Technical Solutions Guide

Wavelength for the Modern Enterprise

A high-capacity optical transport and resilient design.

© 2026 Lumen Technologies. All Rights Reserved.

LUMEN®

Table of contents

Chapter 1: Executive summary & the rise of east-west traffic……..…..3

Chapter 2: What a wavelength service provides ……..……5

Chapter 3: Optical networking fundamentals……..……..7

Chapter 4: Wavelength architecture for Data Center Interconnect…….....9

Chapter 5: Designing for resiliency, and protection vs. diversity…...…..…11

Chapter 6: Capacity, latency, and traffic engineering……17

Chapter 7: Comparing connectivity options……...19

Chapter 8: Accelerated deployment with pre-engineered routes...……21

Chapter 9: Getting to service – what to expect and how to prepare…..23

Glossary……26

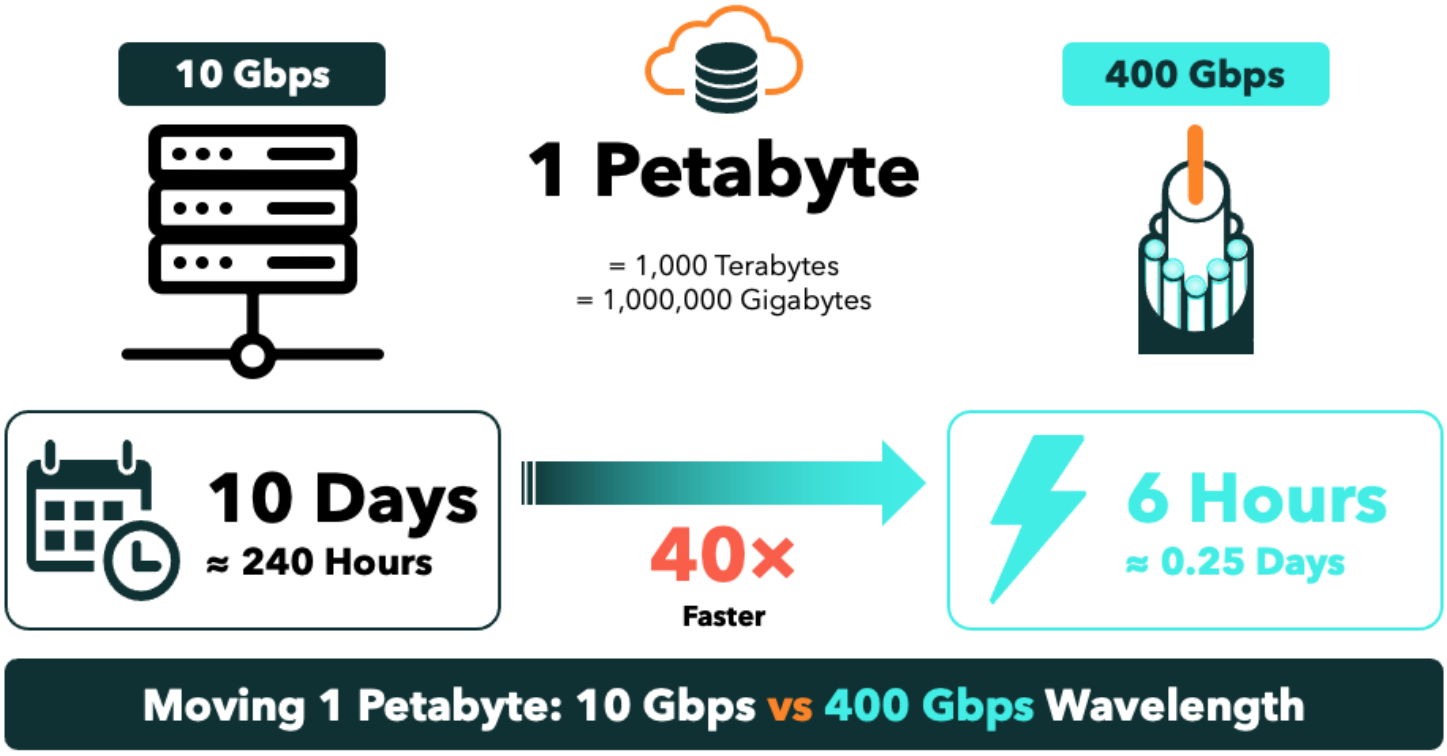
© 2026 Lumen Technologies. All Rights Reserved.

