



NVIDIA NetQ

Holistic insight and real-time visibility, troubleshooting, and monitoring for AI factories.



As AI factories become the norm, complexity increases. Network operators face constant change within multiple network layers, and legacy polling-based network management tools simply can't adapt to the dynamic demands of AI infrastructure. With the advent of generative AI, maintaining a healthy network is critical for maximizing the utilization of GPU infrastructure.

Network operators often struggle with operational challenges, such as network disruptions from maintenance, configuration changes, and simple misconfigurations, all of which increase operator workloads. As networks grow larger and more complex, the tasks a network administrator needs to perform can quickly overwhelm manual efforts. This shift demands modern networking and operational tools ready for the era of AI.

NVIDIA NetQ is a highly scalable, modern network operations tool that provides actionable visibility for the NVIDIA Spectrum-X Ethernet platform, NVIDIA Cumulus Linux, NVLink Switches and NVIDIA SuperNICs, NVIDIA BlueField DPUs, NVIDIA GPUs and AI workloads.

NetQ is built to optimize NVIDIA's accelerated computing platforms—such as NVIDIA DGX BasePOD™ and DGX SuperPOD—and AI solution stacks like NVIDIA Enterprise. NetQ helps identify network issues that may slow down AI processes. Leveraging fabric-wide telemetry data, NetQ provides real-time visibility and troubleshooting for the entire AI fabric. It offers:

- > Network-outage prevention using validation and functional testing with network-continuous integration/continuous delivery (CI/CD) in the NVIDIA Air platform
- > Holistic visibility across Ethernet and NVLink switches, combined with AI workload insights
- > Prescriptive Grafana dashboards—battle-tested in NVIDIA's own data centers—that accelerate troubleshooting and enable optimization based on proven operational knowledge
- > Comprehensive validation capabilities, now including Cable Validation, to accelerate time to first token and ensure AI factory readiness
- > Rapid root-cause detection using network telemetry data reducing mean time to innocence for NVIDIA NVLink switches and NVIDIA Spectrum-X Ethernet switches including NVIDIA What Just Happened (WJH) data

Key Benefits

- > End-to-End holistic observability
- > Grafana integration
- > Validations
- > GPU partitioning management
- > Embedded cable validation tool
- > Slurm job monitoring
- > Path trace insights
- > NVIDIA What Just Happened (WJH)
- > Flow telemetry analysis
- > RDMA over Converged Ethernet (RoCE) monitoring
- > Events and threshold crossing alerts
- > Notification channels
- > Software upgrade management
- > Snapshot and compare
- > Topology visualization and validation
- > Microservices architecture
- > High-availability clustering
- > APIs for integration
- > Switch process monitoring

- > Multiple-event notification integrations (Syslog, PagerDuty, Slack, email, and Generic Webhook)

NetQ is tested and validated with NVIDIA NVLink switches and the NVIDIA Spectrum-X Ethernet networking platform, providing an end-to-end switch-to-host solution that powers accelerated workloads. It delivers the high performance and innovative feature set needed to supercharge AI applications at scale.

Protect Network Integrity With Validations and CI/CD

Network configuration changes and software upgrades can cause numerous trouble tickets due to the inability to test updates before deploying them in production, which also impacts the continuity required for AI workloads. When this happens, a large amount of data is collected and stored in multiple tools, making it difficult to correlate events to resolve issues—something that's especially problematic for latency-sensitive AI applications. NetQ can be used as the functional test platform for the network CI/CD in conjunction with NVIDIA Air. Customers benefit from testing the new configuration as a digital twin with NetQ in the NVIDIA Air environment, being able to fix errors before deploying to their real-world production network. NetQ validations provide insight into the live state of the network, shorten troubleshooting times, and prevent network issues like maximum transmission unit (MTU) mismatches, virtual local area network (VLAN) misconfigurations, and more.

Rapid Root-Cause Detection

NetQ greatly reduces time to innocence by pinpointing and isolating faults caused by network state changes. Working with Cumulus Linux and NVOS, NetQ enables organizations to validate network state, both during regular operations and for post-mortem diagnostic analysis. NetQ provides both a CLI and a robust GUI, enabling on-box interactions for troubleshooting and high-level visual insights and Grafana dashboards.

With the NetQ trace capability, paths are verified, providing additional information that NetQ uses to discover misconfigurations along all the hops, speeding the time to resolution. NetQ trace allows users to view all of the paths between devices to find potential problems.

