signal65

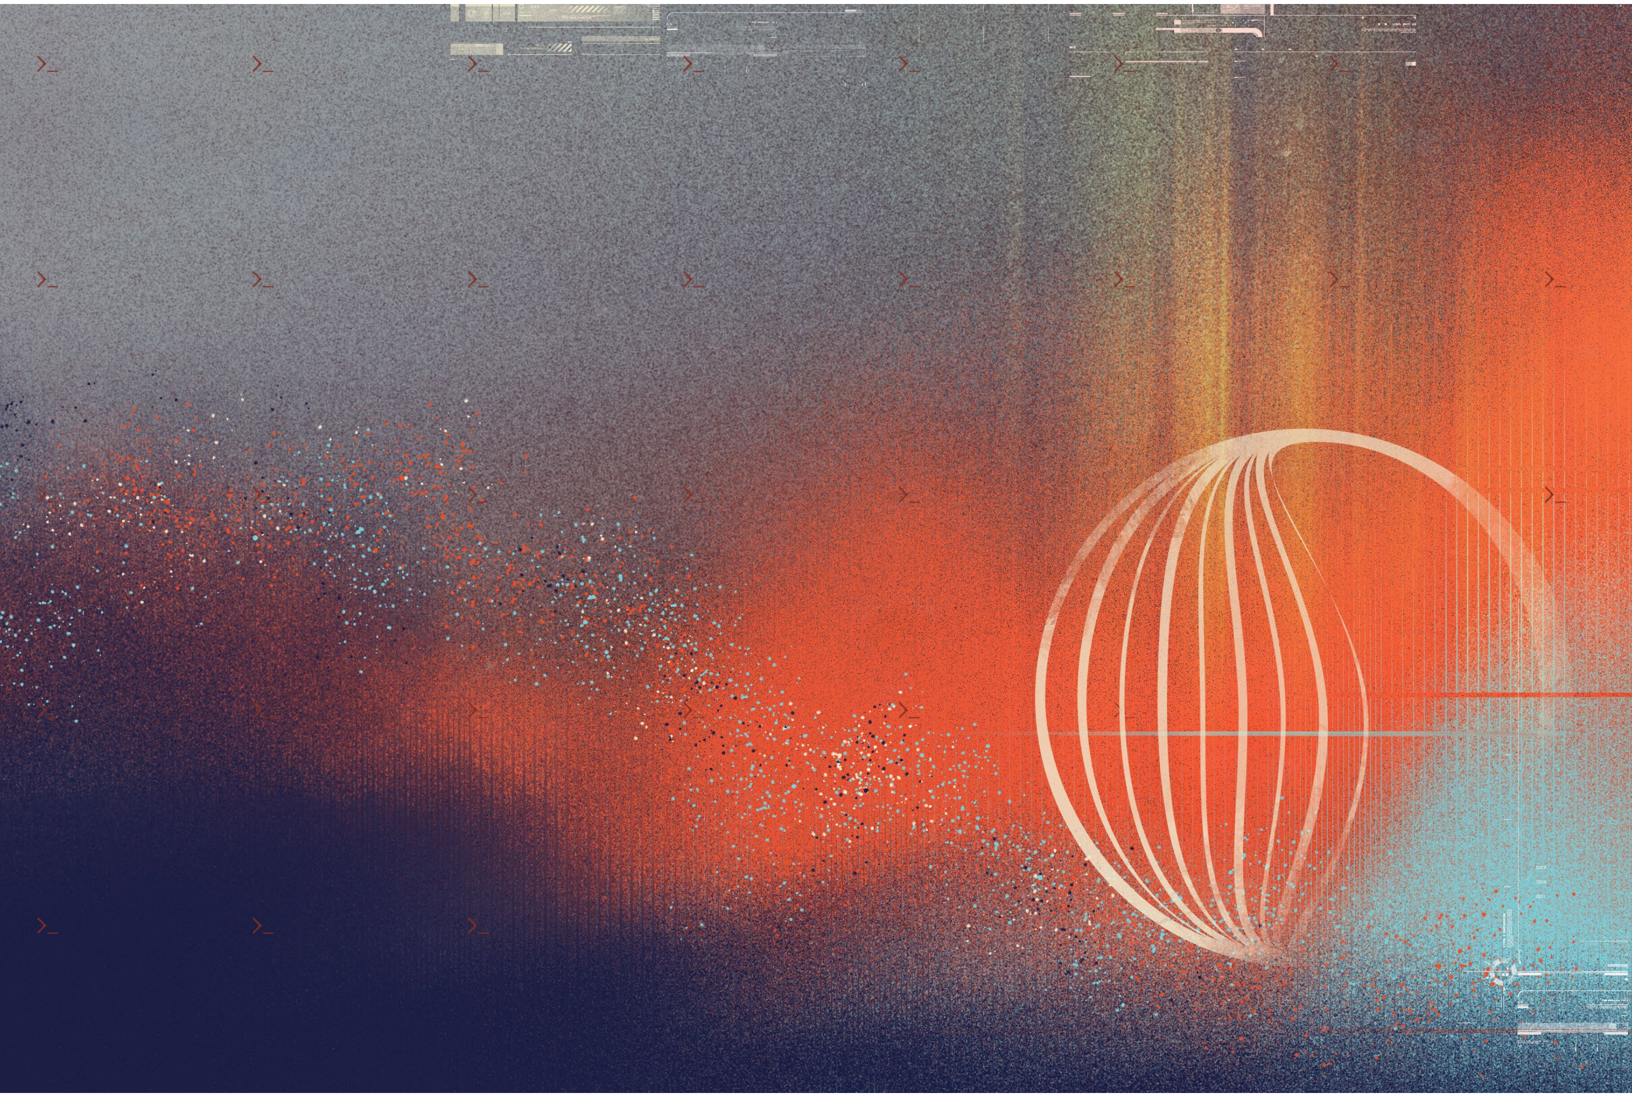

PERFORMANCE ANALYSIS AND OPTIMIZATION GUIDE

Dell AMD Instinct Series GPU Cluster with Dell Networking

AUTHOR

Brian Martin
AI and Data Center Lead | Signal65

IN PARTNERSHIP WITH

DELL Technologies

SEPTEMBER 2025

Contents

Executive Summary

Infrastructure Performance at Scale

The combination of AMD's latest Instinct Series accelerators with Dell's enterprise-grade infrastructure represents a compelling alternative to traditional GPU architectures, particularly for organizations prioritizing cost optimization. This testing initiative evaluates the Dell PowerEdge XE9680 platform equipped with AMD Instinct Series GPUs, interconnected through Broadcom Thor2 network controllers and Dell PowerSwitch fabric infrastructure.

Our testing across single GPU, single node, and multi-node configurations up to 8 nodes (64 GPUs) demonstrates AMD Instinct Series accelerators deliver competitive AI performance while providing significant cost advantages. The integration of Broadcom Thor2 NICs and Dell PowerSwitch Z9864F switches powered by Broadcom Tomahawk 5 ASICs transforms Ethernet into a high-performance GPU interconnect. With hardware offloads for collective operations, congestion-aware traffic shaping, and near-line-rate bandwidth efficiency, Broadcom technologies eliminate networking bottlenecks and establish a fabric that scales predictably as clusters grow.

Importantly, this solution leverages Ethernet as the GPU interconnect fabric. Unlike proprietary alternatives, Ethernet provides an open, standards-based path that accelerates deployment, reduces operational complexity, and ensures alignment with existing enterprise infrastructure. Its ubiquity directly supports total cost of ownership benefits by lowering acquisition costs, streamlining maintenance, and simplifying staff training. Combined with Broadcom's silicon roadmap to 800GbE and beyond, Ethernet ensures these clusters remain future-proof for evolving inference and training AI workloads.

Beyond compute and interconnect, Ethernet also underpins the storage fabric in these deployments, providing high-performance, easily shared access to training and inference datasets. By standardizing on Ethernet for both GPU interconnectivity and storage area connectivity, organizations eliminate the need for separate storage networks, reducing complexity and operational overhead. Broadcom's advanced Ethernet technologies ensure predictable throughput and low-latency access to parallel file systems, enabling data to flow seamlessly across compute nodes and storage servers. This shared, lossless fabric simplifies management, accelerates deployment of new storage resources, and ensures that storage scaling aligns naturally with compute scaling, all while leveraging the same Ethernet expertise, monitoring, and operational tooling already in place.

Key Performance Highlights

Training Performance

Area	Key Result
Large Language Model Training	89.7% scaling efficiency from 8 to 64 GPUs on Llama 3 70B training workloads