Chapter 1: Executive summary and the rise of east-west traffic

Modern applications are increasingly distributed across data centers, cloud environments, colocation facilities, and specialized compute locations. The result is growth in **east-west traffic**—data moving laterally between infrastructure locations rather than simply North-South from a user to an application. Replication, backup, storage synchronization, application clustering, data-pipeline movement, and AI workflows require large amounts of predictable bandwidth between fixed endpoints.

As enterprises rapidly adopt artificial intelligence, machine learning, and hybrid-cloud architectures, the volume of data moving across corporate networks has reached unprecedented levels. Traditional IP and Ethernet networks, while highly capable for general business traffic, can become a constraint when tasked with moving petabytes of training data or supporting ultra-low-latency distributed applications.

The scale of this challenge is stark. By 2030, it is projected that AI models will be trained on datasets comprising up to 162 trillion words [1], approximately 38,500 times more data than all the words currently in Wikipedia (ARK Investment Management LLC. *Big Ideas 2023: Innovation Trends with Exponential Growth Potential*). Moving just one petabyte of data across a standard 10 Gbps connection takes approximately ten days; across a 400 Gbps wavelength, that same transfer takes around six hours [2].



© 2026 Lumen Technologies. All Rights Reserved.

For organizations navigating these new demands, optical networking—specifically wavelength services—has emerged as a practical starting point. Wavelength services provide dedicated, high-capacity, point-to-point fiber connectivity that operates at the physical layer of the network.

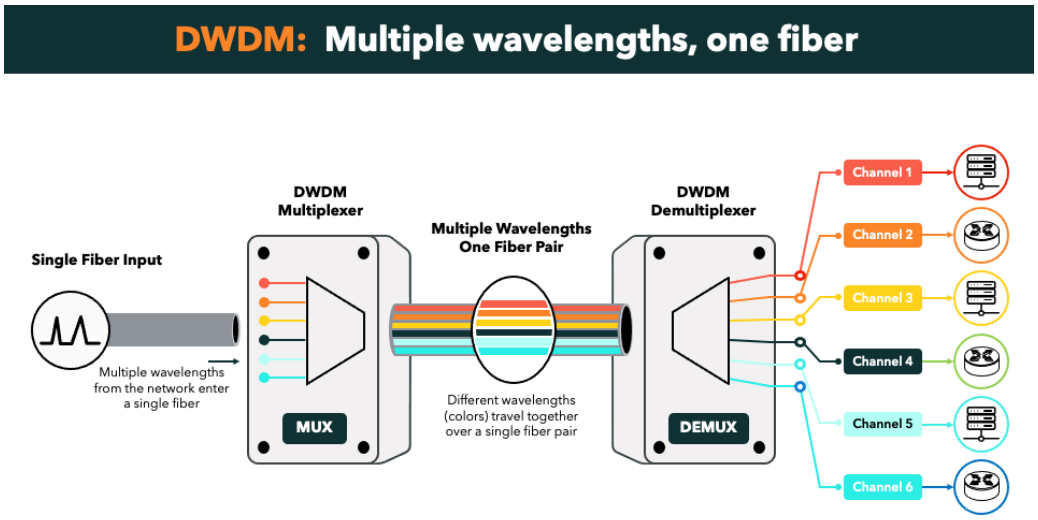
Lumen stands at the forefront of this optical evolution. Recognition from Frost & Sullivan and Vertical Systems Group, combined with broad network investment and next-generation optical expansion, positions Lumen to support large-scale enterprise connectivity requirements. The underlying fiber and optical platform matter: modern fiber with lower loss characteristics can support higher-capacity optical transmission over greater distances and help reduce the need for intermediate signal regeneration. Enterprises should validate current network reach, route availability, technology support, and service terms for their specific locations and capacity requirements.

Key Takeaway: The era of AI and distributed applications demands optical-scale capacity. Traditional WAN architectures may not be sufficient for massive east-west data movement; enterprise leaders should evaluate dedicated, high-capacity wavelength services as part of their long-term connectivity foundation.

© 2026 Lumen Technologies. All Rights Reserved.

Chapter 2: What a wavelength service provides

A wavelength service is a dedicated optical channel between two endpoints. Dense Wavelength Division Multiplexing (DWDM) technology allows multiple optical channels, each using a different wavelength—or color—of light, to travel over a single fiber pair.



Lumen Wavelength Solutions® provide private Layer 1, point-to-point connectivity with low-latency, high-capacity operation. Unlike shared or packet-based services, such as MPLS or SD-WAN, wavelength services provide dedicated Layer 1 connectivity. This means the data travels on its own isolated lane of light, supporting consistent, high-throughput performance without contention from unrelated traffic.

