

Cisco AI Networking for Data Center with NVIDIA Spectrum-X



Contents

Introduction

Expected outcomes from AI networking

How AI networking is driving outcomes

- Meet Cisco Intelligent Packet Flow
- Challenge for AI networking:
Equal-cost multipath (ECMP)
- Cisco Intelligent Packet Flow capability
#1: dynamic load balancing – flowlet
- Cisco Intelligent Packet Flow capability
#2: per-packet load balancing

Cisco data-center switching + NVIDIA Spectrum-X architecture

Seamless integration: Cisco data-center switching and NVIDIA Spectrum-X architecture

References

Introduction

The rapid growth of AI applications demands a network fabric that is scalable, efficient, and purpose-built for performance. Cisco's data-center networking solutions integrate silicon, hardware, software, management tools, and optics into a unified system to meet these needs. AI pipelines, spanning stages such as data ingestion, model training, optimization, and inferencing, generate diverse traffic patterns that require tailored infrastructure to ensure high bandwidth, low latency, and reduced Job Completion Time (JCT). For example, model training relies on east-west traffic with both high bandwidth and low latency as GPUs or accelerators exchange gradients and parameters. Cisco's data-center switching enables AI network fabrics to scale seamlessly from a few dozen to thousands of GPU clusters, delivering the bandwidth, latency optimization, congestion management, telemetry, and power efficiency required to support all stages of the AI pipeline with lossless data transmission.




Figure 1. Cisco AI infrastructure

Contents

Introduction

Expected outcomes from AI networking

How AI networking is driving outcomes

- Meet Cisco Intelligent Packet Flow
- Challenge for AI networking: Equal-cost multipath (ECMP)
- Cisco Intelligent Packet Flow capability #1: dynamic load balancing – flowlet
- Cisco Intelligent Packet Flow capability #2: per-packet load balancing

Cisco data-center switching + NVIDIA Spectrum-X architecture

Seamless integration: Cisco data-center switching and NVIDIA Spectrum-X architecture

References

Expected outcomes from AI networking

A robust AI network relies on several critical capabilities to ensure optimal performance and efficiency across the different stages of an AI pipeline.

- 1. Uniform traffic distribution in the network fabric** ensures that data flows are evenly balanced across available paths, preventing overutilization of specific links while underutilizing others. This balance optimizes resource usage, minimizes latency, and helps sustain peak performance even as AI workloads scale.
- 2. Congestion avoidance and management** are essential for maintaining smooth network operations, especially under the heavy data-load characteristics of AI workloads. By proactively identifying and mitigating congestion points, the network can prevent delays, reduce packet loss, and improve overall throughput, enabling maximum GPU utilization and faster job completion times.
- 3. Hardware link failure management** is vital for ensuring high availability and reliability. By quickly detecting and rerouting traffic around failed links, the network can maintain uninterrupted operations, minimize disruptions, and ensure that critical AI processes continue without downtime.

Together, these capabilities form the backbone of a resilient and efficient AI-driven network environment.

Contents

Introduction

Expected outcomes from AI networking

How AI networking is driving outcomes

- Meet Cisco Intelligent Packet Flow
- Challenge for AI networking: Equal-cost multipath (ECMP)
- Cisco Intelligent Packet Flow capability #1: dynamic load balancing – flowlet
- Cisco Intelligent Packet Flow capability #2: per-packet load balancing

Cisco data-center switching + NVIDIA Spectrum-X architecture

Seamless integration: Cisco data-center switching and NVIDIA Spectrum-X architecture

References

How AI networking is driving outcomes

Meet Cisco Intelligent Packet Flow

Cisco Intelligent Packet Flow optimizes AI traffic for seamless performance, even under demanding workloads, reducing Job Completion Time (JCT) through:

- 1. Optimal path utilization:** efficient traffic distribution across the fabric by leveraging advanced load-balancing techniques, including flowlet-based dynamic load balancing, per-packet load balancing, weighted cost multipath, policy-based load balancing, etc.
- 2. Congestion-aware traffic management:** real-time traffic visibility with telemetry features such as microburst detection, congestion signaling, tail timestamping, and in-band Network Telemetry (INT).
- 3. Autonomous recovery:** fault-aware rerouting and fast convergence to minimize network hotspots and prevent disruptions.

Let's take a closer look at a few key capabilities for AI networking with Cisco Intelligent Packet Flow, especially around providing optimal path utilization such as dynamic load balancing (flowlet) and per-packet load balancing.




Figure 2. Challenge for AI networking and Cisco's capabilities to address

Contents

Introduction

Expected outcomes from AI networking

How AI networking is driving outcomes

- Meet Cisco Intelligent Packet Flow
- Challenge for AI networking: Equal-cost multipath (ECMP)
- Cisco Intelligent Packet Flow capability #1: dynamic load balancing – flowlet
- Cisco Intelligent Packet Flow capability #2: per-packet load balancing

Cisco data-center switching + NVIDIA Spectrum-X architecture

Seamless integration: Cisco data-center switching and NVIDIA Spectrum-X architecture

References

Challenge for AI networking: Equal-cost multi-path (ECMP)

What it means in simple language

Imagine a toll plaza with multiple lanes, but every car with the same license-plate pattern (for example, odd or even numbers) is always sent to the same lane. This ensures that cars of the same group stay in order, avoiding confusion. However, if one lane becomes overcrowded with cars while others remain underused, traffic flow becomes unbalanced.