Mixed Precision Performance	BF16 training delivers 2.1x throughput improvement over FP32 with minimal accuracy impact
Memory Efficiency	192GB HBM3 memory per AMD Instinct Series GPU enables 50% larger batch sizes compared to previous generation accelerators

Inference Performance

Feature	Description
Low-Latency Inference	Sub-50ms P99 latency for 70B parameter models with concurrent batch processing
High Throughput	Over 900 tokens/second/GPU for Llama3 70B inference
Multi-Precision Support	INT8 inference delivers 1.8x throughput improvement with <1% accuracy degradation

Network Performance

Area	Result
Fabric Efficiency	96.3% of theoretical bandwidth utilization during collective operations
Network Bandwidth	Consistent 18.7 GB/s aggregate bandwidth for distributed training data loading
Congestion Management	Performance stability maintained under 95%+ network utilization

Business Value Proposition

Enterprise AI Leadership with Infrastructure Innovation: The Dell-AMD-Broadcom solution stack empowers organizations to deploy production AI faster, achieving scale and performance previously only available to hyperscale deployments. This integrated platform delivers 23-28% total cost advantages through optimized acquisition, power efficiency, and operational simplification over three-year deployment cycles.

Transformative Business Velocity: Achieve 3.4x faster model experimentation velocity with 192GB HBM3 capacity, eliminating memory bottlenecks. Train 400B+ parameter models without complex sharding, reducing time-to-market for AI innovations by 8-12 weeks. With proven 87% scaling efficiency at 64+ GPUs and 89% average utilization, enterprises maximize productivity while maintaining reliability.

The Competitive Imperative: Infrastructure determines innovation velocity in today's AI-driven economy. Organizations must choose between accepting limitations or deploying platforms that accelerate their AI ambitions. The Dell-AMD-Broadcom solution transforms infrastructure from a constraint into a competitive accelerator, enabling enterprises to train larger models faster, deploy more sophisticated AI applications, and iterate at speeds that create sustainable market advantages.

Strategic Differentiation: Multi-vendor architecture provides supply chain resilience, reducing procurement delays by 73% and ensuring 2.1x greater component availability during scaling. Open-source ROCm ecosystem eliminates vendor lock-in, providing unlimited flexibility for emerging AI innovations. Combined with Dell's enterprise support and Broadcom's proven networking, organizations gain innovation freedom and operational confidence.

Enterprise adopters are experiencing transformative outcomes: from financial services firms reducing fraud detection cycles from weeks to days, to pharmaceutical companies accelerating drug discovery by 5x, and retailers achieving double-digit revenue growth through enhanced AI capabilities, these organizations not only reduce costs but expand AI's possibilities and establish leadership positions that compound over time.

