

Intel® Omni-Path Architecture

Product Update


Ed Groden

Intel DCG Connectivity Group

November 2016


The Interconnect Landscape: Why Intel® OPA?

Performance




I/O struggling to keep up with CPU innovation

Increasing Scale



Existing solutions reaching limits

Fabric: Cluster Budget¹



Fabric an increasing % of HPC hardware costs

Goal: Keep cluster costs in check → maximize COMPUTE power per dollar

¹ Source: Internal analysis based on a 256-node to 2048-node clusters configured with Mellanox FDR and EDR InfiniBand products. Mellanox component pricing from www.kernelsoftware.com Prices as of November 3, 2015. Compute node pricing based on Dell PowerEdge R730 server from www.dell.com. Prices as of May 26, 2015. Intel® OPA (x8) utilizes a 2-1 over-subscribed Fabric. Intel® OPA pricing based on estimated reseller pricing using projected Intel MSRP pricing on day of launch.

Intel® Omni-Path Architecture Rapid Ramp



It's been quite a year since launch at Supercomputing 2015



Major wins across the globe:



50% MSS of 100Gb systems on the June '16 Top500 list¹

Expect dramatic increase for November '16 list next week

Gained significant momentum in the market²

¹ Source Top500.org; **Other names and brands may be claimed as the property of others

² As measured by 100Gbps High Perf. Fabrics revenue

Intel® Omni-Path Architecture

Evolutionary Approach, Revolutionary Features, End-to-End Solution

Building on the industry's best technologies

- Highly leverage existing Aries and Intel® True Scale fabric
- Adds innovative new features and capabilities to improve performance, reliability, and QoS
- Re-use of existing OpenFabrics Alliance* software

Robust product offerings and ecosystem

- End-to-end Intel product line
- >100 OEM designs¹
- Strong ecosystem with 70+ Fabric Builders members

HFI Adapters

Single port
x8 and x16



x16
Adapter
(100 Gb/s)

x8 Adapter
(58 Gb/s)

Edge Switches

1U Form Factor
24 and 48 port



48-port
Edge Switch



24-port
Edge Switch

Director Switches

QSFP-based
192 and 768 port



768-port
Director Switch
(20U chassis)



192-port
Director Switch
(7U chassis)

Silicon

OEM custom designs
HFI and Switch ASICs



HFI silicon
Up to 2 ports
(50 GB/s
total b/w)



Switch silicon
up to 48 ports
(1200 GB/s
total b/w)

Software

Open Source
Host Software and
Fabric Manager



Cables

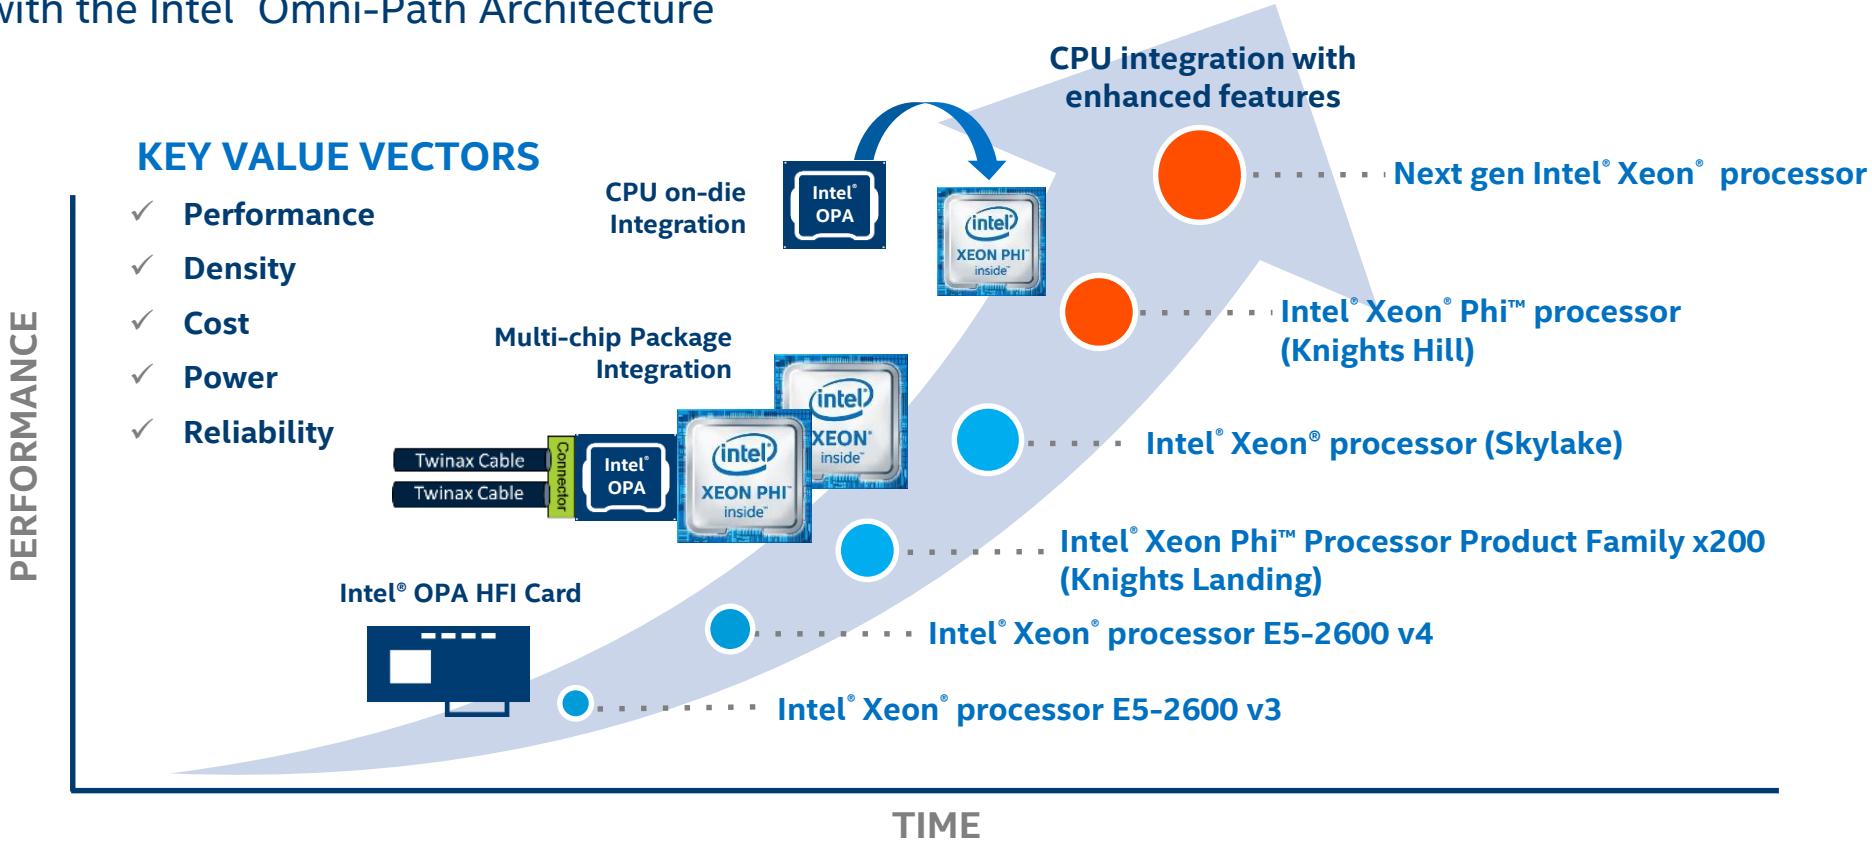
Third Party Vendors
Passive Copper
Active Optical



¹ Source: Intel internal information. Design win count based on OEM and HPC storage vendors who are planning to offer either Intel-branded or custom switch products, along with the total number of OEM platforms that are currently planned to support custom and/or standard Intel® OPA adapters. Design win count as of November 1, 2015 and subject to change without notice based on vendor product plans. *Other names and brands may be claimed as property of others.

CPU-Fabric Integration

with the Intel® Omni-Path Architecture



Latency, Bandwidth, and Message Rate

Intel® Xeon® processor E5-2699 v3 & E5-2699 v4

Intel® Omni-Path Architecture (Intel® OPA)

| Metric | E5-2699 v3 ¹ | E5-2699 v4 ² |
|--|-------------------------|-------------------------|
| Latency (one-way, 1 switch, 8B) [ns] | 910 | 910 |
| Bandwidth (1 rank per node, 1 port, uni-dir, 1MB) [GB/s] | 12.3 | 12.3 |
| Bandwidth (1 rank per node, 1 port, bi-dir, 1MB) [GB/s] | 24.5 | 24.5 |
| Message Rate (max ranks per node, uni-dir, 8B) [Mmps] | 112.0 | 141.1 |
| Message Rate (max ranks per node, bi-dir, 8B) [Mmps] | 137.8 | 172.5 |

Near linear scaling of message rate with added cores on successive Intel® Xeon® processors

Dual socket servers. Intel® Turbo Boost Technology enabled, Intel® Hyper-Threading Technology disabled. OSU OMB 5.1. Intel® OPA: Open MPI 1.10.0-hfi as packaged with IFS 10.0.0.0.697. Benchmark processes pinned to the cores on the socket that is local to the Intel® OP Host Fabric Interface (HFI) before using the remote socket. RHEL 7.2. Bi-directional message rate measured with osu_mbw_mr, modified for bi-directional measurement. We can provide a description of the code modification if requested. BIOS settings: IOU non-posted prefetch disabled. Snoop timer for posted prefetch=9. Early snoop disabled. Cluster on Die disabled.

1. Intel® Xeon® processor E5-2699 v3 2.30 GHz 18 cores, 36 ranks per node for message rate test


2. Intel® Xeon® processor E5-2699 v4 2.20 GHz 22 cores, 44 ranks per node for message rate test

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, 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.


