

Overview

HPE Slingshot

HPE Slingshot is a modern high-performance interconnect for HPC and AI clusters that delivers industry leading performance, bandwidth, and low latency for HPC, AI/ML, and data analytics applications. It is Cray's 8th generation interconnect architecture with unique and powerful innovations engineered from the silicon up to combine the best of a high-performance supercomputing-class interconnect with Ethernet interoperability. This enables a unified high-performing cluster that delivers great performance on both traditional modeling and simulation codes alongside native sockets-based applications, and with direct Ethernet connectivity to storage without requiring gateway nodes.

Beyond raw speed (25.6 Tbps of bidirectional bandwidth per switch and 200 Gbps endpoints), HPE Slingshot dramatically controls loaded latency, the true determinant of realized performance on tightly coupled HPC and AI workloads. Innovative congestion management reduces tail latency and run-to-run execution time variability, and fine-grained adaptive routing ensures high utilization of the bandwidth available.

With HPE Slingshot, the same interconnect in the world's most powerful supercomputers can now be deployed across a wide spectrum of HPC and AI/ML clusters based on HPE Apollo, HPE ProLiant, and HPE Cray XD server systems. This document focuses on standalone HPE Slingshot components, systems, and software for such clusters. For information about HPE Slingshot blades and elements used in HPE Cray Supercomputing EX systems, please refer to the [Cray EX QuickSpecs](#). In addition, please refer to the [Supplemental Configuration and Integration Guide](#) for in-depth guidance on designing and implementing HPC and AI/ML clusters with HPE Slingshot.

HPE Slingshot Highlights

High Bandwidth and Low Latency

HPE Slingshot switches provide 25.6 Tbps of bi-directional bandwidth with 64 ports each capable of 200 Gbps, delivering high bandwidth and low latency simultaneously, overcoming traditional Ethernet approaches to buffering and error handling that result in queuing latency and hinder low-latency RDMA applications.

Automatic Congestion Management

The HPE Slingshot network uniquely addresses the congestion challenges of short-lived, small packet flows typical in HPC applications. In hardware, each switch detects congestion, identifies its causes, and provides real-time feedback to its peers. The system distinguishes perpetrators from victims of injection and limits the injection rate from the perpetrator until congestion clears.

Fine-Grained Adaptive Routing

HPE Slingshot tracks real-time information on load across each switch-to-switch path, and dynamically re-routes traffic to balance load (rather than sticking only to a predetermined path) to improve network performance and bandwidth utilization. Uniquely, HPE Slingshot's adaptive routing decisions are made on a packet-by-packet basis for unordered traffic, or a flow-by-flow basis for protocols (like TCP/IP) requiring ordered delivery.

Quality of Service

HPE Slingshot supports highly configurable quality-of-service (QoS) that provides the ability to control how network bandwidth is allocated to different classes of traffic and applications.

Low-Diameter Network

The high-radix 64-port HPE Slingshot switches enable low-diameter topologies to reduce the quantity of network equipment, cabling, power and cooling costs. The initial topology offered is a 3-hop dragonfly topology based on a mesh of all-to-all that further reduces latency and improves sustained bandwidth and reliability. Fine-grained adaptive routing and congestion management virtually eliminate any locality concerns related to job placement and storage-connectivity in the topology.

Ethernet Standards

The HPE Slingshot system interconnect implements a full range of Ethernet standards, with hardware support focusing on the data movement protocols, IPv4 and IPv6 in particular.

Standard Features

HPE Slingshot for HPC Clusters consists of the following elements:

HPE Slingshot Top-of-Rack Switch

Based on HPE-designed silicon, the high-performance HPE Slingshot switch provides 64-ports at 200 Gbps each for 25.6 Tbps of bidirectional bandwidth. Ports are configured either as edge ports for connectivity to supported Ethernet NICs or external routers at 100GE or 200GE, or as 200 Gbps fabric ports that link the switches together with enhanced resiliency and performance.

HPE Slingshot Network Interface Card

The HPE Slingshot Network Interface Card provides a single port of 200 Gbps connectivity to PCI-Express Gen 4-based host systems. It is based on HPE-designed silicon and is qualified in select HPE Apollo, HPE ProLiant, and HPE Cray XD servers. The HPE Slingshot NIC features high performance RDMA and extensive hardware acceleration for MPI and SHMEM based software. These HPC optimized features are exposed to software applications using the Libfabric API interface for RDMA software. Additionally, the NIC supports a broad range of host software natively using the Linux “sockets” interface for IP software.

