

RIVO MULTITRADE OPTICAL SYSTEMS

Optical module used in B200



Overview

800G Optical Interconnect for NVIDIA B200 GPU Clusters Explore 800G and 400G optical interconnect solutions for NVIDIA B200 GPU clusters. Learn architecture.


Key Takeaways

The NVIDIA B200 does not strictly require an optical module for operation; it can use copper interconnects for single-cabinet setups, while optical modules are primarily used for multi-cabinet or large-scale cluster interconnections. Interconnect Options

The B200 platform, including the HGX B200 and GB200 variants, supports high-speed NVLink connectivity between GPUs. Internally, the GPUs and NVSwitches can be connected via copper backplane connections, which are sufficient for single-cabinet deployments and smaller clusters. This copper solution provides high bandwidth and energy efficiency without the need for optical modules.

Role of Optical Modules

OSFP optical modules are used in B200 systems primarily for external interconnects when scaling beyond a single cabinet. Each OSFP module can support multiple NVLink ports, enabling high-bandwidth connections across multiple cabinets or racks. Optical modules are advantageous for long-distance connections and ultra-large-scale clusters, where copper cabling would be impractical due to distance limitations and signal integrity concerns.



Practical Implications

- Single-cabinet deployments: Copper interconnects are sufficient, and optical modules are not required.
- Multi-cabinet or large-scale clusters: Optical modules are recommended to maintain high bandwidth and low latency across racks.
- Network adapters: The B200 supports NVIDIA ConnectX-7 OSFP400 InfiniBand adapters, which can utilize optical modules for external connectivity, but these are optional depending on the deployment scale .

Summary

While optical modules are part of the B200 ecosystem for high-speed, long-distance interconnects, they are not mandatory for all configurations. For standard single-cabinet setups, the system can operate entirely with copper connections, and optical modules are primarily used when scaling to multi-cabinet clusters or when ultra-high bandwidth is required .

Article Content

800G Optical Interconnect for NVIDIA B200 GPU Clusters

Explore 800G and 400G optical interconnect solutions for NVIDIA B200 GPU clusters. Learn architecture, scale-out

NVIDIA GB200 Interconnect Architecture Analysis: NVLink, InfiniBand ...

Expansion interfaces employ OSFP optical modules. As depicted in the following diagram, an OSFP optical module

Unveiling NVIDIA GB200 Interconnect Architecture

Expansion interfaces employ OSFP optical modules. As depicted in the following diagram, an OSFP optical module

From Node to SuperPod: Interconnect and Optical Design ...

This article provides a structured overview of Blackwell interconnect architectures and explains how optical module requirements

Nvidia's Optical Boogeyman - NVL72, Infiniband Scale Out, 800G

This optical boogeyman dramatically reduces the number of transceivers by a significant volume, and it's shipping in

NVIDIA HGX B200 and Thoughts on Its Liquid Cooling Solution

H3C R5500 G6 H100 Module Liquid Cooling 4-parallel-3-series (2 GPUs in Parallel + 1 Switch in Series) Based on the

NVIDIA GB200 Interconnect Architecture Analysis-NVLink, InfiniBand,

In this process, high-speed interconnect devices such as optical modules will continue to play a crucial role in

Contact Us

Continue your research online:

Website: <https://rivomultitrade.co.za>

Technical inquiry: <https://rivomultitrade.co.za/contact.html>

This document is for preliminary research. Verify specifications against current standards, product documentation and project requirements.

