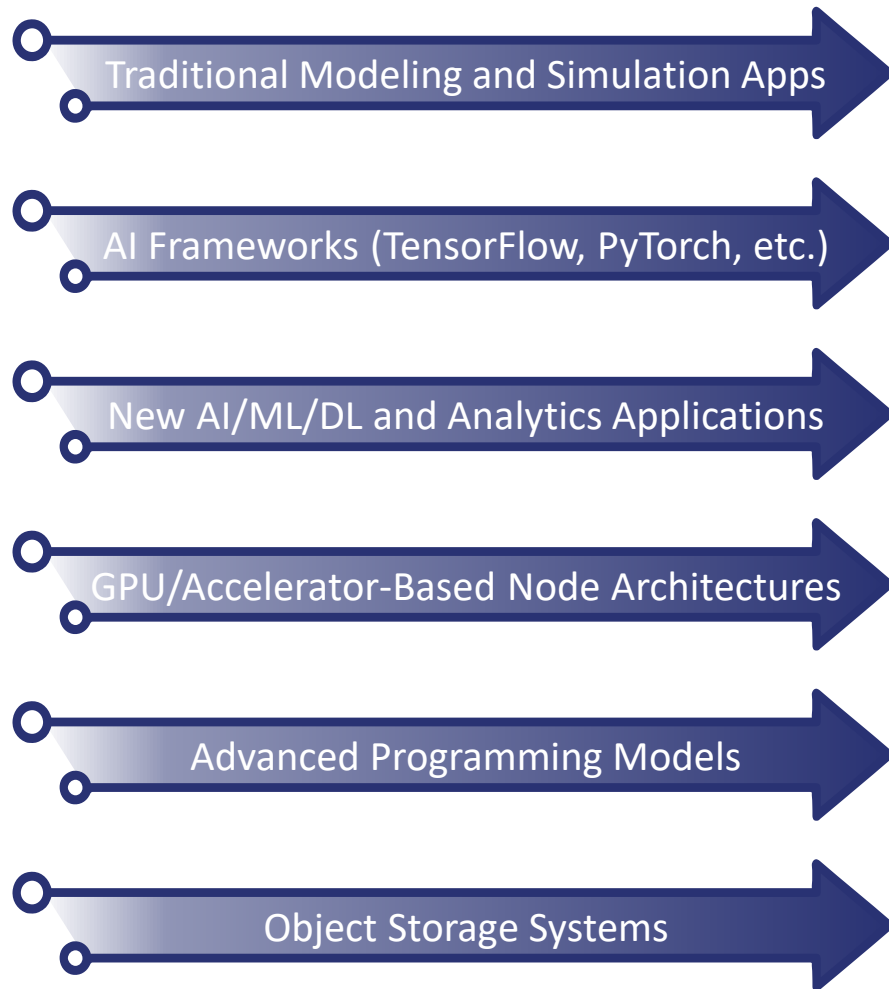
400 Gbps

Increasing Bandwidth and Network Intelligence to target HPC, AI and HPDA workloads

# Omni-Path Express: Leveraging OpenFabrics Interfaces Today's Performance and Price/Performance Leader



**OpenFabrics Interfaces:** Next generation software stack for high performance converged infrastructures



**Delivering unmatched performance, optimizing message rate, latency, overlap, and scalability for today's converged infrastructures**

**Industry-leading performance and extremely compelling price/performance leveraging today's proven hardware components**