HPE Slingshot also supports select 100 Gbps and 200 Gbps standard-based Ethernet adapters as an alternative option.

Cables

HPE Slingshot cables are designed and validated to reliably support the demanding signaling speed used in 200 Gbps links. Copper cables support shorter lengths, and Active Optical Cables provide connectivity up to 100m.


HPE Slingshot Fabric Management Suite software

The fabric manager provides configuration, monitoring, control, and networking functions for the entire HPE Slingshot network. This software runs on a standard ProLiant server (or a redundant pair of servers) and connects to the switches using the cluster’s out-of-band management network.

HPE Slingshot Top-of-Rack Switch

The HPE Slingshot Switch is streamlined from the silicon-level up for the highest throughput to minimize *queueing* latency, which is the main contributor to latency in practice. Optimizations ensure performance and reliability for the small granularity messages common in latency-sensitive HPC and AI workloads. For example, optimizations on links between switches reduce packet overheads and eliminate end-to-end timeouts and retransmissions that otherwise hinder use of Ethernet for latency sensitive applications. Meanwhile, adaptive routing and congestion control between switches ensure bandwidth is used extremely efficiently.

Each HPE Slingshot top of rack switch (TOR) contains 1 HPE Slingshot switch ASIC and supplies sixty-four 200 Gbps ports via 32 QSFP double density (QSFP-DD) connectors in a 1U 19” standard rack-mount form factor chassis. Each QSFP-DD physical port can provide either 2 x 100 Gbps or 2 x 200 Gbps network links. Ports can be configured to connect to HPE Slingshot switches or to edge devices including select HPE or 3rd party enterprise Ethernet switches or supported NICs. Redundant N+1 hot swappable power supplies and fans ensure availability. Management and monitoring operations utilize the cluster’s out-of-band management network.



Standard Features



Ordering Information

HPE Slingshot 200Gb 64-port QSFP-DD 1U Switch

R4K41B

HPE Slingshot 200 Gbps NIC

The HPE Slingshot SA210S Ethernet 200Gb 1-port PCIe NIC works with the HPE Slingshot Top-of-Rack Switches in supported HPE Apollo, ProLiant, and Cray XD rack servers. It delivers supercomputing-class performance on challenging latency-sensitive HPC codes without compromise while also running standard IP applications using native Ethernet IP protocols. The card supports either QSFP-56 200 Gbps (4x50 Gbps PAM 4) and QSFP-28 100 Gbps (4 x 25 Gbps NRZ) Ethernet standard connectivity.


Many features accelerate HPC and AI performance. MPI message matching offload overlaps communications and compute and frees up host memory bandwidth and CPU resources. The NIC reduces compute overhead for GPU initiated communications by allowing the host CPU to set up operations that the GPU can trigger asynchronously. The efficient command launch mechanism reduces latency by receiving and pipelining commands without reading host memory. And a reliable delivery mechanism detects packet loss and enables end-to-end retry in hardware to reduce software overhead.

Some HPE Slingshot fabric capabilities are extended to the NIC-to-switch port link when using the HPE Slingshot NIC versus industry standard NICs including enhanced link resiliency and fine-grained flow control. The latter helps the automatic congestion apply stiff backpressure to the specific flow causing congestion which delivers granularity over backpressure to the entire port.

The HPE Slingshot NIC exposes RDMA and HPC optimized features to software using the Libfabric software interface.



Standard Features



Ordering Information:

Consult HPE Slingshot for HPC Clusters reference architecture information for supported host systems.

HPE Slingshot SA210S Ethernet 200Gb 1-port PCIe NIC

R4K46A

HPE Slingshot Fabric Management Suite

The HPE Slingshot Fabric Management Suite provides the management, control, and network functions and is required to configure and operate the HPE Slingshot network. The software is installed on a standard server that communicates with the switches via a 1 or 10 Gbps management network. This separation of data and control planes allows for independent reliable data transfers and management operations. All network management traffic is encrypted with switch-based certificates.

The Fabric Management Suite supports comprehensive health checking, monitoring and logging, performance profiling, and diagnostics that assure high performance with high availability and effective operational processes. Event, error, and counter data is gathered from all the switches in the network. Open documented APIs and a scriptable command line interface allow for integration with cluster management and network monitoring tools. All of these APIs support AuthN based on client certificates.

Scaling, efficiency, and performance of management and operations are achieved via a hierarchical distributed control plane design. Multiple Fabric Management Suite nodes can optionally operate in an Active/Standby mode to deliver a high-availability configuration.