What it means technically

The traditional method of Equal-Cost Multipath (ECMP) enables packets with the same 5-tuple (source IP, destination IP, source port, destination port, and protocol) to be distributed over the best path, with subsequent frames maintained on the same link, avoiding out-of-order packet delivery.




Figure 3. Equal-cost multi-path (ECMP) – traditional approach

Contents

Introduction

Expected outcomes from AI networking

How AI networking is driving outcomes

- Meet Cisco Intelligent Packet Flow
- Challenge for AI networking:
Equal-cost multipath (ECMP)
- Cisco Intelligent Packet Flow capability
#1: dynamic load balancing – flowlet
- Cisco Intelligent Packet Flow capability
#2: per-packet load balancing

Cisco data-center switching + NVIDIA Spectrum-X architecture

Seamless integration: Cisco data-center switching and NVIDIA Spectrum-X architecture

References

This results in under-utilization or over-utilization of links due to low hashing entropy and is not adaptive to bursty traffic. Based on performance and benchmark testing, this leads to unbalanced traffic distribution across different links.




Figure 4. Unbalanced traffic distribution across various links

The results in Figure 4 show how Tx % and OutUcastPkts across eight Ethernet interfaces have uneven amounts of traffic distributed amongst them. This can lead to link congestion, increased packet loss, and suboptimal AI application performance.

Because this degrades the overall throughput efficiency, traditional ECMP is typically not suitable for AI fabrics.

Cisco Intelligent Packet Flow capability #1: dynamic load balancing – flowlet

What it means in simple language

Dynamic Load Balancing (DLB) is like a smart traffic-management system on a multilane highway. Imagine a highway where traffic is monitored in real time, and instead of assigning entire groups of cars to a single lane based on their license plate pattern (as in the ECMP toll booth analogy above), the system breaks down the traffic into smaller groups, like individual car convoys (representing flowlets).

Each convoy is directed to the lane with the least congestion at that moment, ensuring that all lanes are utilized efficiently. If one lane becomes blocked or slows down (such as during a link failure), the system quickly reroutes the convoys to other available lanes without causing a delay. This intelligent system also ensures that cars within a convoy always arrive in the correct order at their destination, avoiding confusion and maintaining smooth travel.

What it means technically

Dynamic Load Balancing – flowlet (DLB – flowlet) is a sophisticated mechanism that optimizes network performance by intelligently distributing traffic across multiple paths. It identifies flowlets, which are subsets of flows within a five-tuple flow, by analyzing packet burst intervals. By breaking down traffic at the flowlet level, DLB allows each flowlet to be transmitted over a distinct path based on real-time transmission rates, ensuring efficient utilization of available links. This approach not only maximizes link efficiency but also enables fast convergence in the event of link failures, because traffic can be quickly rerouted without disruption.




Figure 5. Dynamic load balancing – flowlet

Additionally, DLB ensures that packets arrive in order, eliminating the need for reordering and preserving the integrity of data flows. Based on performance and benchmark testing, this leads to near perfect traffic distribution across various links.

Contents

Introduction

Expected outcomes from AI networking

How AI networking is driving outcomes

- Meet Cisco Intelligent Packet Flow
- Challenge for AI networking:
Equal-cost multipath (ECMP)
- Cisco Intelligent Packet Flow capability
#1: dynamic load balancing – flowlet
- Cisco Intelligent Packet Flow capability
#2: per-packet load balancing

Cisco data-center switching + NVIDIA Spectrum-X architecture

Seamless integration: Cisco data-center switching and NVIDIA Spectrum-X architecture

References




Figure 6. Optimized traffic distribution for AI networks with DLB - flowlet

The results in Figure 6 show how Tx % and OutUcastPkts across eight Ethernet interfaces have optimized traffic distribution. This load-balancing strategy is ideal for distributing traffic in real-time for AI workloads, delivering high throughput efficiency and shorter job completion times.

Cisco Intelligent Packet Flow capability #2: per-packet load balancing

What it means in simple language

Visualize a high-tech logistics system managing the delivery of goods across a network of roads. Imagine a central warehouse (representing the sender's data-center network switch) sending packages (packets) to a destination using a fleet of trucks (network links). Instead of assigning all packages for a specific destination to a single truck, the system intelligently divides the packages and distributes them across multiple available trucks to maximize road (link) utilization and ensure faster delivery.

However, because each truck might take a slightly different route (path), some packages might arrive out of order. To address this, the receiving facility (the NIC on the server) is equipped with advanced sorting technology to reorder the packages as they arrive, ensuring that everything is in the correct sequence before processing.