With AI workloads demanding high network throughput, NetQ offers insights into real-time performance with curated dashboards tailored for Spectrum-X Ethernet. The NetQ agents running on switches and hosts monitor various events—such as interface state, Border Gateway Protocol (BGP) neighbors, media access control (MAC) addresses, RoCE and routes—in real time. This provides a single source of truth for fabric-wide events that can be viewed via NetQ's CLI, GUI, Grafana and multiple third-party notification services like PagerDuty and Slack.

Management

NetQ provides centralized management for NVLink fabrics, simplifying administration at AI scale. With its GPU partitioning capability, operators can carve the fabric into logical partitions aligned with specific workloads, tenants, or clusters ensuring traffic isolation, predictable GPU-to-GPU performance and efficient resource utilization.

NetQ also supports **domains**, enabling segmentation of large fabrics into independent administrative units for scoped monitoring, role-based access, and domain-level health validation.

- > Precision Time Protocol (PTP) monitoring
- > Application-specific integrated circuit (ASIC) monitoring

Proof Points

- > Simplifies scaling of NVIDIA Spectrum-X Ethernet platform and NVLink Fabrics
- > Speeds mean time to innocence and deployment
- > Reduces opex
- > Reduces downtime
- > Increases productivity
- > Simplifies upgrades
- > Reduces security risks
- > Maximizes value of network infrastructure

Additionally, NetQ delivers lifecycle management for NVLink and Spectrum switches, NetQ agents covering software deployment, upgrades, and ongoing operational consistency across the fabric.

NetQ integrates with orchestration pipelines through REST APIs and exposes Prometheus metrics for observability. Seamless integration with Grafana and other monitoring platforms delivers real-time KPIs and domain/partition health insights.

Together, these management capabilities give operators a single control point to create, assign, and monitor GPU partitions, manage domains, validate NVLink and Ethernet fabric health, and streamline operations across AI environments.

Deploy Reliable Networking With WJH and Flow Telemetry

WJH is a hardware-accelerated telemetry feature available on NVIDIA Spectrum Ethernet switches that streams detailed and contextual telemetry data for analysis.

WJH provides real-time visibility into problems in the network, such as hardware packet drops due to misconfigurations, buffer congestion, access control list (ACL), or layer-one problems. This visibility is crucial for data centers built for AI, as AI workloads are especially sensitive to performance degradation and outages.

WJH provides telemetry data from the switches collected by NetQ, extending GUI and CLI functionality to WJH as well. When WJH capabilities are combined with NetQ, packet drops can be identified anywhere in the fabric to improve network reliability:

- Viewing current or historical drop information, including the reason for the drop.
- Identify problematic flows or endpoints and pinpoint exactly where communication is failing in the network.
- Get contextual WJH drops information in the output with NetQ trace.

gRPC network management interface (gNMI) can also be used to collect WJH data from the NetQ Agent.

While WJH is always on—detecting packet drops, latency, and congestion events—flow telemetry provides on-demand analysis of specific application flows. NetQ, working with Cumulus Linux, samples packets matching a 4-tuple or 5-tuple application flow, analyzes and reports per-switch latency (maximum, minimum, and average), and provides buffer-occupancy details along the path of the flow. The NetQ GUI reports all the possible paths, paths in use, and per-path details.

By combining WJH with flow telemetry analysis, network operators can proactively identify root-cause server and application issues. This lets them inform server or application administrators about potential outages or performance impacts critical to AI workloads.

NetQ Components and Deployment Options

NetQ Components

- > **NetQ Agents:** run on Cumulus Linux switches and other certified Linux systems and Ubuntu hosts. NetQ Agents capture network data and other state information in real time and transmit the data to the NetQ Server.
- > **NetQ Server:** consists of telemetry data collection software, an on-premises telemetry aggregator (OPTA), data analytics applications, and the database. The NetQ applications and database can be deployed on premises or consumed as a cloud-based service.

NetQ on Customer Premises

In this deployment option, all NetQ components are deployed on customer premises. Deployments can span a single site or multiple sites.

- > **Single-site deployment:** NetQ Agents running on switches and hosts collect and transmit data to the NetQ OPTA, which hosts the NetQ applications and database.
- > **Multi-site deployment:** NetQ Agents at each site collect and transmit data from the switches and hosts to the local OPTA. The OPTAs then transmit the data to a common NetQ application server for processing and storage.

For high availability, OPTAs and applications with storage can be deployed as a cluster.

Ready to Get Started?

To learn more about the NVIDIA NetQ, visit:
nvidia.com/netq