



# NVIDIA Spectrum-4

Unmatched performance, virtualization, and telemetry for AI and cloud.

AI, cloud-native computing, and simulation at scale are the new networking frontiers, and data is the fuel that drives them. AI represents a powerful growth driver that spans every industry. Cloud-native workloads are becoming the norm, scaled out with containerized microservices that need massive bandwidth to access data and communicate across the data center. The new NVIDIA Spectrum™-4 Ethernet switch application-specific integrated circuit (ASIC) drives best-in-class performance for AI and cloud-native computing workloads—from core to cloud to edge.

As NVIDIA's fifth-generation Ethernet switch ASIC, Spectrum-4 powers the NVIDIA Spectrum™-X platform, the world's first accelerated Ethernet solution that meets the needs of AI infrastructures. Spectrum-4 delivers the highest performance for bandwidth-hungry accelerated applications based on generative AI, large language models, recommenders, video analytics, and more. With enhanced remote direct-memory access (RDMA) over Converged Ethernet (RoCE) built on adaptive routing, performance isolation, and auto-path configuration, the Spectrum-4 switch ASIC is unmatched for workloads running on Ethernet.

In addition to delivering accelerated Ethernet for AI workloads, Spectrum-4 provides clouds with 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.

## NVIDIA Spectrum-X Networking Platform

Consistency and fair bandwidth sharing are critical for AI superclouds built with NVIDIA platforms such as HGX™, DGX™, and OVX™. Accelerated AI applications require a network designed to enable:

- > Tightly coupled processes
- > Bursty network capacity with predictive performance and no statistical averaging
- > Highly efficient RDMA bandwidth
- > Low jitter tolerance with deterministic short tails
- > Non-blocking, rail-optimized fabric support


## Key Features

- > 51.2Tb/s switching and routing
- > 12.8Tb/s crypto engine traffic
- > 100Gb/s PAM4 SerDes
- > TSMC 4N Process
- > RoCE dynamic datapath extensions
- > Fully shared buffer architecture
- > 4X packet rate
- > 30 percent lower power consumption
- > Flexible port configurations
  - 64x 800GbE ports
  - 128x 400GbE ports
  - 256x 200/100GbE ports

## Proof Points

### Consistent Performance for Generative AI via RoCE Extensions

- > RoCE Adaptive Routing
- > RoCE Congestion Control
- > RoCE Performance Isolation
- > RoCE port-to-port (RP2P) Auto-path Configuration
- > Synchronized Collectives
- > Industry-low Latency



To meet these needs, NVIDIA has introduced the Spectrum-X networking platform, the world's first Ethernet platform for AI. Spectrum-X tightly coordinates the Spectrum-4 ASIC and NVIDIA® BlueField®-3 DPUs to optimize hyperscale generative AI infrastructures using a suite of end-to-end innovations. Central to these innovations are the new RoCE dynamic datapath extensions:

- > RoCE Adaptive Routing
- > RoCE Congestion Control
- > RoCE Performance Isolation
- > RoCE Port-to-Port (RP2P) Auto-path Configuration
- > Synchronized Collectives

## Robust Design

Designed in-house using an NVIDIA 4 nanometer (4N) process and leveraging NVIDIA technology shared from other NVIDIA platform innovations, Spectrum-4 features the most robust and stable serializer/deserializer (SerDes) yet. This enables NVIDIA direct drive technology to simplify solution design, leading to best-in-class power efficiency and a lower TCO.

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.

## 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, Spectrum-4 can deliver industry-leading packet-processing rates at scale, concurrent with programmability and virtualization.

Spectrum-4 also supports single-pass VXLAN routing and provides 10X better VXLAN VTEP and tunnel scale compared to the market. In addition to VXLAN, 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.

### Advanced Virtualization

- > Single-pass VXLAN bridging and routing
- > Centralized VXLAN routing
- > Comprehensive and programmable tunneling technology support
- > Network address translation (NAT)

### Cloud-Tier Scale and Security

- > Half a million flexible entries shared across routes, MAC, ACL, MPLS, etc.
- > 10X better VXLAN VTEP and tunnel scale
- > 12.8Tb/s line rate encryption
- > Class C PTP and SyncE time synchronization
- > Secured switch offerings via hardware root-of-trust

### Accelerated Event-Based Telemetry

- > Actionable telemetry with WJH
- > FabricView histograms
- > Streaming telemetry
- > In-band network telemetry
- > Telemetry CPU offloads

## Actionable Visibility and Telemetry

NVIDIA Spectrum-4 provides packet-granular, network-based visibility data that seamlessly integrates with open APIs at the SDK layer. Referred to as NVIDIA 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, 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, VXLANsec, and secure boot as the default via hardware root of trust (HROt).