Ordering Information:

Consult HPE Slingshot for HPC Clusters reference architecture information for current details on supported host systems and recommended configurations.

Models

HPE Slingshot Fabric Manager E-RTU

R6M77AAE

HPE Slingshot Fabric Manager FIO Software

R6M78A

HPE Slingshot Cables

HPE validates HPE Slingshot cables to ensure reliability and performance running up to 200 Gbps with the switch and supported adapters. Because HPE Slingshot switches use double density ports, each QSFP-DD cable-end carries two 200 Gbps links, often reducing the total cable count to deploy. HPE offers both copper and Active Optical (AOC) cables with various lengths. However, certain restriction may apply for using copper cables for switch-to-switch or switch-to-adapter connections. Consult reference architectures for specific wiring guidelines that must be followed for proper operation.



Service and Support

HPE Services

No matter where you are in your digital transformation journey, you can count on HPE Services to deliver the expertise you need when, where and how you need it. From planning to deployment, ongoing operations and beyond, our experts can help you realize your digital ambitions.

<https://www.hpe.com/services>

Consulting Services

No matter where you are in your journey to hybrid cloud, experts can help you map out your next steps. From determining what workloads should live where, to handling governance and compliance, to managing costs, our experts can help you optimize your operations.

<https://www.hpe.com/services/consulting>

HPE Managed Services

HPE runs your IT operations, providing services that monitor, operate, and optimize your infrastructure and applications, delivered consistently and globally to give you unified control and let you focus on innovation.

[HPE Managed Services | HPE](#)

Operational services

Optimize your entire IT environment and drive innovation. Manage day-to-day IT operational tasks while freeing up valuable time and resources. Meet service-level targets and business objectives with features designed to drive better business outcomes.

<https://www.hpe.com/services/operational>

HPE Complete Care Service

HPE Complete Care Service is a modular, edge-to-cloud IT environment service designed to help optimize your entire IT environment and achieve agreed upon IT outcomes and business goals through a personalized experience. All delivered by an assigned team of HPE Services experts. HPE Complete Care Service provides:

- A complete coverage approach -- edge to cloud
- An assigned HPE team
- Modular and fully personalized engagement
- Enhanced Incident Management experience with priority access
- Digitally enabled and AI driven customer experience

<https://www.hpe.com/services/complecare>

HPE Tech Care Service

HPE Tech Care Service is the operational support service experience for HPE products. The service goes beyond traditional support by providing access to product specific experts, an AI driven digital experience, and general technical guidance to not only reduce risk but constantly search for ways to do things better. HPE Tech Care Service delivers a customer-centric, AI driven, and digitally enabled customer experience to move your business forward. HPE Tech Care Service is available in three response levels. Basic, which provides 9x5 business hour availability and a 2-hour response time. Essential which provides a 15-minute response time 24x7 for most enterprise level customers, and Critical which includes a 6-hour repair commitment where available and outage management response for severity 1 incidents.

<https://www.hpe.com/services/techcare>



Service and Support

HPE Lifecycle Services

HPE Lifecycle Services provide a variety of options to help maintain your HPE systems and solutions at all stages of the product lifecycle. A few popular examples include:

- Lifecycle Install and Startup Services: Various levels for physical installation and power on, remote access setup, installation and startup, and enhanced installation services with the operating system.
- HPE Firmware Update Analysis Service: Recommendations for firmware revision levels for selected HPE products, taking into account the relevant revision dependencies within your IT environment.
- HPE Firmware Update Implementation Service: Implementation of firmware updates for selected HPE server, storage, and solution products, taking into account the relevant revision dependencies within your IT environment.
- Implementation assistance services: Highly trained technical service specialists to assist you with a variety of activities, ranging from design, implementation, and platform deployment to consolidation, migration, project management, and onsite technical forums.
- HPE Service Credits: Access to prepaid services for flexibility to choose from a variety of specialized service activities, including assessments, performance maintenance reviews, firmware management, professional services, and operational best practices.

Notes: To review the list of Lifecycle Services available for your product go to:

<https://www.hpe.com/services/lifecycle>

For a list of the most frequently purchased services using service credits, see the [HPE Service Credits Menu](#)

Other Related Services from HPE Services:

HPE Education Services

Training and certification designed for IT and business professionals across all industries. Broad catalogue of course offerings to expand skills and proficiencies in topics ranging from cloud and cybersecurity to AI and DevOps. Create learning paths to expand proficiency in a specific subject. Schedule training in a way that works best for your business with flexible continuous learning options.