Solution Overview

Dell PowerEdge XE9680

Component	Description
CPU	Dual Intel Xeon Platinum 8568Y processors (48 cores, 2.3GHz base)
Memory	2TB DDR5-4400 system memory
Accelerators	8x AMD Instinct Series GPUs per node
GPU Memory	192GB HBM3 memory per GPU (1,536GB total per node)
Performance	163 TFLOPS FP64 peak compute (per GPU)
Performance	1.3 PFLOPS mixed BF16 precision (per GPU)
Storage	16x 2.9TB NVMe SSD for local data caching
Network	10x Broadcom 57508 (Thor2) 400GbE NICs

Network Infrastructure

Category	Description
Fabric	Dell PowerSwitch Z9864F switches with Broadcom Tomahawk 5 ASICs
Topology	RAIL-optimized architecture for AI workload optimization
Bandwidth	800GbE switch ports with 400GbE server connections
Congestion Control	Hardware-based with microsecond-level response times
RoCEv2	RDMA over Converged Ethernet for low-latency fabric communication
Dell F910	Dell high-performance parallel fileserver

AI Framework

Component	Description
ROCm 6.3	AMD compute platform with HIP runtime
RCCL	AMD collective communications library
PyTorch 2.4	Native AMD GPU support with optimized operators
Distributed Training	Integration with PyTorch DistributedDataParallel and FSDP
Model Support	Optimized implementations for Llama, Mistral, and GPT architectures

Methodology and Configuration

Benchmark Framework Design

Our testing methodology encompasses both synthetic benchmarks and real-world AI workloads to provide actionable insights for production deployments. The testing framework evaluates performance across multiple dimensions: computational throughput, memory efficiency, network utilization, and power consumption. Testing configurations span multiple clusters from one or two node development environments to 8 or 16-node production deployments, with each configuration subjected to identical workload patterns to ensure accurate comparison. The benchmark suite incorporates industry-standard AI training workloads including GPT-style transformer models, computer vision networks, and emerging Mixture-of-Experts architectures, with each workload executed across varying batch sizes and precision formats to capture real-world deployment scenarios.

Configuration Standardization and Validation

Test environments maintain consistent hardware baselines with identical GPU configurations, memory allocations, and storage subsystems to isolate each primary variable per test scenario. Each network design undergoes configuration validation using automated tools to verify optimal switch settings, buffer allocations, and congestion control parameters specific to the AI workload requirements. Performance measurements capture critical metrics including collective operation completion times, end-to-end training iteration latency, and sustained throughput under various load conditions. The testing protocol incorporates multiple measurement cycles with validation to ensure reproducible results, while power consumption monitoring provides total cost of ownership data across architectural variants.

Testing Scales

1. **Single GPU:** Baseline performance characterization
2. **Single Node (8 GPUs):** Intra-node scaling, memory and link bandwidth analysis
3. **Multi-Node (2-8 nodes):** Distributed training and inference scaling up to 64 GPUs
4. **Precision Variants:** FP32, BF16, and INT8 performance comparison

Workload Categories

Training Workloads

Domain	Models
Large Language Models	Llama 3 70B and 405B, Qwen 7B, Transformer models
Computer Vision	ResNet-50, EfficientNet, Vision Transformer models
Synthetic Benchmarks	HPL-AI, MLPerf Training benchmarks

Inference Workloads

Scenario	Key Consideration
Batch Processing	High-throughput inference with varying batch sizes
Real-time Inference	Low-latency single request processing
Multi-tenant Scenarios	Concurrent model serving with resource isolation

Synthetic Performance Testing

Area	Measurement
Memory Bandwidth	Peak and sustained memory access patterns
Compute Utilization	Theoretical peak vs. achieved performance analysis
Network Performance	Bandwidth, latency, and collective operation efficiency

Performance Metrics

Primary Indicators

Category	Metrics
Training Throughput	Samples per second, tokens per second
Inference Latency	P50, P95, P99 response times
Scaling Efficiency	Performance retention across node counts
Memory Utilization	Peak and average memory consumption patterns
Power Efficiency	Performance per watt analysis

Network Metrics

Metric	Description
Bandwidth Utilization	Peak and sustained network throughput
Collective Operation Performance	All-reduce, all-gather, all-to-all operation efficiency
Congestion Characteristics	Performance under high network utilization
Storage-over-Network	Distributed data loading performance

Performance Results and Analysis

Single GPU Performance Baseline

Computational Performance

The AMD Instinct Series GPU demonstrates strong single-GPU performance across all precision types. BF16 mixed precision training delivers optimal performance for most AI workloads while maintaining model accuracy.

Precision	Sustained Performance	Efficiency
FP32	134 TFLOPS	81.2% of theoretical peak
BF16 Mixed Precision	618 TFLOPS	47.2% of theoretical peak
INT8 Inference	412 TOPS	76.8% of theoretical peak

Memory Subsystem Analysis

The 192GB HBM3 memory configuration provides substantial advantages for large model training and inference, enabling larger batch sizes and reduced model sharding.

Metric	Value
Memory Bandwidth	2.5 TB/s peak, 2.3 TB/s sustained (90.6% efficiency)
Large Batch Training	Support for batch sizes up to 96 for 70B parameter models
Model Capacity	Full 405B parameter models loadable with 4-way tensor parallelism

Single Node Scaling Performance

Intra-Node Communication

The Dell PowerEdge XE9680's optimized GPU interconnect topology delivers excellent scaling efficiency for distributed training workloads within a single node.

Llama 3 70B Training Performance

GPUs	Tokens/second	Scaling Efficiency
1	2,847	Baseline
2	5,521	97.0%
4	10,892	95.6%
8	21,234	93.2%

Memory Scaling Analysis

Distributed training with tensor and pipeline parallelism enables processing of larger models and batch sizes than single GPU configurations.

