

Intel® Gaudi® 3 AI Accelerator HLB-325 Baseboard



The Intel® Gaudi® 3 baseboard (HLB-325) accommodates 8 Intel® Gaudi® 3 AI accelerator OAM mezzanine cards and provides customers a module subsystem that is easy to migrate into their AI server designs. The HLB-325 provides direct all-to-all connectivity with 4.2 Tera-Bytes (TB) per second of bi-directional bandwidth between the 8 Intel® Gaudi® 3 accelerators, without requiring a separate switching IC. In addition, the HLB-325 provides an additional 1.2TB per second of bi-directional scale-out bandwidth through 6 OSFP connectors, enabling massive scale-out in data centers using Ethernet switches.

Designed to enable efficient and simple integration into AI servers, the HLB-325 derives the majority of its power from 54V power input, only requiring a separate 12V input for standby. On-board PCIe retimers and scale-out retimers provide robust signal integrity of the high-speed signals coming to and from the baseboard. The 417mm width and 585mm length of the HLB-325 allow it to be incorporated into a variety of new and existing 19" AI accelerator server designs to provide state-of-the-art performance, scalability and power efficiency.

HLFB-325L Integrated Subsystem

The Intel® Gaudi® 3 baseboard and 8 Intel® Gaudi® 3 AI accelerator mezzanine cards are offered as an integrated subsystem to simplify customers' server manufacturing operations. The HLFB-325L is the combination of one HLB-325 baseboard and 8 HL-325L Intel® Gaudi® 3 AI accelerator mezzanine cards, tested together to enhance quality and shipped as a whole subsystem to ease high-volume server assembly and manufacturing.

Intel® Gaudi® 3 AI Accelerator HL-325L Mezzanine Card

| Host Interface | Scale-Out Interconnect | TDP | System Dimension | Power Input |
|----------------------|--|---|------------------|---------------------------|
| 8x PCIe Gen 5.0 x 16 | RDMA (RoCE v2) 24x200 Gbps via 6 OSFP connectors | 7.6 KW (including 8 HL-325L OAMs) | 19" | 54V main (12V standby) |