



# NVIDIA SPECTRUM-3

## 400G Ethernet Switch Silicon



NVIDIA® Spectrum®-3 is our 4th generation of Ethernet switch ASIC and provides an unmatched combination of performance, virtualization, and telemetry capabilities in a 12.8 Tbps Ethernet switch silicon with 256 x 50 Gb/s PAM4 SerDes. The SerDes can act independently, in pairs, quads or groups of 8, to operate at speeds between 1 GbE and 400vGbE for maximum flexibility. The switch ASIC is optimized for cloud, storage and Machine Learning solutions. The 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. Spectrum-3 features a 64 MB fully-shared and monolithic packet buffer that is dynamically shared across all ports. This provides excellent burst absorption with true, port-to-port, cut-through latency.

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, Spectrum-3 supports intelligent congestion management mechanisms, including Explicit Congestion Notification (ECN), that enable robust RoCEbased, hardware-accelerated, data transport and load balancing, such as adaptive routing and Priority-based Flow Control (PFC).

### Advanced Network Visualization

Spectrum-3 is designed with a highly-flexible programmable packet 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, Spectrum-3 can deliver industry-leading packet-processing rates at scale, concurrently with programmability and Advanced Network Virtualization.

Spectrum-3 supports single-pass VXLAN routing and provides 10X better VXLAN VTEP and tunnel scale. In addition to VXLAN, Spectrum-3 supports a myriad of virtualization technologies including SRV6, Network Address Translation (NAT), and more.

### HIGHLIGHTS

- > 12.8 Tb/s of Ethernet switching and IP routing
- > Integrated 256 x 50 Gb/s PAM4 SerDes
- > Flexible port configurations:
  - > Up to 32 x 400 GbE ports
  - > Up to 64 x 200 GbE ports
  - > Up to 128

### CONSISTENT PERFORMANCE

- > Industry leading, true, cut-through latency
- > Low power consumption per port
- > 64 MB dynamically-shared, flexible, packet buffering
- > Robust RoCE data path for NVMe-oF and GPUDirect®
- > Adaptive Routing congestion avoidance

### 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
- > Network Address Translation (NAT)

### EVENT-BASED TELEMETRY

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

## Industry-Leading Visibility with What Just Happened (WJH)

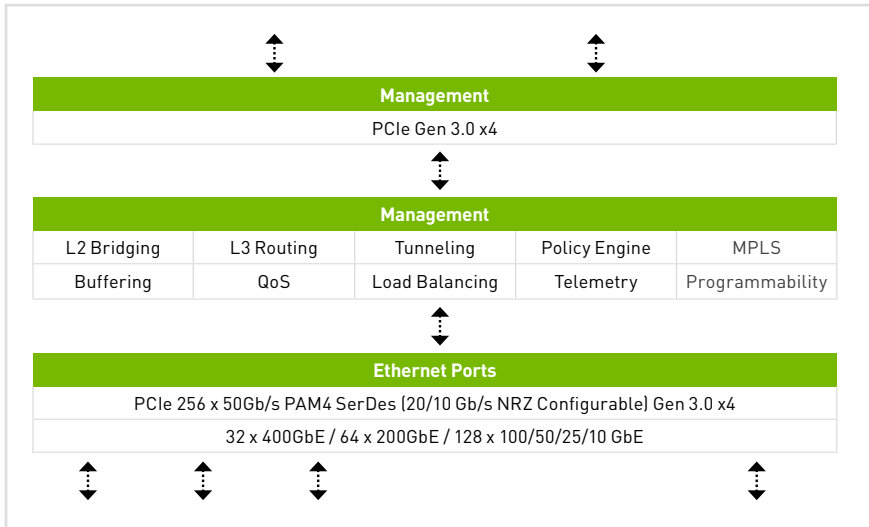
Spectrum-3 provides rich, contextual, and actionable network-based visibility data that seamlessly integrates with open APIs at the SDK layer. Referred to as 'What Just Happened' (WJH), this feature allows network operators to finally get instant answers to the 5 W's—When, What, Who, Where and Why, as 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 thirdparty network monitoring and analysis tools. For those wanting a simple turn-key solution, NVIDIA NetQ® 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-3 also supports traditional hardware-based visibility features such as sFlow, in-band and streaming telemetry.

## Choice of Software

NVIDIA 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 Spectrum-3 at low-cost, with short time-to-market. Users have the flexibility to choose the model best suited to their needs, including Spectrum-3 SDK, OCP SAI (Switch Abstraction Interface), or Linux® Switch driver, which enables native Linux operating systems and applications on top of Spectrum-3.



High Level System Block Overview

## Ordering Information

For ordering information, please contact [gopny@pny.com](mailto:gopny@pny.com)

### CLOUD-SCALE

- > Half a million flexible entries shared across routes, MAC, MPLS, etc.
- > Up to 1 M on-chip IPv4 routes; extendable up to 4 M
- > Half a million flow tracking counters
- > 100X more ACLs
- > 10X better VXLAN VTEP and tunnel scale

### SYNCHRONIZATION

- > Precision Time Protocol (PTP) network synchronization

### COMPATIBILITY

Ethernet	<b>32 x 400 GbE ports, 64 x 200 GbE ports, or 128 x 100/50/25/10/1 GbE ports</b> <b>DCB (PFC, ETS, DCBX)</b>
CPU Interoperability	<b>PowerPC, Intel® x86, and AMD x86</b>
Host Interface	<b>PCIe Gen 3.0; 4 lanes of 8 GT/s, 5 GT/s or 2.5 GT/s</b>
Interoperability	<b>Full interoperability with standard Ethernet NICs and switches</b>
SerDes	<b>256 x 50 Gb/s PAM4 (802.3bs, 802.3cd)</b> <b>10 Gb/s (802.3ba, 802.3ae) and 25 GbE (802.3bj, 802.3bm) NRZ operational modes</b>
Connectivity	<b>Drives active/passive copper cables, fiber optics, PCB, or backplanes</b>
I/O Specifications	<b>SPI flash interface</b> <b>I<sup>2</sup>C, SMBus, MDIO</b> <b>IEEE 1149.1 boundary-scan JTAG</b> <b>IEEE 1588 1-step and 2-step</b>
Package	<b>70mm x 70mm HFCBGA</b>

[Learn more](#)

Learn more at: [pny.com/networking](http://pny.com/networking)