

OcNOS ETHERNET-BASED 800G LOSSLESS AI FABRIC



OcNOS Traffic Engineering for AI Fabric:

- **PFC (Priority-based Flow Control)** over Layer 3 for **guaranteed lossless** transport.
- **ETS (Enhanced Transmission Selection)** for intelligent bandwidth allocation and **guaranteed** prioritization.
- **DCBX (Data Center Bridging Capabilities Exchange Protocol)** for consistent QoS configuration and **guaranteed** network-wide policy enforcement.
- **ECN (Explicit Congestion Notification)** for proactive congestion control to **prevent packet loss**.
- **DLB (Dynamic Load Balancing)** for optimal link utilization and **guaranteed** traffic distribution.

These features work together to create a high-throughput, low-latency, and resilient network foundation, guaranteeing zero packet loss for AI computations.

Elevating AI Workloads: The OcNOS Advantage

Modern AI training demands exceptional network performance. With modern ultra-low latency transport technologies, like RoCEv2, even a single packet loss can drastically degrade Job Completion Time (JCT) and GPU utilization. IP Infusion's OcNOS Data Center (OcNOS-DC) provides the **guaranteed lossless networking** solution essential for maximizing your AI infrastructure's potential.

The AI Networking Challenge

Modern AI/ML workloads, particularly distributed GPU training utilizing RoCEv2 (RDMA over Converged Ethernet v2), present unique network demands:

- **Extreme Loss Sensitivity:** RoCEv2 is highly intolerant to packet loss; even minor drops trigger costly retransmissions, leading to significant Job Completion Time (JCT) delays and reduced GPU efficiency.
- **High Bandwidth & Low Latency:** GPU-to-GPU and GPU-to-storage communication requires massive throughput with minimal delay for optimal performance.
- **Congestion Management:** Dynamic and bursty AI traffic patterns can easily lead to network congestion, necessitating advanced mechanisms to prevent packet loss.
- **Telemetry:** Building and managing large AI clusters requires a network that scales effortlessly and provides real-time visibility for proactive issue resolution.

The OcNOS Solution for AI Fabrics

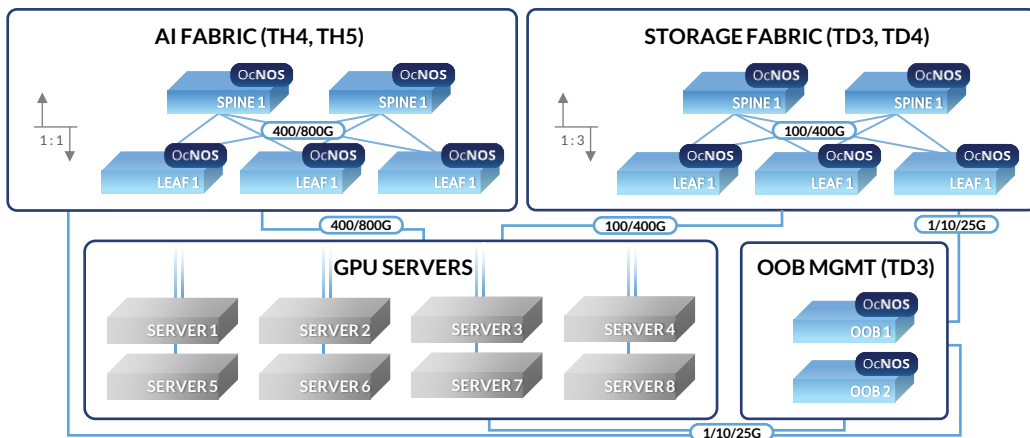
IP Infusion OcNOS-DC transforms open, high-performance Broadcom ASICs (including Tomahawk 5, Tomahawk 4, Trident 4, Trident 3) into a purpose-built AI fabric. It ensures **zero packet loss** and uninterrupted data flow for GPU workloads, leading to:

- **Optimized Job Completion Time (JCT)**
- **Maximized GPU Utilization**
- **Guaranteed Training Stability**

AI Fabric Design: Purpose-Built for Peak Performance

OcNOS enables you to segment your network, guaranteeing optimal performance for each traffic type:

- **Compute Fabric (GPU-to-GPU):** Ultra-low latency, high bandwidth, and **guaranteed lossless transport** for critical training traffic (e.g., gradient synchronization). Powered by OcNOS on Broadcom Tomahawk 5 (800G) and Tomahawk 4 (400G) based switches.
- **Storage Fabric (CPU-to-Storage):** High-throughput data transfer for checkpointing and bulk I/O. Enabled by OcNOS on Broadcom Trident 4 (400G) and Trident 3 (100G) based switches.
- **Management Fabric:** Reliable out-of-band connectivity for operations and monitoring. Supported by OcNOS on Broadcom Trident 3 (1/10/25G) based switches.



