

# PRODUCT BRIEF

High Performance Computing (HPC), Data Analytics and Artificial Intelligence  
Intel® Omni-Path Director Class Switch 100 Series  
26-slot and 6-slot Director Switched (100 Gbps per port)



# Delivering High-Performance HPC Networking with Attractive Cost Models

## The Right Fabric for HPC and AI

### Benefits

- Highly integrated design reduces space, power and cost
  - Up to 1,152 ports in 20U
  - Up to 288 ports in 7U
  - Expandable in 32/48 port
- Optimized for high message rates and low end-to-end MPI latency at scale
- Simple generational upgrades with:
  - Binary compatible applications
    - OFA Compliant, Same as InfiniBand\*
  - FastFabric Installation/Tests Toolset
  - All open-source software

### Key Features

- Up to 28.8 Terabytes of aggregate BW
- Fully Redundant Subsystems
  - Spine/Management/Power/Cooling

### Next Generation Fabric Innovations

- Link Protection
  - Packet Integrity Protection (PIP)
    - “No-Load” Error Detection and Correction
  - Dynamic Lane Scaling (DLS)
    - Uptime Protection
- Quality of Service (QoS)
  - Traffic Fabric Optimization (TFO)
  - Larger MTUs (8/10K): Traffic Efficiency eg. Storage traffic

## Overview

Intel® Omni-Path Director Class Switches are the cornerstone of an end-to-end product family for HPC fabrics that delivers high performance with highly attractive cost models<sup>1</sup>. Intel® Omni-Path Architecture (Intel® OPA) builds on proven technologies including the previous generation Intel® True Scale Architecture and the Cray\* Aries interconnect and open source software to provide an evolutionary on-ramp to revolutionary new fabric capabilities.

## Higher Performance at Lower Cost

Intel Omni-Path Director Class Switches deliver 100 Gbps port bandwidth with latency that stays low even at extreme scale. Based on Intel's 48-port switch silicon, these switches can lower fabric acquisition costs<sup>1</sup> versus lower radix switches while simultaneously reducing space and power requirements. With these savings, you can achieve higher total cluster performance within the same hardware budget to expand and accelerate your research.

## Flexible Fabrics at Every Scale

Intel Omni-Path Director Class Switches support mid-level clusters to supercomputers with 10s of thousands of servers. These switches can be used in combination with Intel Omni-Path Edge Switches to build low-latency, multi-tier fabrics with an exceptional set of features for high-speed networking.

## Quality of Service (QoS)

All Intel OPA products feature excellent Quality of Service features including high availability that includes power, cooling and management module redundancy, component-level diagnostics and alarming, and out-of-band management. Beyond hardware availability, innovative Intel OPA features provide additional QoS options to the fabric. These include:

- **Packet Integrity Protection (PIP)** brings the ability to monitor for link errors with no added performance load, then resolves them.
- **Traffic Fabric Optimization (TFO)** can add traffic prioritization of data on an inter-switch link in 64-bit increments.
- **Dynamic Lane Scaling (DLS)** enables continuing operation of a link in a degraded mode, allowing the running job to complete, in the case of a lane failure.



