WHITE PAPER

Zero-trust optical transport: A tiered architecture for metro network resilience

Abstract

Modern networks demand more than simple redundancy; they demand true resilience engineered into every layer. Microsoft's zero-trust optical business continuity and disaster recovery (optical BCDR) architecture embodies this shift: It moves beyond traditional backup paths and instead combines two fully independent optical systems designed to sustain uninterrupted services even during systemic failures.

Disruptions in complex optical networks result not just from hardware faults but from human error, misconfigurations, automation glitches, or unexpected maintenance impacts, any of which can disable even the most robust transport infrastructure. To address this issue, Microsoft's tiered zero-trust optical architecture couples two completely independent optical domains: a flexible ROADM-based transport network and an optical BCDR layer, tied together only at the routed Ethernet edge.

Built on proven Ciena platforms such as the 6500 Packet-Optical Platform, 6500 Reconfigurable Line System (RLS), and the Waveserver® family equipped with fixed-grid channel multiplexer/demultiplexer (CMD) modules, this approach ensures operational continuity during planned or unplanned outages. By logically bonding these domains at the Ethernet layer, it guarantees seamless failover and uninterrupted client services.

With this blueprint, service providers, cloud operators, and enterprise architects can move beyond simple redundancy and design networks that guarantee continuity, shaping a new benchmark for resilient optical transport.

1. Evolving risks in optical transport

The new reality of metro optical networks

Optical networks are more intelligent, faster, and programmable than ever. Automation and software control have provided unmatched flexibility and scale. But with this progress comes added risk. One misconfigured automation rollout, a fiber cut and equipment failure, unintended policy change, or mismatched software version can take down an entire optical domain. Events like these do not happen often. However, their effects can be severe. Thus, making sure that network designs consider such events is critical. ROADM systems might be equipped with layer 0 (L0) control plane restoration to improve resiliency, but restoration can take seconds or more depending upon network configuration.

2. Tiered metro optical architecture: Overview

To overcome these concerns, the optical BCDR model introduces a second, independent optical system that runs in parallel with the ROADM network. Unlike the reactive approach of L0 control plane restoration, the optical BCDR link is integrated alongside the ROADM paths at the Ethernet edge, enabling active-active forwarding. If a ROADM degree, or even the entire system, goes down, traffic continues flowing uninterrupted through the BCDR link, without any delay or dependence on reactive control plane actions. With a separate overlay network, the optical BCDR architecture enhances traditional ROADM systems, providing reliable protection and always-on connectivity for critical traffic, AI workloads, and latency-sensitive applications.

This tiered architecture doesn't replace the inherent resiliency of ROADM systems; instead, it adds a layer of separation that keeps services running even during rare, high-impact events. It changes the mindset from responding to outages to prioritizing uninterrupted service.

The proposed design consists of:

- A primary ROADM-based transport layer, built with 6500 or RLS photonic line systems
- A completely independent, secondary optical BCDR transport layer, built using Waveserver with fixed-filter CMD modules
- Layer 2 (L2) Ethernet bonding at the client interface, unifying both systems for seamless failover

Both optical systems are always live and active, each delivering production services through independent fibers, shelves, racks, power, and software stacks.

Photonic layer	Hardware
ROADM transport 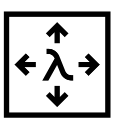	Flexible photonic layer (6500, RLS, or Coherent ELS) with compact modular transport (Waveserver) powering 400G/800G/1.6T wavelengths  6500  RLS  Coherent ELS
Optical BCDR layer (CMD-based) 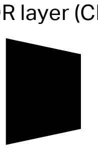	Completely separate Waveserver platforms with integrated fixed-grid CMD modules (wavelengths and photonic layer all in a single platform)  CMD10 Waveserver  Waveserver CMD12
Client integration 10G 100G 400G	Client interfaces from the ROADM transport and optical BCDR networks connect into layer 3 (L3) routers—multiple ports are bundled into a single logical interface with L2 Ethernet bonding

3. Logical coupling at the Ethernet edge: How it works and why it matters

A foundational element of the zero-trust optical BCDR architecture is the logical bundling of client-facing services across two physically and operationally independent optical domains.