Key characteristics

Characteristic	Business Benefit
Dedicated Bandwidth	A dedicated channel supports predictable performance between defined endpoints without unrelated traffic sharing the service capacity.
Ultra-Low Latency	Direct optical routes can reduce network hops and packet-processing delays. Actual latency depends on physical route, distance, regeneration, and the applicable service design.
Protocol Transparency	Layer 1 transport is generally transparent to the customer's higher-layer protocol design.
High-Capacity Scalability	Wavelength options can scale from 10 Gbps up to 400 Gbps, subject to route and service availability, to support growing DCI and backbone requirements.
Architectural Flexibility	Unprotected, protected, diverse, and dual path architectures address different failure-response, capacity, and operating requirements.

© 2026 Lumen Technologies. All Rights Reserved.

What the customer still owns

A wavelength provides the transport foundation, but it does not replace the customer's LAN, storage, compute, application, or security architecture. The customer typically remains responsible for endpoint routers or switches, interface configuration, application recovery, data protection, and any required higher-layer encryption, unless an encrypted wavelength option is explicitly selected and validated for the route.

Because wavelengths operate at Layer 1, they do not include native IP routing logic. An **unprotected wavelength** cannot switch or route around a problem the way a Layer 2 or Layer 3 network can; if the service path fails, the service is interrupted until the path is restored or the customer moves traffic through another design element. A **protected wavelength**, by contrast, includes provider-managed optical protection switching within one logical service. In every architecture, the customer remains responsible for the logical network and application behavior above the optical transport layer.

Key Takeaway: Wavelengths enable predictable economics, high-capacity transport, and a dedicated foundation for modern DCI. A successful deployment depends on understanding the demarcation between the carrier's optical service, the customer's logical network, and the specific failure response the selected architecture provides.

© 2026 Lumen Technologies. All Rights Reserved.

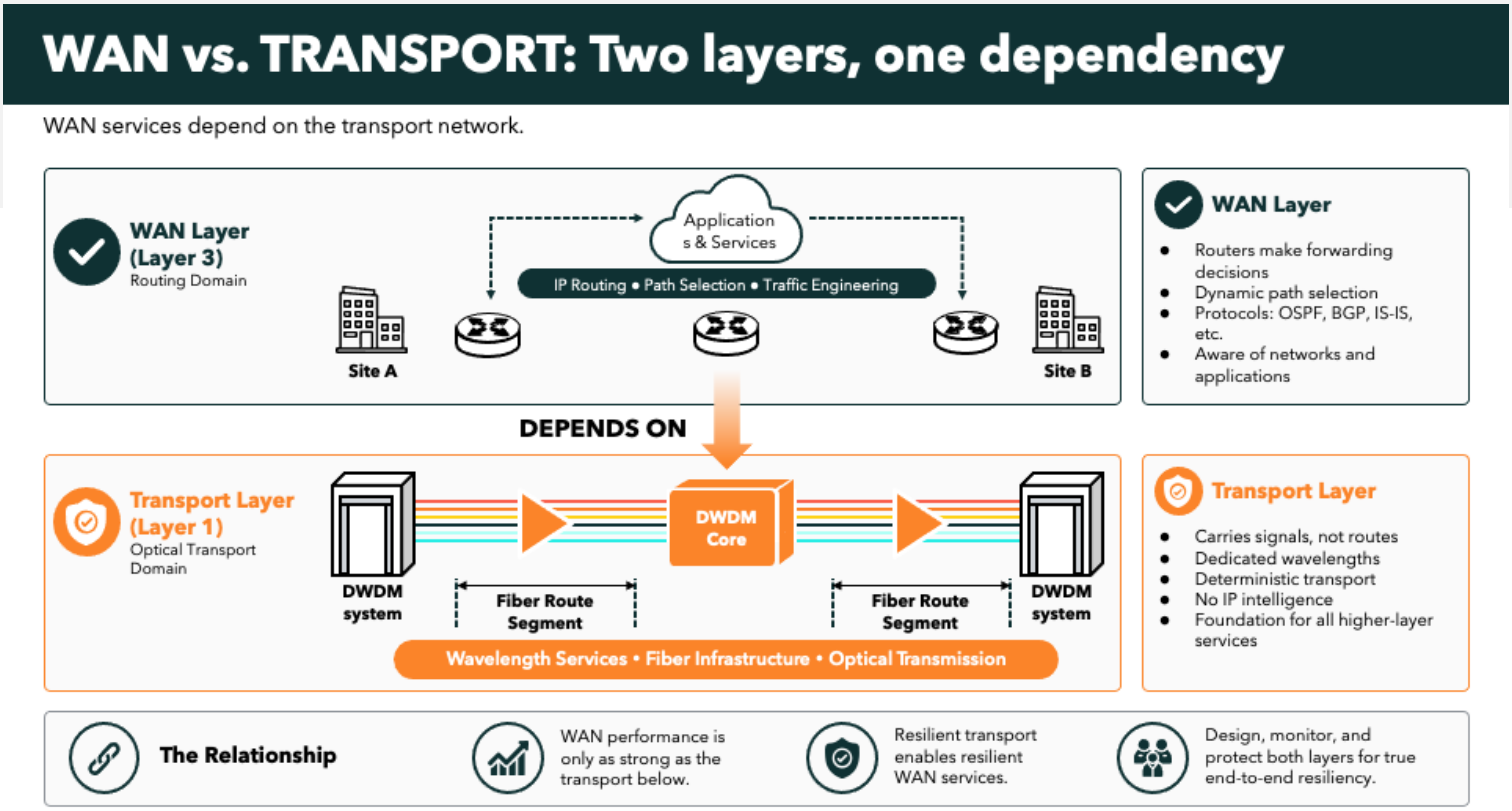
Chapter 3: Optical networking fundamentals

