


NVLink Bridge Support

NVIDIA NVLink is a high-speed point-to-point (P2P) peer transfer connection. Where one GPU can transfer data to and receive data from one other GPU. The NVIDIA H100 NVL card supports NVLink bridge connection with a single adjacent NVIDIA H100 NVL card.

Each of the three attached bridges spans two PCIe slots. To function correctly as well as to provide peak bridge bandwidth, bridge connection with an adjacent NVIDIA H100 NVL card must incorporate all three NVLink bridges. Wherever an adjacent pair of the NVIDIA H100 NVL cards exists in the server, for best bridging performance and balanced bridge topology, the NVIDIA H100 NVL pair should be bridged. Figure 3 illustrates correct and incorrect NVIDIA H100 NVL NVLink connection topologies.

Figure 3. NVLink Topology – Top Views



For systems that feature multiple CPUs, both NVIDIA H100 NVL cards of a bridged card pair, should be within the same CPU domain. That is, under the same CPU's topology, and ensuring this benefits workload application performance. There are exceptions, for example, in a system with dual CPUs wherein each CPU has a single NVIDIA H100 NVL card under it. In that case, the two NVIDIA H100 NVL cards in the system may be bridged together. See the “PCIe and NVLink Topology” section.

NVIDIA H100 NVL card, NVLink speed, and bandwidth are given in the following table.

Table 6. H100 NVL Card NVLink Speed and Bandwidth

Parameter	Value
Total NVLink bridges supported by NVIDIA H100	3
Total NVLink Rx and Tx lanes supported	48
Data rate per NVIDIA H100 NVL NVLink lane (each direction)	100 Gbps
Total maximum NVLink bandwidth	600 Gbytes per second


NVLink Bridge

The 2-slot NVLink bridge for the NVIDIA H100 NVL card (the same NVLink bridge used in the NVIDIA Ampere architecture generation, including the NVIDIA A100 PCIe card), has the following NVIDIA part number: 900-53651-0000-000.

NVLink Connector Placement

Figure 4 shows the connector keep-out area for the NVLink bridge support of the NVIDIA H100.

Figure 4. NVLink Connector Placement – Top View



Sufficient clearance must be provided both above the card's north edge and behind the backside of the card's PCB to accommodate NVIDIA H100 NVL NVLink bridges. The clearance above the north edge should meet or exceed 2.5 mm. The backside clearance (from the rear card's rear PCB surface) should meet or exceed 2.67 mm. Consult the *NVIDIA Form Factor 5.5 Specification for Enterprise PCIe Products Specification* (NVOnline: 1063377) for more detailed information.

NVLink bridge interfaces of the H100 NVL card include removable caps to protect the interfaces in non-bridged system configurations.

PCIe and NVLink Topology

As stated, it is strongly recommended that both NVIDIA H100 NVL cards of a bridged card pair should be within the same CPU topology domain. Unless a dual CPU system has only two H100 NVL cards each of which is under its own CPU. Full NVLink connection topology guidance is as follows:

- > Best NVLink Topology (Recommended):
 - Bridge two GPUs under the same CPU or PCIe switch
 - GPU count in a system should be in powers of two (1, 2, 4, 8, and so on)
 - Locate the same (even) number of GPUs under each CPU socket
 - Maintain a balanced configuration: same count of CPU:GPU:NIC for each grouping
- > Good NVLink Topology:
 - Bridge two GPUs under different PCIe switches but under the same CPU
 - Same number of GPUs and NICs under each CPU socket, but not powers of two
- > Allowed but Not Recommended:
 - Bridge two GPUs under two different CPUs
 - Odd number of GPUs under each CPU
 - Unbalanced configurations: Different ratios of CPU:GPU:NIC for each grouping