<https://www.hpe.com/services/training>

Defective Media Retention

An option available with HPE Complete Care Service and HPE Tech Care Service and applies only to Disk or eligible SSD/Flash Drives replaced by HPE due to malfunction.

Consult your HPE Sales Representative or Authorized Channel Partner of choice for any additional questions and services options.

Parts and Materials

HPE will provide HPE-supported replacement parts and materials necessary to maintain the covered hardware product in operating condition, including parts and materials for available and recommended engineering improvements.

Parts and components that have reached their maximum supported lifetime and/or the maximum usage limitations as set forth in the manufacturer's operating manual, product quick-specs, or the technical product data sheet will not be provided, repaired, or replaced as part of these services.

How to Purchase Services

Services are sold by Hewlett Packard Enterprise and Hewlett Packard Enterprise Authorized Service Partners:

- Services for customers purchasing from HPE or an enterprise reseller are quoted using HPE order configuration tools.
- Customers purchasing from a commercial reseller can find services at <https://ssc.hpe.com/portal/site/ssc/>



Service and Support

AI Powered and Digitally Enabled Support Experience

Achieve faster time to resolution with access to product-specific resources and expertise through a digital and data driven customer experience

Sign into the HPE Support Center experience, featuring streamlined self-serve case creation and management capabilities with inline knowledge recommendations. You will also find personalized task alerts and powerful troubleshooting support through an intelligent virtual agent with seamless transition when needed to a live support agent.

<https://support.hpe.com/hpesc/public/home/signin>

Consume IT On Your Terms

HPE GreenLake edge-to-cloud platform brings the cloud experience directly to your apps and data wherever they are—the edge, colocations, or your data center. It delivers cloud services for on-premises IT infrastructure specifically tailored to your most demanding workloads. With a pay-per-use, scalable, point-and-click self-service experience that is managed for you, HPE GreenLake edge-to-cloud platform accelerates digital transformation in a distributed, edge-to-cloud world.

- Get faster time to market
- Save on TCO, align costs to business
- Scale quickly, meet unpredictable demand
- Simplify IT operations across your data centers and clouds

To learn more about HPE Services, please contact your Hewlett Packard Enterprise sales representative or Hewlett Packard Enterprise Authorized Channel Partner. Contact information for a representative in your area can be found at "Contact HPE"

<https://www.hpe.com/us/en/contact-hpe.html>

For more information

<http://www.hpe.com/services>



Configuration Information

Each HPE Slingshot network is a custom tuned network configuration to achieve the best price/performance. Please contact your Hewlett Packard Enterprise representative for requirements to configure the solution correctly and for recent additions and details on supported configurations.

HPE strongly recommends HPE supplied rack-level integration of HPE Slingshot-based clusters.

HPE Slingshot Switch

HPE Slingshot 200Gb 64-port QSFP-DD 1U Switch

R4K41B

- 64 QSFP-DD ports externally managed switch
- 2 (1+1) Power supplies
- 2 Power cords, 1.5M, C13-C14
- 5 (N+1) Fan Units
- Sliding rail rackmount kit

Network Interface Cards

HPE Slingshot for Apollo and DL Clusters currently supports the following Network Interface cards:

HPE Slingshot SA210S Ethernet 200Gb 1-port PCIe NIC

R4K46A

HPE Ethernet 100Gb 1-port QSFP28 PCIe3 x16 MCX515A-CCAT Adapter

P31246-B21

HPE Ethernet 100Gb 1-port QSFP28 PCIe3 x16 MCX515A-CCAT Adapter

P31246-H21

HPE Slingshot for Apollo, ProLiant, and Cray XD HPC Clusters currently supports the following compute node servers:

- HPE Apollo 2000 and 6500 Gen 10+ based on AMD processors
- HPE ProLiant DL 325 Gen 10+ based on AMD processors
- HPE ProLiant DL 385 Gen10+ based on AMD processors
- HPE ProLiant DL345 Gen11 based on AMD processors
- HPE Cray XD2000 models with choice of Intel or AMD processors

Topologies

HPE Slingshot supports “dragonfly” network topologies that ensure great bandwidth using low-cost electrical links for most of the connections. Ports on every switch are configured either as downlinks that connect to server endpoints, “intra-group” links that connect to other switches in the dragonfly group, and “inter-group” links (or “global” links) that connect each group to all the other groups. Group size is chosen based on factors such as the maximum system size, the desired scalable unit(s) for building and expanding the cluster, density and rack size impact to the economic tradeoffs between switches and optical cables, and preferences such as physical grouping, placement of resources, and cable routing policies.

