

CORNING



How to Build a Multi-POD Cluster with a Leaf-Spine-Core Architecture

How to build a Multi-POD Cluster with a Leaf-Spine-Core Architecture: GB200 16 POD Example

Scaling from a single SU, using 72 x 128F Core-Trunks at the Spine-to-Core layer and 64 x 144F Core-Trunks at both the Leaf-to-Spine and Server-to-Leaf layers, a fully built-out 16-POD architecture comprises:

Level C

Spine-to-Core Cabling

(choose between Structured Cabling or Point-to-Point for up to 500 m, but keep it simple!)

288 x 128F CORE-Trunks per SU or > 4,600 individual patch cords

Level B

Leaf-to-Spine Cabling

(Point-to-Point)

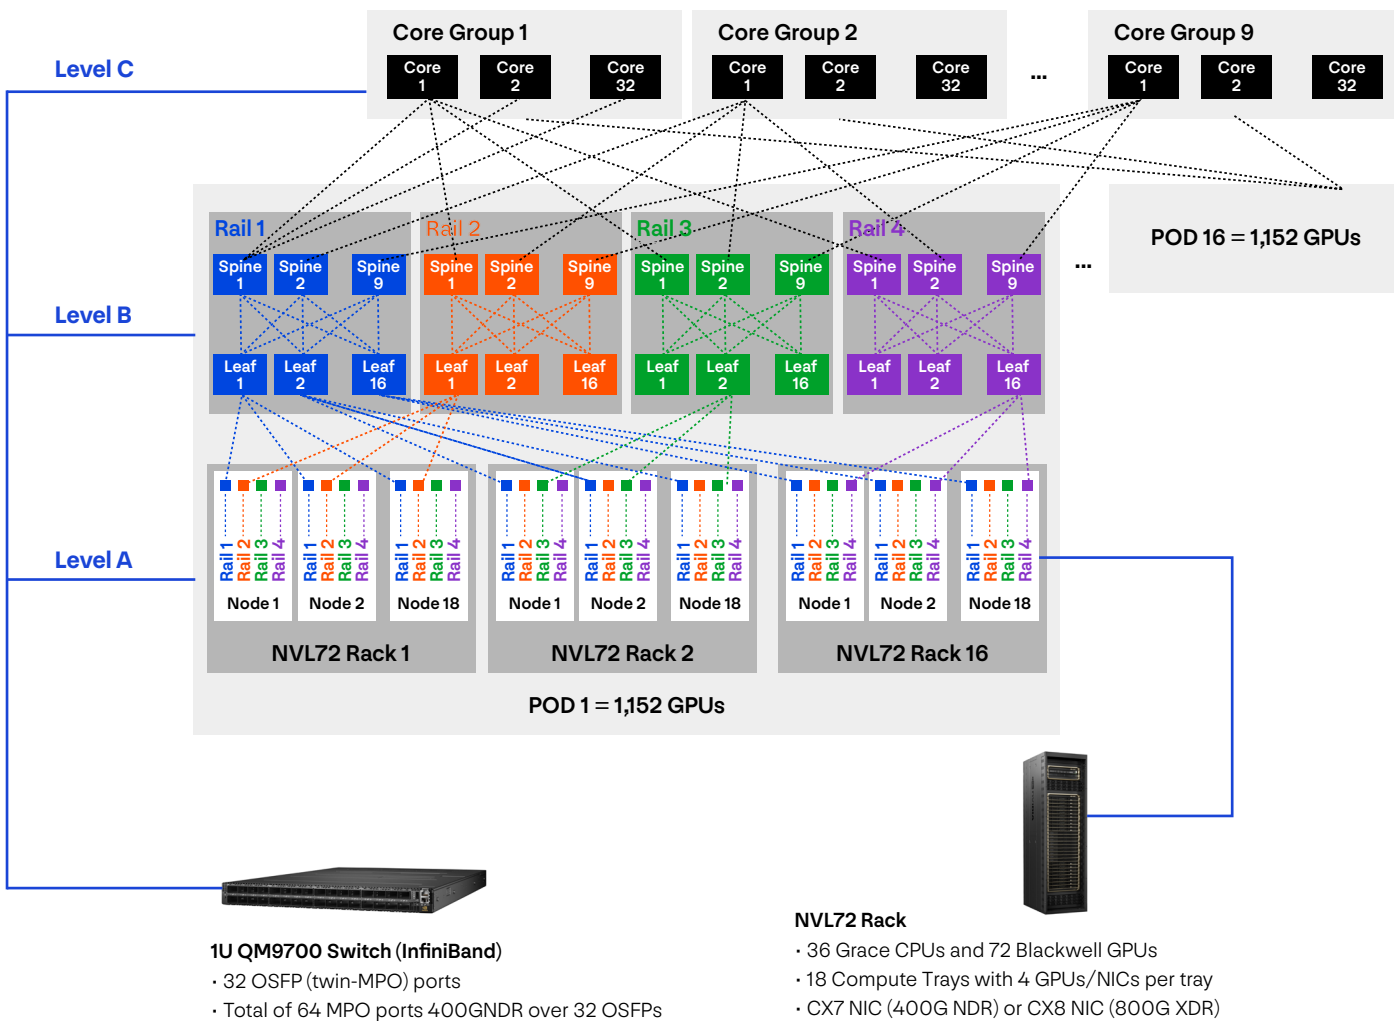
256 x 144F CORE-Trunks per SU or > 4,000 individual patch cords

Level A

Server-to-Leaf Cabling

(Point-to-Point)

256 x 144F CORE-Trunks per SU or > 4,000 individual patch cords



Getting ready for NVIDIA Blackwell/Grace-Blackwell Ultra (B300, GB300) or Vera Rubin (VR200)? We are happy to support your network build with our [Corning® GlassWorks AI™](#) portfolio.

Contact us for more information or help on network design.

CORNING

Corning Optical Communications GmbH & Co. KG • Leipziger Strasse 121 • 10117 Berlin, GERMANY
+00 800 2676 4641 • FAX: +49 30 5303 2335 • www.corning.com/opcomm/emea

Corning Optical Communications reserves the right to improve, enhance, and modify the features and specifications of Corning Optical Communications products without prior notification. A complete listing of the trademarks of Corning Optical Communications is available at www.corning.com/opcomm/trademarks. All other trademarks are the properties of their respective owners. Corning Optical Communications is ISO 9001 certified. © 2026 Corning Optical Communications. All rights reserved. LAN-3571-A4-BEN / April 2026