



Naddod N9570-128QC Datasheet

AI Datacenter 128*400G RoCE Switch

NVIDIA Spectrum-4 Solution

NADDOD Pte.Ltd.

All rights reserved.

1. Product Overview

In the era of rapid advancements in generative AI and exponential growth in data traffic, network infrastructure faces unprecedented challenges. Recognizing these industry trends, Naddod has developed the **N9570-128QC AI RoCE Switch**, a cutting-edge solution designed to meet the demands of modern AI-driven networks.

The N9570-128QC integrates **128*400G QSFP112 ports**, significantly enhancing port density while minimizing physical footprint. This design ensures seamless connectivity for large-scale data center deployments, addressing the growing need for high-bandwidth applications. With its robust port configuration, the switch reduces the risk of network bottlenecks caused by insufficient ports, ensuring smooth and efficient data transmission.


Equipped with AI-driven network functionalities, such as **intelligent traffic scheduling** and **adaptive congestion control**, the N9570-128QC optimizes network performance for AI applications. These features reduce latency, accelerate AI model training, and improve overall network efficiency.

The switch features an **independent management plane**, enabling administrators to perform comprehensive and precise management operations. Supported by the **Naddod Network Operating System**, the N9570-128QC offers exceptional flexibility and scalability, allowing users to customize network configurations to meet specific business needs.

The N9570-128QC is versatile and can be deployed as an **AIGC network switch**, providing reliable network support for AI research and development, or as a **data center switch**, meeting the daily data processing and transmission requirements of modern data centers. This versatility helps businesses build high-performance **AI Fabric** and data center networks, confidently addressing the network challenges of the AI era.

2. Product Pictures

N9570-128QC



Front View



Rear View



Isometric View

3. Product Features

Data Center AI Switch Powered by Spectrum-X High-Performance Switching Chip.

This high-performance AI switch leverages the robust computational power and advanced architecture of NVIDIA's Spectrum-4 chip, meticulously optimized from the hardware layer up to the software layer. Specifically designed for complex computations of large AI models and massive data interaction scenarios, it excels in key performance metrics such as network bandwidth, low latency, and high concurrency. It provides a solid, reliable, and efficient network foundation for both training and inference of large AI models, empowering enterprises to achieve breakthroughs and innovation in the AI domain.

Seamlessly compatible with the NVIDIA Spectrum-X networking platform, its AI networking performance rivals that of InfiniBand, **delivering 1.6x the performance of traditional Ethernet.**

This high-performance AI switch leverages the robust computational power and advanced architecture of NVIDIA's Spectrum-4 chip, meticulously optimized from the hardware layer up to the software layer. Specifically designed for complex computations of large AI models and massive data interaction scenarios, it excels in key performance metrics such as network bandwidth, low latency, and high concurrency. It provides a solid, reliable, and efficient network foundation for both training and inference of large AI models, empowering enterprises to achieve breakthroughs and innovation in the AI domain.

Ultra-Fast Connectivity, Scalable Expansion, Building Ultra-Large-Scale AI Networks.

With an ultra-high switching capacity of 51.2Tbps, combined with cutting-edge intelligent buffering technology and load balancing algorithms, it ensures zero packet loss during end-to-end 400G data transmission. Even in the face of sudden surges in massive data, its unique queue management mechanism precisely schedules data flows. This not only guarantees stable network performance but also significantly enhances data transmission efficiency, meeting the demands of high-performance application scenarios.

Featuring 128*400G ports, it offers an ultra-large-scale configuration, enabling the seamless deployment of AI computing clusters with up to 524K GPUs, empowering large-scale AI networks.

End-to-Network Collaboration, Ultimate Efficiency, Accelerating AI Computing Power

This device deeply integrates adaptive routing technology and fine-grained load balancing mechanisms. It intelligently plans optimal data transmission paths and precisely allocates network loads based on real-time network topology and traffic dynamics. With this exceptional performance, the device maximizes network bandwidth potential, **achieving up to 95% bandwidth utilization**. In large model training scenarios, efficient bandwidth utilization significantly reduces data transmission time, minimizes waiting periods and latency during training, and greatly enhances the training efficiency of large models.

The device fully supports advanced congestion control technologies such as Round-Trip Time Congestion Control (RTTCC). These technologies enable high-precision real-time monitoring of network traffic, dynamically adjust transmission rates, and optimize data paths through intelligent algorithms, minimizing network congestion. In AI computing scenarios, they can reduce job completion time (JCT) by up to 30%, significantly improving the execution efficiency and timeliness of AI tasks, ensuring smooth and efficient business operations.

In complex multi-tenant and multi-task operational environments, it effectively isolates the performance of different AI tasks. Using innovative resource allocation and precise traffic control, it ensures that AI tasks run without interference.

Open Ecosystem, Network Innovation, Driving Continuous Evolution of AI Networks

Compliant with the S3IP-UNP standard, it demonstrates exceptional openness and compatibility, seamlessly integrating with third-party network operating systems such as SONiC. This feature not only provides users with diverse system options to meet various network management needs but also breaks down technical barriers, promoting collaboration between different systems. It enables users to build flexible, efficient, and customizable AI network architectures, providing robust support for stable network operations and continuous innovation.

