# Juniper Artificial Intelligence Data Center: Comparison of InfiniBand and RDMA over Converged Ethernet

Peter Fetterolf, Ph.D.

# EXECUTIVE SUMMARY

Data centers are the engines behind artificial intelligence (AI). Their design reflects cutting-edge technology, integrating high-performance AI accelerators, such as Graphics Processing Units (GPUs), low-latency networking, and high-speed storage solutions. The need for such dedicated infrastructure has grown as AI and machine learning (ML) applications have evolved rapidly, introducing challenging new requirements into the data center. Although the core of these AI data center architectures is state-of-the-art GPU servers, the network is critical as AI models are essentially massive parallel processing problems. Data center networks play a vital role in optimizing the utilization of these expensive GPUs. GPU and storage nodes are interconnected using high-speed 400G and 800G networking. Additionally, a dedicated 100 Gbps management network provides cluster management and operation. This report examines the economics of AI data centers, with an emphasis on AI training.

A comparison between the two leading network technologies, InfiniBand and RDMA over Converged Ethernet (RoCE), is presented, outlining the technical and economic advantages of RoCE, particularly in terms of availability, familiarity among engineers, and the faster development trajectory of Ethernet technologies. The paper explores the benefits of automated network operations, which significantly reduce operational expense (OpEx) through intent-based automation, as exemplified by Juniper Apstra. Our analysis shows that deployments of Juniper Ethernet with RoCE and Apstra result in 55% total cost of ownership (TCO) savings, including 56% OpEx savings and 55% capital expense (CapEx) savings over the first three years versus InfiniBand networks.

The paper also contrasts the economics of public cloud GPU instances for AI data centers, such as the AWS p5.48xlarge, with private cloud solutions. The TCO model demonstrates a 46% TCO savings over three years for private AI data centers compared to public cloud offerings, primarily due to the high recurring costs associated with public cloud services. This report underscores the economic and performance advantages of private AI data centers over public cloud solutions for AI training, with a detailed focus on network architecture and the impactful reduction of TCO offered by technologies such as RoCE and automated data center network management systems.

## Overview of AI-Optimized Data Center Architectures

AI-optimized data centers represent the pinnacle of computational power, designed to meet the intensive demands of AI model training and operations. The architecture of these data centers is an integration of state-of-the-art hardware, including AI accelerator/GPU systems and advanced networking technologies that ensure seamless communication and data transfer across computing nodes. With the exponential growth of data volumes and the iterative nature of deep learning, these clusters are optimized for parallel processing and high-speed data accessibility, making them indispensable for cutting-edge AI research and development.

Key uses cases in AI-optimized data centers include:

- Large language model (LLM) training and inference
- High-performance compute (HPC) and handling of massive data sets
- Other use cases requiring substantial parallel operations, such as image recognition

Rigorous AI training, such as LLMs, has become a high-priority use case and is the primary driver of compute resources. AI training requires many parallel mathematical operations facilitated by GPUs interconnected with a high-performance network.

The key requirements for AI optimized data centers:

- Provide a large number of GPUs for parallel high-performance computing
- Provide high-performance, low-latency, lossless connectivity between GPUs to accelerate parallel mathematical operations
- Provide high-performance storage infrastructure to support the GPU clusters

The AI data center architecture is represented in Figure 1. GPUs are interconnected with 400G/800G networking and GPUs are connected to storage nodes with 400G networking.

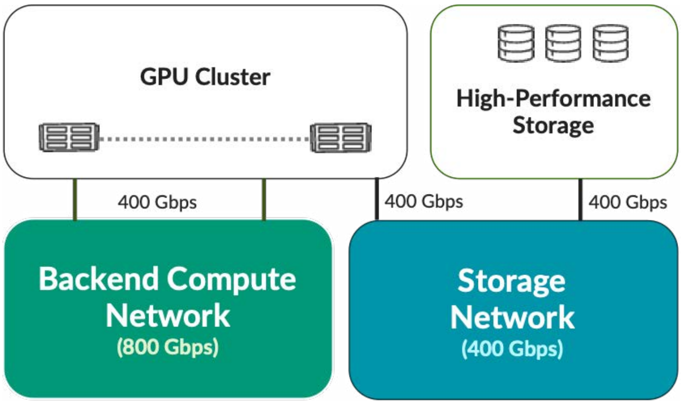


**Figure 1. AI-Optimized Data Center Architecture**

The key components of a state-of-the-art AI data center that we are modeling are:

