

Juniper's Validated Design (JVD) for Building an AI Data Center

A proven, repeatable data center network design to support your AI aspirations

Get help designing and deploying your AI data center.

[Learn more →](#)

The challenge


Unprecedented AI network demands

Artificial intelligence (AI) and machine learning (ML) training is a massive parallel processing problem requiring the world's most sophisticated data center networks to handle the demands of AI cluster workloads. AI clusters pose unique demands on network infrastructure due to their high-density traffic patterns, characterized by frequent elephant flows with minimal flow variation. As with any new technology, the AI/ML cluster learning curve can be steep, resulting in time and cost delays that are exceedingly expensive for AI/ML architectures.

AI data center networks are comprised of frontend, backend, and storage, and they must be able to scale up and scale out. The backend and storage components must also seamlessly work together in order to minimize job completion time (JCT). This is important because even a slight loss of efficiency can add significant training time to models.

The traffic in data centers generated from distributed ML workloads dwarfs that of most other applications. AI requirements to communicate large datasets and process billions or trillions of model parameters put heavy stress on the network.

The key objectives in designing your network for an AI cluster are to provide **maximum throughput, minimal latency, and minimal network interference** for AI traffic flows over a lossless fabric.



The capabilities you need

A fully tested, well-documented network design

Optimizing AI/ML performance requires maximizing Graphics Processing Unit (GPU) utilization for speed and efficiency.

To simplify the process and economize your AI investment, Juniper Validated Designs (JVDs) offer frontend and backend data center blueprints to speed up time to stability, accelerate job completion time (JCT), and optimize AI inference performance.

- **Repeatability**
Prescriptive designs where all JVD users benefit from lessons learned in worldwide deployments
- **Reliability**
Integrated, best practice designs tested with real-world traffic and described with measured results
- **Velocity**
Streamlined deployments and upgrades with step-by-step guidance, automation, and prebuilt integrations
- **Consistency**
JVDs result in networks that are far more rigorous, stable, and maintainable

The answer

Juniper’s Validated Design for an AI Data Center

The [JVD for an AI Data Center](#) delivers turnkey designs for a comprehensive end-to-end template. It’s rigorously pretested and validated, and sharply reduces complexity and risk.

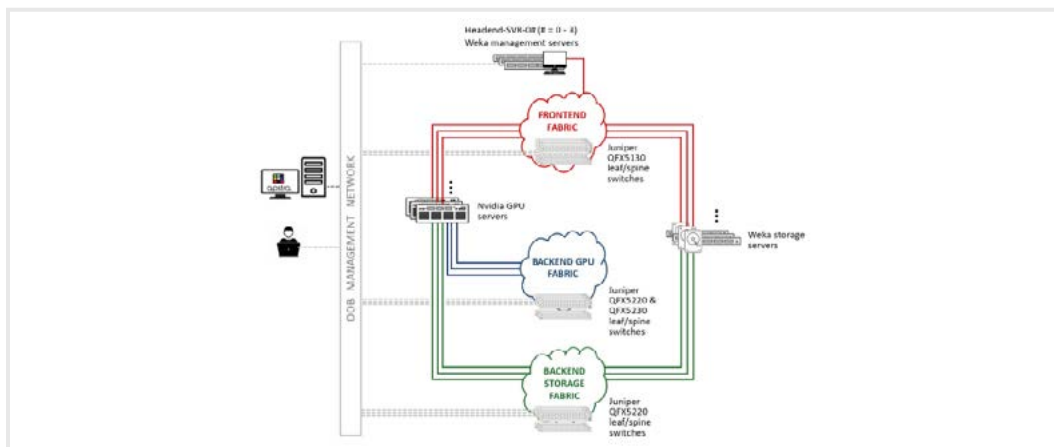


FIGURE 1
Design for AI data center

This JVD covers a complete end-to-end Ethernet-based AI infrastructure, including three separate fabrics:

- **Frontend Fabric:** This fabric is the gateway network to the GPU nodes and storage nodes. It is where users and systems access the GPU clusters.
- **Backend GPU Fabric:** This fabric connects the GPUs that consume the large models into memory from the stored datasets and facilitate the requests coming from the frontend network. The nodes (servers) in the backend GPU fabric pass high-speed information that requires a lossless network between GPUs during learning and training events.
- **Backend Storage Fabric:** This fabric connects the high-availability storage systems that hold the large model training data. The backend storage fabric must ensure seamless and reliable delivery of data to the GPUs for model learning.

These three fabrics (Figure 1) have a symbiotic relationship, and each provides unique functions in training and inference.

The fabrics are designed to ensure that the network is never a bottleneck in AI model training and GPU job completion. NVIDIA prescribes a network design called “rail-optimization,” which minimizes communication latency.

In a rail-optimized design, each server has eight network interfaces (one per GPU). Each GPU interface is cabled to a separate leaf switch. The frontend, storage, and backend GPU compute interfaces are kept on separate, dedicated networks. [Juniper Apstra](#) is ideal for managing all three fabrics as blueprints in the same Apstra instance.