MPI Performance - Ohio State Microbenchmarks

Intel® Omni-Path Architecture (Intel® OPA) vs. InfiniBand* EDR - Intel® MPI


MPI Latency¹



MPI Bandwidth²



MPI Message Rate³

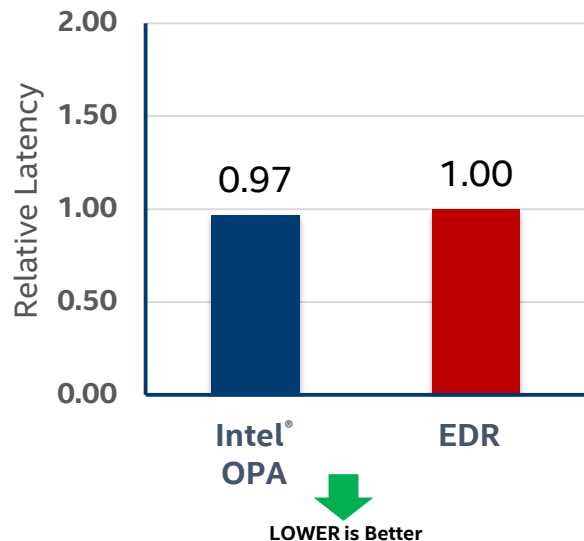


Tests performed on Intel® Xeon® Processor E5-2697A v4 dual-socket servers with 2133 MHz DDR4 memory. Intel® Turbo Boost Technology and Intel® Hyper-Thread Technology enabled. Ohio State Micro Benchmarks v. 5.0. Intel MPI 5.1.3, RHEL7.2. Intel® OPA: tmi fabric, I_MPI_TMI_DRECV=1. Intel Corporation Device 24f0 – Series 100 HFI ASIC (B0 silicon). OPA Switch: Series 100 Edge Switch – 48 port (B0 silicon). IOU Non-posted Prefetch disabled in BIOS. Snooper hold-off timer = 9. EDR based on internal testing: shm:dapl fabric. `-genv I_MPI_DAPL_EAGER_MESSAGE_AGGREGATION off`. Mellanox EDR ConnectX-4 Single Port Rev 3 MCX455A HCA. Mellanox SB7700 - 36 Port EDR Infiniband switch. MLNX_OFED_LINUX-3.2-2.0.0.0 (OFED-3.2-2.0.0). 1. osu_latency 8 B message. 2. osu_bw 1 MB message. 3. osu_mbw_mr, 8 B message (uni-directional), 32 MPI rank pairs. Maximum rank pair communication time used instead of average time, average time introduced into Ohio State Micro Benchmarks as of v3.9 (2/28/13).

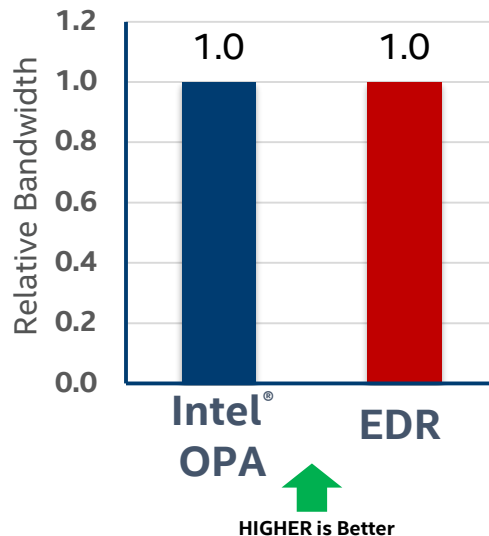
MPI Performance - Ohio State Microbenchmarks

Intel® Omni-Path Architecture (Intel® OPA) vs. InfiniBand* EDR - Open MPI

MPI Latency¹




MPI Uni-directional Bandwidth²



MPI Message Rate³

[Non-coalesced]



Tests performed on Intel® Xeon® Processor E5-2697A v4 dual-socket servers with 2133 MHz DDR4 memory. Intel® Turbo Boost Technology and Intel® Hyper Threading Technology enabled. Ohio State Micro Benchmarks v. 5.0. RHEL 7.2 Intel OPA: Open MPI 1.10.0 with PSM2 as packaged with IFS 10.0.1.0.50. Intel Corporation Device 24f0 – Series 100 HFI ASIC (B0 silicon). OPA Switch: Series 100 Edge Switch – 48 port (B0 silicon). IOU Non-posted Prefetch disabled in BIOS. EDR: Open MPI 1.10-mellanox released with hpcx-v1.5.370-gcc-MLNX_OFED_LINUX-3.2-1.0.1.1-redhat7.2-x86_64. MLNX_OFED_LINUX-3.2-2.0.0.0 (OFED-3.2-2.0.0). Best of default, MXM_TLS=self,rc, and -mca pml yalla tunings. Mellanox EDR ConnectX-4 Single Port Rev 3 MCX455A HCA. Mellanox SB7700 - 36 Port EDR Infiniband switch. All measurements include one switch hop. 1. osu_latency 8 B message. 2. osu_bw 1 MB message. 3. osu_mbw_mr, 8 B message (uni-directional), 32 MPI rank pairs. Maximum rank pair communication time used instead of average time, average timing introduced into Ohio State Micro Benchmarks as of v3.9 (2/28/13).


Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, 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.




Ohio State Microbenchmarks - 8 Byte MPI latency with a Switch Hop

Intel® Omni-Path Architecture (Intel® OPA) PSM vs. InfiniBand® EDR MXM

3 M Copper Cables



2 M Copper Cables



FEC - Forward Error Correction

LOWER is Better

Tests performed on Intel® Xeon® Processor E5-2697A v4 dual-socket servers with 2133 MHz DDR4 memory. Intel® Turbo Boost Technology and Intel® Hyper-Thread Technology enabled. IOU Non-posted Prefetch disabled in BIOS. Snoop hold-off timer = 9. Ohio State Micro Benchmarks v. 5.3. RHEL7.2. Intel® OPA: Intel Corporation Device 24f0 - Series 100 HFI ASIC (B0 silicon). OPA Switch: Series 100 Edge Switch - 48 port (B0 silicon). Open MPI 1.10.2-hfi as packaged in IFS 10.2.0.0.158. EDR based on internal testing: Mellanox EDR ConnectX-4 Single Port Rev 3 MCX455A HCA. Mellanox SB7700 - 36 Port EDR Infiniband switch. MLNX_OFED_LINUX-3.3-1.0.4. Open MPI 1.10.5a1 as packaged in hpcx-v1.7.405-gcc-MLNX_OFED_LINUX-3.4-1.0.0.0-redhat7.2-x86_64. Older switch firmware: 11.0100.0112. Upgraded switch firmware: 11.1200.0102. 1. SwitchIB-FW-11_1200_0102-release_nodes.pdf: "Removed out-of-the-box FEC, reaching 90ns latency, on Mellanox GA level copper cables equal to or shorter than 2m."

Latency Impact of InfiniBand EDR FEC Implementation

- IB EDR FEC latency penalty is driving a recommendation to operate clusters with FEC DISABLED
- Disabling FEC requires the raw link bit error rate to be extremely low for AOCs and DAC cables

Expected IB EDR End-to-End Retries per Second with FEC Disabled

| Fat Nodes | Link Average BER | | | |
|--------------|------------------|-------|-------|-------|
| | 1e-15 | 1e-14 | 1e-13 | 1e-12 |
| 36 | 0.007 | .07 | .7 | 7 |
| 48 | 0.02 | .20 | 2 | 20 |
| 648 | 0.27 | 2.7 | 27 | 267 |
| 1152 | 0.7 | 7 | 71 | 713 |
| 11664 | 7.2 | 72 | 722 | 7217 |
| 27648 | 22.8 | 228 | 2281 | 22810 |

Mellanox BER claim



InfiniBand specification

For a 1k node cluster, expect between ~1 to 700 end-to-end retries per second

Intel® OPA Packet Integrity Protection is always on providing link BER < 2.99e-29
27K node Intel® OPA cluster expected to generate an end-to-end retry **once per 62K years**

MPI Latency at Scale

Intel® Omni-Path Architecture (Intel® OPA) vs. InfiniBand® EDR - Open MPI




16 Nodes, 32 MPI ranks per node

Tests performed on Intel® Xeon® Processor E5-2697A v4 dual-socket servers with 2133 MHz DDR4 memory. Intel® Turbo Boost Technology and Intel® Hyper-Thread Technology enabled. HPCC 1.4.3. RHEL7.2. Intel OPA: Open MPI 1.10.0 with PSM2 as packaged with IFS 10.0.1.0.50. Intel Corporation Device 24f0 – Series 100 HFI ASIC (B0 silicon), OPA Switch: Series 100 Edge Switch – 48 port (B0 silicon). IOU Non-posted Prefetch disabled in BIOS. EDR: Open MPI 1.10-mellanox released with hpcx-v1.5.370-gcc-MLNX_OFED_LINUX-3.2-1.0.1.1-redhat7.2-x86_64. MLNX_OFED_LINUX-3.2-2.0.0.0 (OFED-3.2-2.0.0). Mellanox EDR ConnectX-4 Single Port Rev 3 MCX455A HCA. Mellanox SB7700 - 36 Port EDR Infiniband switch.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, 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.

Intel® Omni-Path Architecture (Intel® OPA)

Application Performance - Intel® MPI - 16 Nodes



*Spec MPI2007 results estimates until published

**see following slide for system configurations

No Intel® OPA or EDR specific optimizations applied to any workloads except LS-DYNA and ANSYS Fluent: Intel® OPA HFI driver parameter: eager_buffer_size=8388608
WIEN2k comparison is for 8 nodes because EDR IB measurements did not scale above 8 nodes*

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, 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.

Intel Confidential.




Quantum Espresso - Intel® Xeon® processor E5-2697A v4

Intel® Omni-Path Architecture (Intel® OPA) vs EDR

Mellanox Performance Data

ausurf111 benchmark

Colors adjusted to match internal chart



<http://www.hpcwire.com/2016/04/12/interconnect-offloading-versus-onloading/>


- No system configurations
- No details on test methodology
- No way to repeat measurements

Intel® Xeon® Processor E5-2697A v4 dual socket servers. 64 GB DDR4 memory per node, 2133 MHz. RHEL 7.2. BIOS settings: Snoop hold-off timer = 9, Early snoop disabled, Cluster on die disabled. Intel Fabric Suite 10.0.1.0.50. Intel Corporation Device 24f0 – Series 100 HFI ASIC (B0 silicon). OPA Switch: Series 100 Edge Switch – 48 port (B0 silicon). IOU Non-posted prefetch disabled. MLNX_OFED_LINUX-3.2-2.0.0.0 (OFED-3.2-2.0.0). Mellanox EDR ConnectX-4 Single Port Rev 3 MCX455A HCA. Mellanox SB7700 - 36 Port EDR Infiniband switch.

Quantum Espresso version 5.3.0. Intel Compiler 2016 Update 2. ELPA 2015.11.001 (<http://elpa.mpcdf.mpg.de/elpa-tar-archive>). Minor patch set for QE to accommodate latest ELPA. Most optimal NPOOL, NDIAG, and NTG settings reported for both OPA and EDR. Intel MPI 5.1.3, shm:tmi fabric for Intel® OPA and shm:dapl fabric for EDR, all default settings.

Intel Performance Data

ausurf112 benchmark




Our absolute performance for EDR is better than what they posted for themselves and Intel® OPA is still better!

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, 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.

Intel® Omni-Path Architecture (Intel® OPA)

Application Performance Per Fabric Dollar* - Intel® MPI - 16 Nodes




**HIGHER
is Better**

*Spec MPI2007 results estimates until published

**see following slide for system configurations

*All pricing data obtained from www.kernelsoftware.com May 4, 2016. All cluster configurations estimated via internal Intel configuration tool. Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction. **Fabric hardware assumes one edge switch, 16 network adapters and 16 cables.**

*No Intel® OPA or EDR specific optimizations applied to any workloads except LS-DYNA and ANSYS Fluent: Intel® OPA HFI driver parameter: eager_buffer_size=8388608
WIEN2k comparison is for 8 nodes because EDR IB* measurements did not scale above 8 nodes*

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, 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.

Intel Confidential.