OcNOS Solution Coverage for AI and DC Fabric

OcNOS Capabilities: Guaranteed Lossless AI Transport

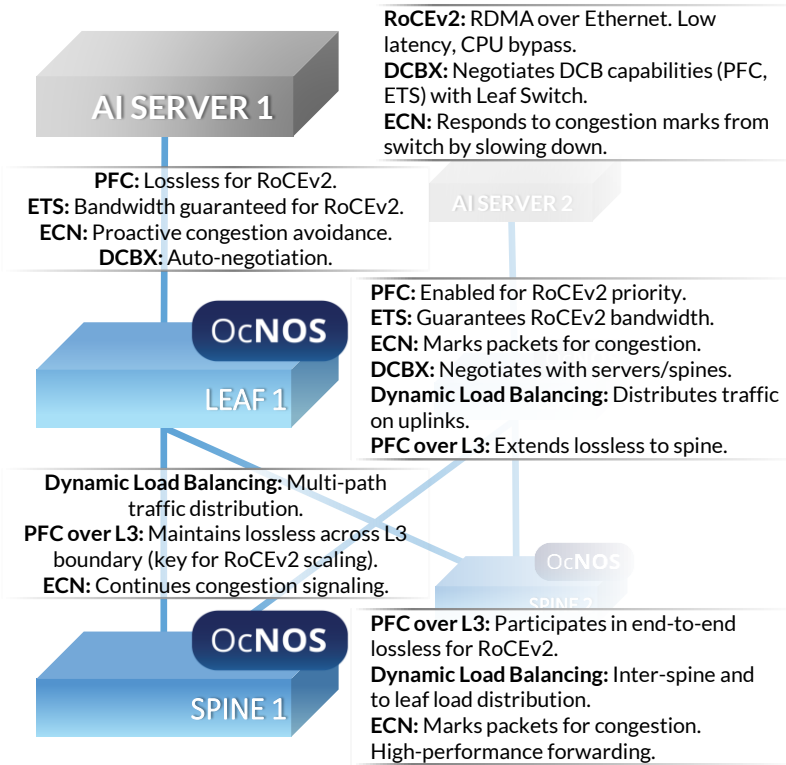
AI workloads demand a network that precisely handles sensitive traffic patterns. OcNOS provides a **unified suite of traffic engineering features** that work in concert to transport these critical flows flawlessly, guaranteeing zero packet loss, and directly improving GPU utilization.

Achieving Guaranteed Lossless AI Transport Through Coordinated Mechanisms

The foundation of high-performance AI networking, especially for RoCEv2 traffic, lies in guaranteed lossless transport over scalable Layer 3 routed fabrics. OcNOS achieves this through a carefully coordinated sequence of mechanisms designed to **prevent even a single packet loss**:

1. **ECN (Explicit Congestion Notification): Proactive Loss Prevention.** ECN is the **first line of defense**. OcNOS proactively marks packet headers when congestion occurs in ques, signaling the RoCEv2 receiver to notify the sender (via CNP) to reduce its rate. This non-disruptive feedback loop prevents critical congestion and guarantees the prevention of packet drops, maintaining high throughput.
2. **PFC (Priority-based Flow Control) over Layer 3: The Ultimate Lossless Guarantee.** Should congestion persist, OcNOS with PFC (IEEE 802.1Qbb) will send pause frames requests to the GPU server to guarantee packet loss avoidance. Even a single packet loss can cripple AI training. OcNOS extends this Layer 2 pause mechanism across your entire Layer 3 routed fabric. At every hop, OcNOS ensures pause frames are honored for high-priority AI traffic. This **guarantees** that AI data never experiences packet loss due to congestion, ensuring uninterrupted and high-fidelity flow from source to destination.
3. **ETS (Enhanced Transmission Selection): Intelligent Bandwidth Prioritization.** ETS (IEEE 802.1Qaz) complements PFC by intelligently allocating bandwidth on shared links. OcNOS uses ETS to classify traffic into priority groups, precisely assigning a percentage of link bandwidth to each. This **guarantees** that critical AI traffic is always prioritized and assured its necessary bandwidth, preventing starvation.

