

Architecting the Next-Generation Data Center Interconnect

Disaggregated DCI Solutions with OcNOS

OcNOS KEY ADVANTAGES

- **Reduce TCO, Power, and Footprint:** Eliminate standalone optical transponders by enabling a converged IP-over-DWDM (IPoDWDM) architecture, deploying 400G ZR/ZR+ coherent optics directly in routers.
- **Accelerate Multi-Service Deployment:** Rapidly deploy sophisticated Layer 2 and Layer 3 services over the DCI fabric using a comprehensive, carrier-grade suite of MPLS, L2VPN/L3VPN, and EVPN technologies.
- **Enable Modern Network Operations:** Stream high-fidelity, real-time telemetry using gNMI and standardized OpenConfig data models, enabling proactive monitoring and automation.
- **De-Risk and Innovate with an Open Ecosystem:** Mitigate supply chain risks and avoid vendor lock-in. The OcNOS disaggregated software model provides the freedom to deploy on a competitive ecosystem of best-in-class whitebox hardware and optics.

The explosive growth of AI and distributed cloud services is forcing a radical re-evaluation of Data Center Interconnect (DCI) architectures. Legacy, monolithic network designs are proving too costly and inflexible to handle the demands for massive bandwidth and ultra-low latency. This solution brief outlines a new architectural paradigm: **disaggregated DCI solutions with OcNOS**.

This approach addresses the challenges of modern AI and enterprise application workloads. Traditional DCI architectures, with their separate IP and optical layers, introduce performance-degrading latency.

The strategic path forward lies in the convergence of these layers through **IP over DWDM (IPoDWDM)** and **network disaggregation**. By decoupling the network operating system (like OcNOS) from the hardware, operators gain the agility to innovate at the speed of software and silicon, breaking free from vendor lock-in and creating a more efficient, scalable, and future-proof network foundation.

Deconstructing the Modern DCI: Roles and Legacy Challenges

The DCI provides high-speed, packet-optical connectivity between two or more data centers over short, medium, or long distances. Its primary purpose is to create a cohesive, virtualized pool of resources, enabling critical business functions such as synchronous data replication, disaster recovery, and dynamic load balancing.

Historically, DCI architectures have relied on a rigid, multi-layered approach. An IP router provides Layer 3 services and connects via short-reach "grey" optics to a dedicated optical transponder shelf. This transponder then performs an (optical-electrical-optical) OEO conversion to map the signal onto a specific Dense Wavelength Division Multiplexing (DWDM) wavelength for long-haul transport. This legacy model presents significant challenges:

- **High Total Cost of Ownership (TCO):** The requirement for redundant hardware, including dedicated transponder chassis and multiple sets of optics, substantially increases capital expenditures (CapEx), ongoing power consumption, and the physical rack space footprint.
- **Operational Silos:** This layered approach fosters organizational silos. Separate IP and optical engineering teams manage their respective domains using different tools, processes, and management systems. This legacy design leads to slow, complex service provisioning and protracted troubleshooting cycles.
- **Increased Latency:** Every optical-electrical-optical (OEO) conversion performed by a transponder introduces a tangible processing delay, typically in the range of 50 to 100 microseconds per site. In a network with multiple hops, this accumulated latency can become a significant performance bottleneck for sensitive applications.

The Architectural Shift: Disaggregation and IP-Optical Convergence

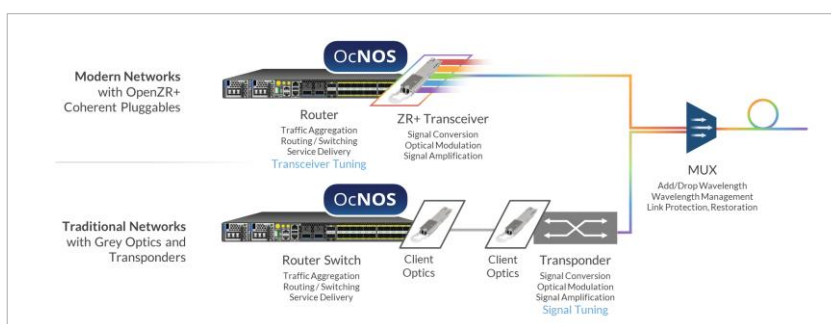
The solution to the challenges posed by legacy DCI architectures is a two-pronged approach: leveraging IP-optical convergence to collapse network layers and embracing network disaggregation to break vendor lock-in. This combination creates a more agile, efficient, and economically sustainable network foundation.