System & Software Configuration [include with previous slides]

Common configuration for bullets 1-11 unless otherwise specified: Intel® Xeon® Processor E5-2697A v4 dual socket servers. 64 GB DDR4 memory per node, 2133 MHz. RHEL 7.2. BIOS settings: Snoop hold-off timer = 9, Early snoop disabled, Cluster on die disabled, IOU Non-posted prefetch disabled. Intel® Omni-Path Architecture (Intel® OPA): Intel Fabric Suite 10.0.1.0.50. Intel Corporation Device 24f0 – Series 100 HFI ASIC (Production silicon). OPA Switch: Series 100 Edge Switch – 48 port (Production silicon). EDR Infiniband: MLNX_OFED_LINUX-3.2-2.0.0.0 (OFED-3.2-2.0.0). Mellanox EDR ConnectX-4 Single Port Rev 3 MCX455A HCA. Mellanox SB7700 – 36 Port EDR Infiniband switch.

1. WIEN2k version 14.2. <http://www.wien2k.at/>. http://www.wien2k.at/reg_user/benchmark/. Run command: "mpirun ... lapw1c_mpi lapw1.def". Intel Fortran Compiler 17.0.0 20160517. Compile flags: -FR -mp1 -w -prec_div -pc80 -pad -ip -DINTEL_VML -traceback -assume buffered_io -DFTW3 -l/opt/intel/compilers_and_libraries_2017.0.064/linux/mkl/include/fftw/ -DParallel.shm:tmi fabric used for Intel® OPA and shm:dapl fabric used for EDR IB®.
2. GROMACS version 5.0.4. Intel Composer XE 2015.1.133. Intel MPI 5.1.3. FFTW-3.3.4. ~/bin/cmake .. -DGMX_BUILD_OWN_FFTW=OFF -DREGRESSIONTEST_DOWNLOAD=OFF -DCMAKE_C_COMPILER=icc -DCMAKE_CXX_COMPILER=icpc -DCMAKE_INSTALL_PREFIX=~/gromacs-5.0.4-installed. Intel® OPA MPI parameters: I_MPI_FABRICS=shm:tmi, EDR MPI parameters: I_MPI_FABRICS=shm:dapl
3. NWChem release 6.6. Binary: nwchem_comex-mpi-pr_mkl with MPI-PR run over MPI-1. Workload: sios3 and sios5. Intel® MPI Library 2017.0.064. 2 ranks per node, 1 rank for computation and 1 rank for communication. shm:tmi fabric for Intel® OPA and shm:dapl fabric for EDR, all default settings. Intel Fabric Suite 10.2.0.0.153. http://www.nwchem-sw.org/index.php/Main_Page
4. LS-DYNA MPP R8.1.0 dynamic link. Intel Fortran Compiler 13.1 AVX2. Intel® OPA - Intel MPI 2017 Library Beta Release Candidate 1. mpi.2017.0.0.BETA.U1.RC1.x86_64.wv20.20160512.143008. MPI parameters: I_MPI_FABRICS=shm:tmi. HFI driver parameter: eager_buffer_size=8388608. EDR MPI parameters: I_MPI_FABRICS=shm:ofa.
5. ANSYS Fluent v17.0, Rotor_3m benchmark. Intel® MPI Library 5.0.3 as included with Fluent 17.0 distribution, and libpsm_infinipath.so.1 added to the Fluent syslib library path for PSM/PSM2 compatibility. Intel® OPA MPI parameters: -piib.infinipath, EDR MPI parameters: -piib.dapl
6. NAMD: Intel Composer XE 2015.1.133. NAMD V2.11, Charm 6.7.0, FFTW 3.3.4. Intel MPI 5.1.3. Intel® OPA MPI parameters: I_MPI_FABRICS=shm:tmi, EDR MPI parameters: I_MPI_FABRICS=shm:dapl
7. Quantum Espresso version 5.3.0. Intel Compiler 2016 Update 2. ELPA 2015.11.001 (<http://elpa.mpcdf.mpg.de/elpa-tar-archive>). Minor patch set for QE to accommodate latest ELPA. Most optimal NPOOL, NDIAG, and NTG settings reported for both OPA and EDR. Intel® OPA MPI parameters: I_MPI_FABRICS=shm:tmi, EDR MPI parameters: I_MPI_FABRICS=shm:dapl
8. CD-adapco STAR-CCM+® version 11.04.010. Workload: lemans_poly_17m.amg.sim benchmark. Intel MPI version 5.0.3.048. 32 ranks per node. OPA command: \$ /starccm+ -ldlibpath /STAR-CCM+11.04.010/mpi/intel/5.0.3.048/linux-x86_64/lib64 -ldpreload /usr/lib64/psm2-compat/libpsm_infinipath.so.1 -mpi intel -mppflags "-env I_MPI_DEBUG 5 -env I_MPI_FABRICS shm:tmi -env I_MPI_TMI_PROVIDER psm" -power -rsh ssh -np 512 -machinefile hosts -benchmark:"-nps 512,256,128,64,32 -nits 20 -preits 40 -tag lemans_opa_n16" lemans_poly_17m.amg.sim. EDR command: \$ /starccm+ -mpi intel -mppflags "-env I_MPI_DEBUG 5" -power -rsh ssh -np 512 -machinefile hosts -benchmark:"-nps 512,256,128,64,32 -nits 20 -preits 40 -tag lemans_edr_n16" lemans_poly_17m.amg.sim
9. LAMMPS (Large-scale Atomic/Molecular Massively Parallel Simulator) Feb 16, 2016 stable version release. MPI: Intel® MPI Library 5.1 Update 3 for Linux. Workload: Rhodopsin protein benchmark. Number of time steps=100, warm up time steps=10 (not timed) Number of copies of the simulation box in each dimension: 8x8x4 and problem size: 8x8x4x32k = 8,192k atoms Intel® OPA: MPI parameters: I_MPI_FABRICS=shm:tmi, I_MPI_PIN_DOMAIN=core EDR: MPI parameters: I_MPI_FABRICS=shm:dapl, I_MPI_PIN_DOMAIN=core
10. WRF version 3.5.1, Intel Composer XE 2015.1.133. Intel MPI 5.1.3. NetCDF version 4.4.2. FCBASEOPTS=-w -ftz -align all -fno-alias -fp-model precise. CFLAGS_LOCAL = -w -O3 -ip. Intel® OPA MPI parameters: I_MPI_FABRICS=shm:tmi, EDR MPI parameters: I_MPI_FABRICS=shm:dapl
11. Spec MPI 2007: 16 nodes, 32 MPI ranks/node. SPEC MPI2007, Large suite, <https://www.spec.org/mpi/>. *Intel Internal measurements marked estimates until published. Intel MPI 5.1.3. Intel® OPA MPI parameters: I_MPI_FABRICS=shm:tmi, EDR MPI parameters: I_MPI_FABRICS=shm:dapl

Common configuration for bullets 12-13: Intel® Xeon® Processor E5-2697 v4 dual socket servers. 128 GB DDR4 memory per node, 2400 MHz. RHEL 6.5. BIOS settings: Snoop hold-off timer = 9. Intel® OPA: Intel Fabric Suite 10.0.1.0.50. Intel Corporation Device 24f0 – Series 100 HFI ASIC (Production silicon). OPA Switch: Series 100 Edge Switch – 48 port (Production silicon). IOU Non-posted prefetch disabled. 2). Mellanox EDR based on internal measurements: Mellanox EDR ConnectX-4 Single Port Rev 3 MCX455A HCA. Mellanox SB7700 – 36 Port EDR Infiniband switch. IOU Non-posted prefetch enabled.

12. MiniFE 2.0, Intel compiler 16.0.2. Intel® MPI Library version 5.1.3. Build settings: -O3 -xCORE-AVX2 -DMINIFE_CSR_MATRIX -DMINIFE_GLOBAL_ORDINAL="long long int", mpirun -bootstrap ssh -env OMP_NUM_THREADS 1 - perhost 36 minife_nx=200 ny=200 nz=200, 200x200x200 grid using 36 MPI ranks pinned to 36 cores per node. Intel® OPA MPI parameters: I_MPI_FABRICS=shm:tmi, EDR MPI parameters: I_MPI_FABRICS=shm:dapl. Intel® Turbo Mode technology and Intel® Hyper threading technology disabled.
13. VASP (developer branch). MKL: 11.3 Update 3 Product build 20160413. Compiler: 2016u3. Intel MPI-2017 Build 20160718 . elpa-2016.05.002. Intel® OPA MPI parameters: I_MPI_FABRICS=shm:tmi, EDR MPI parameters: I_MPI_FABRICS=shm:dapl, I_MPI_PLATFORM=BDW, I_MPI_DAPL_PROVIDER=ofa-v2-mlx5_0-1u, I_MPI_DAPL_DIRECT_COPY_THRESHOLD=331072. Intel® Turbo Mode technology disabled. Intel Hyper Threading technology enabled.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, 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.

LINK LEVEL INNOVATIONS

New Intel® OPA Fabric Features: Fine-grained Control Improves Resiliency and Optimizes Traffic Movement



Traffic Flow Optimization

- Optimizes Quality of Service (QoS) in mixed traffic environments, such as storage and MPI
- Transmission of lower-priority packets can be paused so higher priority packets can be transmitted

- Ensures high priority traffic is not delayed → Faster time to solution
- Deterministic latency → Lowers run-to-run timing inconsistencies



Packet Integrity Protection

- Allows for rapid and transparent recovery of transmission errors on an Intel® OPA link without additional latency
- Resends 1056-bit bundle w/errors only instead of entire packet (based on MTU size)

- Fixes happen at the link level rather than end-to-end level
- Much lower latency than Forward Error Correction (FEC) defined in the InfiniBand* specification¹



Dynamic Lane Scaling

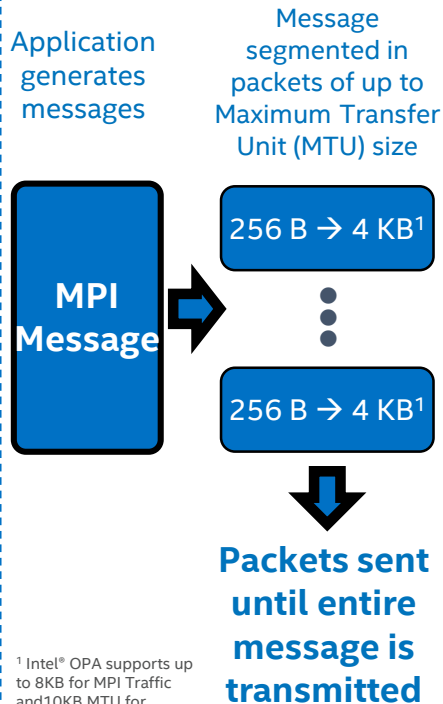
- Maintain link continuity in the event of a failure of one of more physical lanes
- Operates with the remaining lanes until the failure can be corrected at a later time

- Enables a workload to continue to completion. **Note:** InfiniBand will shut down the entire link in the event of a physical lane failure

¹ Lower latency based on the use of InfiniBand with Forward Error Correction (FEC) Mode A or C in the public presentation titled "Option to Bypass Error Marking (supporting comment #205)," authored by Adeel Ran (Intel) and Oran Sela (Mellanox), January 2013. Mode A modeled to add as much as 140ns latency above baseline, and Mode C can add up to 90ns latency above baseline. Link: www.ieee802.org/3/bj/public/jan13/ran_3bj_01a_0113.pdf