Unified client interface

Both the primary ROADM-based system and the optical BCDR system deliver services at multiple bandwidth tiers (10G/100G/400G). Each system's client interfaces terminate at L3 routers or switches. These diverse physical links are then logically bonded using standard EtherChannel (port-channel) techniques.

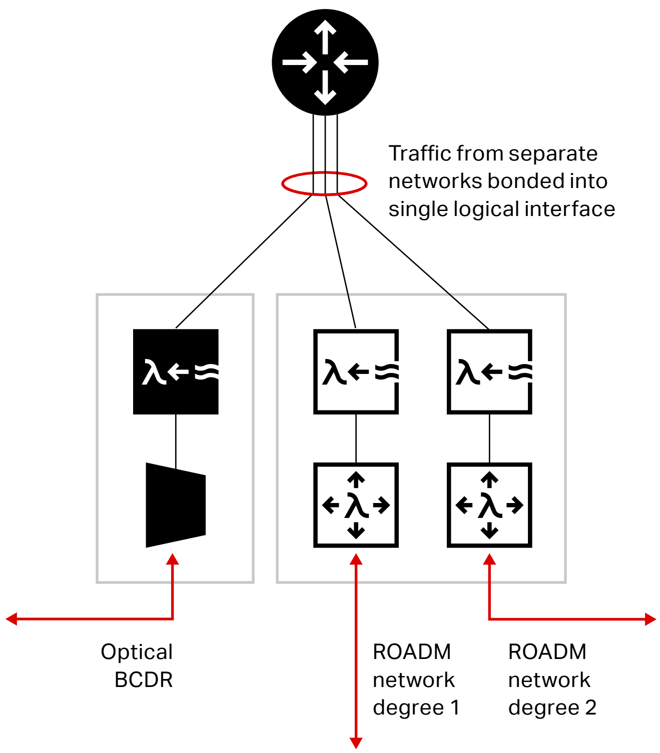

Figure 1. Multiple physical client interfaces bonded into a single logical port

From the client or application perspective, this means:

- One logical interface, regardless of which optical system is used to carry the bits
- Automatic load balancing and traffic distribution, improving overall link utilization
- Seamless failover with enough capacity to ensure zero impact if either the primary ROADM-based system or optical BCDR system fails

Because this coupling happens at the Ethernet layer, there is no need for optical-level reconvergence or OTN switching. The port channel automatically handles forwarding continuity, allowing WAN traffic to shift between independent optical paths in real time. This approach empowers clients to maintain deterministic performance with active-active optical layers underneath, reducing operational complexity while increasing overall resilience.

This design principle is a key takeaway for any enterprise or service provider with mission-critical network traffic: Logical coupling at the client edge makes independent systems act as one, combining their strengths while avoiding single points of failure.

4. Why zero trust at the infrastructure level?

The term "zero trust" is often applied to identity and security. Here, it refers to infrastructure isolation. Not a distrust of platform integrity, but a recognition that complex systems can be unintentionally taken offline through external factors.

By assuming such incidents can occur, and architecting for continuity, the network shifts from reactive recovery to proactive survivability.

Examples of zero-trust failures:

- An automation rollout accidentally zeroes out a ROADM configuration
- A rollback process inadvertently affects active channels
- A provisioning task disables optical services across multiple nodes
- Fiber assignments are overwritten or misrouted at patch panels

5. Understanding how to build a resilient metro ROADM system

A resilient metro ROADM network starts with proven platforms. Ciena's 6500 is a robust, widely deployed, and trusted optical platform, and the next-generation RLS adds advanced photonic features and scalability. Both systems support flexible ring or mesh topologies.

In a typical metro deployment, a ring topology with colorless direct attach (CDA) ROADM nodes ensures each site has at least two degrees for incoming and outgoing channel routing. For maximum agility, core or major edge sites can be built with colorless-directionless-contentionless (CDC) ROADM nodes, while tail or client sites use simpler CDA configurations. This configuration balances flexibility and cost.

