NVIDIA MELLANOX SPECTRUM-2
400G ETHERNET SWITCH SILICON

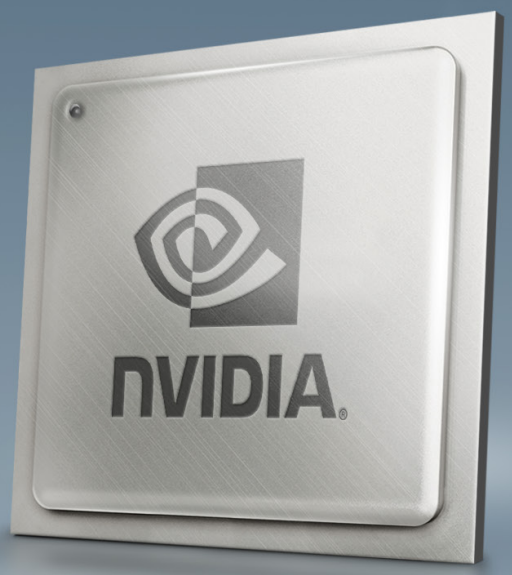

NVIDIA® Mellanox® Spectrum®-2 provides an unmatched combination of performance, virtualization and telemetry capabilities in a 6.4 Tbps Ethernet switch silicon with 128 x 50 Gb/s PAM4 SerDes. The SerDes and can act independently, in pairs or groups of 4 or 8, to operate at speeds between 100 Mbps and 400 GbE for maximum flexibility. The switch ASIC is optimized for cloud, storage and Machine Learning solutions. The Mellanox Spectrum® family of products enjoy a strong software ecosystem that can be leveraged by partners to develop and deploy solutions with a rapid time-to-market.

CONSISTENT PERFORMANCE

Consistent performance and fair bandwidth sharing are critical for multi-tenant cloud, distributed storage workloads as well as hardware-accelerated data transfers using technologies such as GPUDirect® or NVMe-OF. Mellanox Spectrum-2 features a 42 MB fully-shared and monolithic packet buffer that is dynamically shared across all ports. This provides excellent burst absorption and low latency cut-through performance.

The on-chip packet buffer supports high-bandwidth packet reads/writes, enabling consistent and predictable cut-through performance. The monolithic packet buffer architecture simplifies buffer management and traffic scheduling, while also enabling fair resource sharing. Additionally, Mellanox Spectrum-2 supports intelligent congestion management mechanisms, including Explicit Congestion Notification (ECN) that enable robust RoCE-based, hardware-accelerated data transport.

ADVANCED NETWORK VIRTUALIZATION

Mellanox Spectrum-2 is designed with a highly-flexible programmable pipeline that enables advanced network virtualization. The hybrid packet-forwarding pipeline comprises optimized blocks, which are streamlined in hardware to deliver the highest performing traditional functionality. The pipeline also is made up of programmable blocks for supporting new functionality. Leveraging this capability, Mellanox Spectrum-2 can deliver industry-leading packet-processing rates at scale, concurrently with programmability and Advanced Network Virtualization.

Mellanox Spectrum-2 supports single-pass VXLAN routing, while providing 10X better VXLAN VTEP and tunnel scale. In addition to VXLAN, Mellanox Spectrum-2 supports a myriad of virtualization technologies including MPLS, SRV6, and more.

Highlights

- 6.4 Tb/s of Ethernet switching and IP routing
- Integrated 128 x 50 Gb/s PAM4 SerDes
- Flexible port configurations:
  - Up to 16 x 400 GbE ports
  - Up to 32 x 200 GbE ports
  - Up to 64 x 100 GbE ports
  - Up to 128 x 50/25/10 GbE ports

Consistent Performance

- 400 ns, industry leading, true, cut-through latency
- Low power consumption of 7 watts per 200 GbE port
- 42 MB dynamically-shared, flexible, packet buffering
- Robust RoCE data path for NVMe-oF and GPUDirect® hardware-accelerated data transfers

Advanced Network Virtualization

- Single-pass VXLAN bridging and routing
- Centralized VXLAN routing
- Comprehensive overlay and tunneling support: VXLAN-GPE, NVGRE, Geneve, NSH and MPLS/IPv6 based Segment Routing

Open and Actionable Telemetry

- Reduced Mean Time to Recovery/ Innocence
- Actionable and contextual telemetry with What Just Happened (WJH)
- Hardware-based buffer histograms
- Streaming telemetry
- In-band network telemetry

INDUSTRY-LEADING VISIBILITY WITH WHAT JUST HAPPENED (WJH)

Mellanox Spectrum-2 provides rich, contextual, and actionable network-based visibility data that seamlessly integrates with open APIs at the SDK layer. Referred to as 'What Just Happened' (WJH), this feature allows network operators to finally get instant answers to the 5 W's—When, What, Who, Where and Why, as soon as an issue arises. WJH can easily be extended and customized to support both third-party and open-source tools.

