



NVIDIA HGX™ Rubin NVL8

Next Generation 8-GPU Server for
Large-scale AI Model Training & Inference

XA NR1I-E12L Hybrid Cooled

- Powered by dual Intel® Xeon® 6 processors, supporting a 350W TDP
- NVIDIA HGX™ Rubin NVL8, providing 800 Gb/s bandwidth for E/W bandwidth per GPU (OSFP)
- Independent CPU and GPU sled for easy maintenance

XA NR1I-E12LR Fully liquid-Cooled

- Powered by dual Intel® Xeon® 6 processors, supporting a 350W TDP
- NVIDIA HGX™ Rubin NVL8, providing 800 Gb/s bandwidth for E/W bandwidth per GPU (OSFP)
- Extreme Compute Density: 400 PFLOPS in a 2U Chassis
- HBM4 Memory Technology: A Generational Leap in Throughput
- Advanced Full Liquid-Cooled Solutions: Adopting a full liquid-cooled design to ensure peak performance and stability in 2U density while significantly optimizing data center PUE.



Get a quote

servers.asus.com/