= Colorless-directionless-contentionless ROADM

= Colorless direct attach ROADM

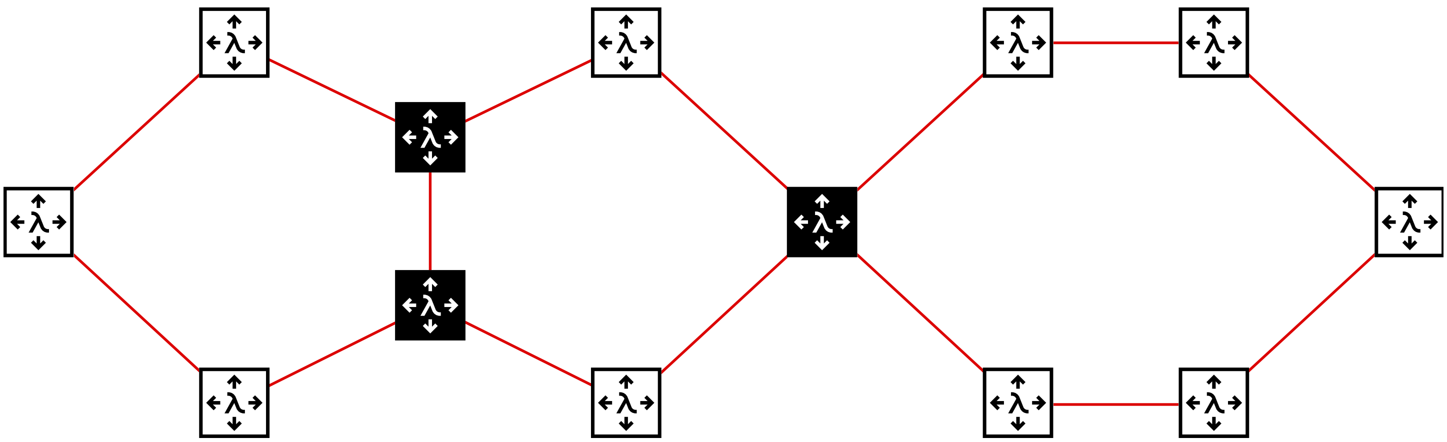

Figure 2. Example metro ROADM–based photonic layer network

Microsoft's approach:

Vinoth Elangovan, Microsoft tiered optical network architect, designed a hybrid architecture with core sites using CDC and tail sites using CDA. This gives the network future-ready scalability and maximizes efficiency for high-bandwidth services.

6. Building the optical BCDR network

The optical BCDR layer adds another dimension of resilience. Using Ciena's Waveserver family of platforms with CMD10 or CMD12 modules, this CMD-based system operates independently of the ROADM system. The optical BCDR network uses a simple, point-to-point line system that is integrated into the Waveserver platform, and, as such, the optical BCDR nodes use a different software stream than the RLS ROADM nodes. This separation protects network traffic from issues that might arise during maintenance of the ROADM network, such as failed automated software upgrades or misconfigured provisioning scripts, as these will not impact the optical BCDR network.

Each CMD module provides an optical path with fixed grid add/drop capability and can be tailored to match the required Ethernet capacity. Options include single or multiple 400G lambdas or an 800G configuration.

To further isolate the optical BCDR network from the ROADM network, its equipment is installed in different racks, using different power feeds. The fiber routes are also independent, so fiber cuts in the ROADM network will not affect the optical BCDR system and vice versa. By maintaining diversity of equipment and facilities, the optical BCDR is unaffected by issues that might arise during maintenance and operation of the ROADM network.

This is an innovative application of Ciena's compact, modular Waveserver and CMD portfolio. By coupling the optical BCDR and ROADM systems logically at the Ethernet layer, customers reduce the operational risk of catastrophic failures and ensure that their high-criticality metro sites remain always on.