The Naddod Network Operating System (NOS) is meticulously constructed on the basis of SONiC, deeply integrating its advanced features and open source ethos. With strong support for container technology, Naddod NOS exhibits remarkable flexibility and scalability. Users can easily expand network functionalities by adding containers based on actual needs. This container-based flexible expansion approach not only significantly reduces the complexity and cost of network function upgrades but also

provides strong support for building personalized high-performance network environments. It helps enterprises maintain a leading position in the ever-evolving landscape of network technology.

4. Product Specifications

Hardware Specifications

System Specifications	N9570-128QC
Ports	128 ×400GE ports (QSFP112)
CPU	Intel Ice lake-D 1734NT(8Core)
memory	16GB
Expansion Modules	Four power module slots Five fan module slots
Management Port	One management port(100/1000Mbps), one console port, and one USB port, compliant with the USB2.0 standard
Switching Capacity	51.2Tbps
Packet Forwarding Rate	37.6bpps
Dimensions and Weight	
Dimensions (W × D × H)	446 mm x 800 mm x 173.6mm
Weight	≤37.6 kg (including five fan modules and four power modules)
Power Supply and Consumption	
Maximum Power Consumption	3,500 W
Environment and Reliability	
Operating Temperature	5°C to 35°C
Operating Humidity	-40°C ~ 70°C


Software Specifications

Software Specifications	N9570-128QC
Layer 2 Protocol	VLAN 802.1Q
	MAC:Dynamic and static address entries MAC address drift suppression
	LLDP (802.1ab)
	Storm suppression:broadcast, Automatic MAC address learning and aging, unknown unicast,unknown multicast
	Loop detection
IP	ARP: Static ARP, Dynamic ARP, ARP advertises a direct route
IP unicast routing	BGP
	Static route
	ECMP
QOS	Queue scheduling:SP, DWRR, SP+DWRR
	CoS: 802.1p, DSCP
	WRED
Security	Copp
AI network features	Adaptive Routing
	RDMA: ECN, PFC, PFC WD
O&M management	Zero Touch Provisioning (ZTP)
	SNMP v1/v2c/v3
	DHCP client
	NTP
	Syslog
	SSH v2.0/Telnet
	License
	gRPC
	PING/Traceroute

Software Specifications	N9570-128QC
	Software and configuration file transfer: FTP、TFTP
	Optical Transceiver monitoring
	Temperature monitoring
	Fan/power supply monitoring

5. Typical Applications

AI Fabric solution



6. Configuration Guide

Take the following steps to order an N9570-128QC switch:

- Select the chassis.
- Select optical transceivers based on port requirements.

7. Ordering Information

Chassis, Fan Modules, and Power Modules

Product Model	Description
N9570-128QC	128 × 400GE ports, four power module slots, and five fan module slots

100GBASE Series Optical Transceivers

Model	Description
QSFP-100G-SR4	100G SR4 transceiver, QSFP28 form factor, MPO-12, 850 nm, 100 m (328.08 ft.) over MMF
QSFP-100G-CWDM4	100G CWDM4 transceiver, QSFP28 form factor, Duplex LC, 1310 nm, 2 km (6,561.68 ft.) over SMF
QSFP-100G-LR4	100G LR4 transceiver, QSFP28 form factor, Duplex LC, 1310 nm, 10 km (32,808.40 ft.) over SMF
QSFP-100G-ER4	100G ER4 transceiver, QSFP28 form factor, Duplex LC, 1310 nm, 40 km (131,233.59 ft.) over SMF
QSFP-100G-AXX	100G AOC cable, QSFP28 form factor, ≤100m (328.08 ft.)
QSFP-100G-CUX	100G DAC cable, QSFP28 form factor, ≤5 m (16.4 ft.)

200GBASE Series Optical Transceivers

Model	Description
Q56-200G-SR4	200G SR4 transceiver, QSFP56 form factor, MPO-12 UPC, 850nm, 100 m (328.08 ft.) over MMF
Q56-200G-FR4	200G FR4 transceiver, QSFP56 form factor, Duplex LC, 1310nm, 2km (6561.68 ft.) over SMF
Q56-200G-AXX	200G AOC cable, QSFP56 form factor, ≤100m (328.08 ft.)
Q56-200G-CUX	200G DAC cable, QSFP56 form factor, ≤3 m (10 ft.)

400GBASE Series Optical Transceivers

Model	Description
Q112-400G-SR4	400G SR4 transceiver, QSFP112 form factor, MPO-12 APC, 850 nm, 100 m (328.08 ft.) over MMF
Q112-400G-VR4	400G VR4 transceiver, QSFP112 form factor, MPO-12 APC, 850 nm, 50 m (1,64.04 ft.) over MMF

Model	Description
Q112-400G-DR4	400G DR4 transceiver, QSFP112 form factor, MPO-12APC, 1310 nm, 500 m (1,640.42 ft.) over SMF
Q112-400G-FR4	400G FR4 transceiver, QSFP-112 form factor, Duplex LC, 1310 nm, 2km (6561.68 ft.) over SMF
Q112-400G-CUX	400G DAC cable, QSFP112 form factor, ≤2 m (7 ft.)

8. Warranty

For more information about warranty terms and period, contact your local sales agency:

- Warranty terms: <https://www.naddod.com/support/>
- Warranty period: <https://www.naddod.com/support/>

Note: The warranty terms are subject to the terms of different countries and distributors.

9. More Information

For more information about Naddod, visit the official Naddod website or contact us:

- Naddod official website: <https://www.Naddod.com/>
- Online support: <https://www.Naddod.com/support>
- Email support: support@Naddod.com