



NVIDIA NETQ

Holistic and Real-Time Visibility,
Troubleshooting, and Monitoring



As cloud-scale networking becomes the enterprise norm, so does complexity. Network operators must manage constant change within multiple network layers and polling-based legacy network management tools simply cannot adapt.

Network operators often struggle with operational challenges, such as network disruption caused by maintenance and configuration changes, and simple misconfigurations can have a significant impact on operator workloads. Furthermore, business networks are usually large and complex, which means the set of tasks a network administrator needs to perform can quickly overwhelm manual efforts. This requires a shift, not only to modern networking, but also to modern operational tools.

NVIDIA® NetQ™ is a highly-scalable, modern network operations tool that provides actionable visibility for the NVIDIA Spectrum™ Ethernet platform, including NVIDIA® Cumulus® Linux and software for open networking in the cloud (SONiC), as well as NVIDIA DPUs.

NetQ is built to accelerate NVIDIA platforms, including NVIDIA EGX™, NVIDIA DGX POD™ and NVIDIA OVX™ SUPERPOD, and AI solution stacks such as NVIDIA AI Enterprise and NVIDIA LaunchPad. NetQ leverages fabric-wide telemetry data to provide visibility and troubleshooting of the overlay and underlay network in real time, delivering the following benefits:

- > Network outage prevention using validation and functional testing with network continuous integration/continuous delivery (CI/CD), in conjunction with NVIDIA Air.
- > Rapid root cause detection using network telemetry data, including NVIDIA What Just Happened™ (WJH) data from NVIDIA switches, reducing mean-time-to-innocence.
- > Fabric-wide latency and buffer occupancy analysis of all the paths of a 4-tuple or 5-tuple flow to identify congestion points impacting application performance.
- > Network-wide telemetry database to optimize network usage supporting GUI, CLI, API, and plugins (Grafana, etc.).
- > Multiple event notification integrations (Syslog, PagerDuty, Slack, Email, and Generic Webhook).

As part of the Spectrum platform, NetQ is tested and validated with NVIDIA's full portfolio of Ethernet networking technology, including NVIDIA® BlueField® DPUs and NVIDIA® ConnectX® SmartNICs. This end-to-end switch to host solution is critical for powering accelerated workloads, and delivers the high performance and innovative feature set needed to supercharge cloud-native applications at scale.

KEY FEATURES

- > Validations
- > Trace
- > WJH
- > Flow Telemetry Analysis
- > RoCE Monitoring
- > Events and Threshold Crossing Alerts
- > Notification Channels
- > Software Upgrade Management
- > Snapshot and Compare
- > Topology
- > Microservices Architecture
- > HA Clustering
- > APIs for Integration

PROOF POINTS

- > Simplified Scaling of Cumulus Linux
- > Speeds Mean-Time-to-Innocence
- > Reduces OpEx
- > Cuts Downtime
- > Increases Productivity
- > Simplifies Upgrades
- > Reduces Security Risks
- > Maximizes Value of Network Infrastructure