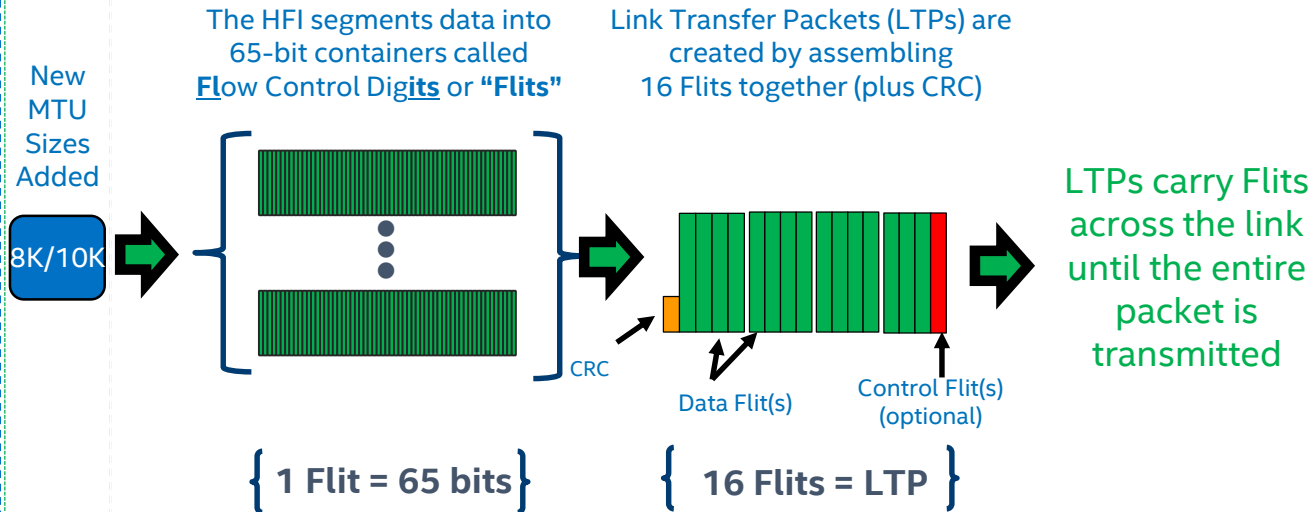
Intel® OPA Link Level Innovation Starts Here

InfiniBand



¹ Intel® OPA supports up to 8KB for MPI Traffic and 10KB MTU for Storage

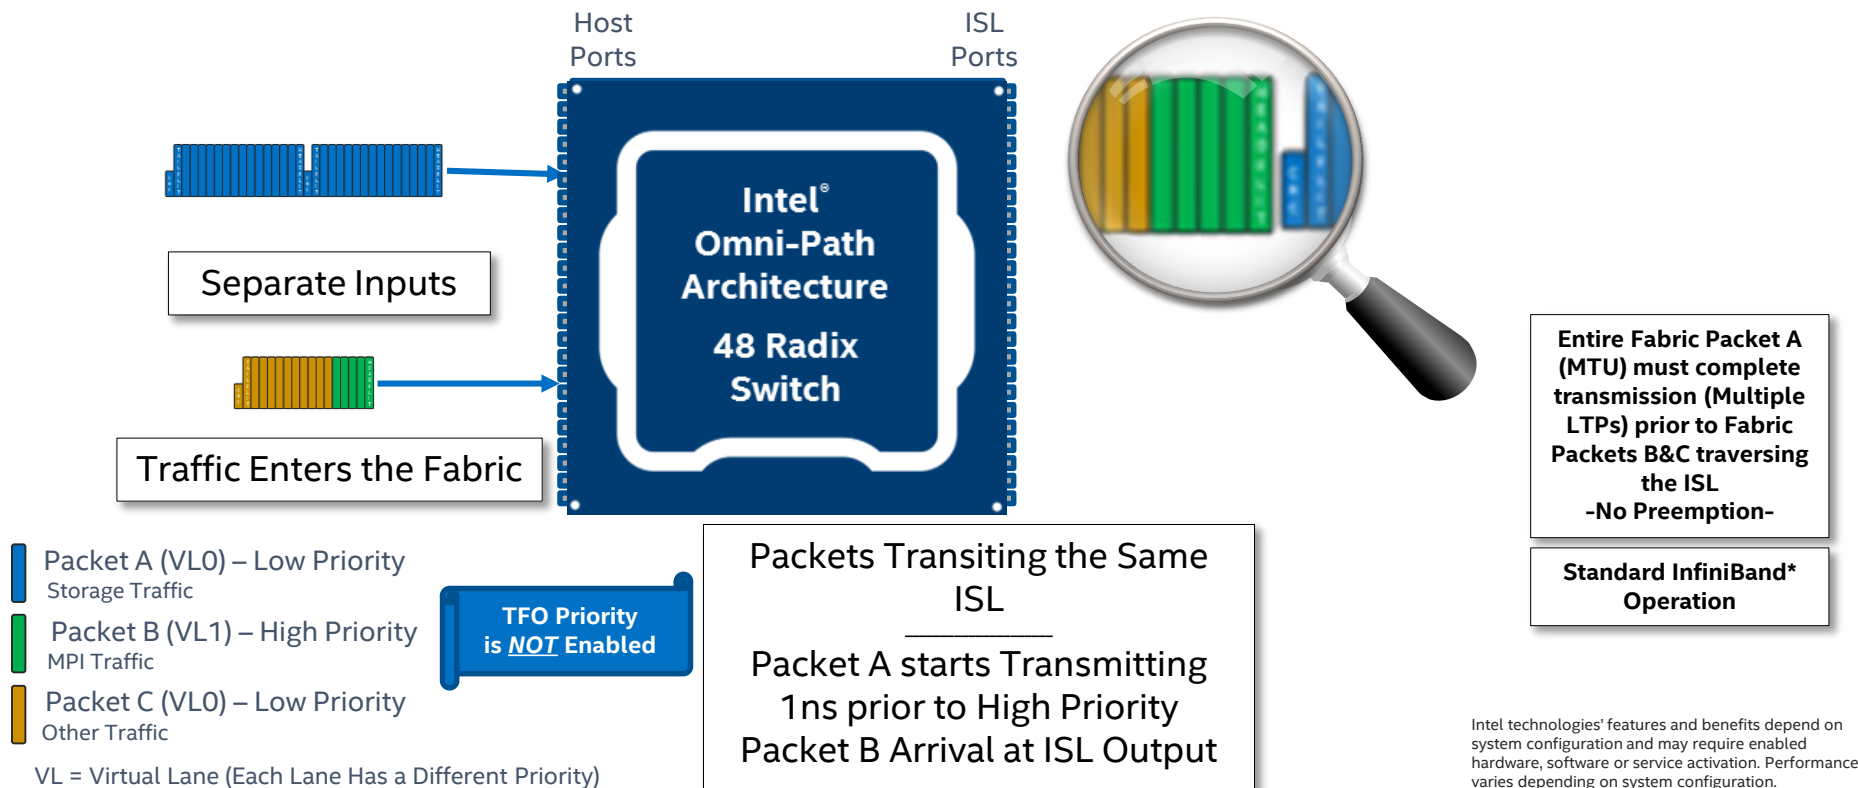
Intel® Omni-Path Fabric



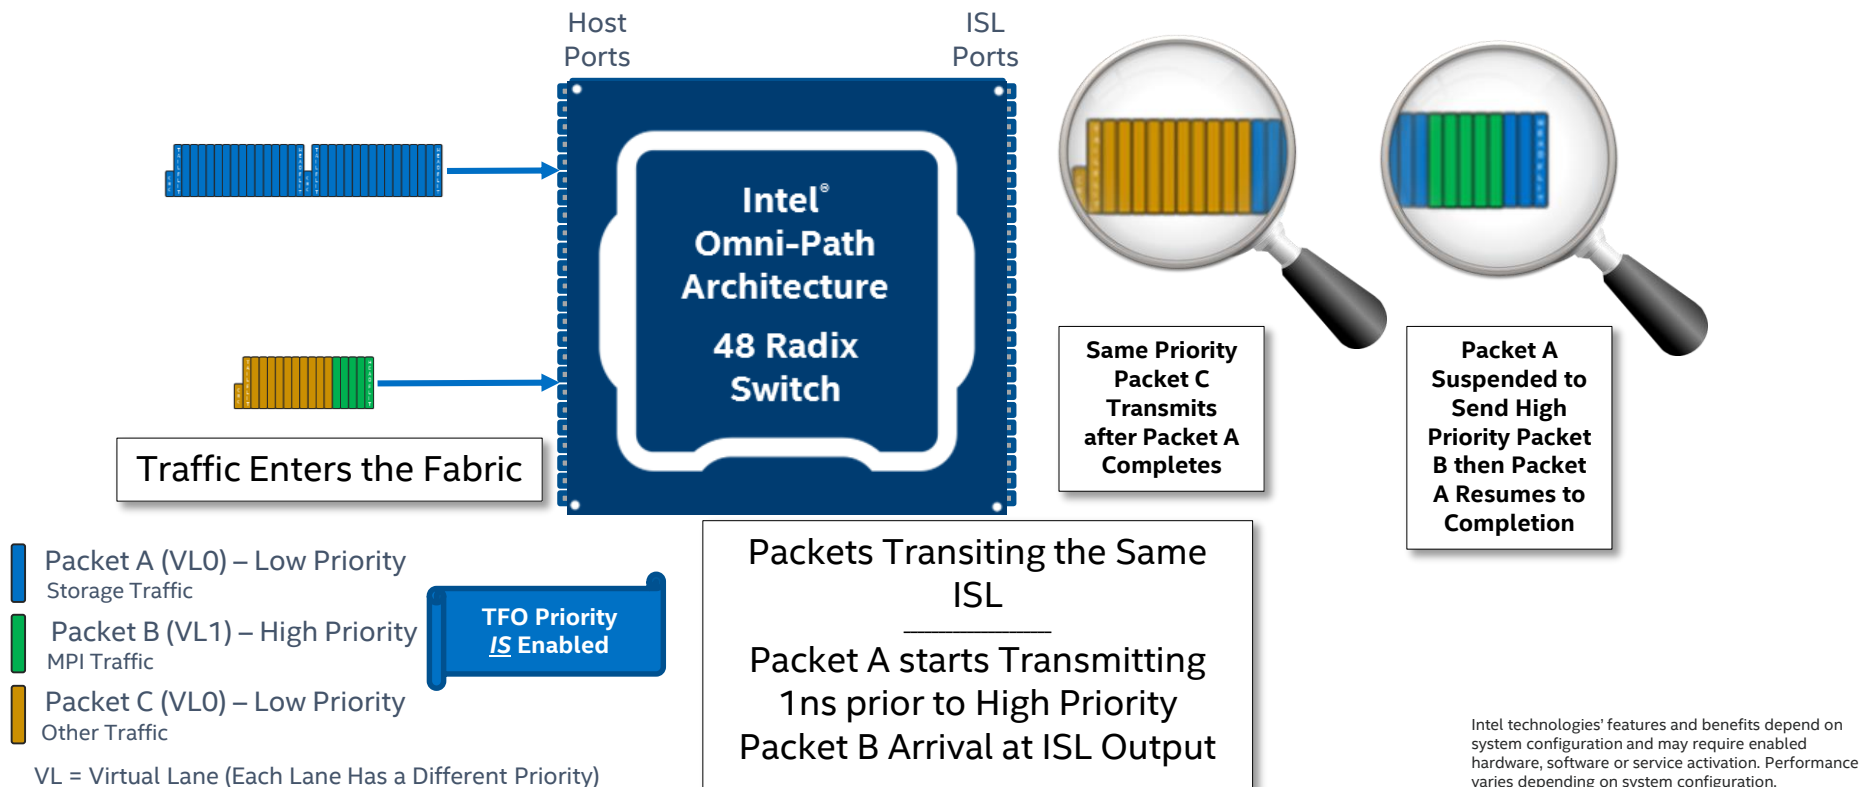
Goals: Improved resiliency, performance and consistent traffic movement