Note that groups of different sizes can be assembled into a single cluster. In addition, clusters with 64 or fewer endpoints can be supported with a single switch.

# of switches per group	Scalable Unit	Maximum Size	Max hops between nodes
1 switch	32 endpoints	544 endpoints	2
2 switches	32 endpoints	1056 endpoints	3
4 switches	64 endpoints	3136 endpoints	3
8 switches	128 endpoints	10,368 endpoints	3
16 switches	256 endpoints	37,120 endpoints	3

Cabling

Validated Direct Attach Copper (DAC) and Active Optical (AOC) cables are available in a wide range of lengths. While copper cables are offered in lengths up to 3 meters, consult [the Supplemental Configuration and Integration Guide](#) for specific switch port wiring guidelines that must be followed in order to incorporate supported lengths for passive copper cables.

When connecting switches to NIC cards in compute nodes or to edge routers that have QSFP connectors, edge (or “Y”) cables are used to split the QSFP-DD connector on the switch into two independent QSFP endpoints. Copper “Y” cables will work at for both

Configuration Information

200 Gbps and 100 Gbps QSFP endpoints. For AOC “Y” cables, specific SKUs support either 100 Gbps (QSFP-28) or 200 Gbps (QSFP-56) endpoints. (See reference architectures for tested edge router cabling compatibility).

In some topologies, bifurcated (or “H”) cables connect two double-density physical connections to two other double-density physical connections while crossing over individual 200 Gbps links.

HPE Slingshot Switch to NIC (QSFP-DD to QSFP “Y”) Cables

Copper Cables

HPE Slingshot L0 QSFP-DD 2x QSFP56 1m Splitter Cable	R4K54A
HPE Slingshot L0 QSFP-DD 2x QSFP56 1.5m Splitter Cable	R4K55A
HPE Slingshot L0 QSFP-DD 2x QSFP56 2m Splitter Cable	R4K56A
HPE Slingshot L0 QSFP-DD 2x QSFP56 2.5m Splitter Cable	R4K57A
HPE Slingshot L0 QSFP-DD 2x QSFP56 3m Splitter Cable	R6P07A

Active Optical Cables

HPE Slingshot L0 QSFP-DD 2x QSFP56 4m Splitter Active Optical Cable	R6F19A
HPE Slingshot L0 QSFP-DD 2x QSFP56 10m Splitter Active Optical Cable	R6F20A
HPE Slingshot L0 QSFP-DD 2x QSFP56 4m Analog Splitter Active Optical Cable	R7N25A
HPE Slingshot L0 QSFP-DD 2x QSFP56 10m Analog Splitter Active Optical Cable	R7N26A
HPE Slingshot L0 QSFP-DD 2x QSFP56 30m Analog Splitter Active Optical Cable	R7N27A
HPE Slingshot L0 QSFP-DD to 2x QSFP56 50m Analog Splitter Active Optical Cable	R8G02A
HPE Slingshot L0 2x100Gb QSFP-DD to 2x QSFP28 5m Splitter Active Optical Cable	R7T63A
HPE Slingshot L0 2x100Gb QSFP-DD to 2x QSFP28 10m Splitter Active Optical Cable	R7T64A
HPE Slingshot L0 2x100Gb QSFP-DD to 2x QSFP28 30m Splitter Active Optical Cable	R7T65A
HPE Slingshot L0 2x100Gb QSFP-DD to 2x QSFP28 50m Splitter Active Optical Cable	R8G03A

HPE Slingshot Straight-Through (QSFP-DD to QSFP-DD) Cables

Copper Cables

HPE Slingshot L1 QSFP-DD 0.4m Straight Cable	R4K58A
HPE Slingshot L1 QSFP-DD 0.8m Straight Cable	R4K59A
HPE Slingshot L1 QSFP-DD 1.2m Straight Cable	R4K60A
HPE Slingshot L1 QSFP-DD 1.4m Straight Cable	R4K61A
HPE Slingshot L1 QSFP-DD 1.6m Straight Cable	R4K62A
HPE Slingshot L1 QSFP-DD 1.9m Straight Cable	R4K63A
HPE Slingshot L1 QSFP-DD 2.1m Straight Cable	R4K64A
HPE Slingshot L1 QSFP-DD 2.4m Straight Cable	R4K65A
HPE Slingshot L1 QSFP-DD 2.6m Straight Cable	R4K66A
HPE Slingshot L1 QSFP-DD 2.9m Straight Cable	R4K67A