To effectively design and procure wavelength services, it is helpful to understand the underlying technology. Modern optical networking is built on DWDM, a technology that revolutionized telecommunications by allowing multiple data streams to share a single physical fiber.

WAN vs. Transport

A critical distinction in network architecture is the difference between the Wide Area Network (WAN) and the transport layer. In practical terms, the WAN refers to the router network, generally operating at Layer 3, while transport denotes the underlying DWDM optical network, operating at Layer 1.

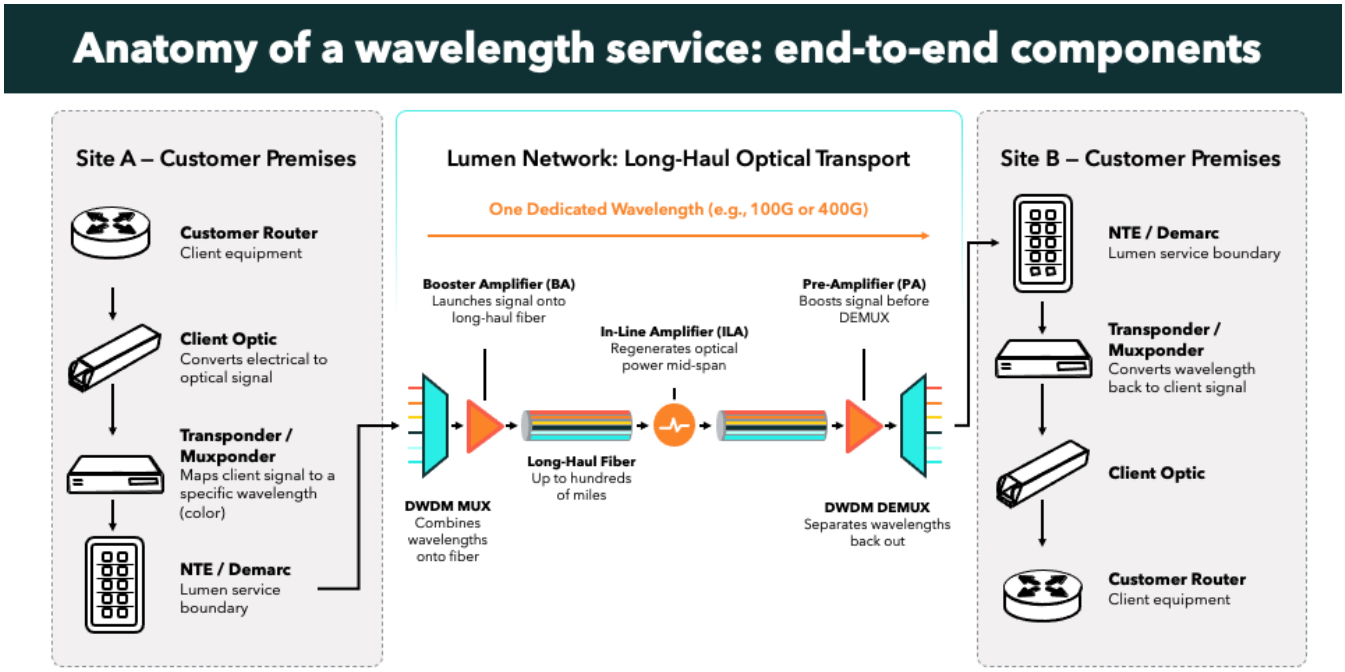
Wavelength services are transport services. They provide the foundational optical paths that connect routers across different markets. This distinction is vital because transport infrastructure must be built, tested, and operational before higher-layer WAN services can be delivered over those paths. Because the transport layer lacks native routing logic, any self-healing behavior must be explicitly engineered either into the wavelength design through optical protection or into the customer's WAN, storage, or application architecture.