CRC: Cyclic Redundancy Check

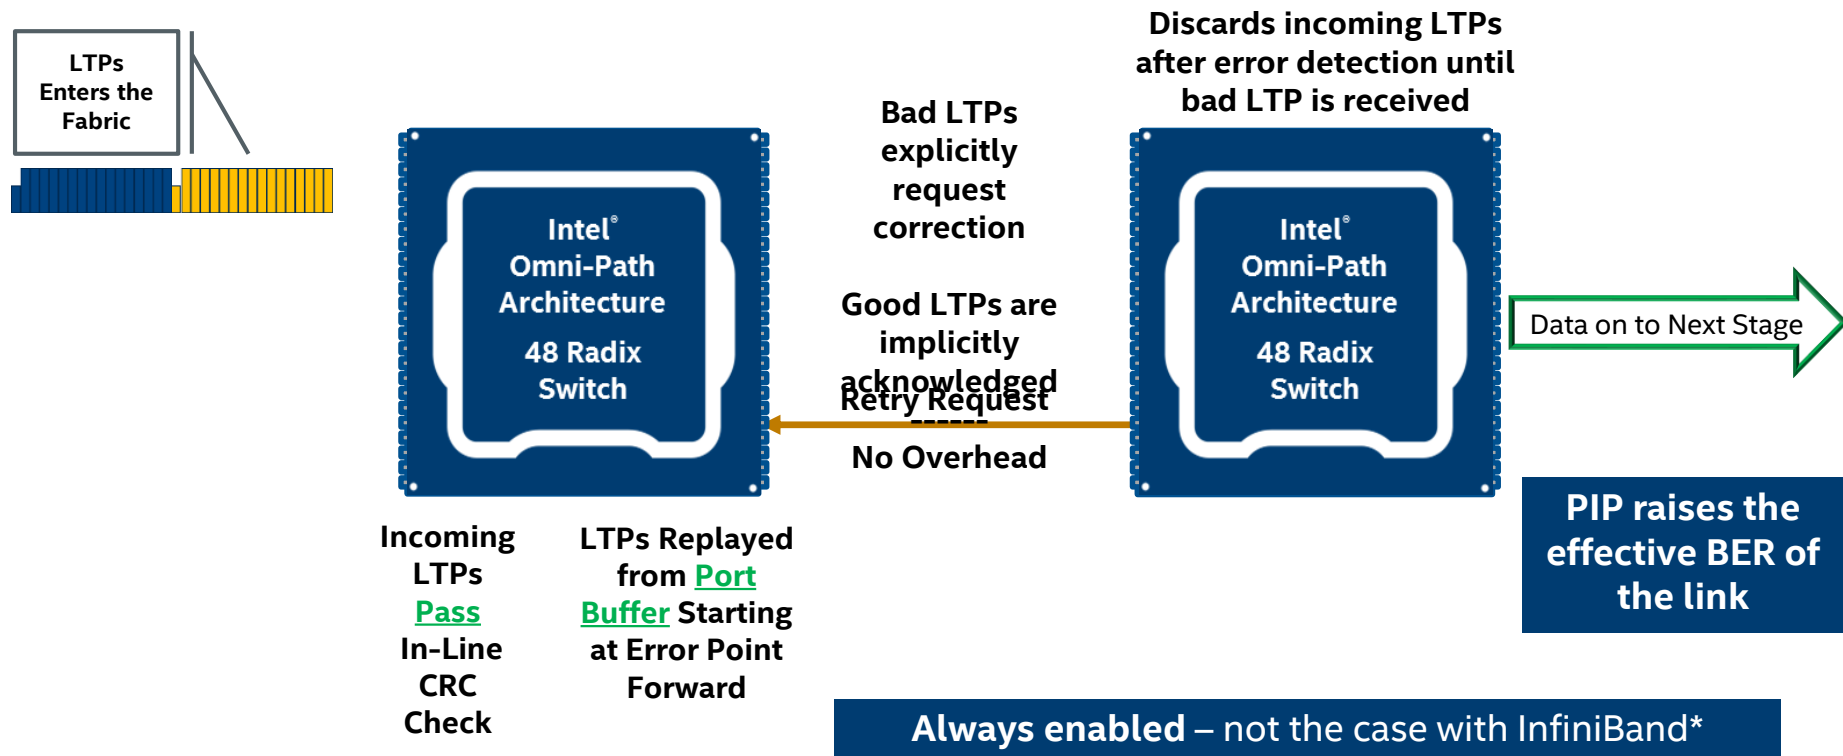
Traffic Flow Optimization (TFO) – Disabled



Traffic Flow Optimization (TFO) – Enabled



Intel® Omni-Path Fabric Link Level Innovation: Packet Integrity Protection (PIP)



Intel® Omni-Path Fabric Link Level Innovation:

Dynamic Lane Scaling (DLS) Traffic Protection



InfiniBand* – 25Gbps

Intel® OPA 75Gbps

User Setting (per Fabric):

- Set maximum degrade option allowable
 - 4x – Any lane failure would cause link reset or take down
 - 3x – Still operates at degraded bandwidth (75 Gbps) with 1 lane failure
 - 2x – Still operates at degraded bandwidth (50 Gbps) with 2 lane failures
 - 1x – Still operates at degraded bandwidth (25 Gbps) with 3 lane failures

Link Recovery:

- PIP is used to recover link without reset – An Intel® OPA innovation

Intel® OPA will continue to pass data at reduced bandwidth with link recovery via PIP

InfiniBand* may close entire link or reinitialize @1x introducing fabric delays or routing issues

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

INTEL[®] OMNI-PATH ARCHITECTURE DATA TRANSPORT CAPABILITIES

Intel® Omni-Path Architecture (Intel® OPA)

Data Transport Summary

Onload, offload, and RDMA:

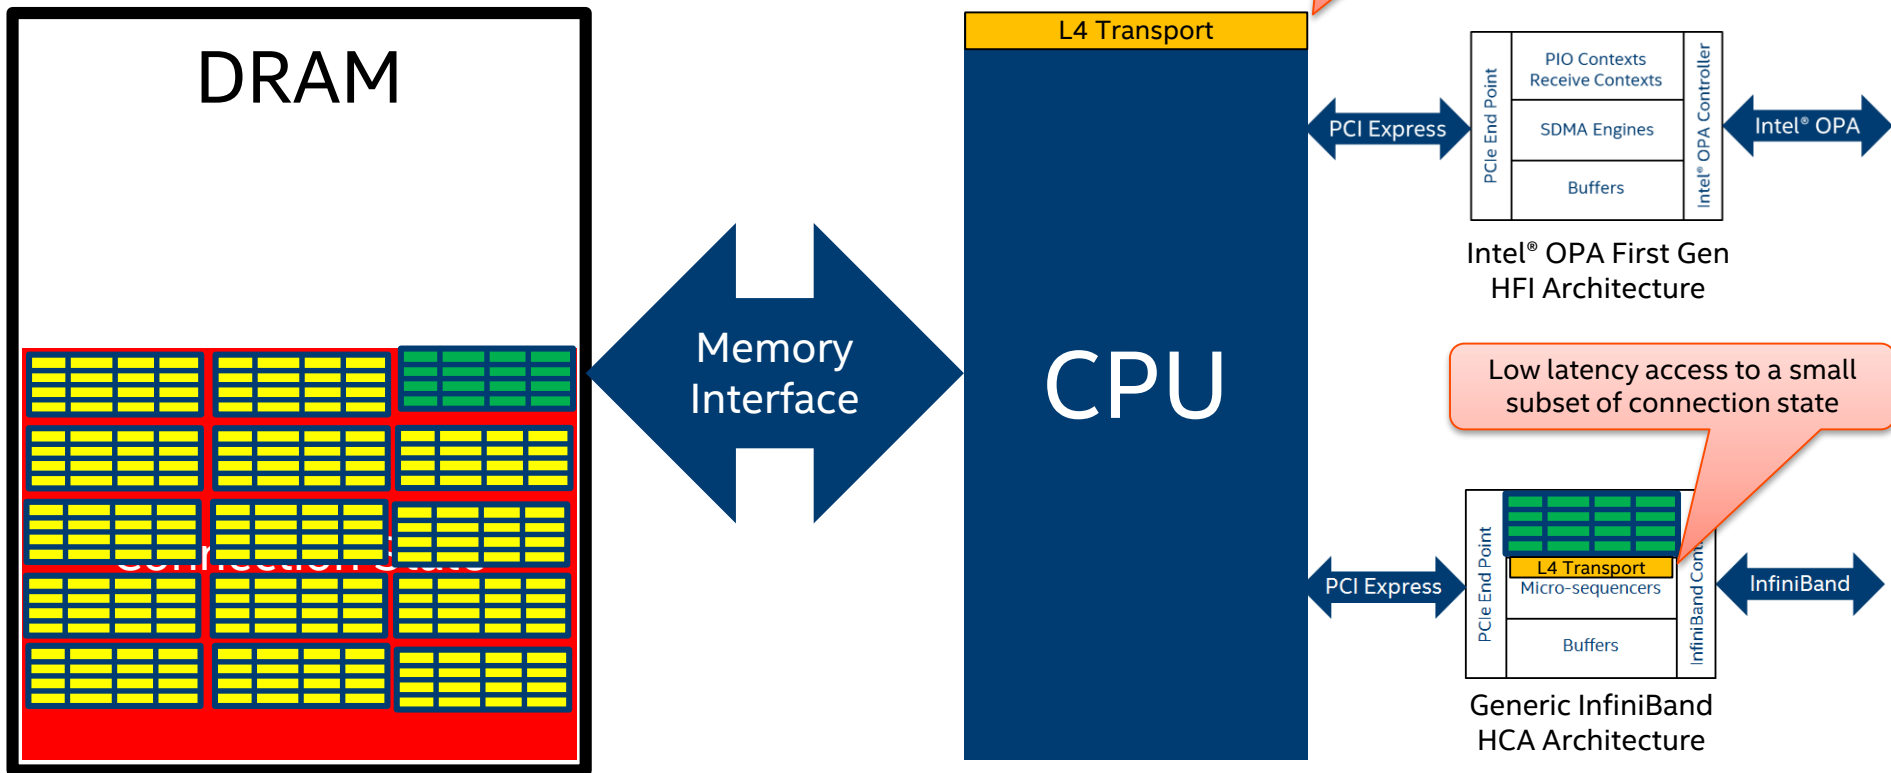
- Onload and offload are approaches for transport layer implementation
 - Choice influences achievable latency, message rate, scalability, data transfer options
- RDMA is an underlying data transfer protocol
 - Intel® OPA absolutely supports it