Analog Active Optical Cables

HPE Slingshot L2 QSFP-DD 5m Analog Active Optical Cable	R4K78A
HPE Slingshot L2 QSFP-DD 7m Analog Active Optical Cable	R4K79A
HPE Slingshot L2 QSFP-DD 10m Analog Active Optical Cable	R4K80A
HPE Slingshot L2 QSFP-DD 14m Analog Active Optical Cable	R4K81A
HPE Slingshot L2 QSFP-DD 19m Analog Active Optical Cable	R4K82A
HPE Slingshot L2 QSFP-DD 25m Analog Active Optical Cable	R4K83A
HPE Slingshot L2 QSFP-DD 35m Analog Active Optical Cable	R4K84A
HPE Slingshot L2 QSFP-DD 50m Analog Active Optical Cable	R4K85A
HPE Slingshot L2 QSFP-DD 75m Analog Active Optical Cable	R4K86A
HPE Slingshot L2 QSFP-DD 100m Analog Active Optical Cable	R4K87A

Configuration Information

Digital Active Optical Cables

HPE Slingshot L2 QSFP-DD 5m Digital Active Optical Cable	R4K88A
HPE Slingshot L2 QSFP-DD 7m Digital Active Optical Cable	R4K89A
HPE Slingshot L2 QSFP-DD 10m Digital Active Optical Cable	R4K90A
HPE Slingshot L2 QSFP-DD 14m Digital Active Optical Cable	R4K91A
HPE Slingshot L2 QSFP-DD 19m Digital Active Optical Cable	R4K92A
HPE Slingshot L2 QSFP-DD 25m Digital Active Optical Cable	R4K93A
HPE Slingshot L2 QSFP-DD 35m Digital Active Optical Cable	R4K94A
HPE Slingshot L2 QSFP-DD 50m Digital Active Optical Cable	R4K95A
HPE Slingshot L2 QSFP-DD 75m Digital Active Optical Cable	R4K96A
HPE Slingshot L2 QSFP-DD 100m Digital Active Optical Cable	R4K97A

HPE Slingshot Bifurcated (QSFP-DD to QSFP-DD “H”) Cables

Copper Cables

HPE Slingshot L1 QSFP-DD 0.8m Bifurcated Cable	R4K68A
HPE Slingshot L1 QSFP-DD 1.2m Bifurcated Cable	R4K69A
HPE Slingshot L1 QSFP-DD 1.4m Bifurcated Cable	R4K70A
HPE Slingshot L1 QSFP-DD 1.6m Bifurcated Cable	R4K71A
HPE Slingshot L1 QSFP-DD 1.9m Bifurcated Cable	R4K72A
HPE Slingshot L1 QSFP-DD 2.1m Bifurcated Cable	R4K73A
HPE Slingshot L1 QSFP-DD 2.3m Bifurcated Cable	R4K74A
HPE Slingshot L1 QSFP-DD 2.6m Bifurcated Cable	R4K75A
HPE Slingshot L1 QSFP-DD 2.9m Bifurcated Cable	R4K77A

HPE Slingshot Fabric Management Suite

The HPE Slingshot Fabric Management Suite runs on management nodes using SuSE, Red Hat, or CentOS Linux. The management node connects to each switch using a management network that connects to each switch using the out-of-band 1 Gbps or 10 Gbps management network. The fabric management node should be configured with two 10 Gbps connections to the management network. There is no current requirement to connect the management node to the Slingshot high speed network itself though future software release may enable in-band functionality from the fabric management suite.

One software license per HPE Slingshot cluster is required.

Consult reference architectures for supported OS versions, hardware reference architectures, and sizing guides.

HPE Slingshot Fabric Manager E-RTU	R6M77AAE
HPE Slingshot Fabric Manager FIO Software	R6M78A

A pre-configured fabric management node hardware SKU is optionally offered to simplify ordering and integration. (Software must still be ordered separately).

Edge Routers

HPE Slingshot switches are validated with select HPE Aruba Networking and other enterprise Ethernet routing switches for connectivity between the HPE Slingshot cluster and external Ethernet networks without use of gateway nodes. This support includes aggregating links from across HPE Slingshot switches into a single wide link aggregation group (LAG). Connections are supported at 100 Gbps (4 lanes of NRZ) or 200 Gbps (4 lanes of PAM 4 encoding). Consult [reference architectures](#) for the current list of validated switches and cables.