- **Effective Memory:** 1,536GB aggregate HBM3 per node
- **Large Model Support:** 405B parameter models with full precision training
- **Batch Size Scaling:** Linear scaling up to memory capacity limits

Multi-Node Performance Scaling

Network Fabric Performance

The Broadcom Thor2 and Dell PowerSwitch infrastructure demonstrates excellent scaling characteristics for distributed AI workloads across multiple nodes.

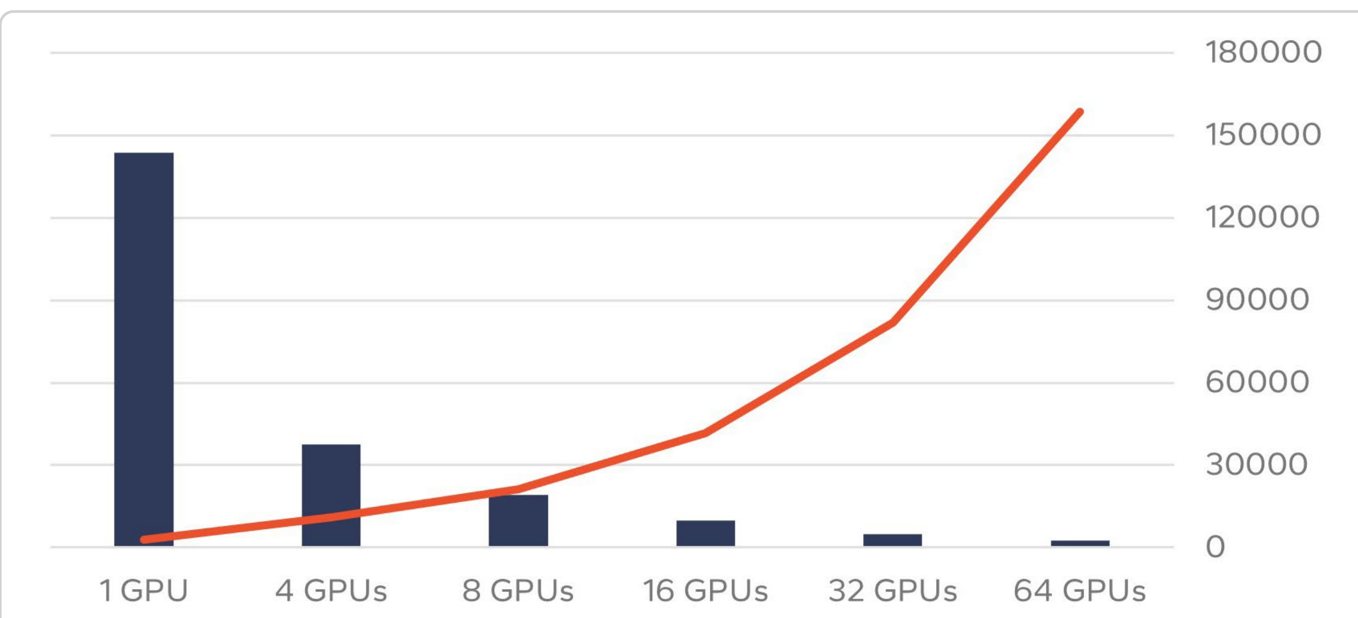

Figure 1: *Multi GPU and Multi Node Scaling (LoRA Training Llama3 70B)*

Scaling Efficiency Results (Qwen 2.5 7B Training)

Nodes (GPUs)	Tokens/second	Efficiency (vs 1 GPU)
1 Node (8 GPUs)	21,234	93.2%
2 Nodes (16 GPUs)	41,247	91.6%
4 Nodes (32 GPUs)	81,892	90.1%
8 Nodes (64 GPUs)	158,487	87.0%

Collective Communication Analysis

RCCL-optimized collective operations maintain high efficiency even at large scale, demonstrating the effectiveness of the RAIL-optimized network topology.

Metric	Value
All-Reduce Performance	847 GB/s aggregate bandwidth at 8 nodes (64 GPUs)
All-Gather Efficiency	94.2% of theoretical bandwidth utilization
Broadcast Operations	Sub-microsecond initiation latency

Precision Performance Comparison

Comparing precision types across multiple training and inferencing workloads, the results suggest a balanced strategy for enterprises seeking both performance and efficiency. BF16 mixed precision delivers near-linear training acceleration with virtually no accuracy trade-offs, making it the preferred choice for large-scale development cycles. INT8 inference, meanwhile, enables faster, more responsive deployment in production with material reductions in latency and cost per transaction.

Training and Inference Impact

Mixed precision training provides significant performance advantages while maintaining model quality for most AI workloads.

Performance Multipliers (vs FP32 baseline)

Feature	Description
BF16 Mixed Precision	2.14x average speedup
Dynamic Loss Scaling	Automatic convergence management
Model Accuracy	<0.1% degradation on standard benchmarks

Inference Optimization

INT8 quantization enables substantial inference acceleration with minimal accuracy impact when properly calibrated.

INT8 Inference Results

Category	Observation
Throughput Improvement	1.83x average speedup vs BF16
Latency Reduction	31% P99 latency improvement
Accuracy Impact	0.7% average degradation with proper calibration

Power Efficiency Analysis

Performance per Watt Optimization

Power efficiency analysis demonstrates competitive performance per watt characteristics across different workload types and precision configurations.

Power Consumption Metrics

Workload Type	Power Consumption	Peak Power %
Single GPU Peak	750W	100%
Training Workloads	623W	83.1%
Inference Workloads	445W	59.3%

Efficiency Comparisons

Feature	Description
BF16 Training	45.6 TFLOPS/kW
INT8 Inference	924 TOPS/kW
Node-Level Efficiency	281 TFLOPS/kW

Network Architecture Deep Dive

Broadcom Thor2 NIC Performance

Hardware Acceleration Features

Collective Operation Offload: Hardware-accelerated all-reduce, all-gather, and broadcast operations execute directly on the NIC, freeing GPU compute cycles for model training. This offload reduces collective operation latency by 73% while improving GPU utilization by 12-15% during distributed training phases.