## Physical Specifications

### CHASSIS CONFIGURATIONS: 100SWD24 AND 100SWD06

FEATURE	100SWD24		100SWD06	
100 Gb port equivalents (min)	32 (add in 32 or 48 port increments)		32 (add in 32 or 48 port increments)	
100 Gb port equivalents (max)	1,152		288	
Total System Bandwidth (bi-di)	28.8 TB/s		7.2 TB/s	
Chassis Height (+ Support Tray)	20U (+0U)		7U	
Dimensions (w x h x d) (in/cm)	17.6 x 35.0 x 29.5/44.7 x 88.9 x 74.9		17.6 x 12.2 x 29.5/44.7 x 31.0 x 74.9	
Weight (Empty FRU Chassis)	256 lb/116.2 kg		103.5 lb/47 kg	
Weight (Base 1N: All Fans, 1 Mgmt Module, min PSUs)	311.5 lb/141.3 kg		123.1 lb/55.9 kg	
Weight (fully configured)	659.7 lb/299.2 kg		214.8 lb/97.4 kg	
Packaged Dimensions (w x h x d) (in/cm)	36.0 x 48.0 x 40.0/91.4 x 121.9 x 101.6		36.0 x 25.0 x 40.0/91.4 x 63.5 x 101.6	
Packaged Weight (fully configured/palletized)	716 lb/324.8 kg		278.0 lb/126.1 kg	
Leaf Modules (max)	24		6	
Cables	32-Port Leaf QSFP28   QSFP28	48-Port Leaf QSFP-DD   Dual QSFP28	32-Port Leaf QSFP28   QSFP28	48-Port Leaf QSFP-DD   Dual QSFP28
Spine Modules (max)	12 (48-port Leaf) 8 (32-port Leaf) 12 (48/32 port mixed)		3 (48-port leaf) 2 (32-port Leaf) 3 (48/32 port mixed)	
Fan Modules	9		3	
Management Modules (standard/redundant)	1 / 2		1 / 2	
Power Supplies (min/DC/AC redundancy)	6 / 7 / 12		2 / 3 / 4	
Power (KW AC, typical-fully loaded, copper cables)†	8.5		1.98	
Power (KW AC, max-fully loaded, copper cables)††	11.2		2.6	
Optical Transceiver—per port QSFP28 QSFP-DD	Class 2/3 (3.5W Max) 3.5W Max		Class 2/3 (3.5W Max) 3.5W Max	
Cooling (max required)	1100 CFM at 40°C		390 CFM at 40°C	
Status LEDs††† (Ethernet/DC_On)	2 / 1		2 / 1	

† Typical: Fully populated chassis (max leaves, max spines, 2MM, AC Redundant PSUs), fully populated copper cables, running looptest (estimated)

†† Max, Fully populated chassis (max leaves, max spines, 2MM, AC Redundant PSUs), fully populated copper cables, worst process/temp/voltage (estimated)

††† Status LEDs - Ethernet Activity (Green), Ethernet Speed (Green/Orange) / DC\_On (Green on push button)

## MODULES

MODULE	DIMENSIONS (W x H x D) (in/cm)	WEIGHT	PACKAGE DIM (W x H x D) (in/cm)	PACKAGED WEIGHT	POWER (W, typical)†	POWER (W, max)††
QSFP 32p Leaf Module	6.9 x 2.2 x 12.3/ 17.5 x 5.6 x 31.2	8.5 lb/3.86 kg	11.5 x 6.5 x 17.0/ 29.2 x 16.5 x 43.2	10.0 lb/4.54 kg	175 W DC 189 W AC	229 W DC 247.5 W AC
QSFP-DD 48p Leaf Module	6.9 x 2.2 x 12.4/ 17.5 x 5.6 x 31.4	8.5 lb/3.86 kg	11.5 x 6.5 x 17.0/ 29.2 x 16.5 x 43.2	10.0 lb/4.54 kg	207 W DC 223.5 W AC	271 W DC 293.5 W AC
Spine Module	17.4 x 2.2 x 7.5/ 44.2 x 5.6 x 19.1	8.6 lb/3.9 kg	22.5 x 6.5 x 11.5/ 57.2 x 16.5 x 29.2	10.5 lb/4.76 kg	203 W DC 219 W AC	262 W DC 283 W AC
Fan Module	3.2 x 3.3 x 8.7/ 8.1 x 8.4 x 22.1	1.6 lb/0.73 kg	7.5 x 7.5 x 11.0/ 19.1 x 19.1 x 27.9	2.40 lb/1.09 kg	45 W DC 48.5 W AC	64 W DC 69 W AC
Power Supply Module†††	4.0 x 1.6 x 15.8/ 10.2 x 4.1 x 40.1	6.55 lb/2.98 kg	8.0 x 5.5 x 19.5/ 20.3 x 14.0 x 49.5	7.55 lb/3.42 kg	94% eff. at 230 V/50% load	91% eff. at 230 V/max load
Management Module	8.6 x 1.7 x 7.2/ 21.8 x 4.3 x 18.3	1.75 lb/0.79 kg	11.0 x 4.9 x 9.5/ 28.0 x 12.5 x 24.1	2.65 lb/1.2 kg	17W DC 18.5W AC	25 W DC 27 W AC

† Typical: copper cables for Leaf (all estimated).