DCBX (Data Center Bridging Capabilities Exchange Protocol): Fabric-Wide QoS Automation. DCBX is the essential automation layer that binds PFC and ETS, ensuring consistent QoS configurations across the entire fabric. It enables switches and NICs to automatically negotiate precise lossless and bandwidth settings. Without DCBX, manual errors could jeopardize the lossless guarantee. OcNOS leverages DCBX for unwavering QoS, robust workload isolation, and precise prioritization end-to-end, **making complex coordination seamless, scalable, and guaranteed to prevent packet loss.**



OcNOS Guaranteed Lossless Networking for AI Fabric

Dynamic Load Balancing (DLB): Maximizing Network Utilization

DLB is critical for maximizing network utilization. Unlike static ECMP hashing, OcNOS's DLB intelligently distributes traffic across all available ECMP and LAG paths, proactively preventing congestion hotspots. DLB dynamically adjusts forwarding based on real-time metrics (utilization, queue depth), rebalancing new "flowlets" while ensuring strict in-order delivery. This **guarantees** efficient use of all available paths, preventing scenarios where even lossless traffic could stall due to an artificial bottleneck.

Seamless Automation & Orchestration for AI Fabrics

Managing a high-performance AI/ML fabric demands effortless control and deep visibility. OcNOS transforms network

What IP Infusion Offers: Your AI Networking Advantage

IP Infusion provides the network operating system (NOS) and solutions to build high-performance, cost-effective, and adaptable AI data centers:

OcNOS Software: The Intelligent Core

- **OcNOS-DC:** Engineered for AI workloads, guaranteeing uninterrupted flow.
- **Advanced TE:** PFC over L3, ETS, DCBX, ECN, DLB – guaranteed lossless delivery.
- **AI Orchestration :** gNMI telemetry & Ansible automation for consistent network states.
- **Perpetual Licensing:** OcNOS-DC is offered in 4 SKUs w/o add-on or subscription fees.

Open Networking Benefits:

- **Open Whitebox Ecosystem:** wider platform selection, shorter lead times.
- **Broadcom ASIC Optimization:** Full support for Tomahawk 5/4 (800G/400G) and Trident 4/3 (400G/25G/100G) for unified fabrics, guaranteeing hardware capabilities.
- **Significant CAPEX/OPEX savings,** vendor diversity, and agility, freedom of choice.
- **Open Optics Ecosystem:** IP Infusion offers you with a wide selection of rigorously qualified transceivers.

Expert Support:

- Comprehensive technical support and professional services, **guaranteeing your success.**

operations into a streamlined, automated experience, freeing your teams to innovate.

- **Open APIs & Real-time Telemetry:** OcNOS integrates deeply with AI orchestrators via Netconf, gNMI, and OpenConfig. It provides streaming telemetry for granular, real-time insights into network performance, enabling proactive bottleneck detection and intelligent resource allocation.
- **Effortless Automation with Ansible:** The IP Infusion OcNOS [Ansible Collection](#) automates EVPN-VXLAN provisioning and namespace-based configurations, drastically reducing human error and accelerating deployment cycles.
- **Intuitive CLI:** For direct interaction and troubleshooting, a powerful yet user-friendly Command Line Interface (CLI) ensures quick verification and on-the-fly adjustments.

By blending open standards, real-time insights, and powerful automation, OcNOS ensures your AI fabric is easy to deploy, monitor, and evolve, making it a truly future-proof investment.

Summary

IP Infusion OcNOS-DC is the essential network operating system for next-generation AI/ML data centers. It delivers **guaranteed lossless networking** through a comprehensive suite of advanced traffic engineering features, ensuring peak GPU utilization and accelerated job completion. By embracing a disaggregated model on Broadcom ASICs, OcNOS provides significant cost savings, enhanced agility, and robust observability and security. Choose OcNOS to build a high-performance, future-proof AI fabric that maximizes your AI investment.

Contact for More Information:

For more information on the OcNOS software, please contact sales@ipinfusion.com

ABOUT IP INFUSION

IP Infusion is a leading provider of open network software and solutions for carriers, service providers and data center operators. Our solutions enable network operators to disaggregate their networks to accelerate innovation, streamline operations, and reduce Total Cost of Ownership (TCO). Network OEMs may also disaggregate network devices to expedite time to market, offer comprehensive services, and achieve carrier grade robustness. IP Infusion network software platforms have a proven track record in carrier-grade open networking with over 500 customers and over 10,000 deployments. IP Infusion is headquartered in Santa Clara, Calif., and is a wholly owned and independently operated subsidiary of ACCESS CO., LTD. Additional information can be found at <http://www.ipinfusion.com>

© 2025 IP Infusion, Inc. All rights reserved. IP Infusion is a registered trademark and the IP Infusion logo and OcNOS are trademarks of IP Infusion, Inc. All other trademarks and logos are the property of their respective owners. IP Infusion assumes no responsibility for any inaccuracies in this document. IP Infusion reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Phone | +1-877-699-3267 Email | sales@ipinfusion.com Web | www.ipinfusion.com