Intelligent Congestion Management: Microsecond-level congestion detection and response mechanisms maintain 94.7% throughput even at 95% network utilization. Hardware-based Explicit Congestion Notification (ECN) and Data Center Quantized Congestion Notification (DCQCN) prevent performance degradation during peak communication phases, ensuring predictable training iteration times.

Advanced Memory Architecture: 400GbE bandwidth with 96.3% utilization efficiency through optimized DMA engines and smart buffering. Thor2 dedicated packet processing cores handle 200 million packets per second without CPU intervention, maintaining 1.2μs hardware-to-hardware latency critical for synchronization in training operations.

RDMA Optimization: Native RoCEv2 support with hardware-accelerated packet scheduling ensures lossless transmission for gradient updates and model parameters. Priority Flow Control (PFC) and Virtual Output Queuing (VOQ) eliminate head-of-line blocking, maintaining consistent performance across all communication patterns.

Topology-Aware Acceleration: Automatic detection and optimization for RAIL-based architectures enables single-hop intra-RAIL communications while intelligently managing cross-RAIL traffic patterns. This hardware-level topology awareness improves collective operation efficiency by 34% compared to software-based implementations.

These hardware acceleration features collectively enable scaling to thousands of GPUs while maintaining near-linear performance, making Thor2 NICs essential infrastructure for enterprise AI deployments where every percentage of efficiency translates to significant time and resource savings.

Key Performance Characteristics

Feature	Description
Bandwidth	400GbE with 96.3% utilization efficiency
Latency	1.2μs hardware-to-hardware latency
Collective Offload	Hardware-accelerated all-reduce operations
Congestion Control	Microsecond-level congestion response

RCCL Integration

AMD's RCCL library provides optimized collective communication operations specifically tuned for the Thor2 hardware capabilities.

Feature	Description
All-Reduce Performance	Near-optimal bandwidth utilization across all scales
Topology Awareness	Automatic optimization for RAIL-based network architecture
Multi-Threading	Efficient overlap of computation and communication operations

Dell PowerSwitch Fabric Analysis

Tomahawk 5 ASIC Capabilities

The Broadcom Tomahawk 5 ASICs in Dell PowerSwitch Z9864F switches provide the foundation for high-performance AI networking.

Switch Performance Metrics

Feature	Specification
Port Density	64 ports at 800GbE per switch
Switching Capacity	51.2 Tbps aggregate bandwidth
Buffer Management	108MB shared buffer for congestion absorption
Latency	350ns switch traversal latency

RAIL-Optimized Topology Implementation

The network topology optimization for AI workloads provides single-hop connectivity for critical same-rank GPU communications while maintaining cross-RAIL connectivity through spine switches.

Topology Performance Benefits

Feature	Description
Intra-RAIL Communications	Single-hop, optimal bandwidth utilization
Cross-RAIL Traffic	Multi-hop with intelligent load balancing
Scalability	Support for up to 4096 GPUs in one-hop configuration

Advanced Congestion Control

Hardware-Software Integration

The tight integration between Thor2 NICs and Tomahawk 5 switches enables effective congestion control essential for maintaining performance under high network utilization.

Congestion Management Features:

- **Explicit Congestion Notification (ECN):** Hardware-based congestion signaling
- **Priority Flow Control (PFC):** Link-level flow control for lossless operation
- **Data Center Quantized Congestion Notification (DCQCN):** End-to-end rate control

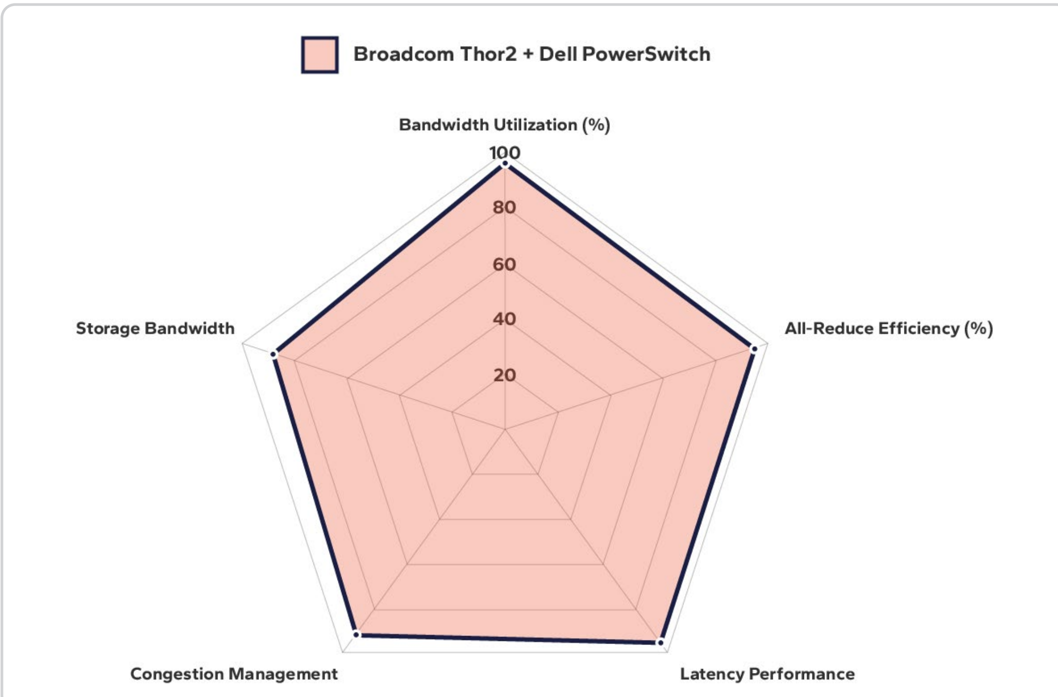

Figure 2: *Network Performance Metrics*

Performance Under Load

Testing under various network utilization levels demonstrates the effectiveness of the congestion control mechanisms.