## Choice of Software

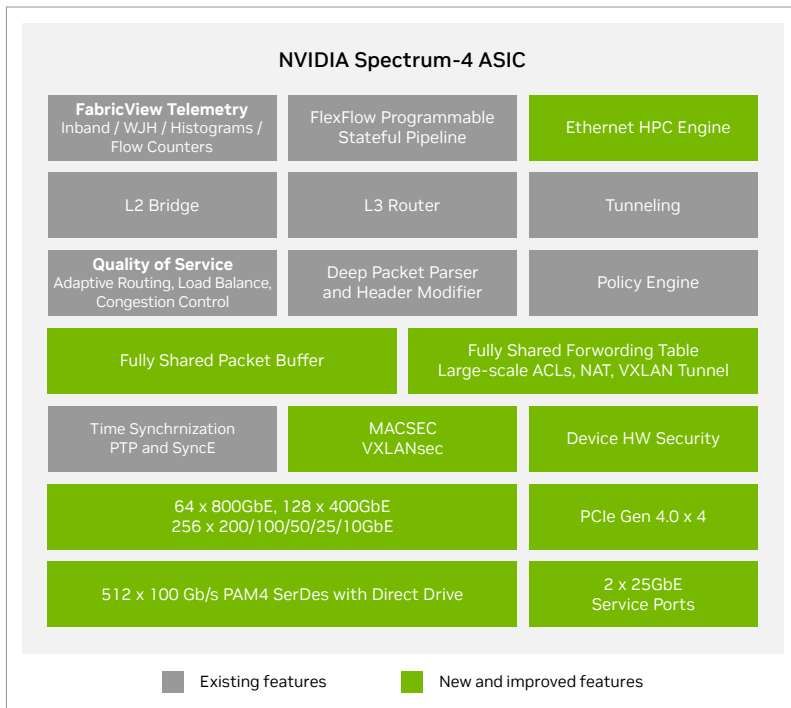
NVIDIA offers a rich set of software development tools and open APIs that enable users and partners to develop fully functional switching solutions on top of Spectrum-4 at low TCO, with a short time to market. Users have the flexibility to choose the network operating system that's best suited to their needs, including the flagship NVIDIA 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 Spectrum-4. Combined with the Spectrum SDK and choice of building on custom operating systems on top of the Switch Abstraction Interface (SAI), the Spectrum-4 ASIC is purpose-built to support full network customization and programmability.

## End-to-End Synergy

As part of the NVIDIA Spectrum platform, the Spectrum-4 Ethernet switch ASIC is pretested and prevalidated with NVIDIA's full portfolio of Ethernet networking technology, including BlueField-3 DPUs and ConnectX-7 smart network interface cards (SmartNICs), based on the same 100G SerDes channels used with Spectrum-4. This end-to-end solution is critical for powering NVIDIA accelerated AI and simulation platforms, and delivers the high performance and innovative feature set required to supercharge cloud-native applications at scale.

Compatibility	
<b>Ethernet</b>	<b>Switch ports</b> 800GbE 400GbE 200GbE 10/100GbE DCB (PFC, ETS, DCBX)
<b>CPU Options</b>	<b>Intel x86, AMD x86, Arm®</b>
<b>AI Solutions</b>	<b>Optimized for NVIDIA:</b> Spectrum-X DGX HGX OVX BlueField ConnectX® NVIDIA Magnum IO™ NVIDIA RAPIDS™ NVMe-oF
<b>IO Specs</b>	<b>Standards:</b> SPI flash I2C, SMBus, MDIO IEEE 1149.1, 1588

## New and Improved Features



## Ordering Information

OPN	Description
<b>SPC4-E0256EC11C-A0</b>	NVIDIA Spectrum-4 Ethernet switch ASIC 51.2Tb/s with 512x 100Gb/s PAM4 and 64x 800GbE interfaces
<b>SPC4-E0128DC11C-A0</b>	NVIDIA Spectrum-4 Ethernet switch ASIC 25.6Tb/s with 256x 100Gb/s PAM4 and 32x 800GbE interfaces
<b>SPC4-E0256EG11C-A0</b>	NVIDIA Spectrum-4 Ethernet switch ASIC 25.6Tb/s with 512x 50Gb/s PAM4 and 64x 400GbE interfaces

## Ready to Get Started?

To learn more about NVIDIA Spectrum-4, visit:  
[nvidia.com/en-us/networking/ethernet-switching](https://nvidia.com/en-us/networking/ethernet-switching)