'What Just Happened' is open and extendable, making it ideal for integration into third-party network monitoring and analysis tools. For those wanting a simple turn-key solution, Mellanox NEO® is an easy way to utilize the benefits of What Just Happened.

Beyond providing hardware-based mechanisms that measure and summarize packet buffer dynamics, and over half a million purpose-built flow tracking counters, Mellanox Spectrum-2 also supports traditional hardware-based visibility features such as sFlow, in-band and streaming telemetry.

CHOICE OF SOFTWARE

Mellanox offers a rich set of software development tools and open APIs that enable users and partners to develop fully-functional switching solutions on top of Mellanox Spectrum-2 at low-cost, with short time-to-market. Users have the flexibility to choose the model best suited to their needs, including Mellanox Spectrum-2 SDK, OCP SAI (Switch Abstraction Interface), or Linux® Switch driver, which enables native Linux operating systems and applications on top of Mellanox Spectrum-2.

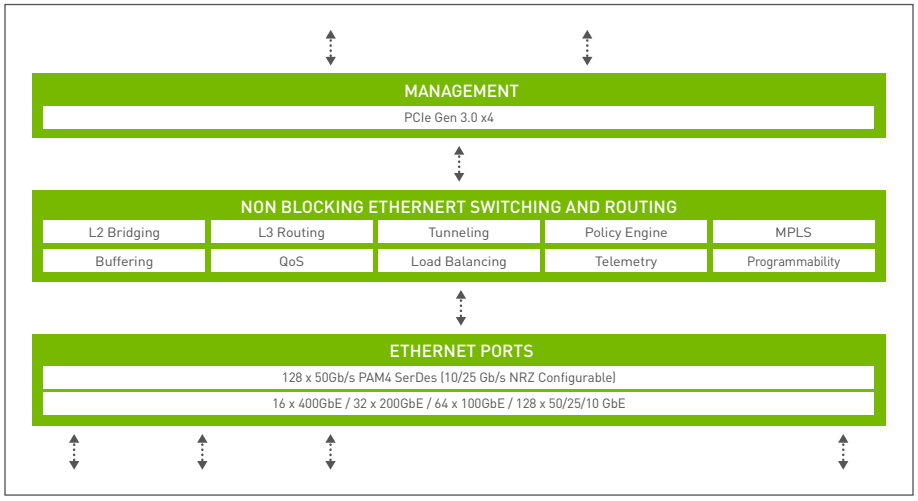

High Level System Block Overview

Cloud-Scale

- Half a million flexible forwarding entries shared across IPv4/IPv6 routes, MAC, MPLS, etc.
- Half a million flow tracking counters
- 100X more ACLs
- 10X better VXLAN VTEP and Tunnel Scale

Compatibility

Ethernet

- 16 x 400 GbE ports,
  32 x 200 GbE ports,
  64 x 100 GbE ports, or
  128 x 1/10/25/50 GbE ports
- DCB (PFC, ETS, DCBX)

CPU Interoperability

- PowerPC, Intel® x86, AMD x86, and MIPS

Host Interface

- PCIe Gen 3.0; 4 lanes of 8 GT/s, 5 GT/s or 2.5 GT/s

Interoperability

- Full interoperability with standard Ethernet NICs and switches

SerDes

- 128 x 50 Gb/s PAM4 (802.3bs, 802.3cd)
- 10 Gb/s (802.3ba, 802.3ae) and 25GbE (802.3bj, 802.3bm) NRZ operational modes

Connectivity

- Drives active/passive copper cables, fiber optics, PCB, or backplanes

I/O Specifications

- SPI flash interface
- I²C, SMBus, MDIO
- IEEE 1149.1 boundary-scan JTAG
- IEEE 1588 1-step and 2-step

Package

- 57.5mm x 57.5mm HFCBGA

ORDERING INFORMATION

OPN	DESCRIPTION
MT53132-NCCR-V	Mellanox Spectrum®-2, 128x50Gb/s PAM4, 32 Port 200GbE Switch IC

Learn more at **www.mellanox.com/products/ethernet-switch-ic**

Sept20/52460PB-R2