### Nvidia DGX H100 GPU Servers

The GPU servers provide the compute power required for high-performance computing and AI model training. The configuration of servers is:

- 8x H100 GPUs
- Each GPU is interconnected inside the DGX H100 with a NVLink Switch high-speed interconnection network
- Each of the 8x GPUs is connected to the compute network with individual 400 Gbps network interfaces
- The network interface is either InfiniBand or RoCE

### GPU Compute Network

- InfiniBand: Nvidia 144 port 800 Gbps switches
- Ethernet RoCE: Juniper Networks 64 port 800 Gbps switches
- All interfaces to the DGX H100 GPUs are 400G
- Spine-leaf network architecture with no oversubscription to receive the full network throughput between any GPU in the network

### Storage Network

- High-performance 400 Gbps network interfaces to each server
- InfiniBand: Nvidia 64 port 400 Gbps switches
- Ethernet RoCE: Juniper Networks 64 port 400 Gbps switches
- Spine-leaf network architecture with an oversubscription of 4:1

Deep learning models leverage large datasets for training and high-performance storage is required for frequent and iterative access from the GPUs .

## High-Level Overview of Data Center Expenses

The economics of modern AI data centers is unique, with a disproportionate investment geared toward compute capacity predominantly due to the prohibitive cost of GPUs. In AI training data centers, the cost allocation shifts significantly toward compute as opposed to traditional data centers where storage and network components command a larger percentage of investment. Despite its smaller financial footprint, the network infrastructure plays an outsized role in the overall performance and efficiency of AI training. A suboptimal network can lead to GPU underutilization, spawning inefficiencies that swell operational costs and inflate the TCO. A breakdown of expenses in traditional data centers and newer AI Training data centers is presented in Table 1.

| Component | Traditional Data Center | AI Training Data Center |
|---|---|---|
| Compute[1] | 55% | 80% |
| Storage | 35% | 14% |
| Network[2] | 10% | 6% |

**Table 1. Data Center Expense Breakdowns, ACG Estimates**

## The Compute Network: A Comparison of Remote Direct Memory Access over Converged Ethernet (RoCE) with InfiniBand

The compute network demands exceptional performance and minimal latency. Remote Direct Memory Access (RDMA) is a technology that enables one computer to access the memory of another computer in a network directly without involving the operating systems or processors of either machine, allowing for more efficient data transfers between GPU memories. It significantly accelerates AI training and job completion times by enabling direct memory access between

[1] Compute category includes network interface cards.

[2] Network category includes cables and optics.

networked GPUs. There are two alternatives for the compute network: InfiniBand and RDMA over Converged Ethernet (RoCE). Although InfiniBand delivers the performance for AI training workload, there are advantages of RoCE:

- Ethernet networks are open, standards-based, GPU-agnostic, proven, and ubiquitous
- Leveraged expertise as most network engineers are familiar with Ethernet while far fewer engineers are familiar with InfiniBand
- Ethernet has traditionally developed at a faster pace than other networking technologies due to breadth of deployments and economies of scale
- More robust and efficient operations: the distributed architecture of Ethernet, where each switch autonomously makes local, swift decisions, enhances resiliency by reducing single points of failure and facilitating faster path selection compared to the centralized decision-making process of InfiniBand
- Over time Ethernet will accelerate at a faster rate in both price declines and performance improvements

Almost all frontend AI data centers are based on Ethernet. If a customer deploys InfiniBand in a backend AI training data center this forces a separate operational and automation paradigm and generates additional complexity.

## Benefits of Automated Network Operations

In addition to the general advantages of Ethernet, automation significantly benefits network deployment and operations, reducing OpEx. Juniper Apstra is an example of intent-based network automation software. Apstra automates and continuously validates data center network design, deployment, and operation from Day 0 through Day 2+. With multivendor support, Apstra empowers organizations to automate and manage their networks across virtually any data center design, vendor, and topology.