†† Max: copper cables for Leaf, worst process/temp/voltage (all estimated).

††† Weight Includes AC Cord; When a form/fit option exists, use of a DC/DC supply may be possible.

# Environmental Specifications

## (planned for 48p Leaf Module#)

### Temperature

Operating: 5° to 40° C, supporting AOC up to 2 W (Temperature de-rating 5°C/W up to 3 W AOC) supporting QSFP-DD AOC up to 3.5 W

Storage: -40° to 70° C

### Humidity

Operating: 5% to 85% non-condensing

Storage: 5% to 95% non-condensing

### Altitude

Operating: 0 – 10,000 feet (Temperature de-rating 1°C / 175 M above 900 M)

Storage: 0 – 40,000 feet

### Shock

Unpackaged: Half-sine, 2g 11ms 300 pulses total

Packaged: 9" vertical and rotational drop

### Vibration

Unpackaged: 5-500 Hz, 2.2 g RMS random

Packaged: 5-500 Hz, 1.09 g RMS random

### Airflow

Front-to-back

1100 CFM maximum for 24 slot

390 CFM maximum for 6 slot

### RoHS/REACH

Complies with RoHS II Directive 2011/65/EU of the European Parliament

Complies with REACH Regulation (EC) No 1907/2006

### Acoustics

Less than 7.0 Bels

### Agency Approvals - EMI and EMC

#### US/Canada

FCC Part 15, Subpart B, Class A

CAN ICES-3 (A)

#### Europe/International

CISPR22

CISPR32/EN55032

EN55024

EN61000-3-2

EN61000-3-3

#### Japan

VCCI, Class A

#### New Zealand/Australia

AS/NZS CISPR 22, Class A

#### Korea

RRA/KC (KN32, KN35), Class A

#### Taiwan

BSMI (CNS 13438), Class A

#### Customs Union: Russian Federation, Belarus and Kazakhstan

TR CU 020/2011 "Electromagnetic compatibility of technical equipment"

#### Agency Approvals – Safety (planned for 48p Leaf Module#)

#### US/Canada

TUV NRTL: UL 60950-1, CSA 22.1.No. 60950-1

#### Europe

TUV: EN60950-1

#### International

CB Scheme: IEC 60950-1

#### Customs Union: Russian Federation, Belarus and Kazakhstan

TR CU 004/2011 "On Safety of Low-Voltage Equipment"

#### Country of Origin

USA

#### Miscellaneous

#### Non-volatile memories

A link will be provided to the statement of volatility document

<sup>1</sup>A sample configuration assumes a 750-node cluster, and number of switch chips required is based on a full bisectonal bandwidth (FBB) Fat-Tree configuration. Intel® OPA uses one fully-populated 768-port director switch, and Mellanox EDR solution uses a combination of 648-port director switches and 36-port edge switches. Mellanox component pricing from [www.kernelsoftware.com](http://www.kernelsoftware.com), with prices as of December 30, 2017.



You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

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

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

The products described 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.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting: <http://www.intel.com/design/literature.htm> Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. No computer is absolutely secure. Learn more at <http://www.intel.com/> or from the OEM or retailer.

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. Power and cooling data is based on this Product Brief and datasheets available on Mellanox.com

QoS features are described in detail in internal Intel documents compared to published InfiniBand standards

Intel and the Intel logo 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.

Copyright © 2018, Intel Corporation. All rights reserved

052018/SY/RJM/J/PDF

337582-001

Please Recycle

H30385 Rev. 1.5