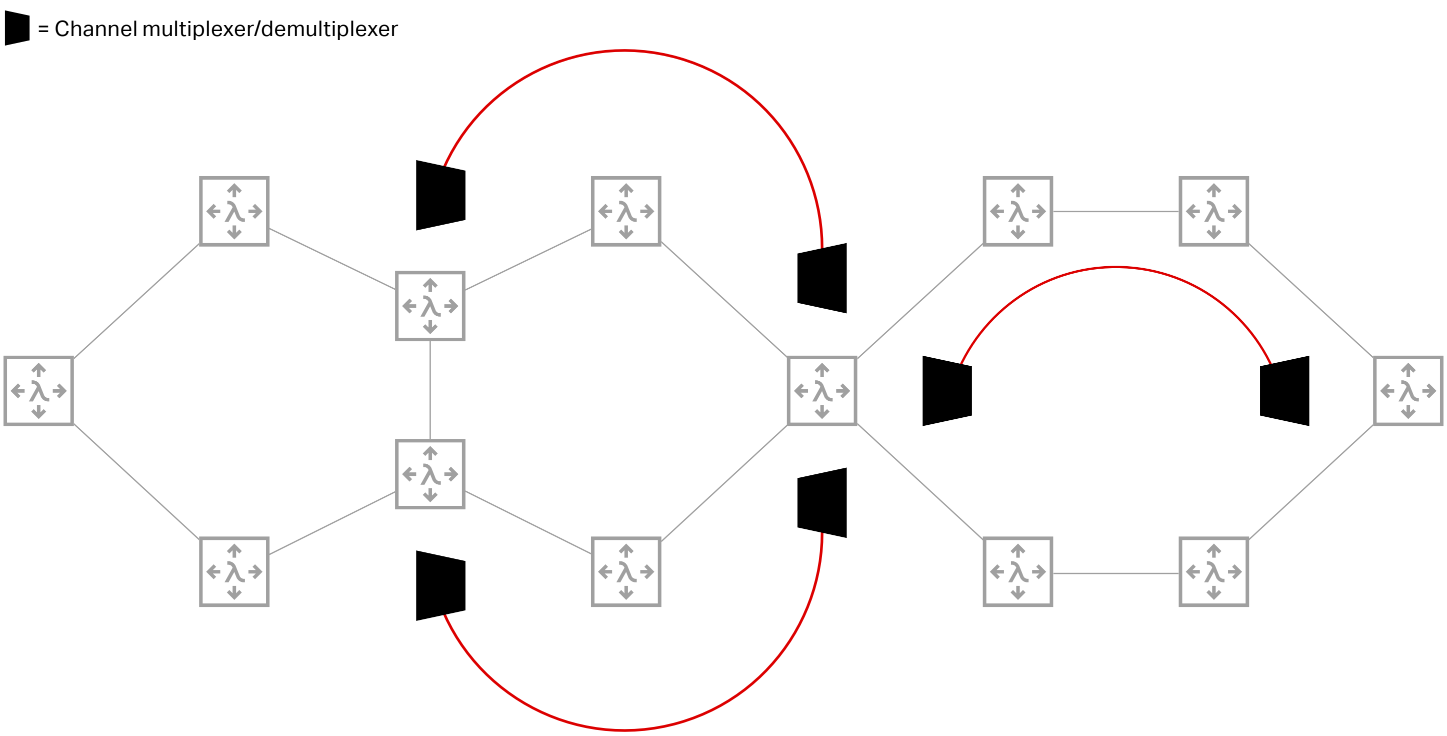

Figure 3. Example of an optical BCDR (CMD-based photonic) layer

How Microsoft implemented tiered optical architecture:

Vinoth designed the architecture with a fixed-filter CMD network to overlay critical segments of the metro ROADM system. Critical sites were carefully selected so that, even in the event of a major ROADM failure, production traffic continues uninterrupted, and clients experience no disruption. This active-active design gives network engineers the breathing room to restore the ROADM layer without business interruption.

7. Key architectural benefits

Site migration and metro expansion

A metro optical network is further safeguarded by an additional protection system through the use of optical BCDR. It is always active-active in case something happens so services can continue. When a site migration occurs and an optical node needs to be moved elsewhere, migration can be complicated and usually requires planned downtime for the ROADM network. With the optical BCDR layer in place, traffic continues flowing, as the primary ROADM system is prepared and migrated from the old site to the new site. Business operations and user services continue without interruption.

Similarly, when the metro design needs to expand and a new CDA or CDC optical node is added, planning, configuration, and execution can be complex and could take days to complete. This process includes integrating the new node with core or tail sites and bringing services to production. Here, the optical BCDR layer provides an extra active leg for the participating sites, ensuring service continuity while the new addition is executed.