Intel® OPA data transfer overview:

- Host Fabric Interface (HFI) automatically selects the most efficient resources and data path to send and receive data over an Intel® OPA fabric based on the data type and message size of the data transfer
- RDMA transfers when it makes sense, typically LARGE MESSAGES where setup time is a relatively small part of the overall transfer time
- Programmed I/O sends, coupled with fast mem copy receives, for SMALLER MESSAGES for higher efficiency than setting up an RDMA connection

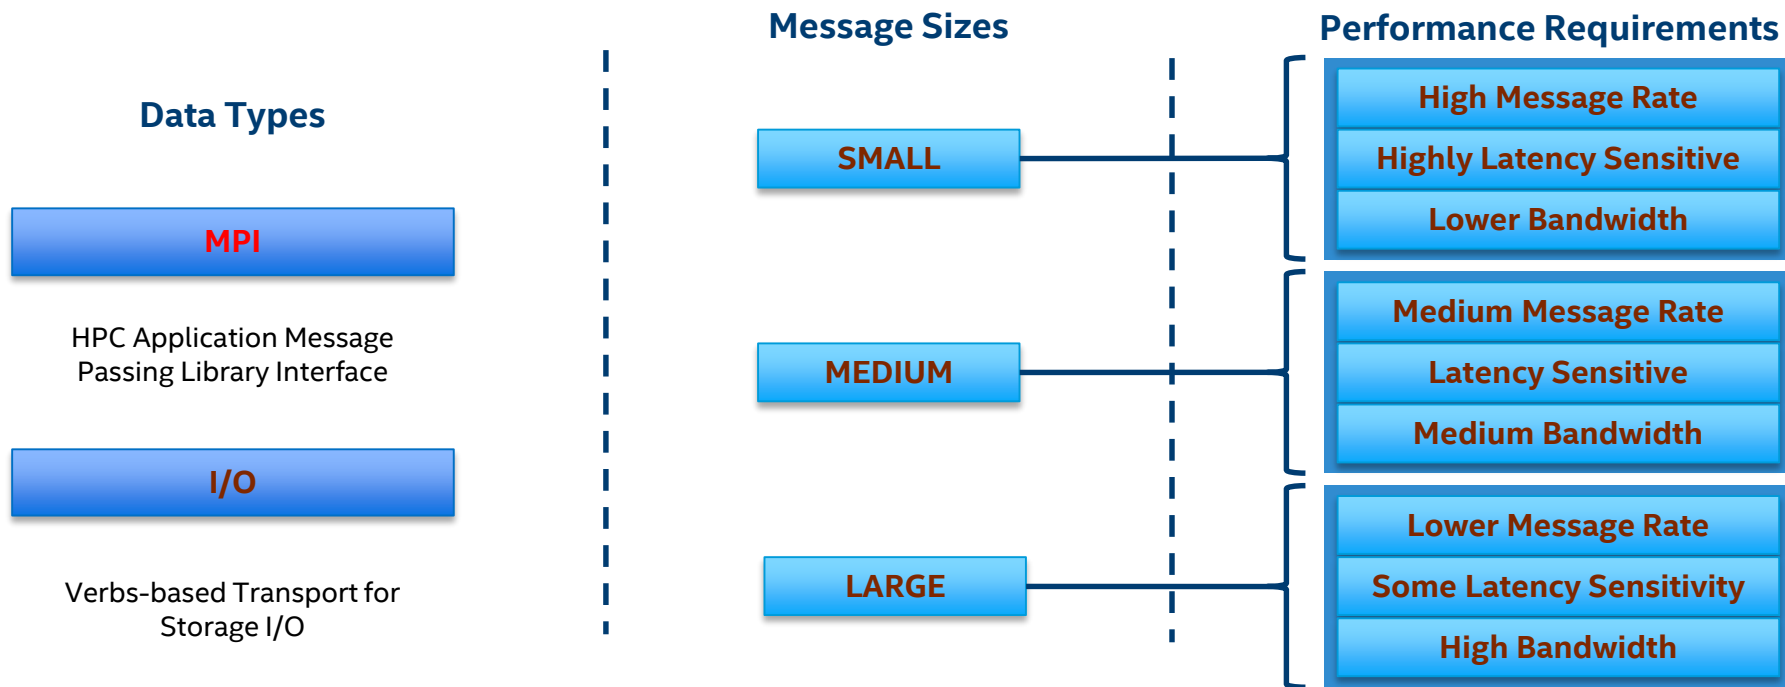
Onload vs. Offload

InfiniBand HCA vs. Intel® OPA HFI



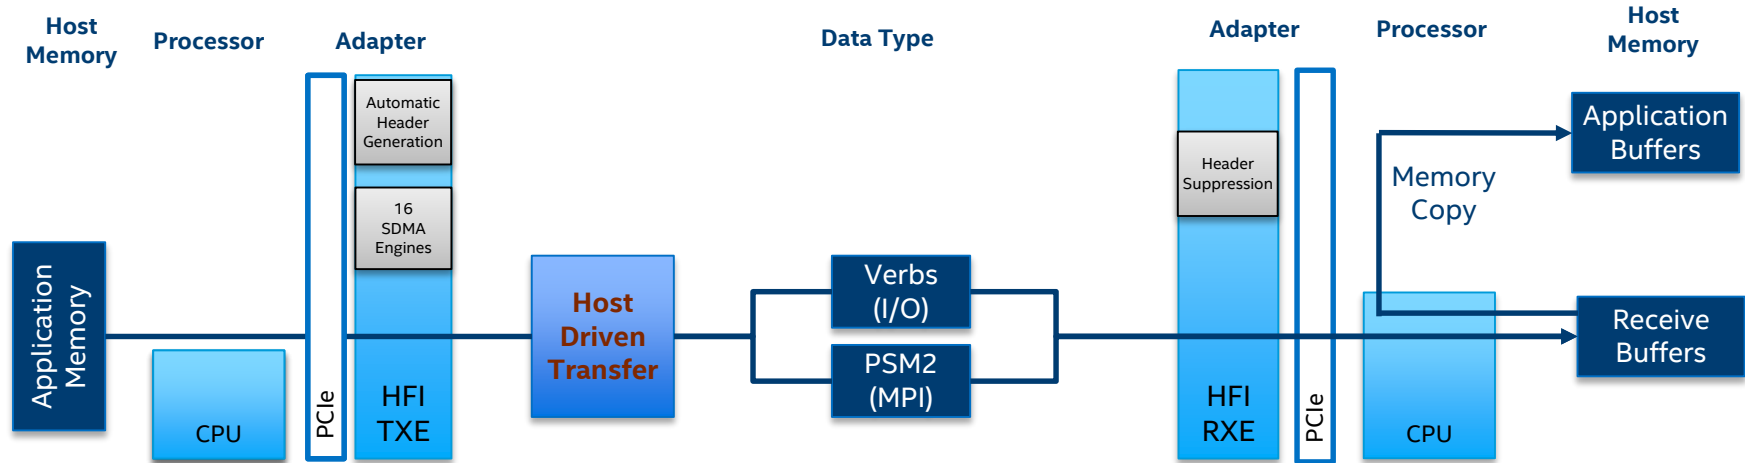
Choosing the Right Data Transfer Approach

MPI and Storage Traffic Performance Needs



Multimodal Data Acceleration

Highest performance small message transfer



Host Driven Send

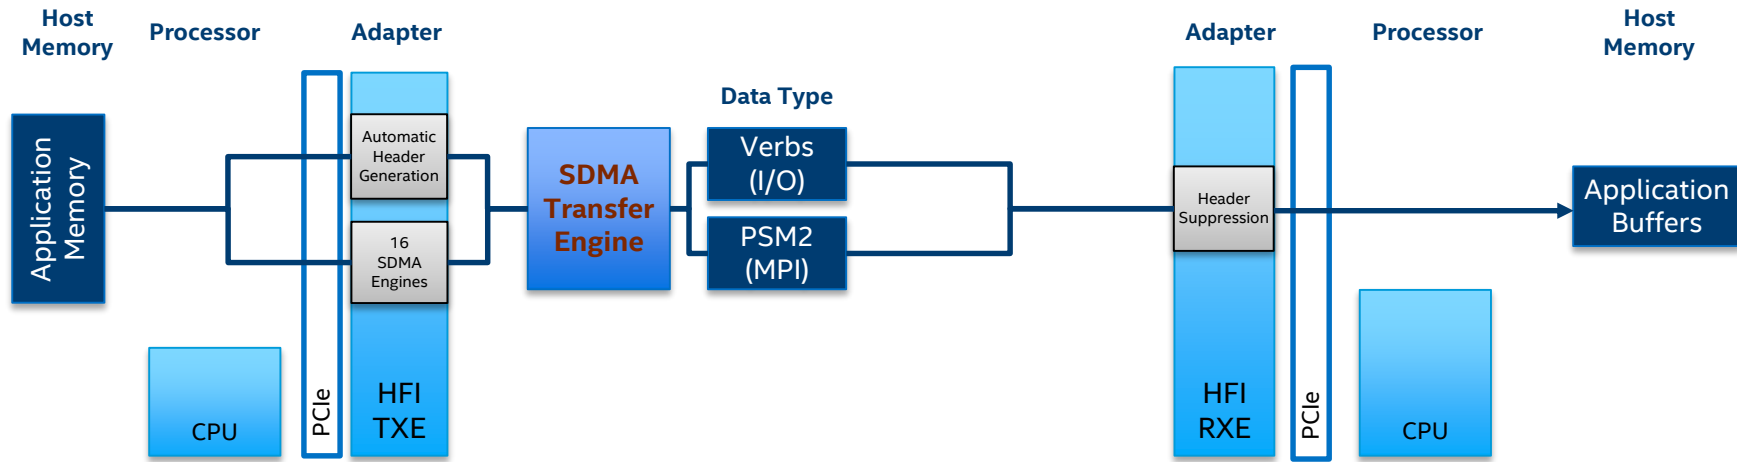
- ❖ Optimizes latency and message rate for high priority messages
- ❖ Transfer time lower than memory handle exchange, memory registration

Receive Buffer Placement

- ❖ Data placed in kernel receive buffers
- ❖ Buffers copied to application buffer

Multimodal Data Acceleration

Lowest overhead RDMA-based large message transfer



Send DMA (SDMA) Engine

- ❖ Stateless offloads on send side
- ❖ Connection setup required

Direct Data Placement

- ❖ Direct data placement on receive side
- ❖ Eliminates memory copy

Proven Technology Required for Today's Bids:

Intel® OPA is **the Future** of High Performance Fabrics



Aries

Highly Leverages
existing Aries and Intel®
True Scale technologies



Open Source software and supports
standards like the **OpenFabrics
Alliance***



Innovative Features
for high fabric performance,
resiliency, and QoS



Leading Edge Integration
with Intel® Xeon® processor
and Intel® Xeon Phi™ processor



Robust Ecosystem
of trusted computing
partners and providers

*Other names and brands may be claimed as property of others.