The Power of Disaggregation: Breaking Vendor Lock-in and Accelerating Innovation

IP over DWDM is a network architecture that achieves IP-optical convergence by eliminating the separate optical transponder layer. It accomplishes this by integrating coherent DWDM optics directly into the line-card ports of an IP router or switch. This architectural leap is made possible by the standardization and miniaturization of powerful coherent Digital Signal Processors (DSPs) into small, low-power pluggable optical modules, such as the QSFP-DD form factor. Industry standards like 400G ZR and the more capable 400G ZR+ are the key enabling technologies for this shift.

The architectural simplification:

- Traditional Model: Router → Grey Optic → Transponder (OEO Conversion) → Colored Optic → DWDM Line System
- IPoDWDM Model: Router with Coherent Pluggable Optic → DWDM Line System



Modern networks with Open ZR/ZR+ Pluggables vs traditional networks with optical transponders.

The Power of Disaggregation: Breaking Vendor Lock-in and Accelerating Innovation

Network disaggregation is the fundamental principle of decoupling the network operating system (software) from the underlying packet-forwarding hardware. This model allows network operators to run a single, consistent NOS like OcNOS across a wide variety of whitebox switches and routers from different vendors, all built on high-performance merchant silicon. This delivers three transformative benefits:

- **Freedom of Choice:** Operators can mix and match best-in-class hardware and software components to precisely meet their technical and commercial requirements, avoiding vendor lock-in and optimizing for price-performance.

- **Accelerated Innovation:** This model allows for the rapid adoption of new silicon capabilities and higher-density hardware platforms as soon as they become available on the market, rather than being constrained by a monolithic vendor's slower, integrated development cycle.
- **Operational Efficiency:** Deploying a single NOS across diverse hardware platforms enables a unified automation, orchestration, and management framework. This consistency reduces operational complexity and the cost of training and maintenance.

Quantifiable Impact of the Modern Architecture

The shift to a disaggregated, IPoDWDM-based DCI architecture delivers measurable and significant improvements across TCO, power consumption, and network performance.

- **Total Cost of Ownership (TCO) Reduction:** By eliminating an entire class of network hardware (transponder shelves) and the associated client-side optics, IPoDWDM delivers significant cost savings. Industry analysis and reports from early adopters consistently show CapEx reductions ranging from 40% to 65% and OpEx reductions between 50% and over 76%.
- **Power and Footprint Savings:** The architectural simplification yields substantial environmental and operational benefits. Eliminating transponder chassis can reduce power consumption by up to 70% and save as much as 80% of the required physical rack space. This is driven not only by removing hardware but also by the superior efficiency of modern pluggable optics. For example, a standard 400G ZR module consumes only about 15W of power, a fraction of the power required by a fully loaded transponder shelf.
- **Latency Improvement:** Each transponder in a traditional design introduces latency due to the inherent delay of the OEO signal conversion process. This delay, typically in the range of **50 to 100 microseconds** per site, is completely eliminated in an IPoDWDM architecture. While the propagation delay of light through fiber remains the largest component of end-to-end latency, removing multiple OEO hops across a long-haul path can yield a significant improvement in overall network performance.

Technical Deep Dive: Building a Resilient DCI Fabric with OcNOS-SP

A modern DCI fabric built with OcNOS-SP follows a layered, best-practice design that maximizes scalability, resiliency, and service flexibility. This design leverages modern protocols in both the underlay and overlay to create a robust and efficient transport network.