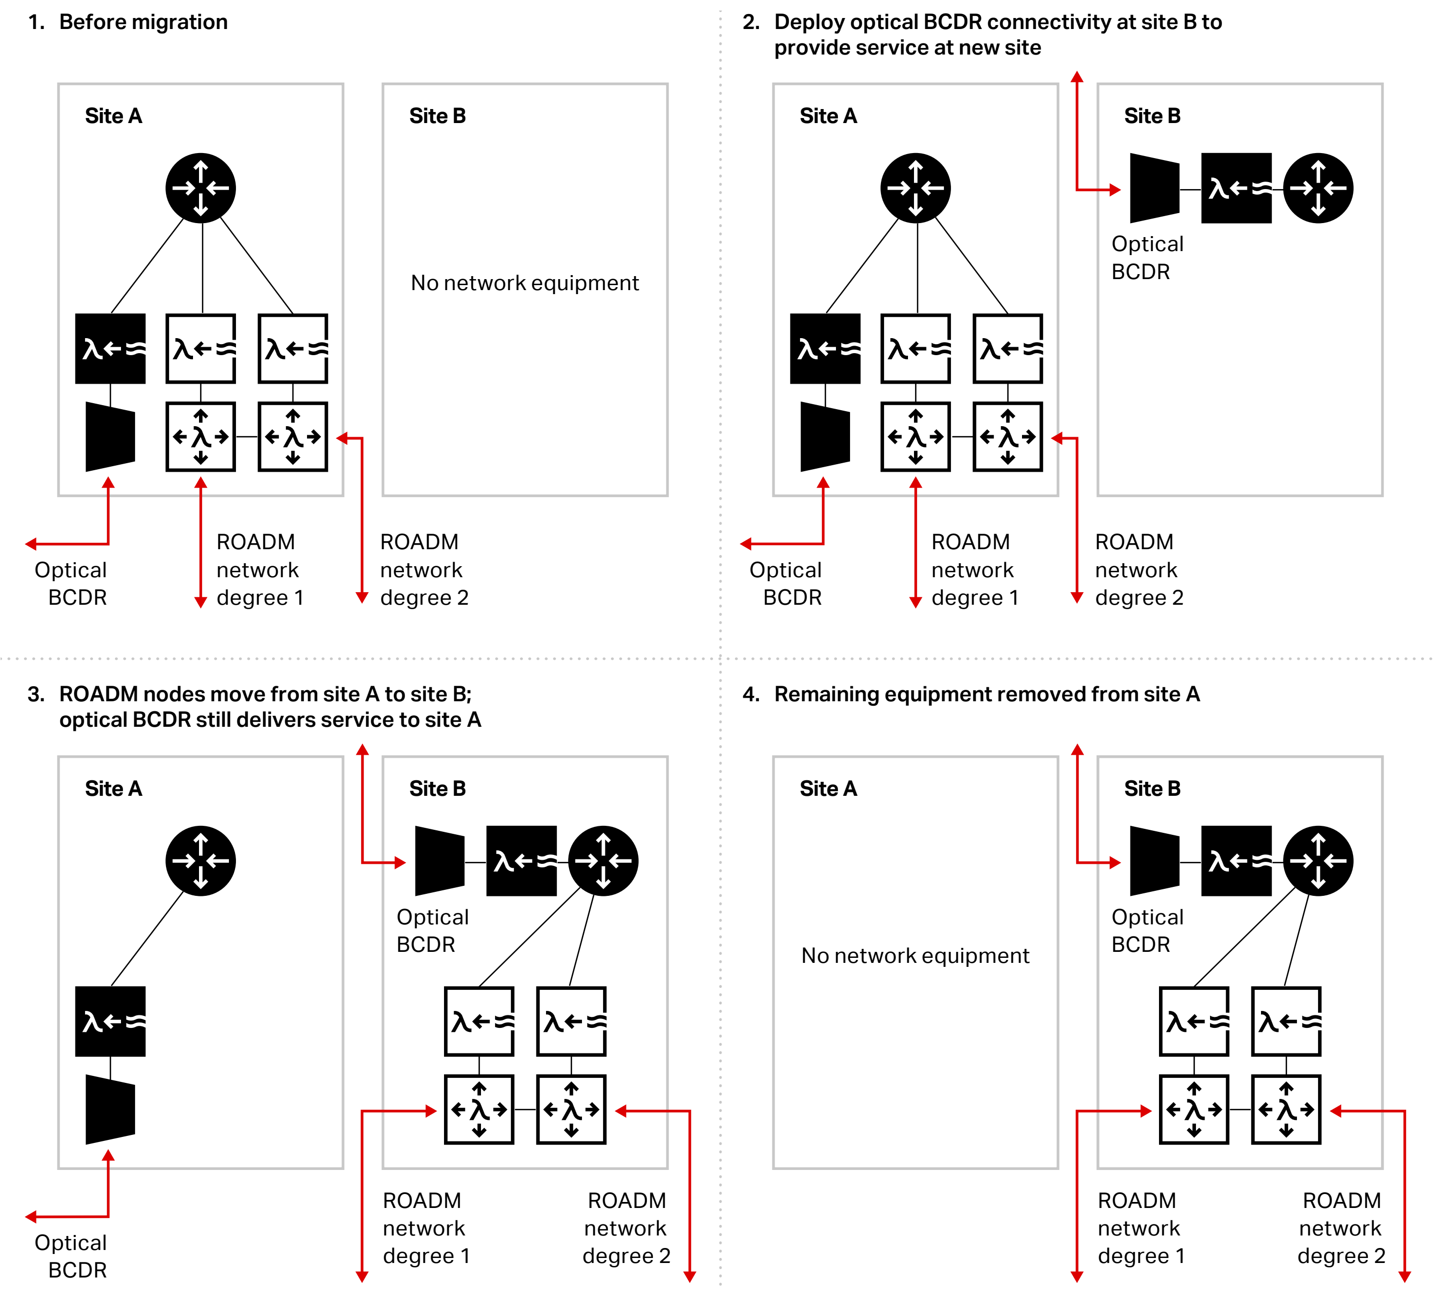

Figure 4. Optical BCDR system maintains connectivity while the ROADM network is migrated from sites A to B

Node upgrades and configuration changes

The optical BCDR layer plays a critical role in maintaining service continuity during node upgrades or configuration changes on the ROADM network. Even with careful planning, unexpected issues such as a failed software rollback, misapplied policy, or other operational errors can impact the primary ROADM system. Because the optical BCDR operates as a fully independent optical layer with its own power and control, traffic is automatically directed to it, ensuring uninterrupted service. By preferring the optical BCDR path during these events, applications, customers, and internal teams continue to experience seamless connectivity even when the primary ROADM network faces disruptions.

This architecture ensures that service continuity and bandwidth parity are maintained across both the transport planes. Clients can terminate multiple bandwidth levels—10G/100G/400G—into a single logical port channel, with one or more links sourced from the ROADM line system and the rest from the optical BCDR system.

This balanced approach ensures that even under outage conditions affecting a single optical system, full-capacity services remain operational with no manual intervention and no traffic loss.