Other ACG financial models confirm that Apstra results in significant reductions in labor time for the following categories:

- Day 0: High-Level Network Design
- Day 0: Detailed Network Design
- Day 0: Scope Requirements

- Day 1: Network Implementation
- Day 1: Testing
- Day 1: Operations and Documentation
- Day 2: Network Operations
- Day 2: Troubleshooting

Unlike InfiniBand, Apstra provides uniform management across training, inference, and storage clusters without requiring a specialized network management platform. It brings consistency in managing all fabrics in the AI cluster with features such as congestion monitoring and management that are unique to AI clusters. Additionally, InfiniBand demands specialized expertise, which organizations must hire to optimize its utilization.

## Economic Benefits of Ethernet versus InfiniBand

ACG developed a TCO model comparing AI data center compute and storage networks. The model compared two scenarios:

- InfiniBand for both the compute and storage networking
- Juniper Ethernet with RoCE v2 for both the compute and storage networking

The model includes the following CapEx components:

- Nvidia InfiniBand 144 port 800G switches, cabling, and optics used for both leaf and spine switches
- Juniper QFX5240 64 port 800G switches, cabling, optics, and Apstra used for both leaf and spine
- For larger scale designs requiring 1000s of GPUs with the highest radix nonblocking Clos, some customers may prefer the Juniper PTX 10008 with its cell-based fabric as the spine and QFX5240 as a leaf

The model includes the following OpEx components:

- Vendor Support and Maintenance
- Day 0: High-Level Network Design
- Day 0: Detailed Network Design
- Day 0: Scope Requirements

- Day 1: Network Implementation
- Day 1: Testing
- Day 1: Operations and Documentation
- Day 2: Network Operations
- Day 2: Troubleshooting
- Day 2: Outage Costs and Remediation

We compared the following system and network configurations depicted in Table 2:

| Item | Value |
|---|---|
| Number of H100 GPUs | 1024 |
| Number of DGX H100 | 128 |
| Number of InfiniBand 800G Compute Leafs | 8 |
| Number of InfiniBand 800G Compute Spines | 4 |
| Number of Ethernet 800G Compute Leafs | 16 |
| Number of Ethernet 800G Compute Spines | 8 |
| Number of Storage Leafs | 8 |
| Number of Storage Spines | 2 |

**Table 2. System and Network Configurations Used in the Three-Year TCO Models**

## Results of Three-Year Network TCO Analysis

The three-year cumulative savings for Juniper RoCE v2 with Apstra are:

- Overall TCO Savings: 55%
- OpEx Savings: 56%
- CapEx Savings: 55%

The details of the three-year cumulative TCO results are provided in Table 3:

| Scenario | CapEx | OpEx | TCO | NPV | TCO Savings |
|---|---|---|---|---|---|
| InfiniBand | $11.67M | $8.12M | $19.75M | ($19.37M) | |
| Juniper Ethernet | $5.29M | $3.55M | $8.84M | ($8.66M) | $10.91M<br>55% |

**Table 3. Three-Year TCO Comparison of InfiniBand and Juniper RoCE v2**

The CapEx advantage is due to the cost difference between InfiniBand and Ethernet switches. InfiniBand switches are approximately double the price of Ethernet switches. InfiniBand hardware is sole-sourced, while Ethernet hardware is produced by a broad, competitive, industry. The OpEx benefits are due to the overall ease of Ethernet operations versus InfiniBand and the reductions in labor due to Juniper Apstra network automation. A breakdown of three-year cumulative OpEx is presented in Figure 2. Vendor Support is the annual vendor support and maintenance contract provided by Nvidia for InfiniBand and Juniper for Ethernet. The other key areas of savings are outage and remediation costs, troubleshooting, and network operations. These savings are due to the benefits of Apstra automation.