Scenario	Typical Distance	Optical Layer	Packet Layer	Business Model & OcNOS Role
Metro DCI	<80 km (same metro)	400G ZR IPoDWDM over dark fiber	SR-MPLS	Enterprise/CSP owns both layers; OcNOS provides a fully converged IP+Optical solution.
Regional DCI	80-500 km (nearby cities)	400G ZR+ IPoDWDM or alien wavelength	SR-MPLS	Mixed ownership; OcNOS provides packet-layer intelligence over a leased optical path.
Long-haul DCI	500-2000+ km	Coherent DWDM w/ amplifiers, ROADMs	SR-MPLS	Carrier owns optical; OcNOS delivers robust L3VPN services over a carrier-provided wavelength.
Intra-campus DCI	<5 km (same campus)	Short-reach optics (SR4, DR4, etc.)	EVPN-VXLAN, SR-MPLS optional	Enterprise owns both; OcNOS extends the DC fabric with EVPN-VXLAN for operational consistency.

Data Center Interconnect (DCI) Scenarios and Architectures

The Underlay: A Foundation of Scalability and Sub-50ms Resiliency

The underlay provides the fundamental IP connectivity and transport paths between all nodes in the DCI fabric.

- **Transport Protocol (SR-MPLS):** For traffic engineering, the fabric employs Segment Routing over an MPLS data plane (SR-MPLS). Unlike its predecessor, RSVP-TE, which requires every node along a path to maintain state for every traffic-engineered tunnel, SR-MPLS is stateless in the network core. Path information is encoded as a stack of labels in the packet header at the ingress node. This source-routing approach simplifies core network devices, enhances scalability, and provides granular, per-application control over traffic paths.
- **Automated Resiliency (TI-LFA):** To ensure carrier-grade availability, OcnOS-SP implements Topology-Independent Loop-Free Alternate (TI-LFA). This advanced Fast Reroute (FRR) mechanism uses a sophisticated algorithm to pre-compute and pre-install a loop-free backup path for every destination in the network. In the event of a link or node failure, traffic is immediately rerouted onto this backup path in under 50 milliseconds, providing automated, sub-50ms protection for all traffic, regardless of network topology.

The Overlay: A Unified Service Delivery Control Plane

Services are decoupled from the underlying transport infrastructure using a BGP-based overlay.

- **Control Plane (MP-BGP EVPN):** OcnOS-SP leverages Multiprotocol BGP (MP-BGP) with the EVPN address family as the single, unified control plane for all DCI services. EVPN provides an integrated framework for delivering both Layer 2 and Layer 3 services, eliminating the need to run and manage separate protocols for different service types. It uses BGP, a protocol already proven to scale to the entire internet, as its foundation, ensuring the overlay can support massive growth in endpoints and services across interconnected data centers.

Domain	Key Supported Technologies & Protocols from OcnOS-SP
Data Plane & Encapsulation	MPLS: Full stack including LDP, RSVP-TE, and MPLS Diffserv for QoS-aware transport. VXLAN: Standards-based overlay networking for massive Layer 2 domain extension.
Control Plane & Services	BGP: Rich feature set including BGP-LU, Add-Path, PIC Edge, and Inter-AS Options A/B/C. EVPN: Unified control plane for modern services including EVPN-MPLS and EVPN-VXLAN with full support for multihoming IRB, and E-TREE. L3VPN / L2VPN: Carrier-grade VPRN and VPLS/VPWS for robust, multi-tenant services.
High Availability & Resiliency	Segment Routing TI-LFA: Topology-Independent Loop-Free Alternate for comprehensive, automated FRR. BGP PIC Edge: Sub-second convergence for edge and core failures. BFD (Bidirectional Forwarding Detection): Low-overhead failure detection for all major routing protocols. ERPS: Carrier Ethernet standard for sub-50ms protection in ring topologies.
Traffic Engineering & OAM	Segment Routing (SR-MPLS): Granular, policy-based traffic control with SR Policy , Flexible Algorithms , and PCEP for path computation. RSVP-TE: Traffic engineering with bandwidth reservation, FRR, and auto-bandwidth. Advanced OAM: Comprehensive toolset including CFM (802.1ag) , Y.1731 (Performance Monitoring), and TWAMP for SLA verification.
Automation	Streaming Telemetry (gNMI): Granular, real-time data export for various platforms. OpenConfig & YANG: Standard model-driven configuration (NETCONF) and state reporting. NETCONF/RESTCONF: Robust APIs for full programmatic control and automation.

OcnOS-SP Core Capabilities

Mastering the Optical Layer: OcnOS Software Capabilities for IPoDWDM

For an IPoDWDM architecture to be truly operational at scale, the NOS must provide deep visibility and control over the underlying optical layer. OcnOS provides the deep software integration required to make the converged architecture manageable.

