H100 optical module requirements

Overview

In summary, the network architecture and optical module requirements are crucial considerations for the NVIDIA DGX H100 server cluster. The compute network utilizes a two-layer switch architecture with NVIDIA QM9700 switches, requiring a significant number of 800G and 400G optical. In the HPC network industry, various methods are employed to determine the quantity of optical modules to GPUs, resulting in inconsistent outcomes. These discrepancies primarily arise from various optical modules utilized in different network architectures. In this article, we delve into these. The NVIDIA DGX™ H100/H200 Systems are the universal systems purpose-built for all AI infrastructure and workloads from analytics to training to inference. The NVIDIA DGX H100 (640. The basic deployment structure for each DGX A100 SuperPOD consists of 140 servers (each server with 8 GPUs) and switches (each switch with 40 ports, each port at 200G). The network topology is an InfiniBand (IB) Fat-Tree structure. 3 400G optic transceivers, meaning each H100 requires 1. Each Leaf switch has 16 interfaces (800G) for upward transmission and 16 interfaces (800G) for downstream transmission. Therefore, a total of 32*16+32*16=1024 pieces optical.

H100 optical module requirements

For a 4SU cluster, assuming an all-optical network and a three-layer Fat-Tree architecture, 400G optical modules are used between servers and leaf switches, while 800G optical ...

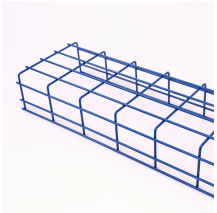

NVIDIA H100 GPUs feature fourth-generation Tensor Cores and the Transformer Engine with FP8 precision, further extending NVIDIA''s market-leading AI leadership with up to 4X faster training and ...

Currently, the ratio of GPUs to optical modules is approximately 1:2.5—meaning that, on average, for one Nvidia H100 GPU, two-and-a-half 800G optical modules are required for training activities within ...

This is driving a surge in the need for optical modules in data center interconnects. GPUs such as the A100, H100, and upcoming GH100 require high-speed optical interconnects to link thousands of GPU ...

Each DGX H100 is connected to eight Leaf switches. Each Leaf switch has 16 uplink ports connected to 16 Spine switches. Optical module usage: A 400 Gbit/s optical module is required for the ...

How Many Optical Transceivers Does an NVIDIA H100 Cluster Really Need? A practical breakdown of transceiver requirements in NVIDIA DGX H100 SuperPOD architectures using 400G ...

In this article, we delve into these factors and explore how they influence the exact quantity of optical modules needed, particularly focusing on the configurations involving A100 and ...

For the 4SU cluster, assuming an all-optical network and a three-layer Fat-Tree architecture, the server and Leaf layer switches use 400G optical modules, and the Leaf-Spine and Spine-Core use 800G ...

Discover the optic transceiver requirements for the NVIDIA DGX H100 server cluster, including computing and storage networks. Learn about the necessary configurations for efficient data transfer.

Here is an image that shows the real panel modules on DGX H100/H200.

Contact Us

For more information, pricing, or custom network solutions, please contact us:

Website: https://www.hashherbcafe.co.za
Email: hello@hashherbcafe.co.za
Phone: +27 63 814 7295
Address: 15 Galaxy Road, Linbro Business Park, Johannesburg, 2065, South Africa

This document is for informational purposes only. Specifications subject to change without notice.