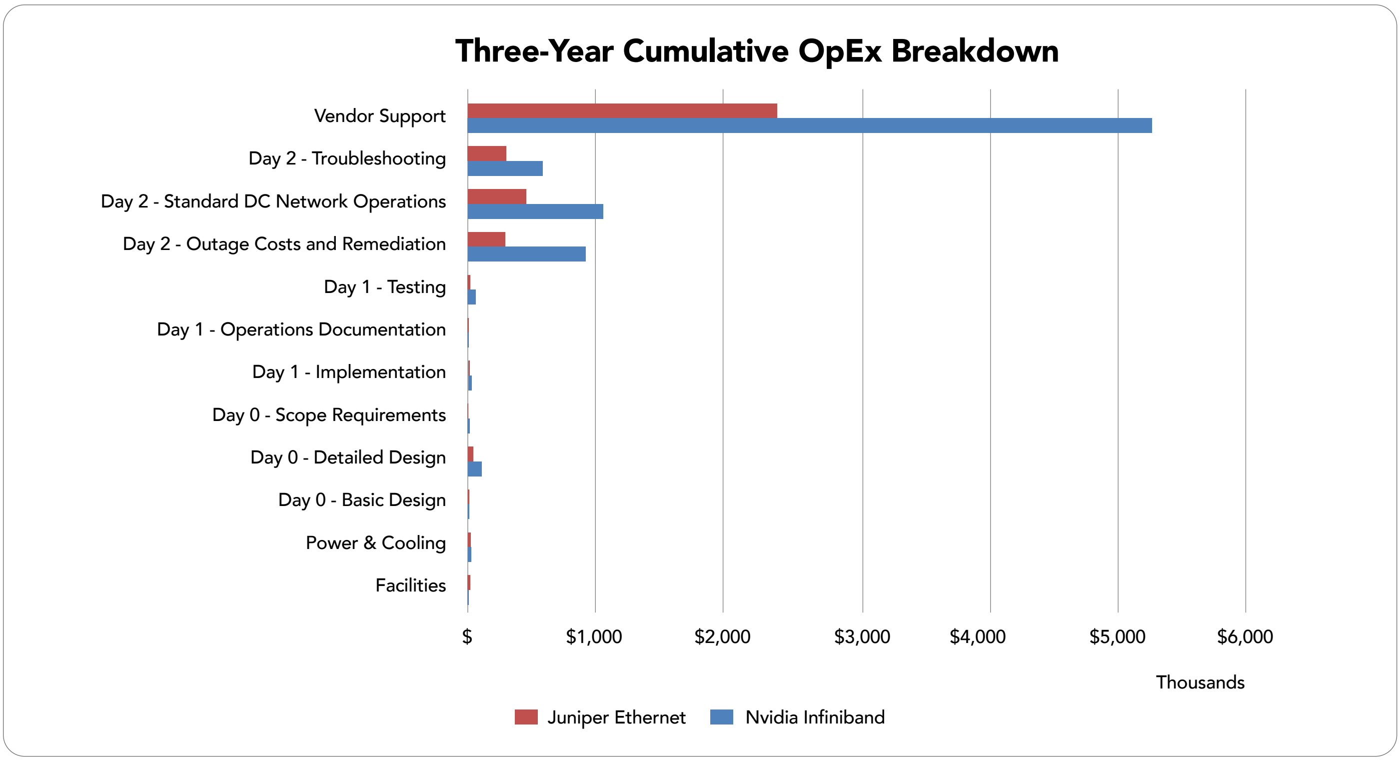


**Figure 2. Breakdown of the Cumulative OpEx Expenses over Three Years**

## Overview of Public Cloud GPU Instances for AI Data Centers

In order to compare the cost of public cloud GPU services with private AI-optimized data centers, we use AWS p5.48xlarge instances, a robust offering from Amazon Web Services, tailored for compute-intensive workloads that offer powerful GPU capabilities. Each p5.48xlarge instance has 8 NVIDIA H100 Tensor Core GPUs, delivering exceptional acceleration for machine learning training and high-performance computing applications. The interconnectivity between these GPUs is facilitated by NVIDIA NVLink Switch, which allows for high-bandwidth, low-latency communication critical for scaling up AI and HPC tasks.

From a networking perspective, these GPU instances offer up to 400 Gbps of network bandwidth for each GPU, ensuring that the data-intensive operations that GPUs perform are not bottlenecked by network throughput.