- **Low-Level Control (CMIS):** OcnOS utilizes the Common Management Interface Specification (CMIS), a standardized interface for managing complex pluggable modules. This allows network operators to query and configure key optical parameters directly from the router's software interface. This includes tuning

the laser frequency to a specific DWDM channel, adjusting the transmit output power, and monitoring critical performance metrics like the pre-Forward Error Correction (FEC) Bit Error Rate (BER).

- **Modern Automation (OpenConfig & gNMI):** OcNOS provides a fully model-driven management interface based on industry-standard OpenConfig YANG models for optical transport. This enables robust, programmatic configuration via NETCONF. Furthermore, it allows for the export of real-time performance data - for both the IP/MPLS layers and the underlying optical parameters - via gNMI (gRPC Network Management Interface) streaming telemetry. This provides a rich, granular dataset that can be fed into modern analytics platforms, allowing a single operator or automation system to correlate IP-layer events with optical-layer events for rapid root-cause analysis and predictive maintenance.

The OcNOS Ecosystem: Open Hardware and End-to-End Consistency

The value of a disaggregated NOS is realized through its broad and robust hardware ecosystem and its ability to provide a consistent operational experience across network domains.

Freedom of Choice: The Disaggregated Hardware Advantage

OcNOS is engineered to operate on a wide and growing variety of open, whitebox hardware platforms from leading Original Design Manufacturers (ODMs). This open ecosystem liberates operators from vendor lock-in, granting them the freedom to select the hardware platform that best meets their specific requirements for port density, throughput, power efficiency, and price-point. The table below provides a representative sample of the diverse, ZR/ZR+-capable platforms qualified to run OcNOS.

Model	Max ZR+ Ports	800GE QSFP-DD	400GE QSFP-DD	100GE QSFP28	Chipset
S9510-28DC	2 (ZR+)	-	2	2	Q2A
AS7946-30XB	2 (ZR+)	-	4	8	Q2C
S9600-56DX	8 (ZR+)	-	8	48	Dual Q2C
AS9947-36XKB	12 (ZR+)	-	12	24	J2C+
AS9726-32DB	16 (ZR+)	-	32	-	TD4
AS9736-64D	16 (ZR+)	-	64	-	TH4
S9321-64E	64 (ZR+)	64	-	-	Tomahawk 5
AS9817-64D	64 (ZR+)	64	-	-	Tomahawk 5

OcNOS Qualified ZR/ZR+ Enabled Whitebox Platforms

Beyond DCI: A Unified Network Operating System

The OcNOS software family extends this philosophy of consistency across the entire network. Operators can deploy OcNOS-DC to build highly scalable EVPN-VXLAN fabrics inside the data center for optimized east-west traffic flow. This internal fabric then seamlessly interconnects with the DCI nodes running OcNOS-SP.

Standardizing on the OcNOS family provides a consistent operational experience - a common Command Line Interface (CLI), a single set of APIs, and unified operational procedures - across all network domains. This reduces administrative overhead, simplifies training, and enables the creation of a single, end-to-end automation pipeline.

Conclusion: Building a Future-Proof DCI Foundation for the AI Era

By embracing a disaggregated model powered by OcNOS, cloud and service providers can strategically move

away from the rigid, expensive, and operationally complex legacy DCI architectures of the past. The convergence of IP and optical layers through IPoDWDM, made viable and manageable by the deep software capabilities within OcNOS, delivers a demonstrably simpler, more powerful, and economically superior network. This modern architecture provides a future-proof foundation, ready to support the most demanding applications of today and tomorrow.

The flexibility to choose best-in-class hardware, the scalability of modern control plane protocols like SR-MPLS and EVPN, and the efficiency of a collapsed IP-optical architecture are the essential building blocks for the AI era. This approach equips operators with the agility and performance necessary to support distributed AI/ML training fabrics, immersive cloud services, and the next generation of data-intensive applications. The comprehensive, carrier-grade feature set within OcNOS-SP provides the technical underpinnings for this transformation, offering a smarter path to building a more powerful and future-ready network.

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
