

NVIDIA SPECTRUM-4

51.2 Tb/s Ethernet Switch ASIC

Cloud-native, accelerated computing, and AI are the new networking frontiers, and data is the fuel that drives these workloads. AI represents a powerful growth driver that spans all industries. All workloads are becoming cloud-native; scaled out with containerized microservices that need massive bandwidth to access data and communicate across the data center. The new NVIDIA Spectrum[®]-4, the 5th generation of ASICs in the NVIDIA Spectrum Ethernet platform, arrives just in time to meet the needs of cloud-native, accelerated computing, and AI workloads. Spectrum accelerates workloads from core, to cloud, to edge.

In addition to delivering accelerated Ethernet for any data center workload, Spectrum-4 provides the most advanced and innovative feature set yet. This includes unique features like cloud security through MACsec over VXLAN, and end-to-end nanosecond level timing precision from switch to host. Spectrum-4 ASIC delivers the following key features:

Consistent Performance

Consistency and fair bandwidth-sharing are critical for multi-tenant clouds, distributed storage workloads, and hardware-accelerated data transfers using technologies like NVIDIA Rapids[™], NVIDIA Magnum IO[™], or NVMe-oF. Designed in-house using an NVIDIA 4N process, leveraging NVIDIA technology shared from other NVIDIA platform innovations, Spectrum-4 features the most robust and stable SerDes yet. This simplifies solution design, leading to best-in-class power efficiency and a lower TCO.

NVIDIA Spectrum-4 features a fully-shared and monolithic packet buffer that's dynamically available to all ports. This provides excellent microburst absorption with true, port-to-port, cut-through latency. Spectrum-4 also supports programmability of the pipeline and packet modifier/parser without impact to latency or packet rate, enabling flexibility without compromising performance.

The on-chip packet buffer supports high-bandwidth packet reads/writes, enabling a consistent and predictable performance. The monolithic packet buffer architecture simplifies buffer management and traffic scheduling, while also enabling fair resource sharing.

KEY FEATURES

- > 51.2 Tb/s switching and routing
- > 12.8 Tb/s crypto engine traffic
- > 100 Gb/s PAM4 SerDes
- > NVIDIA 4N Process
- > Fully shared buffer architecture
- > 4X packet rate
- > 40% lower power consumption
- > Flexible port configurations
 - > 64 x 800GbE ports
 - > 128 x 400GbE ports
 - > 256 x 200/100GbE ports

PROOF POINTS

- > **Consistent Performance:**
 - > Adaptive Routing tuned for AI applications
 - > Industry-low latency
 - > Accelerated RoCE data path for NVMe-oF and Magnum IO
- > **Advanced Network Virtualization:**
 - > Single-pass VXLAN bridging and routing
 - > Centralized VXLAN routing
 - > Comprehensive and programmable tunneling technology support
 - > Network Address Translation (NAT)
- > **Accelerated Event-based Telemetry:**
 - > Actionable telemetry with WJH
 - > FabricView histograms
 - > Streaming telemetry
 - > In-band network telemetry
 - > Telemetry CPU-offloads

Additionally, NVIDIA Spectrum-4 supports intelligent congestion control. This enables accelerated RoCE-based data transport and load balancing through adaptive routing (including adaptive routing notifications for dynamic flow rebalance) as well as high precision congestion control (HPCC) facilitated through in-band network telemetry.

Advanced Network Virtualization and Simulation

NVIDIA Spectrum-4 is designed with a highly-flexible programmable packet pipeline with stateful capability that enables advanced network virtualization. The hybrid packet-forwarding pipeline can be merged into optimized blocks that are streamlined in hardware to deliver the highest performance. Leveraging this capability, NVIDIA Spectrum-4 can deliver industry-leading packet-processing rates at scale, concurrent with programmability and virtualization.

NVIDIA Spectrum-4 supports single-pass VXLAN routing and provides 10X better VXLAN VTEP and tunnel scale. In addition to VXLAN, NVIDIA Spectrum-4 supports multiple virtualization technologies, such as SRV6, NAT, custom virtualization, and more.