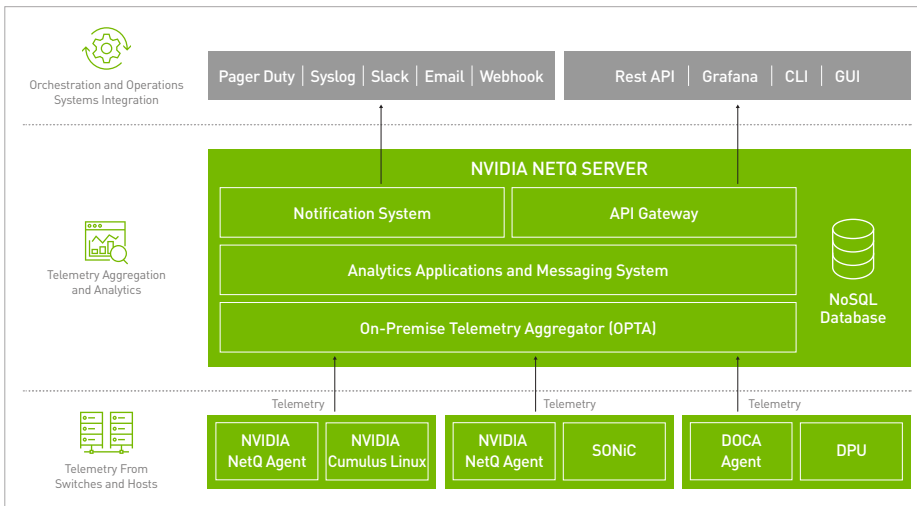


Figure 1: NetQ real-time telemetry data collection and deep analytics.

Protect Network Integrity with Validations and CI/CD

Network configuration changes and software upgrades can cause numerous trouble tickets because of the inability to test before deploying in production. When this happens, a large amount of data is collected and stored in multiple tools, making correlating events to resolve issues difficult at best. 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 with NetQ in the NVIDIA Air environment (“digital twin”) and fix errors before deploying to their production network (“physical twin”). In the physical production network, NetQ validations provide insight into the live state of the network, shorten the troubleshooting time, and prevent network issues like MTU mismatch, 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 hand-in-hand with Cumulus Linux and SONiC, NetQ enables organizations to validate network state, both during regular operations and for post-mortem diagnostic analysis. NetQ provides both a CLI and robust GUI to allow for on-box interactions as part of troubleshooting or visually as a high-level dashboard.

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

The NetQ agents running on switches and hosts monitor various events in real time, like interface state, BGP neighbors, MACs, and routes, providing a single source of truth for data center-wide events. These events can be viewed via NetQ CLI, GUI, and multiple third-party notification services like PagerDuty or Slack.

Deploy Reliable Networking with WJH and Flow Telemetry

WJH is a hardware-accelerated telemetry feature available on NVIDIA Spectrum 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, ACL, or layer 1 problems.

WJH telemetry data from the switches is collected and aggregated 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 by:

- > Viewing current or historic drop information, including the reason for the drop.
- > Identifying problematic flows or endpoints and pinpointing exactly where communication is failing in the network.
- > Including 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 4-tuple/5-tuple application flow, analyzes and reports per switch latency (max, min, avg), and buffers occupancy details along the path of the flow. The NetQ GUI reports all the possible paths, paths in use, and per-path details. On each switch, NetQ shows minimum latency, maximum latency, average latency, and buffer occupancy.

By combining WJH with flow telemetry analysis, network operators can proactively identify server and applications issues, and inform the server/application administrator about the possible outage or performance impact.

NetQ Components and Deployment Options

NetQ Components

- > **NetQ Agents** run on Cumulus Linux and SONiC switches and other certified Linux systems, such as Ubuntu®, Red Hat®, and CentOS 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, “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 (SaaS)

NetQ on Customer Premises

In this deployment, 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:** For the multi-site NetQ implementation, the NetQ agents at each premise collect and transmit data from the switches and hosts to the local OPTA. The OPTAs then transmit the data to a common NetQ applications server for processing and storage.

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

NetQ as a Cloud Service

NetQ as a cloud service is similar to the multi-site deployment, where the OPTAs run on premises at the customer site, securely connecting to the NetQ multi-tenant cloud service operated and maintained by NVIDIA.

NetQ Physical Appliance Specifications

The NetQ physical appliance is a 1RU rack mountable server available in two forms. For deploying NetQ OPTA, applications, and databases, use the NetQ physical appliance with eight cores, 96GB DRAM, 960GB SSD, and a 4TB hard drive. For deploying NetQ OPTA only, use the NetQ physical appliance with four cores, 16GB DRAM, and a 240GB hard drive.

Ordering Information

For ordering information, please contact gopny@pny.com

[Learn more](#)

Learn more at: pny.com/networking