## Economic Comparison of Public Cloud with Private Cloud AI-Optimized Data Centers

ACG created a TCO model comparing a public cloud GPU service with a private cloud AI-optimized data center. The following scenarios were compared:

1. AWS p5.48xlarge EC2 instances
2. Private data center with DGX H100 GPUs and 800 Gbps Juniper Ethernet networks

The following assumptions were used in the cost comparison:

**AWS GPU Instances**

- p5.48xlarge EC2 instances cost per hour of $43.157
- FSX Cluster high-performance storage cost of $0.145 per GB
- EC2 egress expenses cost of $0.05 per GB of Internet egress

**Private AI Data Center**

The private AI data center uses the standard DGX H100 system and network architecture consisting of DGX H100 nodes (described previously), an 800G compute network, a 400G storage network, and a 100G management network. The configuration of this system is described in Table 2.

The CapEx considered in this analysis is:

- 1024 x DGX H100
- Juniper RoCE v2 800G Compute Network
- Juniper 400G Storage Network
- High-Performance Storage
- Rack and Stack
- Cabinets
- Professional Services

The OpEx modeled in the private cloud AI data center is:

- Power and Cooling
- Facilities Expenses
- Management and Operations of Servers

## Public Cloud versus Private AI-Optimized Data Center TCO Results

The economic comparison contrasting AWS Public Cloud GPU services with private data center configurations spanning three years reveals a 46% cost advantage favoring the private AI-optimized data center. This advantage stems from the avoidance of escalating cloud service fees that can accumulate to be a significant portion of the costs in intensive AI training operations. The results of the three-year cumulative TCO comparison are presented in Table 4. Even after three years the TCO of the private AI GPU cloud solution is 46% lower than the AWS public cloud GPUs. Over time the savings of the private data center solution continues to increase because the annual cloud expenses are high and the initial CapEx expenses in the private data center is a one-time charge.

| 3 Years TCO | CapEx | OpEx | TCO | NPV | TCO Savings |
|---|---|---|---|---|---|
| AWS p5.48xlarge | 0 | $148.93M | $148.93M | ($130M) | |
| Private AI Data Center | $67.05M | $12.71M | $79.76M | ($78.25M) | 46% |

**Table 4. Three-Year Cumulative TCO Comparison of Private AI-Optimized Data Center with AWS public Cloud GPUs**

Although the Juniper & DGX H100 3 YR model requires a higher initial investment, it offers lower operational costs and a more favorable NPV and TCO over three years. Therefore, it would be a more financially prudent choice for stakeholders considering long-term cost efficiency.

A breakdown of the OpEx and CapEx for the public cloud and private cloud is provided in Figure 3 and Table 5, respectively. The key driver of expenses in the public cloud is the monthly charge for the p5.48xlarge AWS instance; the key expense in the private cloud is the DGX H100 systems.