Technical Specifications

HPE Slingshot 200Gb 64-port QSFP-DD 1U Switch (R4K41B)	
Form factor	1U enclosure for 19" rack, fixed rails included.
Dimensions	1.7 x 19 x 27in (H x W x D) (4.3cm x 48.3 cm x 68.6 cm)
Weight:	35lbs (15.88 kg)
Physical connectivity	32 QSFP-DD
Port capability – 200 Gbps	<ul style="list-style-type: none"> • IEEE 802.3cd/bd (200 Gbps) Ethernet over 4 x 50 Gb/s (PAM-4) lanes • 200GBASE-SR4, up to 100m multi-mode fiber • 200GBASE-CR4, up to 3m copper cable
Port capability – 100 Gbps	<ul style="list-style-type: none"> • IEEE 802.3 2018 100 Gbps Ethernet over 4 x 25 Gb/s (NRZ) lanes • 100GBASE-CR4, 5m copper cable • 100GBASE-SR4, 100m multi-mode fiber
Ethernet	<ul style="list-style-type: none"> • Optimized Ethernet and HPC Fabric Formats • Link level retry, low latency FEC • Lossy and Lossless delivery • Flow control, 802.1x (PAUS), 802.1p (PFC), credit-based flow control on fabric links, fine grain flow control on host links end edge ports
Quality of Service	<ul style="list-style-type: none"> • IETF DiffServ or IEEE 802.1p/Q • Patented hardware congestion management mechanisms
Routing	<ul style="list-style-type: none"> • Packet-by-packet adaptive routing with real-time feedback on load • Order-preserving adaptive routing of IP flows • Traffic class specific routing
Management ports	1Gb RJ-45, 10Gb RJ-45
Management LEDs	Power and status
Throughput per port	Line rate
Port status	Link and activity
Power and cooling	<ul style="list-style-type: none"> • Two (1+1 redundant) hot plug power supplies • Five hot plug fan modules (N + 1 redundant). • Power 180-264V – inlet connector IEC320 C-14 10A/250Vac • 240VAC typical – No 120VAC or 277 VAC support. Frequency 50-60Hz • Power draw: 481W +12.5W per active optical cable when all fabric ports are attached to compute nodes and running at full bandwidth. Includes fan power, all VR conversion, and copper losses. • Airflow is from the power supply side to the port connector side. • 94.5 cfm at 2.6" of water pressure at maximum fan speed with all QSFP-DD cable ports populated
Accoustic	Sounds pressure while operating at maximum fan speed is 73.5 dBA (ISO 7779 bystander method)
Environmentals	<p>Temperature Ranges:</p> <ul style="list-style-type: none"> • Maximum operating dry bulb temperature range: 5 °C to 35 °C • Recommended operating dry bulb temperature: 18 °C to 25 °C • Maximum non-operating dry bulb temperature range: 5 °C to 45 °C <p>Relative Humidity Specification (for ambient air):</p> <ul style="list-style-type: none"> • Maximum operating relative humidity range: 8% (-12C dew point) to 85% non-condensing • Recommended operating relative humidity range: 30% to 50% non-condensing • Maximum non-operating relative humidity range: 8% to 80% non-condensing • Maximum dew point: 27C

Technical Specifications

Environmentals	Altitude Specifications: <ul style="list-style-type: none"> • The following numbers are specified for altitude. • Maximum operating altitude range: 0ft to 10,000ft or 0m to 3048m • Maximum non-operating altitude range: 0ft to 40,000ft or 0m to 12,192m • Temperature derating with altitude: Max dry bulb must be de-rated by 1° C per 175m above 900m (2°F per 638ft over 2953ft)
Electromagnetic Compliance	<ul style="list-style-type: none"> • EN55032 • EN55035 • Australia/New Zealand: RCM • Canada: ICES-003 Issue 5 • European Union: CE Mark • Israel: IS 961 Part 6.01 and 6.02 • Japan: VCCI • South Korea: KC Mark • Taiwan: BSMI • United States: FCC Part 15, Subpart B, Class A
Safety	<ul style="list-style-type: none"> • IEC 62368-1, Second Edition (as well as the country-specific variants covered by CB report) • NRTL Certified for United States and Canada
Warranty	<ul style="list-style-type: none"> • HPE Slingshot For Apollo, ProLiant, and Cray XD Clusters carry a 3-year warranty, On Site, next business day response • Cables carry a 1-year warranty, parts exchange

HPE Slingshot SA210S Ethernet 200Gb 1-port PCIe NIC (R4K46A)