© 2026 Lumen Technologies. All Rights Reserved.

Client-side vs. Line-side optics

DWDM systems involve two primary interfaces. The client-side interface is where the customer's equipment connects to the DWDM platform. This is often referred to as the *grey* optical interface because it uses standard, uncolored light. The line-side interface is where the DWDM system connects to the long-haul optical fiber. At this point, the grey light has been converted into a specific, precisely tuned wavelength and multiplexed with other precisely tuned wavelengths for transmission across the DWDM core.



The customer owns the router. Lumen owns everything from the NTE outward – the transponders, the fiber, the amplifiers, and the DWDM platform.

Protocol Transparency

One of the advantages of DWDM is protocol transparency. Wavelength systems can transport various data types, including IP, Ethernet, and storage protocols such as Fiber Channel, without altering or interfering with the original framing. This allows enterprises to use existing protocols without modification. Lumen Wavelength Solutions® inherently support a Maximum Transmission Unit (MTU) of 9600 bytes, accommodating the 9000-byte jumbo frames commonly used in high-performance computing and storage environments, subject to applicable service specifications.

Key Takeaway: Wavelengths are the foundational transport layer of the network. Understanding the distinction between transport and WAN is critical because robust DWDM infrastructure must be in place—and appropriately engineered for the intended failure behavior—before higher-layer services can operate as designed.

© 2026 Lumen Technologies. All Rights Reserved.

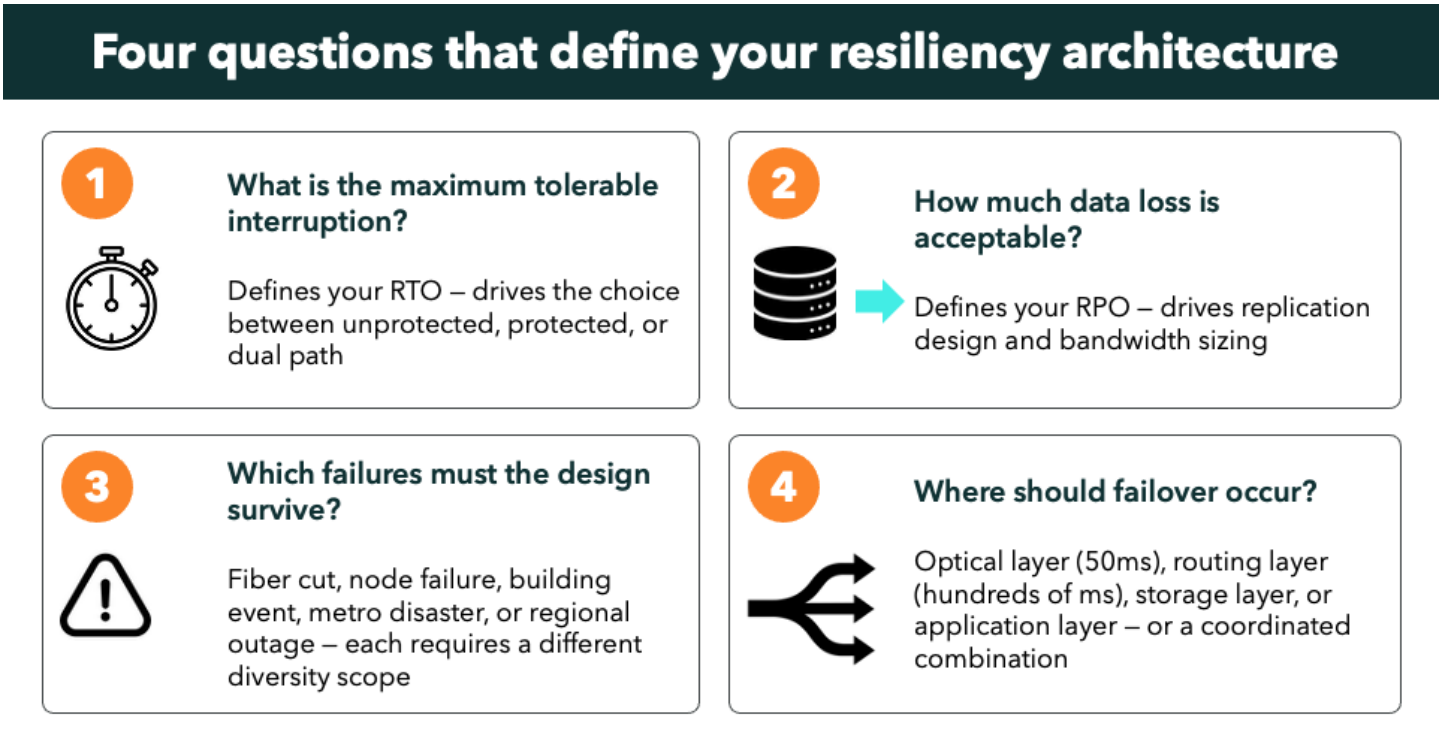
Chapter 4: Wavelength architecture for Data Center Interconnect (DCI)