**Evolving applications and host architectures demand a revolutionary approach to host interfacing**

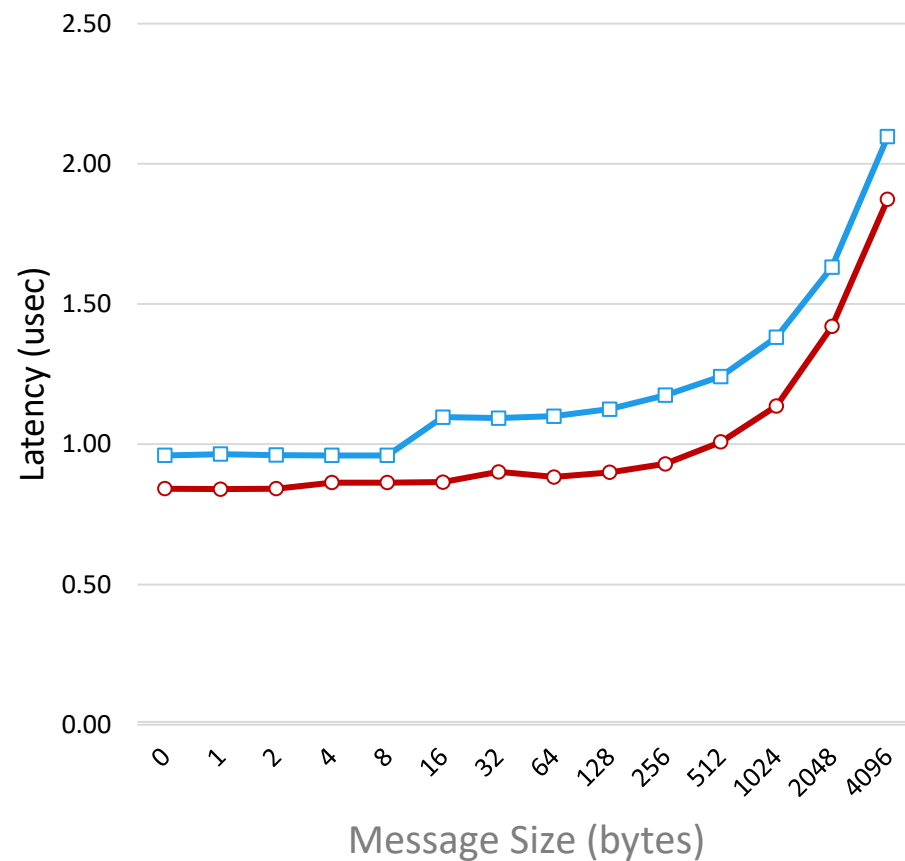


# Omni-Path Express – Significant Performance Improvements




## ✓ Latency

Up to 29% latency reduction



## ✓ Message rate

Over 2.3x more messages/sec/core

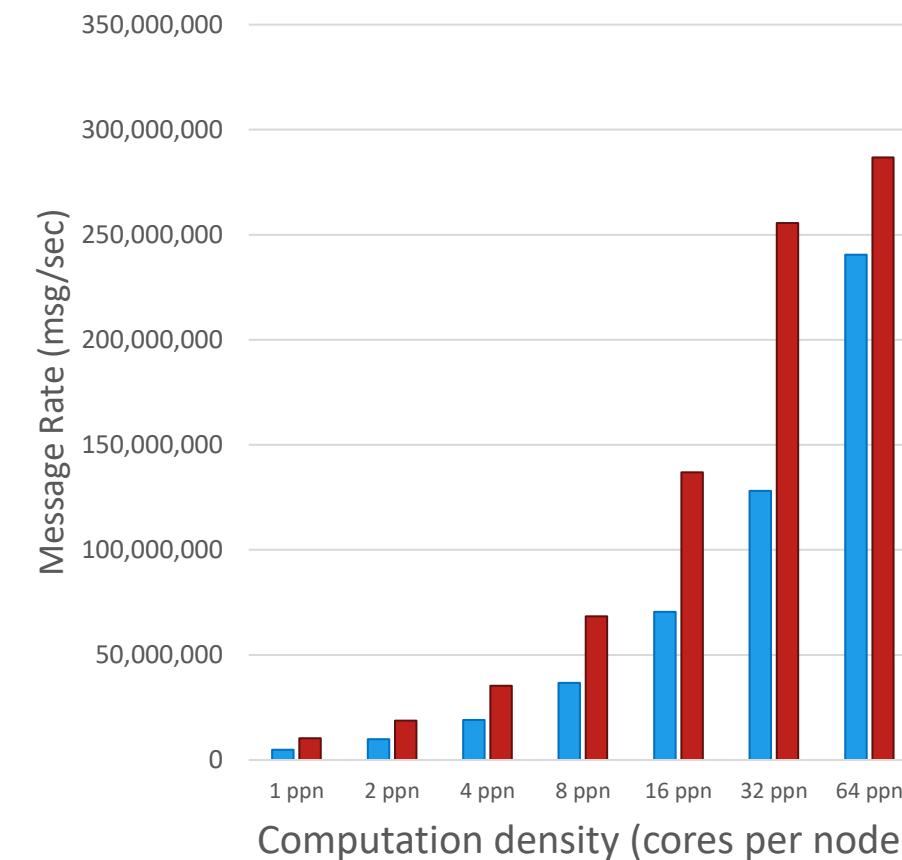


### Legend

■ New software release ■ Current

## ✓ Bandwidth

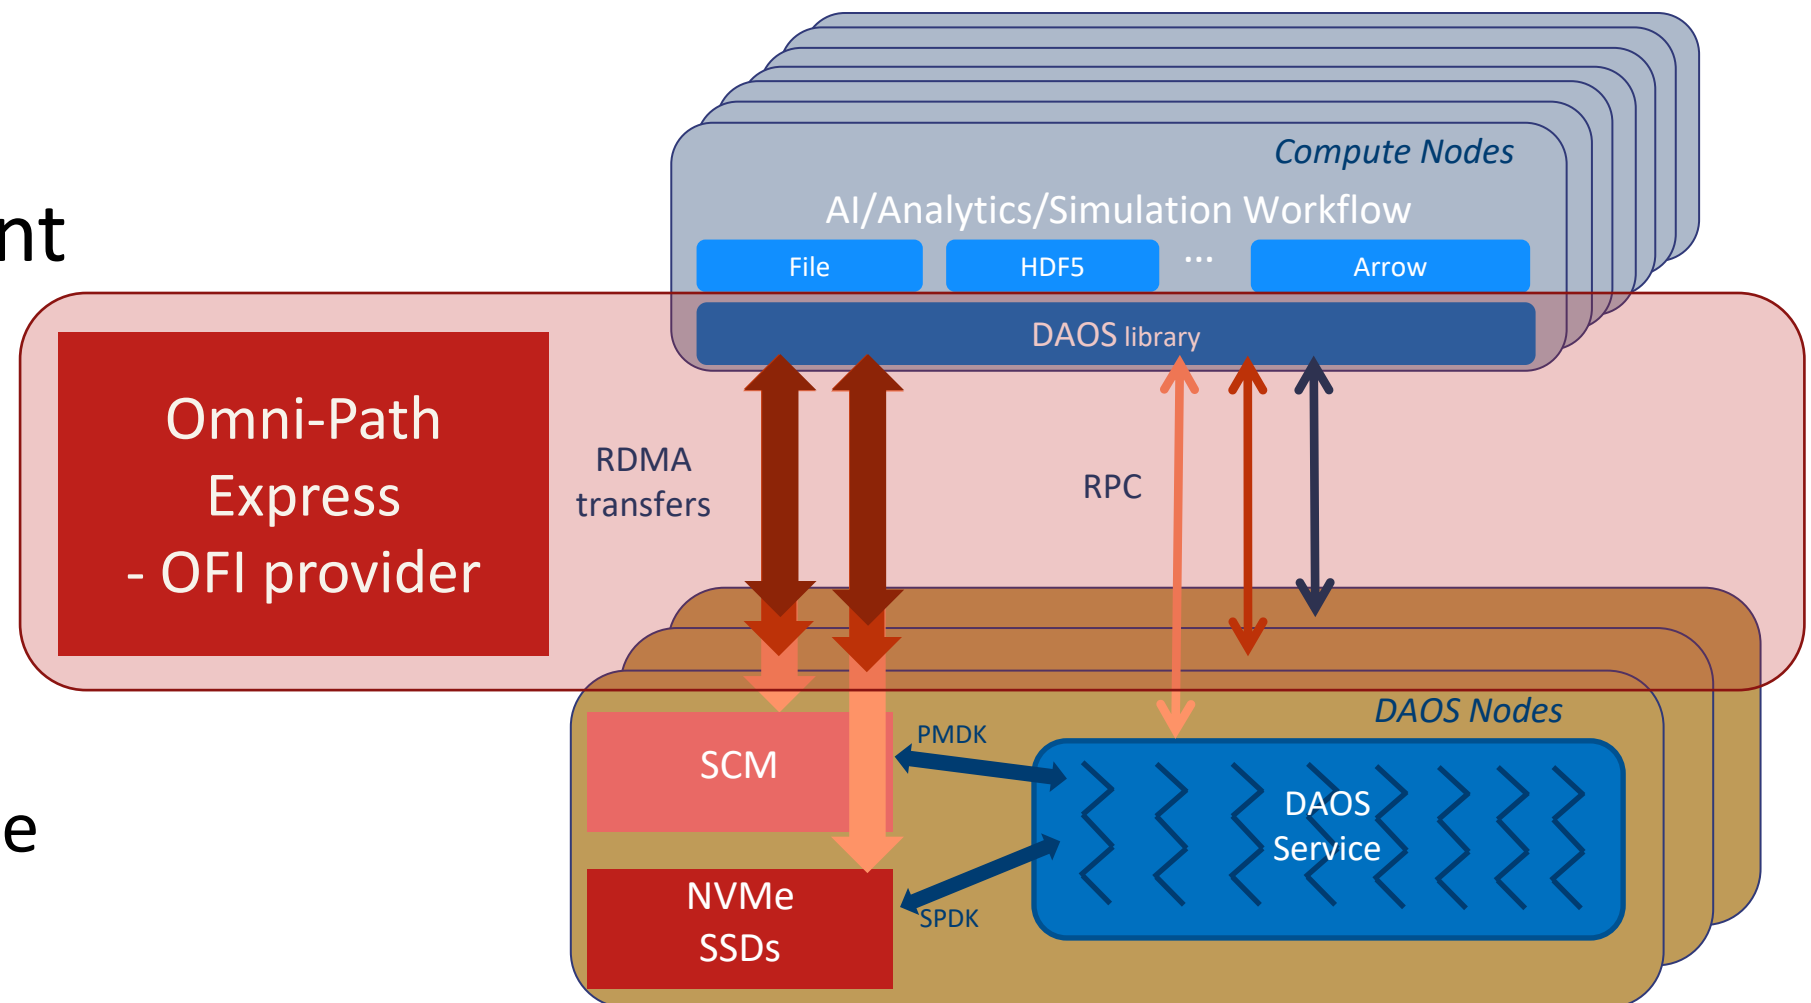
50% core count to max bw





# DAOS on Omni-Path Express

- OFI libfabric
  - Omni-Path Express native provider fully integrated
  - FI\_THREAD\_SAFE operations implemented
- DAOS / Mercury environment
  - Omni-Path Express provider enablement complete
- Features ...
  - Persistent addressing and deterministic endpoints