Contents

Introduction

Expected outcomes from AI networking

How AI networking is driving outcomes

- Meet Cisco Intelligent Packet Flow
- Challenge for AI networking: Equal-cost multipath (ECMP)
- Cisco Intelligent Packet Flow capability #1: dynamic load balancing – flowlet
- Cisco Intelligent Packet Flow capability #2: per-packet load balancing

Cisco data-center switching + NVIDIA Spectrum-X architecture

Seamless integration: Cisco data-center switching and NVIDIA Spectrum-X architecture

References

What it means technically




Figure 7. Cisco data-center switching example

Cisco data-center switches support per-packet load balancing designed to maximize network efficiency and performance. In this mode, packets within the same flow are distributed across multiple available links, enabling AI workloads to evenly balance traffic and fully utilize the network bandwidth. This distribution strategy ensures optimal link utilization and allows for rapid convergence in the event of link failures, because traffic can be quickly redirected without interrupting data flows. However, since packets belonging to the same flow traverse different paths, there is a possibility of out-of-order delivery at the destination. To address this, deployments using per-packet load balancing rely on the receiving host's NICs (for example, NVIDIA BlueField-3 SuperNICs) to perform packet reordering, ensuring that data integrity and sequence are maintained.

Contents

Introduction

Expected outcomes from AI networking

How AI networking is driving outcomes

- Meet Cisco Intelligent Packet Flow
- Challenge for AI networking:
Equal-cost multipath (ECMP)
- Cisco Intelligent Packet Flow capability
#1: dynamic load balancing – flowlet
- Cisco Intelligent Packet Flow capability
#2: per-packet load balancing

Cisco data-center switching + NVIDIA Spectrum-X architecture

Seamless integration: Cisco data-center switching and NVIDIA Spectrum-X architecture

References

Cisco data-center switching + NVIDIA Spectrum-X architecture

As part of the Cisco and NVIDIA partnership, NVIDIA is enhancing communication between its SuperNICs and Cisco silicon technology to create robust and scalable network infrastructure.

Cisco's integration enables the automated deployment of NVIDIA Spectrum-X, which, when combined with Cisco's per-packet load balancing, ensures uniform traffic distribution across the network fabric, preventing congestion and optimizing resource utilization.




Figure 8. Optimized traffic distribution for AI networks with Cisco per-packet load balancing and NVIDIA SuperNICs

The results in Figure 8 show how Tx % and OutUcastPkts across eight Ethernet interfaces have optimized traffic distribution. This load-balancing strategy is ideal for enabling multipath distribution with improved efficiency.

Contents

Introduction

Expected outcomes from AI networking

How AI networking is driving outcomes

- Meet Cisco Intelligent Packet Flow
- Challenge for AI networking: Equal-cost multipath (ECMP)
- Cisco Intelligent Packet Flow capability #1: dynamic load balancing – flowlet
- Cisco Intelligent Packet Flow capability #2: per-packet load balancing

Cisco data-center switching + NVIDIA Spectrum-X architecture

Seamless integration: Cisco data-center switching and NVIDIA Spectrum-X architecture

References

Seamless integration: Cisco data-center switching and NVIDIA Spectrum-X architecture




Figure 9. Cisco + NVIDIA partnership

The integration of the Cisco and NVIDIA ecosystem marks a significant milestone in data center networking innovation to meet the demands of AI workloads. As a stepping stone in this partnership, NVIDIA and Cisco's enhanced communication between NVIDIA BlueField-3 SuperNICs and Cisco silicon-based switches enables tighter integration and exceptional performance through optimized traffic distribution across AI fabric links. On top of this, Cisco Reference Architectures (RAs) are based on NVIDIA Reference Architectures with Spectrum-X. This represents a cutting-edge collaboration to deliver high-performance networking solutions tailored for AI workloads. Additionally, for customers preferring NVIDIA Silicon with Cisco switches, Cisco plans to develop a Nexus switch leveraging NVIDIA Spectrum-X silicon ASIC. This deep integration will ensure that customers benefit from the combined strengths of both Cisco and NVIDIA technologies, delivering a versatile, high-performance solution tailored for modern AI-driven data centers.

References

- Cisco Expands Partnership with NVIDIA to Accelerate AI Adoption in the Enterprise: <https://newsroom.cisco.com/c/r/newsroom/en/us/a/y2025/m02/cisco-expands-partnership-with-nvidia-to-accelerate-ai-adoption-in-the-enterprise.html>.
- AI Infrastructure with Cisco Nexus 9000 Switches – Cisco Enterprise Reference Architecture (ERA): <https://www.cisco.com/c/en/us/products/collateral/data-center-networking/nexus-hyperfabric/nexus-9000-ai-era-ds.html>.
- NVIDIA-Validated Cisco Nexus Hyperfabric AI Enterprise Reference Architecture: <https://www.cisco.com/c/en/us/products/collateral/data-center-networking/nexus-hyperfabric/hyperfabric-ai-era-ds.html>.
- Cisco Secure AI Factory with NVIDIA Solution Overview: <https://www.cisco.com/c/en/us/solutions/collateral/artificial-intelligence/secure-ai-factory-nvidia-so.html>.

