VULTR HPE JUNIPER networking BROADCOM

SOLUTION BRIEF

AI-Ready Ethernet Fabric Networking for GPU Clusters

Vultr + HPE Juniper Networking + Broadcom

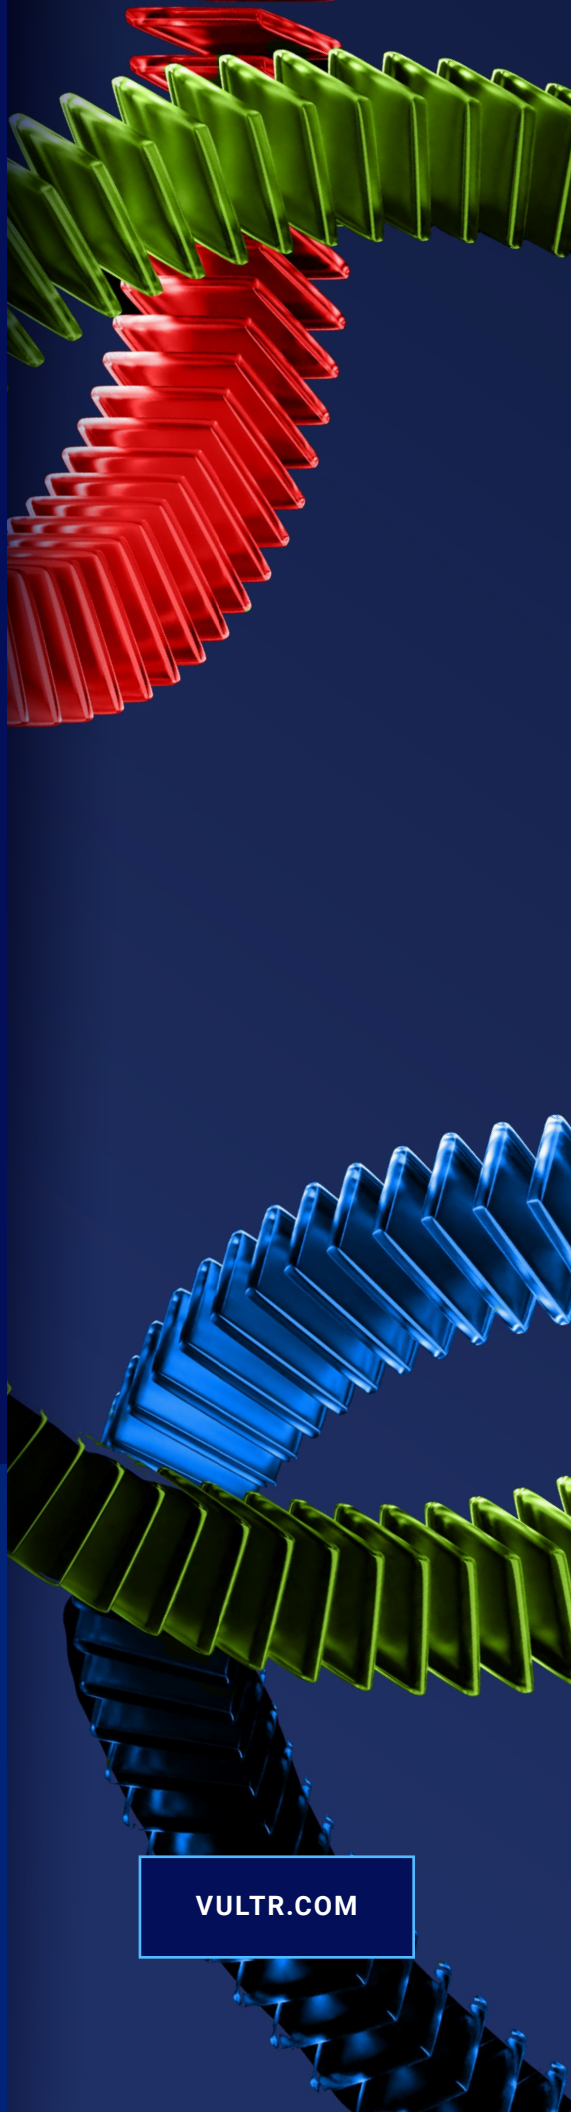

Deliver predictable GPU-to-GPU performance for AI training and inference by keeping throughput stable as clusters scale, improving GPU utilization, shortening job runtimes, and maintaining consistent latency under load on Vultr infrastructure with HPE Juniper Networking and Broadcom Ethernet silicon.