The Spectrum platform enables digital twins of switches based on the Spectrum-4 ASIC (as well as all other Spectrum switches) through the NVIDIA Air infrastructure simulation platform. The digital twin is used for validating security policy compliance, automation process, monitoring tools, interoperability, and upgrade procedures. In addition, NVIDIA Air supports SimX for Spectrum-4, enabling users to simulate the ASIC and test-drive features, such as programmability. With SimX, developers can get hardware-free early access to the ASIC for better debugging.

Actionable Visibility With What Just Happened (WJH)

NVIDIA Spectrum-4 provides packet-granular network-based visibility data that seamlessly integrates with open APIs at the SDK layer. Referred to as What Just Happened (WJH), this always-on feature allows network operators to start troubleshooting as soon as an issue arises. WJH can be extended and customized to support both third-party and open-source tools.

WJH is open and extendable, making it ideal for integration into third-party network monitoring and analysis tools. NetQ, combined with WJH, enables a simple turn-key solution that also provides validation and monitoring of the network fabric.

Beyond providing hardware-accelerated mechanisms that measure and summarize packet buffer dynamics, NVIDIA Spectrum-4 also supports traditional hardware-based visibility tools, such as sFlow, in-band telemetry, and streaming telemetry. In addition to providing actionable visibility, the Spectrum-4 ASIC features state-of-the-art security features like support for MACsec and VXLANsec, and SN5000 switches that support secure boot as default via hardware root of trust (HRoT).

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 NVIDIA Spectrum-4 at low-TCO, with a short time-to-market. Users have the flexibility to choose the network operating system that's best suited to their needs, including the flagship Cumulus Linux, SONiC (Software for Open Networking in the Cloud), or Linux Switch driver, which enables native Linux operating systems and applications on top of NVIDIA Spectrum-4.

> Cloud-tier Scale and Security:

- > Half a million flexible entries shared across routes, MAC, MPLS, etc.
- > Up to one million on-chip IPv4 routes
- > 100X more ACLs
- > 10X better VXLAN VTEP and tunnel scale
- > 12.8 Tb/s line rate encryption
- > Class C PTP and SyncE time synchronization
- > MACsec and VXLANSec encryption
- > Secured switch offerings via hardware root-of-trust

COMPATIBILITY

| | |
|--------------|--|
| Ethernet | Switch ports <ul style="list-style-type: none">> 800GbE> 400GbE> 200GbE> 10/100GbE> DCB (PFC, ETS, DCBX) |
| CPU Options | Intel x86, AMD x86, ARM |
| AI Solutions | Optimized for NVIDIA® <ul style="list-style-type: none">> DGX™> EGX™> OVX> BlueField®> ConnectX® |
| IO Specs | Standards <ul style="list-style-type: none">> SPI flash> I2C, SMBus, MDIO> IEEE 1149.1, 1588 |



NVIDIA Spectrum-4 ASIC

End-to-End Synergy

As part of the NVIDIA Spectrum platform, Spectrum-4 is pre-tested and pre-validated with NVIDIA's full portfolio of Ethernet networking technology, including BlueField-3 DPUs and ConnectX-7 SmartNICs based on the same 100G SerDes channels used with Spectrum-4. This end-to-end solution is critical to powering NVIDIA's accelerated AI and simulation platforms, and delivers the high performance and innovative feature set needed to supercharge cloud-native applications at scale.



Ordering Information

| OPN | DESCRIPTION |
|--------------------|--|
| SPC4-E0256EC11C-A0 | NVIDIA Spectrum-4, 51.2 Tb/s switch IC with 512 x 100 Gb/s PAM4 and 64 x 800GbE interfaces |
| SPC4-E0128DC11C-A0 | NVIDIA Spectrum-4, 25.6 Tb/s switch IC with 256 x 100 Gb/s PAM4 and 32 x 800GbE interfaces |
| SPC4-E0256EG11C-A0 | NVIDIA Spectrum-4, 25.6 Tb/s switch IC with 512 x 50 Gb/s PAM4 and 64 x 400GbE interfaces |