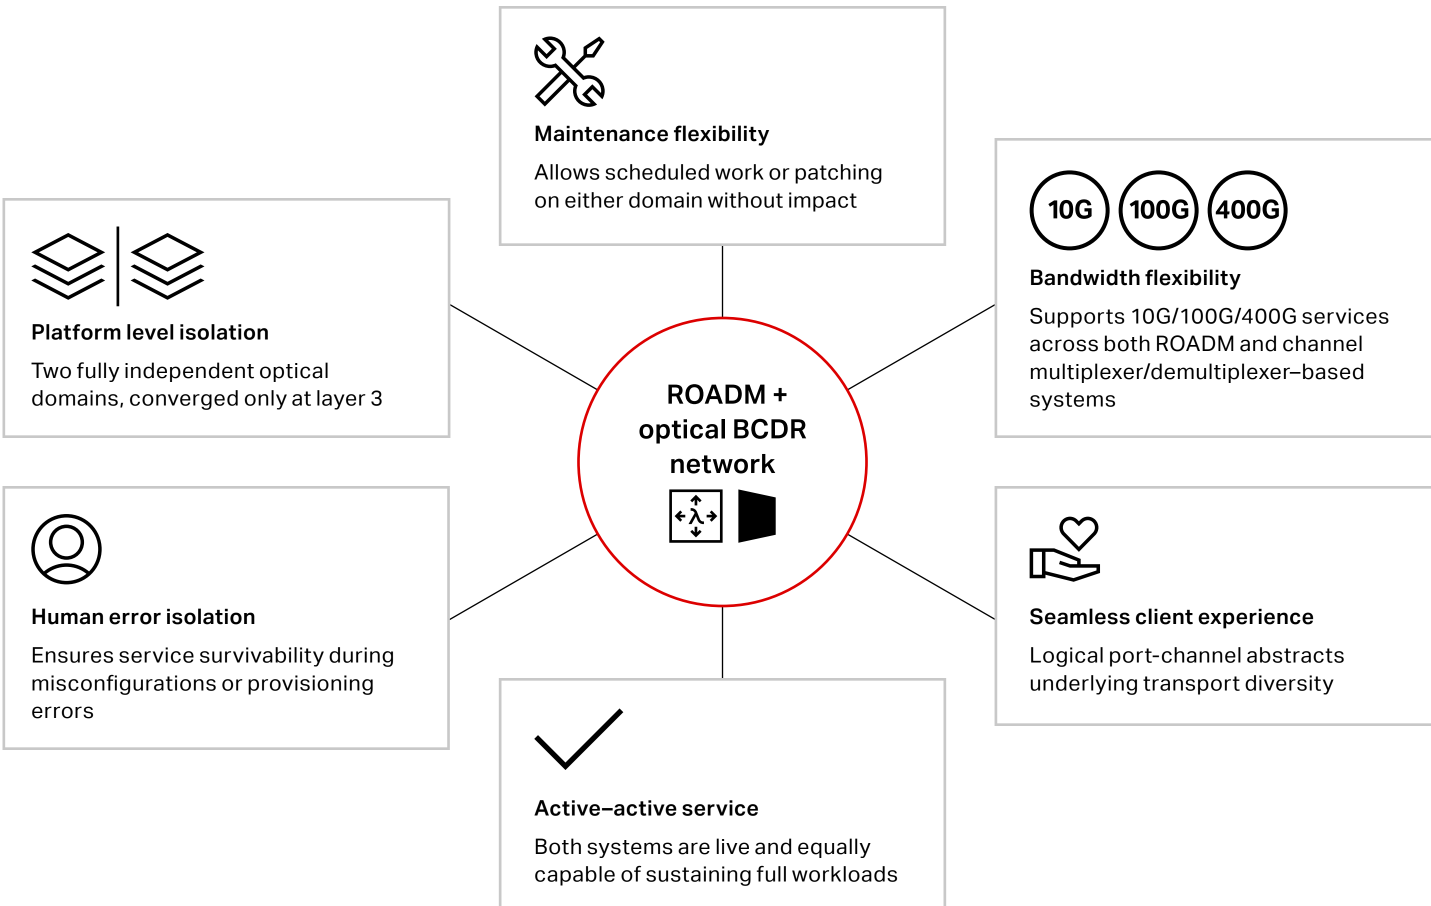

Figure 5. Key benefits of the tiered ROADM + optical BCDR architecture

8. Deployment scenarios

This model is especially suitable for the following networking applications.

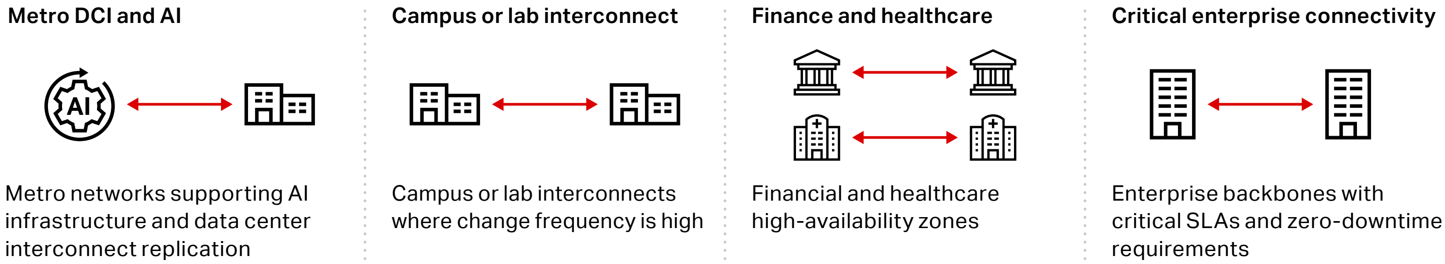

Figure 6. Optical BCDR architecture is ideal for networks requiring high availability

Operators can selectively apply this model to high-priority geographies or customer segments, adapting system configuration and bonding logic to local needs.