The rail-optimized solution architecture is shown in Figure 2. Note the high-performance connectivity in all fabrics, particularly the backend GPU fabric between the GPU servers and the QFX leaf switches.

Juniper data center networks with fabrics built from [QFX Series Switches](#) are high performance with an open Ethernet infrastructure. While the fabrics are complex to meet AI performance demands, [intent-based automation](#) shields the network operator from that complexity.

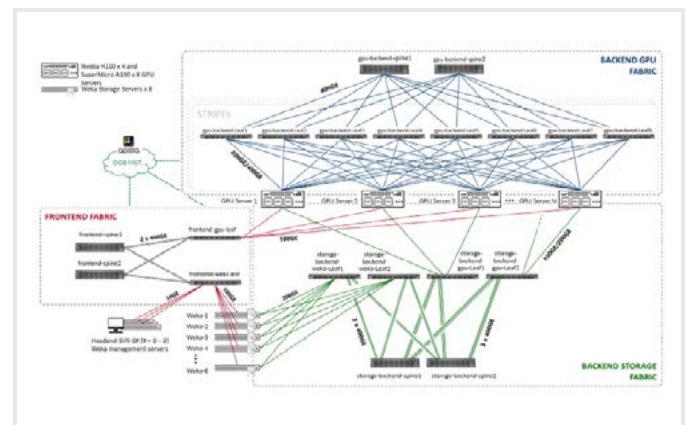



FIGURE 2
Solution architecture

To follow best practice recommendations, a minimum of four spines in each fabric is suggested. The use of traffic load balancing and Class of Service (CoS) should be employed to ensure a lossless fabric in the backend GPU fabric, and possibly in the backend storage fabric, as required per vendor recommendation. The deployment of a rail-optimized fabric will maintain a 1:1 relation with bandwidth subscription and leaf to GPU symmetry.

How it works

Capabilities and benefits

The **JVD for an AI Data Center** is a prescribed best practice approach to deploy high performing AI training and inference networks that minimize job completion time and simplify management with limited IT resources. The architecture meets and exceeds the demands of AI training performance.

 **Integrated prescribed design**

In this design, Juniper has integrated NVIDIA GPU, WEKA storage, and Juniper switches along with Juniper Apstra software and **Terraform** for automation. The prescribed design helps customers easily build high-capacity, easy-to-operate network fabrics that deliver the fastest JCT and maximize GPU utilization

 **Uninterrupted AI/ML traffic flow**

To ensure uninterrupted AI/ML traffic flow, network devices are configured to handle congestion using Data Center Quantized Congestion Notification (DCQCN), which provides Priority Based Flow Control (PFC) and Explicit Congestion Notification (ECN). The design also includes Dynamic Load Balancing (DLB) on the leaf nodes

Core capabilities

No bottlenecks. Ever. The frontend, backend GPU, and backend storage fabrics ensure that the network is never a bottleneck in AI model training and GPU job completion


Rail-optimized design Minimizes communication latency

Intent-based automation Shields the network operator from data center fabric complexity

How we deliver



The **JVD** program develops well-characterized, multidimensional solutions that reduce complexity for networking teams



Physical infrastructure supported by a JVD proves solution viability, with results provided in test reports.

Our advantage

The Juniper advantage

Juniper's AI Data Center Network follows an industry-standard dedicated IP fabric design. Three distinct fabrics (a backend, frontend, and storage fabric) provide maximum efficiency while maintaining focus on AI model scale, expedited completion times, and rapid evolution with the advent of AI technologies.

Configurations are given along with optimal hardware. In all, the best platforms are detailed in terms of features, performance, and the roles that are specified in this JVD.

The AI JVD design enables operators to orchestrate a training cluster systematically without a need for in-depth prerequisite knowledge of the required products and technology.

Why Juniper

The NOW Way to Network

Juniper Networks believes that connectivity is not the same as experiencing a great connection. Juniper's AI-Native Networking Platform is built from the ground up to leverage AI to deliver exceptional, highly secure, and sustainable user experiences from the edge to the data center and cloud. Additional information can be found at Juniper Networks (www.juniper.net) or connect with Juniper on [X](#) (Twitter), [LinkedIn](#), and [Facebook](#).

More information

Next steps

To learn more about AI Data Center Network with Juniper Apstra, NVIDIA GPUs, and WEKA Storage—JVD, visit <https://www.juniper.net/documentation/us/en/software/jvd/jvd-ai-dc-apstra-nvidia-weka/index.html>

For technical data sheets, guides and documentation, visit <https://www.juniper.net/documentation/validated-designs/>

Take the next step

Connect with us

Learn how we can build what's next.

[Contact us](#) →

Explore solutions

Discover Juniper's solution practice.

[Explore more](#) →

Read case studies

See how we help unlock new growth.

[Case study](#) →

More insights

Get the latest news delivered weekly.

[Subscribe](#) →