BACKUP

Intel Fabric Solutions Powered by Intel® Omni-Path Architecture

PCIe Adapters



| Intel Part # | 100HFA018LS 100HFA018FS | 100HFA016LS 100HFA016FS |
|--|--|--|
| Description | Single-port PCIe x8 Adapter, Low Profile and Std Height | Single-port PCIe x16 Adapter, Low Profile and Std Height |
| Availability ¹ | Q3'16 | Shipping |
| Speed | 58 Gbps | 100 Gbps |
| Ports, Media | Single port, QSFP28 | Single port, QSFP28 |
| Form Factor | Low profile PCIe Std Height PCIe | Low profile PCIe Std Height PCIe |
| Features | Passive thermal – QSFP heatsink, supports up to Class 4 max optical transceivers | Passive thermal – QSFP heatsink, supports up to Class 4 max optical transceivers |
| Sandy Bridge | X | X |
| Ivy Bridge | X | X |
| Intel® Xeon® processor E5-2600 v3 (Haswell-EP) | ✓ | ✓ |
| Intel® Xeon® processor E5-2600 v4 (Broadwell-EP) | ✓ | ✓ |

Edge Switches



| Intel Part # | 100SWE48UF2 / R2 100SWE48QF2 / R2 | 100SWE24UF2 / R2 100SWE24QF2 / R2 | 100SWD24B1N 100SWD24B1D 100SWD24B1A | 100SWD06B1N 100SWD06B1D 100SWD06B1A | 100SWDLF32Q | 100SWDSPINE | 100SWDMGTSH |
|---------------------------|--|--|--|---|---------------------------------------|--|---|
| Description | 48 Port Edge Switch (*Q* = mgmt card) | 24 Port Edge Switch (*Q* = mgmt card) | 24-slot Director Class Switch, Base Config | 6-slot Director Class Switch, Base Config | Director Class Switch Leaf Module | Director Class Switch Spine Module | Director Class Switch Management Module |
| Availability ¹ | Shipping | Shipping | Shipping | Shipping | Shipping | Shipping | Shipping |
| Speed | 100 Gbps | 100 Gbps | 100 Gbps | 100 Gbps | 100 Gbps | 100 Gbps | 100 Gbps |
| Max External Ports | 48 | 24 | 768 | 192 | 32 | N/A | N/A |
| Media | QSFP28 | QSFP28 | 10/100/1000 Base-T USB Gen2 | 10/100/1000 Base-T USB Gen2 | QSFP28 | Internal high speed connections | 10/100/1000 Base-T USB Gen2 |
| Form Factor | 1U | 1U | 20U | 7U | Half-width module, 2 modules per leaf | Full width module, 2 boards/module | Half-width module |
| Features | Forward / reverse airflow and mgmt card options, up to 2 PSU | Forward / reverse airflow and mgmt card options, up to 2 PSU | Up to 2 mgmt modules, up to 12 PSUs, AC and DC options | Up to 2 mgmt modules, up to 6 PSUs, AC and DC options | Hot swappable | 96 internal mid-plane connections, hot swappable | N+1 redundancy, hot swappable |

Director Switches



Passive Copper Cables



| 0.5M | 1.0M | 1.5M | 2.0M | 3.0M |
|--|--|-------------------------|-------------------------|-------------------------|
| 100CQQF3005 100CQQH3005 (30 AWG) | 100CQQF3010 100CQQH3010 (30 AWG) | 100CQQH2615 (26 AWG) | 100CQQH2620 (26 AWG) | 100CQQH2630 (26 AWG) |

Active Optical Cables



| 3.0M | 5.0M | 10M | 15M | 20M | 30M | 50M | 100M |
|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
| 100FRRF0030 | 100FRRF0050 | 100FRRF0100 | 100FRRF0150 | 100FRRF0200 | 100FRRF0300 | 100FRRF0500 | 100FRRF1000 |

¹ Production Readiness / General Availability dates

Intel® Omni-Path Host Fabric Interface

100 Series Single Port¹

Low Profile PCIe Card

- 2.71"x 6.6" max. Spec compliant.
- Standard and low profile brackets

Wolf River (WFR-B) HFI ASIC

PCIe Gen3


Single 100 Gb/s Intel® OPA port

- QSFP28 Form Factor
- Supports multiple optical transceivers
- Single Link status LED (Green)


| Power | Copper | | Optical (3W QSFP) | |
|---------|---------|---------|-------------------|---------|
| Model | Typical | Maximum | Typical | Maximum |
| X16 HFI | 7.4W | 11.7W | 10.6W | 14.9W |
| X8 HFI | 6.3W | 8.3W | 9.5W | 11.5W |

Thermal

- Passive thermal - QSFP Port Heatsink
- Standard 55C, 200lfm environment



x16 HFI
(100Gb Throughput)



x8 HFI
(~58Gb Throughput)
PCIe Limited

¹Specifications contained in public Product Briefs.

Intel® Omni-Path Edge Switch

100 Series 24/48 Port: Features¹

Compact Space (1U)

- 1.7"H x 17.3"W x 16.8"L

Switching Capacity

- 4.8/9.6 Tb/s switching capability

Line Speed

- 100Gb/s Link Rate

Standards-based Hardware Connections

- QSFP28

Redundancy

- N+N redundant Power Supplies (optional)
- N+1 Cooling Fans (speed control, customer changeable forward/reverse airflow)

Management Module (optional)

No externally pluggable FRUs

| Power | Copper | | Optical (3W QSFP) | |
|----------|---------|---------|-------------------|---------|
| Model | Typical | Maximum | Typical | Maximum |
| 24-Ports | 146W | 179W | 231W | 264W |
| 48-Ports | 186W | 238W | 356W | 408W |

24-port
Edge Switch



48-port
Edge Switch



¹Specifications contained in public Product Briefs.

Intel® OPA Director Class Systems 100 Series

6-Slot/24-Slot Systems¹

Highly Integrated

- 7U/20U plus 1U Shelf

Switching Capacity

- 38.4/153.6 Tb/s switching capability

Common Features

- Intel® Omni-Path Fabric Switch Silicon 100 Series (100Gb/s)
- Standards-based Hardware Connections – QSFP28
- Up to Full bisectional bandwidth Fat Tree internal topology
- Common Management Card w/Edge Switches
- 32-Port QSFP28-based Leaf Modules
- Air-cooled, front to back (cable side) air cooling
- Hot-Swappable Modules
 - Leaf, Spine, Management, Fan , Power Supply
- Module Redundancy
 - Management (N+1), Fan (N+1, Speed Controlled), PSU (DC, AC/DC)
- System Power : 180-240AC

| Power | Copper | | Optical (3W QSFP) | |
|---------|---------|---------|-------------------|---------|
| Model | Typical | Maximum | Typical | Maximum |
| 6-Slot | 1.6kW | 2.3kW | 2.4kW | 3.0kW |
| 24-Slot | 6.8kW | 8.9kW | 9.5kW | 11.6kW |

**6-Slot
Director Switch**




**24-Slot
Director Switch**

¹Specifications contained in public Product Briefs.

Intel® OPA Director Switch Chassis Configuration

QSFP-based Leaf Module
7U chassis


192 port configuration



16 * 6 leaf modules [12 total switch chips]
= 192 user ports

QSFP-based Leaf Module
20U chassis

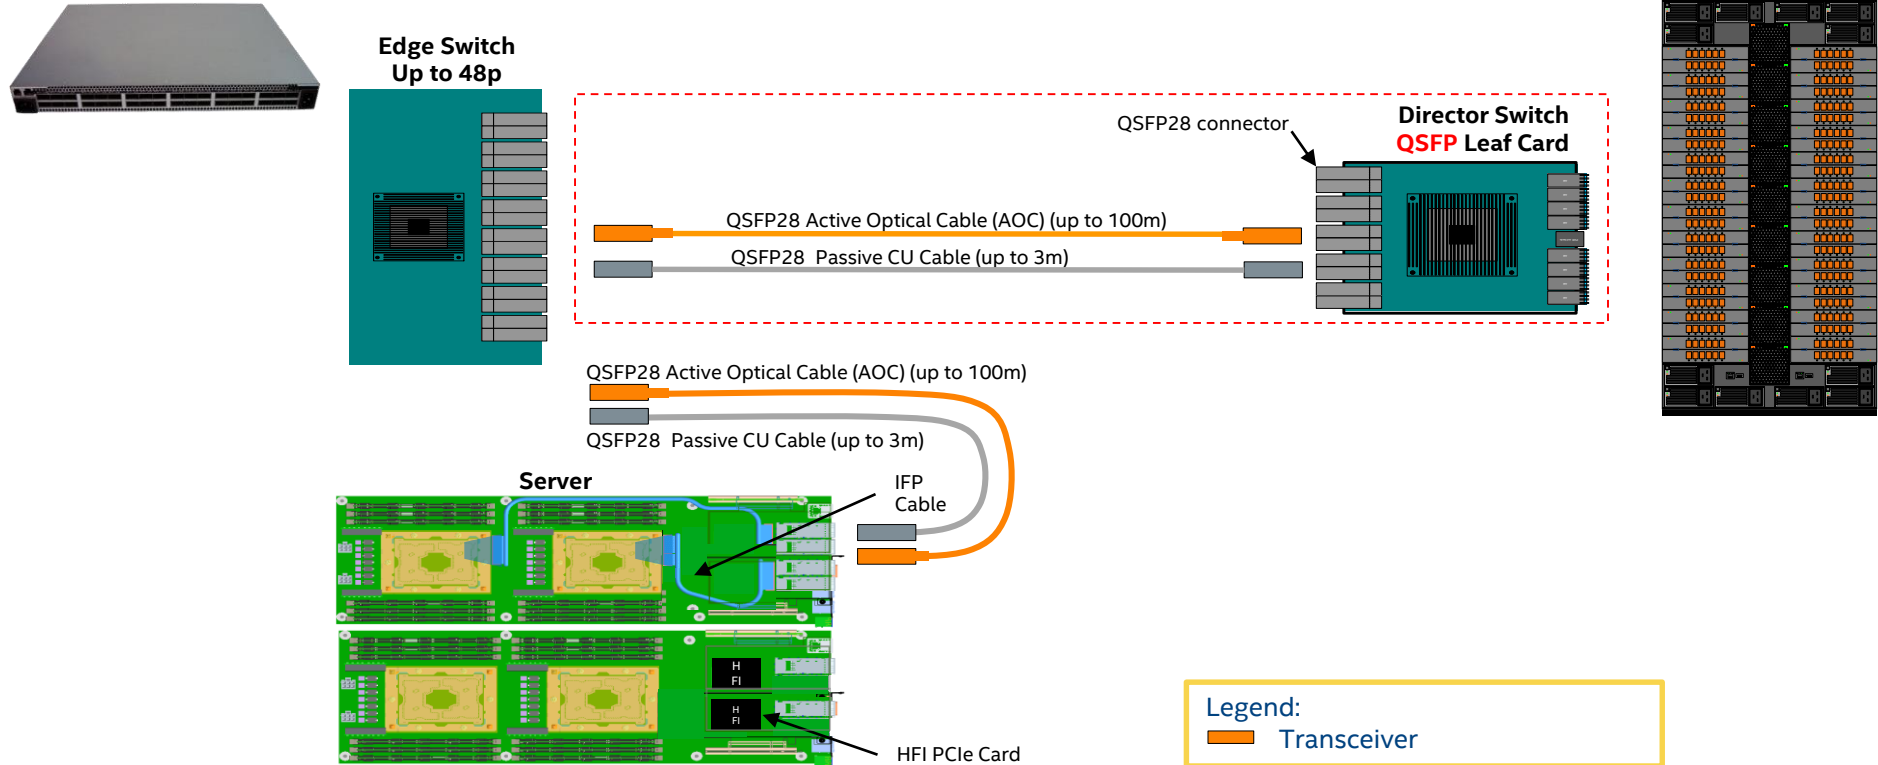
768 port configuration



16 * 24 leaf modules [48 total switch chips]
= 768 user ports

Full bisectional bandwidth for all configurations

Intel® Omni-Path Architecture Fabric Cabling Topology



Intel® OPA Standard Product SKUs:

Edge and Director Class Switches

| CATEGORY | INTEL MM# | PRODUCT CODE | Intel Branding String |
|-----------------|-----------|--------------|---|
| Edge Switch | 948588 | 100SWE48QF2 | Intel® Omni-Path Edge Switch 100 Series 48 Port Managed Forward 2 PSU 100SWE48QF2 |
| Edge Switch | 948678 | 100SWE48UF2 | Intel® Omni-Path Edge Switch 100 Series 48 Port Forward 2 PSU 100SWE48UF2 |
| Edge Switch | 945654 | 100SWE24QF2 | Intel® Omni-Path Edge Switch 100 Series 24 Port Managed Forward 2 PSU 100SWE24QF2 |
| Edge Switch | 945655 | 100SWE24UF2 | Intel® Omni-Path Edge Switch 100 Series 24 Port Forward 2 PSU 100SWE24UF2 |
| Edge Switch | 945662 | 100SWE48QF1 | Intel® Omni-Path Edge Switch 100 Series 48 Port Managed Forward 1 PSU 100SWE48QF1 |
| Edge Switch | 945663 | 100SWE48UF1 | Intel® Omni-Path Edge Switch 100 Series 48 Port Forward 1 PSU 100SWE48UF1 |
| Edge Switch | 945664 | 100SWE24QF1 | Intel® Omni-Path Edge Switch 100 Series 24 Port Managed Forward 1 PSU 100SWE24QF1 |
| Edge Switch | 945665 | 100SWE24UF1 | Intel® Omni-Path Edge Switch 100 Series 24 Port Forward 1 PSU 100SWE24UF1 |
| Edge Switch | 945775 | 100SWEQ7CN1 | Intel® Omni-Path Edge Switch Management Card 100 Series 100SWEQ7CN1 |
| Edge Switch | 945820 | 100SWEKIT1 | Intel® Omni-Path Edge Switch Installation Kit 100 Series 100SWEKIT1 |
| Director Switch | 947192 | 100SWD06CHS | Intel® Omni-Path Director Class Switch 100 Series 6 Slot FRU Chassis 100SWD06CHS |
| Director Switch | 947193 | 100SWD24CHS | Intel® Omni-Path Director Class Switch 100 Series 24 Slot FRU Chassis 100SWD24CHS |
| Director Switch | 945676 | 100SWD06B1N | Intel® Omni-Path Director Class Switch 100 Series 6 Slot Base 1MM 100SWD06B1N |
| Director Switch | 945677 | 100SWD24B1N | Intel® Omni-Path Director Class Switch 100 Series 24 Slot Base 1MM 100SWD24B1N |
| Director Switch | 945776 | 100SWDMGTSH | Intel® Omni-Path Director Switch Management Module 100 Series 100SWDMGTSH |
| Director Switch | 945777 | 100SWDLF32Q | Intel® Omni-Path Director Switch Leaf Module 100 Series 32 port 100SWDLF32Q |
| Director Switch | 945778 | 100SWDSPINE | Intel® Omni-Path Director Switch Spine Module 100 Series 100SWDSPINE |
| Director Switch | 945779 | 100SWDFAN01 | Intel® Omni-Path Director Switch Fan Module 100 Series 100SWDFAN01 |
| Director Switch | 945780 | 100SWDPS001 | Intel® Omni-Path Director Switch Power Supply Module 100 Series 100SWDPS001 |
| Director Switch | 945781 | 100SWDLFFPN | Intel® Omni-Path Director Switch Leaf Filler Panel 100 Series 100SWDLFFPN |
| Director Switch | 945834 | 100SWDSPFPN | Intel® Omni-Path Director Switch Spine Filler Panel 100 Series 100SWDSPFPN |
| Director Switch | 945835 | 100SWDMSFPN | Intel® Omni-Path Director Switch Management Filler Panel 100 Series 100SWDMSFPN |
| Director Switch | 945870 | 100SWDPSFPN | Intel® Omni-Path Director Switch Power Supply Filler Panel 100 Series 100SWDPSFPN |
| Director Switch | 945832 | 100SWDIKT06 | Intel® Omni-Path Director Switch Installation Kit 100 Series 6 Slot 100SWDIKT06 |
| Director Switch | 945833 | 100SWDIKT24 | Intel® Omni-Path Director Switch Installation Kit 100 Series 24 Slot 100SWDIKT24 |

Intel® OPA Standard Product SKUs:

HFI Adapters and Warranty Extensions

| CATEGORY | INTEL MM# | PRODUCT CODE | Intel Branding String |
|------------------|-----------|--------------|---|
| HFI Adapter Card | 948159 | 100HFA016LS | Intel® Omni-Path Host Fabric Interface Adapter 100 Series 1 Port PCIe x16 Low Profile 100HFA016LS |
| HFI Adapter Card | 945670 | 100HFA018LS | Intel® Omni-Path Host Fabric Interface Adapter 100 Series 1 Port PCIe x8 Low Profile 100HFA018LS |
| Warranty Ext. | 945144 | 100SWE24WE1 | Intel® Omni-Path Edge Switch 24 port Warranty Extension 1 year 100SWE24WE1 |
| Warranty Ext. | 945146 | 100SWE48WE1 | Intel® Omni-Path Edge Switch 48 port Warranty Extension 1 year 100SWE48WE1 |
| Warranty Ext. | 945145 | 100SWE24WE3 | Intel® Omni-Path Edge Switch 24 port Warranty Extension 3 year 100SWE24WE3 |
| Warranty Ext. | 945147 | 100SWE48WE3 | Intel® Omni-Path Edge Switch 48 port Warranty Extension 3 year 100SWE48WE3 |
| Warranty Ext. | 945148 | 100SWD06WE1 | Intel® Omni-Path Director Switch 6 slot Warranty Extension 1 year 100SWD06WE1 |
| Warranty Ext. | 945150 | 100SWD24WE1 | Intel® Omni-Path Director Switch 24 slot Warranty Extension 1 year 100SWD24WE1 |
| Warranty Ext. | 945149 | 100SWD06WE3 | Intel® Omni-Path Director Switch 6 slot Warranty Extension 3 year 100SWD06WE3 |
| Warranty Ext. | 945151 | 100SWD24WE3 | Intel® Omni-Path Director Switch 24 slot Warranty Extension 3 year 100SWD24WE3 |

Intel® OPA Standard Product SKUs:

Passive Copper and Active Optical Cables

| CATEGORY | INTEL MM# | PRODUCT CODE | Intel Branding String |
|----------------------|-----------|--------------|--|
| Passive Cu Cable | 947013 | 100CQQF3005 | Intel® Omni-Path Cable Passive Copper Cable QSFP-QSFP F 30AWG 0.5M 100CQQF3005 |
| Passive Cu Cable | 947012 | 100CQQF3010 | Intel® Omni-Path Cable Passive Copper Cable QSFP-QSFP F 30AWG 1.0M 100CQQF3010 |
| Passive Cu Cable | 947001 | 100CQQH3005 | Intel® Omni-Path Cable Passive Copper Cable QSFP-QSFP H 30AWG 0.5M 100CQQH3005 |
| Passive Cu Cable | 947005 | 100CQQH3010 | Intel® Omni-Path Cable Passive Copper Cable QSFP-QSFP H 30AWG 1.0M 100CQQH3010 |
| Passive Cu Cable | 947791 | 100CQQH2615 | Intel® Omni-Path Cable Passive Copper Cable QSFP-QSFP H 26AWG 1.5M 100CQQH2615 |
| Passive Cu Cable | 947010 | 100CQQH2620 | Intel® Omni-Path Cable Passive Copper Cable QSFP-QSFP H 26AWG 2.0M 100CQQH2620 |
| Passive Cu Cable | 947011 | 100CQQH2630 | Intel® Omni-Path Cable Passive Copper Cable QSFP-QSFP H 26AWG 3.0M 100CQQH2630 |
| Active Optical Cable | 947790 | 100FRRF0030 | Intel® Omni-Path Cable Active Optical Cable QSFP-QSFP F 3.0M 100FRRF0030 |
| Active Optical Cable | 947634 | 100FRRF0050 | Intel® Omni-Path Cable Active Optical Cable QSFP-QSFP F 5.0M 100FRRF0050 |
| Active Optical Cable | 947625 | 100FRRF0100 | Intel® Omni-Path Cable Active Optical Cable QSFP-QSFP F 10.0M 100FRRF0100 |
| Active Optical Cable | 947764 | 100FRRF0150 | Intel® Omni-Path Cable Active Optical Cable QSFP-QSFP F 15.0M 100FRRF0150 |
| Active Optical Cable | 947626 | 100FRRF0200 | Intel® Omni-Path Cable Active Optical Cable QSFP-QSFP F 20.0M 100FRRF0200 |
| Active Optical Cable | 947627 | 100FRRF0300 | Intel® Omni-Path Cable Active Optical Cable QSFP-QSFP F 30.0M 100FRRF0300 |
| Active Optical Cable | 947792 | 100FRRF0500 | Intel® Omni-Path Cable Active Optical Cable QSFP-QSFP F 50.0M 100FRRF0200 |
| Active Optical Cable | 947628 | 100FRRF1000 | Intel® Omni-Path Cable Active Optical Cable QSFP-QSFP F 100.0M 100FRRF0300 |

PERFORMANCE BACKUP TEST CONDITIONS

Legal Notices and Disclaimers

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer.

No computer system can be absolutely secure.

Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit <http://www.intel.com/performance>.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Statements in this document that refer to Intel's plans and expectations for the quarter, the year, and the future, are forward-looking statements that involve a number of risks and uncertainties. A detailed discussion of the factors that could affect Intel's results and plans is included in Intel's SEC filings, including the annual report on Form 10-K.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Intel, the Intel logo and others are trademarks of Intel Corporation in the U.S. and/or other countries. *Other names and brands may be claimed as the property of others.

© 2016 Intel Corporation.

Optimization Notice

Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel.

Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice revision #20110804