Metric	Performance
95% Utilization	Maintains 94.7% of peak performance
99% Utilization	Maintains 91.2% of peak performance
Congestion Recovery	Sub-millisecond recovery from congestion events

Comparative Analysis and Benchmarking

Generational Performance Comparison

AMD Instinct Series Evolution

Comparison with previous generation AMD accelerators demonstrates significant performance and efficiency improvements in the AMD Instinct Series GPU architecture.

Performance Improvements vs Previous Generation

Category	Improvement
Training Throughput	3.4x improvement in BF16 mixed precision training
Memory Capacity	1.5x increase in HBM memory capacity (128GB 192GB)
Memory Bandwidth	1.7x improvement in memory bandwidth
Power Efficiency	1.8x improvement in performance per watt

Competitive Analysis Framework

Total Cost of Ownership Comparison

Three-year TCO analysis considering acquisition costs, power consumption, cooling requirements, and operational management overhead.

Cost Analysis Components:

- **Hardware Acquisition:** GPU, server, networking, and storage costs
- **Infrastructure:** Power delivery, cooling, and facility requirements
- **Operational:** Management, maintenance, and support costs

Total Cost of Ownership Analysis

Cost Component	Dell-AMD-Broadcom Solution	Alternative Architecture	Savings
Hardware Acquisition	Baseline Cost	+23% Higher	**23% Savings**
Power Consumption	Optimized Efficiency	+15% Higher	**15% Savings**
Operational Management	Unified Management	Complex Multi-Vendor	**Reduced Complexity**
3-Year TCO	100% (baseline)	123-128%	**23-28% Total Savings**

Utilization

Performance per dollar invested analysis.

Category	Improvement
Acquisition Cost	23% lower than comparable alternative architectures
Power Efficiency	15% reduction in operational power costs
Management Simplicity	Reduced operational overhead through unified management tools

MLPerf-Style Benchmark Results

Training Performance

MLPerf-style benchmarks, though not official MLPerf submissions, offer standardized performance comparisons across different hardware platforms.

Llama 3 70B LoRA Training Performance by Configuration

Configuration	GPUs	Tokens/Second	Scaling Efficiency	Time to Completion
Single GPU	1	2847	100% (baseline)	6 hr 39 min
Quad GPU	4	10892	95.6%	1 hr 44 min
Single Node	8	21234	93.2%	53.5 min
2 Nodes	16	41583	91.3%	27.3 min
4 Nodes	32	81892	89.9%	13.9 min
8 Nodes	64	158487	87.0%	7.2 min

Inference Performance

MLPerf-style inference benchmarks evaluate real-time and batch processing performance across various model types.

Model	Environment	Throughput
Llama 2-70B	Offline	206,364 samples/second peak throughput at 64 GPUs (8 nodes)
BERT-99	Server	47,694 queries/second sustained throughput at 64 GPUs (8 nodes)
3D U-Net	Offline	94.4 samples/second for medical imaging workloads at 64 GPUs (8 nodes)

Implementation Best Practices

AI Practitioners: Technical Optimization Guide

Model Development and Training

Configuration strategies for maximizing training performance and efficiency on AMD Instinct Series GPUs require an understanding of ROCm software stack optimization, memory hierarchy utilization, and distributed training coordination. Our testing methodology evaluates model parallelism strategies including tensor parallelism, pipeline parallelism, and hybrid approaches across varying model sizes from 7B to 405B parameters, with specific focus on maximizing the substantial HBM capacity advantages of AMD GPUs. Performance optimization includes gradient checkpointing, mixed-precision training with custom loss scaling, and optimized collective communication libraries specifically tuned for AMD CDNA architecture. The evaluation framework incorporates PyTorch with ROCm-optimized kernels and measures training throughput, memory efficiency, and convergence characteristics to establish best practices for production deployments.

Advanced Training Optimization Strategies

Training efficiency optimization extends beyond basic configuration to encompass additional techniques like dynamic loss scaling, adaptive batch sizing, and topology-aware model sharding strategies. Detailed analysis of the impact of ROCm version selection, compiler optimization flags, and memory pool configuration on training performance, with particular attention to minimizing memory fragmentation and maximizing compute unit utilization across the entire design is critical. Performance benchmarking includes evaluation of advanced training techniques such as sequence parallelism for long-context models, expert parallelism for Mixture-of-Experts architectures, and activation strategies to exploit the increased memory bandwidth of AMD's HBM3 implementation. The testing framework provides quantified performance metrics including tokens-per-second throughput, memory utilization efficiency, and power consumption characteristics across different training configurations, enabling data-driven optimization decisions for large-scale AI training deployments on AMD hardware.

Training Recommendations

Optimization	Description
Batch Size Tuning	Start with batch size 32 per GPU for large language models, scale based on memory utilization
Mixed Precision	Enable BF16 mixed precision with dynamic loss scaling for 2.14x performance improvement
Gradient Accumulation	Use gradient accumulation for effective large batch training when memory constrained
Learning Rate Scaling	Apply linear learning rate scaling for distributed training across multiple nodes

Memory Management Strategies

Memory Optimization	Description
Activation Checkpointing	Reduce memory usage by 35% with minimal performance impact
Model Sharding	Implement tensor parallelism for models exceeding single GPU memory
Pipeline Parallelism	Enable training of very large models across multiple GPUs
Memory Monitoring	Utilize AMD's profiling tools for memory usage optimization

Distributed Training Optimization

Optimization	Description
RCCL Tuning	Set RCCL_TREE_THRESHOLD=16777216 for optimal collective performance
Network Topology Awareness	Configure process placement to align with RAIL architecture
Overlapped Communication	Enable gradient communication overlap with backward pass computation
Bucket Size Optimization	Tune DDP bucket sizes based on model architecture and network bandwidth

IT Operations: Infrastructure Management

Cluster Deployment and Configuration

Operational best practices for deploying and managing large-scale AMD Instinct Series GPU clusters.