  - Fault Tolerance and Resilience
  - Multi-Tenancy



# Omni-Path Express – DAOS Test Environment



- Local DAOS Cluster
  - 4 DAOS server nodes
    - Thank you to Intel for providing these DAOS servers for this effort!
  - 32 compute nodes
- Initial development base
  - DAOS 2.0 dev, updating to 2.1
  - Mercury v2.1.0.rc1

# Omni-Path Express - DAOS Support Availability



- Upstream Schedule
  - Libfabric on Omni-Path Express – Alpha (Nov-2021)
    - Targeted for inclusion in DAOS 2.2 by freeze date, end of Dec-2021
    - Evaluation support for select customers
  - Libfabric on Omni-Path Express – Beta (Feb-2022)
    - Driver enhancements, performance focused
  - Mercury, DAOS
    - Omni-Path Express support for these environments will be upstreamed as well
- Omni-Path Express, Release 10.12 (Mar-2022)
  - Target release for DAOS 2.2 support

- Cornelis Networks continues to optimize software for performance and efficiency
  - DAOS is an integral part of Omni-Path Express validation for subsequent releases
- Follow on efforts to improve Omni-Path Express and DAOS support – Adding DAOS support is the start
  - Initial release of Omni-Path Express will support DAOS 2.2
  - Subsequent Omni-Path Express releases will address any new requirements, requests



**CORNELIS**<sup>TM</sup>  
NETWORKS

**Thank You!**

[www.cornelisnetworks.com](http://www.cornelisnetworks.com)  
[sales@cornelisnetworks.com](mailto:sales@cornelisnetworks.com)  
[paul.stasurak@cornelisnetworks.com](mailto:paul.stasurak@cornelisnetworks.com)