VULTR.COM

VULTR

Ethernet Fabric for GPU Clusters

Reduce network bottlenecks and deliver predictable GPU-to-GPU performance for AI training and inference, helping teams improve GPU utilization, shorten job runtimes, and scale clusters with consistent latency under load.

AI clusters don't fail because GPUs are slow. They fail when the network becomes the bottleneck under synchronized east–west traffic.

That's why Vultr builds AI infrastructure with a network-first design. For many deployments, Broadcom provides the Ethernet switching silicon foundation (Broadcom Tomahawk 5) for high-radix 400G/800G fabrics, while Broadcom Thor 2 400G NICs provide RoCEv2 hardware offload and ultra-low latency at the server edge. HPE Juniper Networking delivers the data center switching platforms and fabric capabilities optimized for AI, including RoCEv2 and congestion management.

Vultr delivers a production stack on top of bare metal, GPUs, integrated networking, and the operational layer to run AI clusters at scale. The solution uses HPE Juniper Networking QFX5240 built on Broadcom Tomahawk 5 to deliver a rail-optimized backend fabric designed for lossless RoCEv2 performance.

Challenges in networking

Microbursts and congestion

AI training generates bursty, synchronized traffic that can slow GPUs and extend job runtimes. Running on Vultr, HPE Juniper's lossless RoCEv2 fabric (using ECN/PFC with DCQCN), combined with Broadcom Tomahawk 5 load balancing and backend scale, keeps congestion controlled under load. Thor 2 RDMA NICs reduce CPU overhead through RoCEv2 offload and hardware-based congestion control. Together, this architecture delivers faster job completion and higher GPU utilization.

Inconsistent latency and throughput

Mixed AI traffic often leads to unstable latency and uneven throughput. Across Vultr's AI infrastructure, HPE Juniper Networking congestion management stabilizes RDMA traffic end-to-end, supported by Broadcom Tomahawk 5 high-density Ethernet switching and Thor 2 NICs sustaining 400G GPU data transfers. This ensures consistent throughput across distributed training and inference.

Operational complexity at scale

Lossless Ethernet requires careful tuning across ports, queues, and traffic classes to achieve optimal performance. HPE Juniper Validated Designs, combined with Vultr's standardized deployment model, reduce manual configuration and operational overhead. Broadcom-based switching enables repeatable and scalable fabric designs, speeding deployment and simplifying day-to-day operations.

Design risk

New link speeds and GPU generations increase the risk of interoperability and stability issues. In Vultr's 50MW AMD GPU deployment, HPE Juniper switching fabric built on Broadcom Tomahawk 5, together with Broadcom Thor 2 400G NICs, provides sustained host connectivity and stable, high-volume GPU flows for the rail-optimized backend fabric. This enables next-gen adoption with lower risk and predictable performance.

Joint solution overview

Vultr provides the bare metal and GPU infrastructure, Broadcom Thor 2 NICs deliver high-throughput host connectivity with RoCE offload, and HPE Juniper Networking builds the Ethernet fabric across frontend, backend, storage, and out-of-band (OOB) networks. Rail-optimized backend fabrics utilizing HPE Juniper Networking QFX5240 (powered by Broadcom Tomahawk 5 chips) support lossless RoCEv2 on the GPU fabric. Tomahawk 5 shared buffers and scheduling are designed to sustain throughput during heavy AI backend traffic. Together, this delivers a production-ready AI platform that reduces network bottlenecks, enhances GPU utilization, and scales reliably as workloads increase.

Our combined value:

- End-to-end AI networking stack across compute, NICs, and fabric – engineered to work together
- Lossless RoCEv2 backend fabric optimized for GPU-to-GPU communication and distributed training
- 800G-ready scalability with high-density switching for large AI clusters
- Predictable performance through congestion management and validated configurations
- Faster time-to-production with repeatable deployment patterns on Vultr's GPU infrastructure