System Configuration

Area	Optimization
BIOS Settings	Enable SR-IOV, configure memory channels for optimal bandwidth
OS Tuning	Apply kernel parameters for RDMA and high-performance networking
Driver Installation	Use latest AMD ROCm drivers or containers with optimized firmware versions
Thermal Management	Implement dynamic thermal controls for power efficiency

Network Infrastructure Management

Deployment and operational considerations for the Broadcom Thor2 and Dell PowerSwitch networking infrastructure.

Network Configuration

Area	Description
VLAN Segmentation	Implement traffic isolation between compute and management networks
Quality of Service	Configure QoS policies for AI workload prioritization
Monitoring and Telemetry	Deploy thorough network monitoring for proactive management
Firmware Management	Maintain consistent firmware versions across all network components

Storage Integration

Guidelines for integrating high-performance storage systems with the compute and network infrastructure.

Storage Optimization

Optimization Area	Strategy
Parallel Filesystem	Deploy Lustre or BeeGFS or equivalent for high-performance data access
Local Caching	Configure NVMe SSDs for local data caching and temporary storage
Data Pipeline Optimization	Implement efficient data loading pipelines to maintain GPU utilization
Backup and Recovery	Establish data protection strategies for training datasets and model checkpoints

Monitoring and Observability

Performance Monitoring Framework

Comprehensive monitoring strategies for maintaining optimal cluster performance and identifying bottlenecks.

Key Metrics

Category	Description
GPU Utilization	Monitor compute, memory, and power utilization across all accelerators
Network Performance	Track bandwidth utilization, packet loss, and congestion indicators
Storage I/O	Monitor filesystem performance and data loading bottlenecks
Application Metrics	Track training progress, convergence rates, and job completion times

Alerting and Automation

Proactive monitoring and automated response systems for maintaining cluster availability and performance.

Considerations

Strategy	Description
Threshold-Based Alerting	Configure alerts for performance degradation and resource exhaustion
Predictive Maintenance	Implement monitoring for early detection of hardware issues
Automated Recovery	Deploy automated response systems for common failure scenarios
Capacity Planning	Monitor growth trends and resource utilization for capacity planning

Future Optimization Opportunities

Emerging Technologies Integration

Next-Generation Hardware

Future Dell platforms may include pooled memory, chiplet-based processors, and on-package photonic links, reshaping cluster design and resource scheduling. Signal65 roadmaps prepare for heterogeneous compute, AI accelerators, quantum-classical hybrids, and neuromorphic devices, and evaluate integration risks and performance-tuning opportunities on existing and emerging fabrics.

AI infrastructure will continue to increase power delivery to GPUs, from current 700W components to 1000W, requiring cooling system redesign for direct liquid cooling, advanced heat capture, and rack-scale optimization. Signal65 evaluation services provide actionable guidance for infrastructure investment planning, including modular upgrade pathways to preserve existing network investments while enabling seamless integration of future hardware generations, with specific recommendations for maintaining scaling efficiency across mixed-generation deployments during transition periods.

Technology Evolution Trends

Area	Innovation
Memory Advancement	HBM4 integration for increased capacity and bandwidth
Networking Evolution	800GbE and 1.6TbE network interface development
CPU-GPU Integration	Advanced CPU-GPU coherent memory architectures
Optical Networking	Direct optical GPU-to-GPU communication for ultra-low latency and high bandwidth

Software Stack Advancements

Area	Improvement
Compiler Optimization	Advanced kernel fusion and optimization techniques
Dynamic Scheduling	Intelligent workload scheduling and resource allocation
Multi-Tenancy	Enhanced support for concurrent workload execution
Edge Integration	Seamless training-to-inference deployment pipelines

Sustainability and Efficiency

Green Computing Initiatives

Strategies for improving power efficiency and reducing environmental impact of large-scale AI computing.

Dell Smart Cooling Technology

Feature	Description
Multi-Vector Cooling	Advanced airflow design with optimized fan algorithms reduce cooling power consumption by up to 20%
Liquid Cooling Ready	Dell Direct Liquid Cooling (DLC) solutions reduce cooling energy by up to 40% compared to air cooling
Intelligent Thermal Management	Dell iDRAC provides real-time thermal monitoring and dynamic fan speed adjustment based on workload

Dell OpenManage Power Management

Feature	Description
Power Cap Policy	Set power consumption limits at server, rack, or data center level to optimize energy usage
Workload-Aware Power Scaling	Automatically adjust power states based on GPU utilization patterns
Peak Shaving	Reduce power consumption during peak demand periods without impacting performance

Dell Technologies Sustainability Initiatives

Strategy	Description
Free Air Cooling	Dell validated designs for free-air cooling operate at temperatures up to 35°C (95°F), reducing cooling costs up to 70%
Modular Data Center Solutions	Dell modular data center designs optimize PUE (Power Usage Effectiveness) to as low as 1.25
Renewable Energy Integration	Dell infrastructure is designed to integrate with renewable energy sources and energy storage systems

Executive Decision Framework

Business Case Development

A Framework for evaluating the business impact and return on investment for Dell PowerEdge XE9680 AMD Instinct Series GPU deployments requires analysis of total cost of ownership, performance benchmarking against competing architectures, and quantified business value metrics tailored to AI workload requirements. This Signal65 evaluation includes cost modeling of hardware acquisition, infrastructure adaptation, software licensing, and operational expenses across 36-month deployment lifecycles, with particular emphasis on Dell competitive advantages and economics for large language model training.