9. Design principles recap

The zero-trust optical BCDR transport architecture adheres to five core design pillars:

1. Independent optical systems: ROADM and CMD platforms autonomously operate
2. Physically diverse paths: circuits avoid shared conduits, fibers, or splices
3. No shared control software: separate network management system (NMS), automation, and provisioning domains
4. Unified client interface: services terminate into a logical L3 EtherChannel
5. Survivability by design: not assuming systems are immune, but building for outcomes

10. Optimizing metro network resiliency with symmetric and asymmetric paths

To ensure resiliency and performance, it will be necessary to use a combination of symmetric and asymmetric paths to design modern metro optical networks with an optical BCDR overlay. In symmetric designs, where the fiber distances between the tail and core sites are similar, all ROADM paths and the optical BCDR link are configured active-active. Traffic can be distributed equally over the paths, thus achieving the immediate failover without any delays or disturbance.

For designs that are asymmetric, meaning one path is shorter and the backup path is longer, the shortest ROADM path and the optical BCDR are set to active and the longer one to passive. This configuration lowers latency under normal circumstances and allows for seamless redundancy during failures. The optical BCDR layer is designed and configured to be useful at all times because of its low-latency, point-to-point design. As a result, the optical BCDR ensures smooth traffic flow with service continuity. BCDR design becomes especially useful during large ROADM failures.

11. Conclusion: Designing for continuity under uncertainty

Modern networking means planning for the worst while delivering the best. Unexpected interruptions to mission-critical traffic and applications can impact revenue opportunities. Microsoft's zero-trust optical BCDR blueprint shows how to combine proven Ciena ROADM and optical BCDR solutions to achieve unmatched metro resilience. This approach shifts the focus from trusting any single system to trusting a design that anticipates and absorbs failures at any scale.

Modern optical platforms like the 6500 and RLS are foundational to scalable transport. But no system is immune to operational realities: Human mistakes, automation errors, or rollback failures can disrupt even the best-built platforms. In these rare but high-impact scenarios, resilient network design, not just redundant hardware, keeps services running.

The zero-trust optical BCDR tiered architecture is engineered for these moments. By layering fully independent optical domains and logically unifying them at the Ethernet edge, the network can sustain a complete failure of one system while maintaining service continuity.

By combining ROADM and CMD fixed-filter photonic networks based on Ciena's proven photonic layer technology, this architecture offers an innovative approach to improving network resiliency against maintenance and operational outages that might occur due to automation or human error. Organizations with business-critical traffic can adopt this tiered architecture as part of their next-generation transport network strategy to offer highly resilient connectivity or business continuity / data recovery services between locations in metro area networks.

Trust the design—not just the system.

About the author

Vinoth Kumar Saalai Elangovan
Senior Network Engineer, Microsoft
Author and principal architect of the tiered optical network design

Vinoth Elangovan is the architect behind this tiered optical design. His work focuses on building resilient, scalable optical architectures that can withstand both systemic disruptions and operational failures.

Contributors

Ren Hamil
Technical Account Manager, Ciena

Ren Hamil is part of the Ciena account team for Microsoft. His work focuses on building customer networks and architectures for hyperscale customers.

Kent Jordan
Advisor, Product and Technology Marketing, Ciena

Kent Jordan is part of the product and technology marketing team at Ciena with a focus on coherent optical networking and data center interconnect solutions.

Acknowledgments:

Heng Foo
Senior Director, Product Line Management, Ciena

Blaine Martin
Principal Engineering Manager, Microsoft

The author and contributors thank Heng Foo and Blaine Martin for their guidance and technical review of this work.