Vultr + HPE Juniper Networking + Broadcom AI Architecture

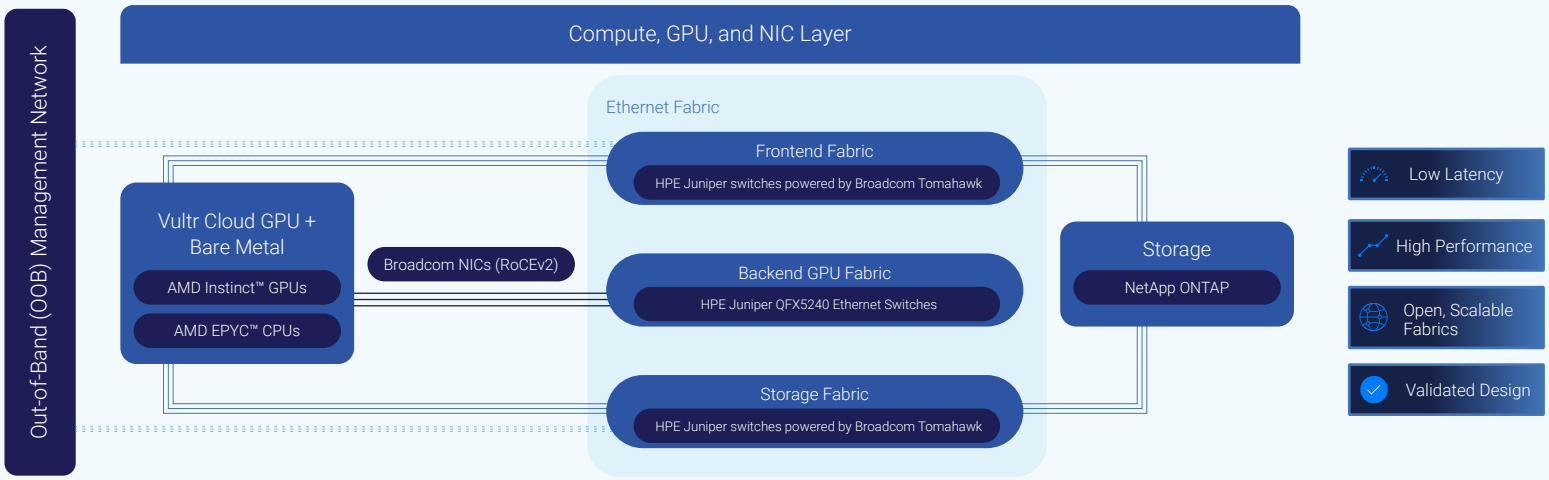

Vultr, HPE Juniper Networking, and Broadcom are members of the Ethernet Alliance

Operational advantages at AI scale

Predictable performance stems from tuned, lossless Ethernet behavior for RoCEv2 backends and silicon-level throughput that prevents GPUs from idling during data transfers. Repeatable outcomes result from HPE Juniper Networking interoperability testing and validated designs for end-to-end AMD GPU systems using QFX5240 components, together with Vultr's standardized deployment patterns.

Built on open standards, the fabric remains aligned with the broader Ethernet ecosystem as AI networking evolves, allowing customers to scale AI with confidence and consistent performance.

Performance outcomes

Low latency under load: Achieved through lossless Ethernet controls (ECN/PFC with DCQCN for RoCEv2) to minimize packet drops and stabilize RDMA performance during congestion

Predictable throughput: Improves distributed training efficiency and reduces variance in synchronized collectives by maintaining stable RDMA traffic across the fabric, supported by Tomahawk 5 shared buffers and Thor 2 RoCEv2 offload for sustained GPU data flows

Scalable fabrics: Aligned with standards-based Ethernet ecosystems and 800G switching enables cluster growth without redesigning the network foundation

Validated configurations: Reduces design and tuning risk with HPE Juniper Validated Designs deployed on Vultr infrastructure, built on Broadcom silicon, and with proven configuration guidance

Built for AI workloads

This architecture is designed for production AI workloads, including distributed training with heavy east–west GPU communication, high-throughput inference across batch and real-time services, RAG pipelines requiring fast storage ingest and efficient checkpointing, and multi-tenant AI platforms where isolation and predictable performance are critical.

Key takeaways for infrastructure leadership

If your goal is to run GPU clusters with consistent, predictable performance, your networking decisions must prioritize lossless behavior on the GPU fabric, operational repeatability, and scalable design. Vultr provides the compute layer and deployment environment, HPE Juniper Networking delivers the high-performant Ethernet fabric and validated designs for AI-scale operation, and Broadcom supplies the silicon foundation that enables high-bandwidth, high-radix Ethernet fabrics with Tomahawk 5 switching and Thor 2 RDMA NICs. The result is high-performance production AI networking on Vultr infrastructure.

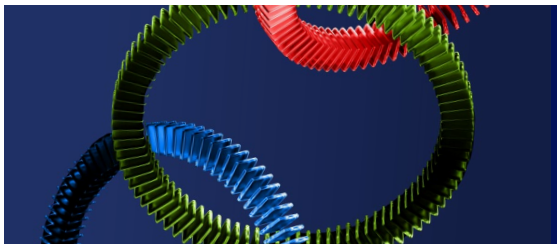

Talk to Vultr about deploying an AI-ready Ethernet fabric.

Contact us at vultr.com to get started.