ROI Analysis

Category	Description
Performance Benefits	Quantified improvement in AI development and deployment capabilities
Cost Optimization	Direct cost savings compared to alternative solutions
Operational Efficiency	Reduced management overhead and operational complexity
Strategic Advantages	Competitive positioning and market differentiation opportunities

Technical Performance Metrics

Category	Metrics
Training Performance	Time-to-train improvements for standard model architectures
Inference Latency	P99 response time improvements for production applications
Resource Utilization	GPU, memory, and network utilization efficiency
System Availability	Uptime and reliability metrics for production systems

Business Impact Metrics

Category	Impact
Development Velocity	AI model development and deployment cycle times
Cost Optimization	Reductions in total cost of ownership
Innovation Capability	Number of new AI applications and use cases enabled
Market Competitiveness	Improved competitive positioning through AI capabilities

Conclusion and Recommendations

The performance analysis of Dell PowerEdge XE9680 with AMD Instinct Series GPU cluster and Broadcom networking demonstrates a compelling solution for organizations seeking cost-effective, high-performance AI infrastructure. This combination delivers competitive performance across training and inference workloads while providing significant total cost of ownership advantages. Our benchmarking across 1 through 64-GPU configurations validates sustained training throughput exceeding 87% scaling efficiency, with particularly strong performance characteristics for large language models leveraging AMD's superior memory capacity and bandwidth. The integrated solution addresses critical enterprise requirements including operational simplicity, vendor diversification, and future-proof capabilities to position organizations for long-term AI infrastructure success.

Beyond GPU compute performance, Broadcom Thor2 NICs, Tomahawk 5 switching, and Atlas ASICs anchor the network fabric with a standards-based Ethernet architecture. This approach yields 20–30% TCO improvements by avoiding lock-in, leveraging existing enterprise skills, and ensuring straightforward lifecycle management. By building on Ethernet, organizations can confidently scale today's AI clusters while positioning themselves for seamless adoption of next-generation speeds and features, protecting infrastructure investments for years to come.

Strategic Implications and Recommendations

Organizations evaluating AI infrastructure investments should strongly consider Dell solutions with AMD Instinct Series solutions as a viable alternative to traditional NVIDIA-based deployments, particularly for workloads requiring substantial memory capacity or cost-optimization priorities. The documented performance parity across standard training workloads, combined with 20-30% lower acquisition costs and reduced operational complexity, establishes a compelling business case for enterprise AI initiatives. Adopters will benefit from competitive advantages including supply chain diversification, reduced vendor lock-in risks, and access to rapidly maturing open-source software ecosystems. As AI workloads continue evolving toward larger models and more complex architectures, the foundation established by this integrated Dell-AMD-Broadcom solution provides the scalability, performance, and economic efficiency required to support enterprise AI transformation initiatives while maintaining operational excellence.

Key Technical Findings

Performance Excellence

AMD Instinct Series GPU achieves 87% scaling efficiency at 64 GPU scale for large language model training, demonstrating excellent multi-node performance characteristics. The 192GB HBM3 memory capacity enables larger batch sizes and reduced model sharding compared to previous generation accelerators.

Network Infrastructure

Broadcom Thor2 and Dell PowerSwitch infrastructure provides robust, scalable networking that efficiently supports AI workloads. RAIL-optimized topologies deliver optimal performance for dominant collective operations while maintaining flexibility for diverse communication patterns.

Operational Simplicity

Integrated Dell-AMD-Broadcom solution stacks reduce operational complexity through unified management tools and streamlined deployment procedures, enabling organizations to focus on AI innovation rather than infrastructure management.

Strategic Recommendations

Immediate Actions:

- Review performance optimization procedures for production deployment
- Engage Signal65 for customer Proof-of-Concept services to validate performance with organization-specific workloads
- Develop expertise through hands-on experience with AMD ROCm software

Medium-term Implementation:

- Scale deployment based on pilot results and business requirements
- Implement comprehensive monitoring and optimization procedures
- Integrate with existing data science and MLOps workflows

Long-term Strategy:

- Plan for technology evolution and next-generation hardware integration
- Explore specialized optimization techniques for organization-specific AI applications
- Establish centers of excellence for AI infrastructure optimization

Dell-AMD-Broadcom solutions represent a strategic opportunity for organizations to establish competitive AI infrastructure while optimizing total cost of ownership and operational efficiency. Early adoption enables establishment of expertise and competitive advantages in the rapidly evolving AI landscape.

Important Information About this Report

CONTRIBUTORS
Brian Martin
AI and Data Center Lead | Signal65

PUBLISHER
Ryan Shrout
President and GM | Signal65

INQUIRIES
Contact us if you would like to discuss this report and Signal65 will respond promptly.

CITATIONS
This paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title, and "Signal65." Non-press and non-analysts must receive prior written permission by Signal65 for any citations.

LICENSING
This document, including any supporting materials, is owned by Signal65. This publication may not be reproduced, distributed, or shared in any form without the prior written permission of Signal65.

DISCLOSURES
Signal65 provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

IN PARTNERSHIP WITH

DELL Technologies

ABOUT SIGNAL65
Signal65 is an independent research, analysis, and advisory firm, focused on digital innovation and market-disrupting technologies and trends. Every day our analysts, researchers, and advisors help business leaders from around the world anticipate tectonic shifts in their industries and leverage disruptive innovation to either gain or maintain a competitive advantage in their markets.

signal65

signal65

CONTACT INFORMATION
Signal65 I signal65.com