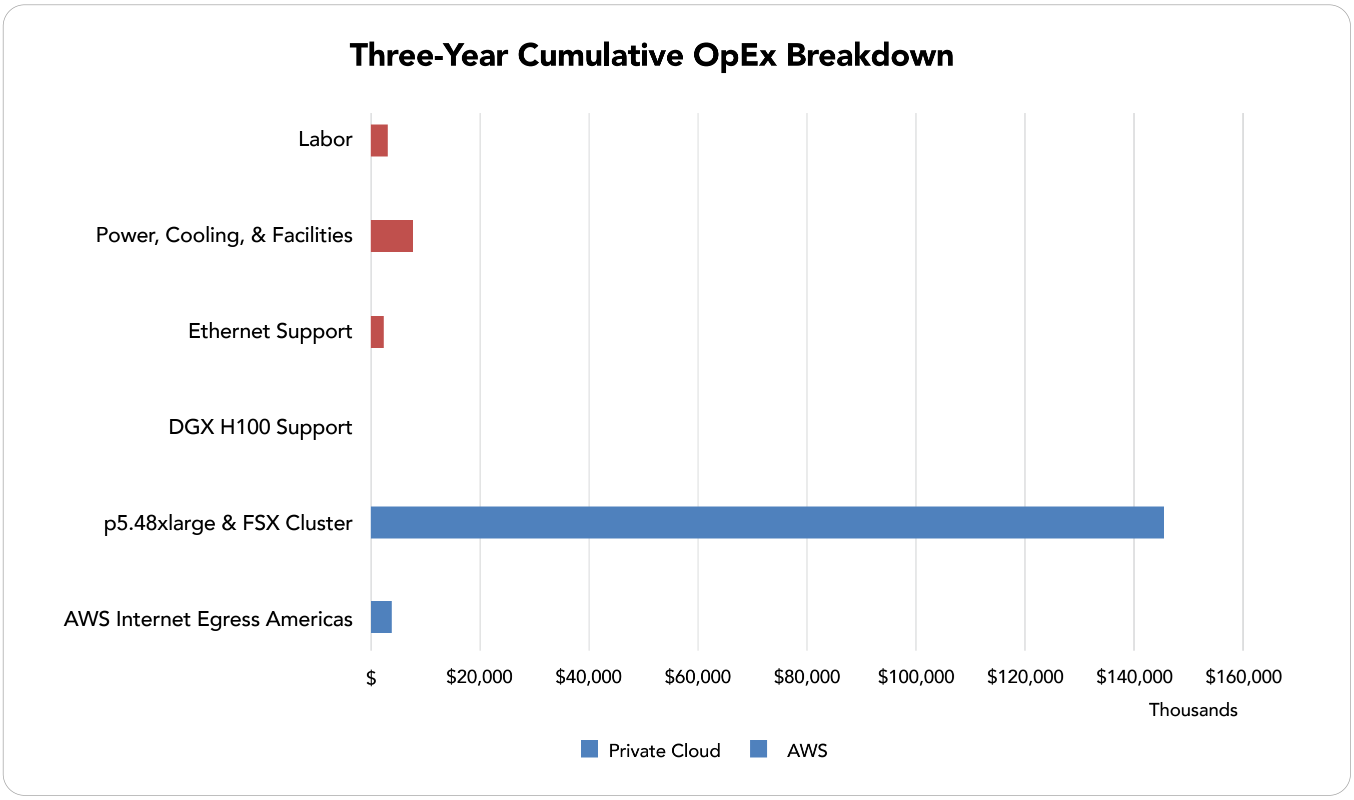


**Figure 3. Three-Year Cumulative Breakdown of Private versus Public Cloud GPU OpEx Savings**

| CapEx | Private Cloud |
|---|---|
| DGX H100 & Cabinet | $58,880,000 |
| Network | $5,290,204 |
| Storage | $2,880,000 |

**Table 5. Three-Year Cumulative CapEx Breakdown for Private Cloud AI Data Center**

# Conclusion

The extensive analysis elucidates the crucial role of AI-optimized data centers in the realm of AI training, particularly for tasks as demanding as training Large Language Models. We have dissected the architectural nuances of such data centers, detailing their reliance on high-performance GPUs, low-latency networking, and the necessity of advanced storage solutions to handle the growing data requirements.

Our economic comparison reveals a paradigm shift in cost structures, with compute expenses constituting the lion's share of investment in AI-optimized data centers. This shift underscores the importance of network optimization, which if overlooked can lead to significant inefficiencies and elevated TCO.

Through comparative analysis, Ethernet emerges as a more economically favorable solution over traditional InfiniBand, not only due to its technical competencies, but also due to its broader adoption and scalability. The advent of intent-based network automation solutions, such as Juniper Apstra, offers additional cost benefits, streamlining operations and significantly reducing OpEx.
The whitepaper also ventures beyond the confines of private data centers to evaluate the economic viability of public cloud services. Despite their initial appeal, the long-term cost analysis favors private AI data centers, which demonstrate substantial cost savings over public cloud alternatives, particularly for organizations with sustained, intensive compute needs.

In light of these findings, stakeholders are armed to make informed decisions regarding the implementation of AI data centers. Organizations can achieve a delicate balance between high-performance and cost-effectiveness with the right architectural and network choices. As AI continues to advance, the investment in private AI infrastructure appears to be not only justified, but also economically prudent, ensuring that AI initiatives are supported by a backbone that is as robust as it is financially viable.

## Peter Fetterolf

Peter Fetterolf, Ph. D. is an expert in network technology, architecture and economic analysis. He is responsible for financial modeling and whitepapers as well as software development of the ACG Research Business Analytics Engine. Dr. Fetterolf has a multidisciplinary background in the networking industry with over thirty years of experience as a management consultant, entrepreneur, executive manager, and academic. He is experienced in economic modeling, business case analysis, engineering management, product definition, market validation, network design, and enterprise, and service provider network strategy.