Form factor	PCIe low profile (includes faceplates for full and half-height)
Dimensions	2.713 x 6.60 (H x W) (68.91mm x 167.64)
Weight:	0.42lbs (191g)
Physical connectivity	1 QSFP-56 port
Host connectivity	PCIe Gen 4 x16
Port capability – 200 Gbps	IEEE 802.3cd/bs (200 Gbps) Ethernet over 4 x 50 Gbps (PAM-4) lanes 100GBASE-CR4 5m copper cable 100GBASE-SR4 100m OM3/OM4 optical cable 50/100/200GBASE-CR4 3m TwinAx copper cable
Port capability – 100 Gbps	IEEE 802.3 2018 100 Gbps Ethernet over 4 x 25 Gbps (NRZ) lanes 100GBASE-CR4 5m copper cable 100GBASE-SR4 100m OM3/OM4 optical cable 50/100/200GBASE-CR4 3m TwinAx copper cable
Ethernet	Optimized Ethernet and HPC Fabric Formats. Support for 32- and 40-byte frame for single word remote memory access operations Link level retry, low latency FEC RoCE CRC offload TCP, checksum
HPC Acceleration	Efficient command launch: Commands with immediate data Low-latency (write-based) commands Multi-threaded DMA engine Packet and connection tracking Hardware message matching:

Technical Specifications

	<p>Full offload of MPI message matching with strong progression</p> <p>Eager and rendezvous transfers</p> <p>Unexpected message handling</p> <p>High rate Put mechanism</p> <p>Increases rate of small Puts and atomic memory operations</p> <p>Triggered operations</p> <p>Synchronization and ordering:</p> <p>Full events, counting events, flush and fence operations</p> <p>Address Translation: Large on-chip translation cache</p>
Quality of Service	<p>IETF DiffServ or IEEE 802.1p/Q</p> <p>Up to four bandwidth shaping classes</p> <p>Fine grain flow control</p>
Software Interfaces	<p>Linux Drivers including IPv4 and IPv6</p> <p>Libfabrics</p> <p>k-fabrics (for LNET)</p>
Port status	<p>Link and activity</p>
Power and cooling	<p>35W max</p> <p>18W typical</p> <p>Passive air-cooled</p>
Electromagnetic Compliance	<p>EN55032 and EN55035</p> <p>Australia/New Zealand: RCM</p> <p>Canada: ICES-003 Issue 5</p> <p>European Union: CE Mark</p> <p>Israel: IS 961 Part 6.01 and 6.02</p> <p>Japan: VCCI</p> <p>South Korea: KC Mark</p> <p>Taiwan: BSMI</p> <p>United States: FCC Part 15, Subpart B</p>
Safety	<p>IEC 60950-1 and IEC62368-1, both Second Edition (as well as the various country-specific variants covered by CB report)</p> <p>NRTL Certified for United States and Canada</p>
Warranty	<p>The HPE Slingshot NIC is covered by warranty of the HPE Apollo, ProLiant, or Cray XD server into which it has been configured.</p>





Summary of Changes

Date	Version History	Action	Description of Change
20-May-2024	Version 12	Changed	Overview, Standard Features, Configuration Information and Technical Specifications sections were updated Added reference to Cray XD and links to related QuickSpecs and configuration guides
19-Feb-2024	Version 11	Changed	Networking product names were updated.
18-Dec-2023	Version 10	Changed	HPE Services Rebranding
17-Apr-2023	Version 9	Changed	Addition of Analog Optical Cable SKUs and 2-switch group. Overview, Standard Features and Configuration Information sections were updated
16-May-2022	Version 8	Changed	Standard Features and Configuration Information sections were updated
20-Dec-2021	Version 7	Changed	Standard Features section was updated
06-Dec-2021	Version 6	Changed	Name of the QS was changed Overview section was updated Service and Support section was updated.
04-Oct-2021	Version 5	Changed	Added information throughout on HPE Slingshot NICs plus overall edits Overview, Standard Features, Configuration Information and Technical Specifications sections were updated
02-Aug-2021	Version 4	Changed	Configuration Information section was updated
07-Jun-2021	Version 3	Changed	Overview, Service and Support and Configuration Information sections were updated
06-Apr-2021	Version 2	Changed	Configuration Information section was updated
07-Dec-2020	Version 1	New	New QuickSpecs

Copyright

Make the right purchase decision.
Contact our presales specialists.



© Copyright 2024 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

To learn more, visit <http://www.hpe.com/>

a50002546enw - 16701 - Worldwide - V12 - 20-May-2024