In the simplest design, customer equipment at Data Center A connects to a Lumen wavelength endpoint. The optical service transports the client signal to Data Center B, where it is handed to the destination equipment. The customer controls the logical network and applications above Layer 1.

The DCI design question

The central question in Data Center Interconnect (DCI) is not simply, how much bandwidth is needed. A complete DCI design reframes the conversation around business outcomes.

- What data must move, in which direction, and at what sustained rate?
- What latency and jitter can the application tolerate?
- How long can the service be unavailable before the business is materially impacted?
- Does the application fail over automatically, or does it depend on the network to switch paths?
- Which physical facilities and network elements can be shared without violating the recovery objective?
- How will demand grow during the service term?



Not every workload needs the same design – match the architecture to the business impact of an interruption.

© 2026 Lumen Technologies. All Rights Reserved.

Translating features to business outcomes

Customers usually ask outcome-focused questions. The table below translates those questions into the architecture conversation required to answer them accurately.

The Buyer Asks...	The Architecture Conversation
"What happens if a fiber cut occurs?"	A protected service can use provider-managed optical protection switching. Diverse and dual path designs rely on customer equipment or an operating process to move traffic between separately usable circuits.
"Can this support active-active data centers?"	Two active circuits can be available to carry traffic at full bandwidth. The customer's routers, switches, load balancers, or applications determine how traffic is distributed and how it moves following a failure.
"Can I separate paths for compliance?"	Diversity can be requested across specific physical and operational domains. The final design should document which domains are separated, which are shared, and which require site-specific validation.
"Can I avoid a particular route or facility?"	Yes. Avoidance documents a specific route constraint, but it does not by itself establish a protected, diverse, or resilient design.
"Do you support zero-touch failover?"	A protected wavelength provides automatic optical-layer switchover within the service. End-to-end application recovery still depends on endpoint behavior, capacity, timers, and application state.

Key Takeaway: Wavelengths are adaptable, serving as direct point-to-point DCI products. A successful architecture answers not only the bandwidth question, but also the application dependency, failure behavior, separation, and recovery requirements that determine business continuity.

© 2026 Lumen Technologies. All Rights Reserved.

Chapter 5: Designing for resiliency, and protection vs. diversity

Designing for uptime requires a clear understanding of concepts that are often conflated. Reliability, redundancy, diversity, protection, and resiliency are related, but they are not interchangeable. Each describes a different part of the customer's architecture and failure-response story.

The five terms that clarify a resiliency design

Term	Working Definition	The Question It Answers
Reliability	How consistently a service performs without interruption over time.	How often does the service operate without failure?
Redundancy	Duplicate capacity, component, circuit, or path.	What additional resources exist?
Diversity	Separation of specified physical or logical dependency domains.	What is actually separate?
Protection	A switching mechanism that responds to a failure.	Who or what moves traffic when a failure occurs?
Resiliency	The ability of the overall service or system to continue operating, limit disruption, or recover within the business requirement.	Can the complete design meet the required recovery outcome?

A credible resiliency design begins with four questions: **What can fail? What is actually separate? Who or what moves traffic? What service level remains after the failure?** The sections that follow apply those questions to unprotected, protected, diverse, and dual path architectures.

Protection vs. Diversity architectures

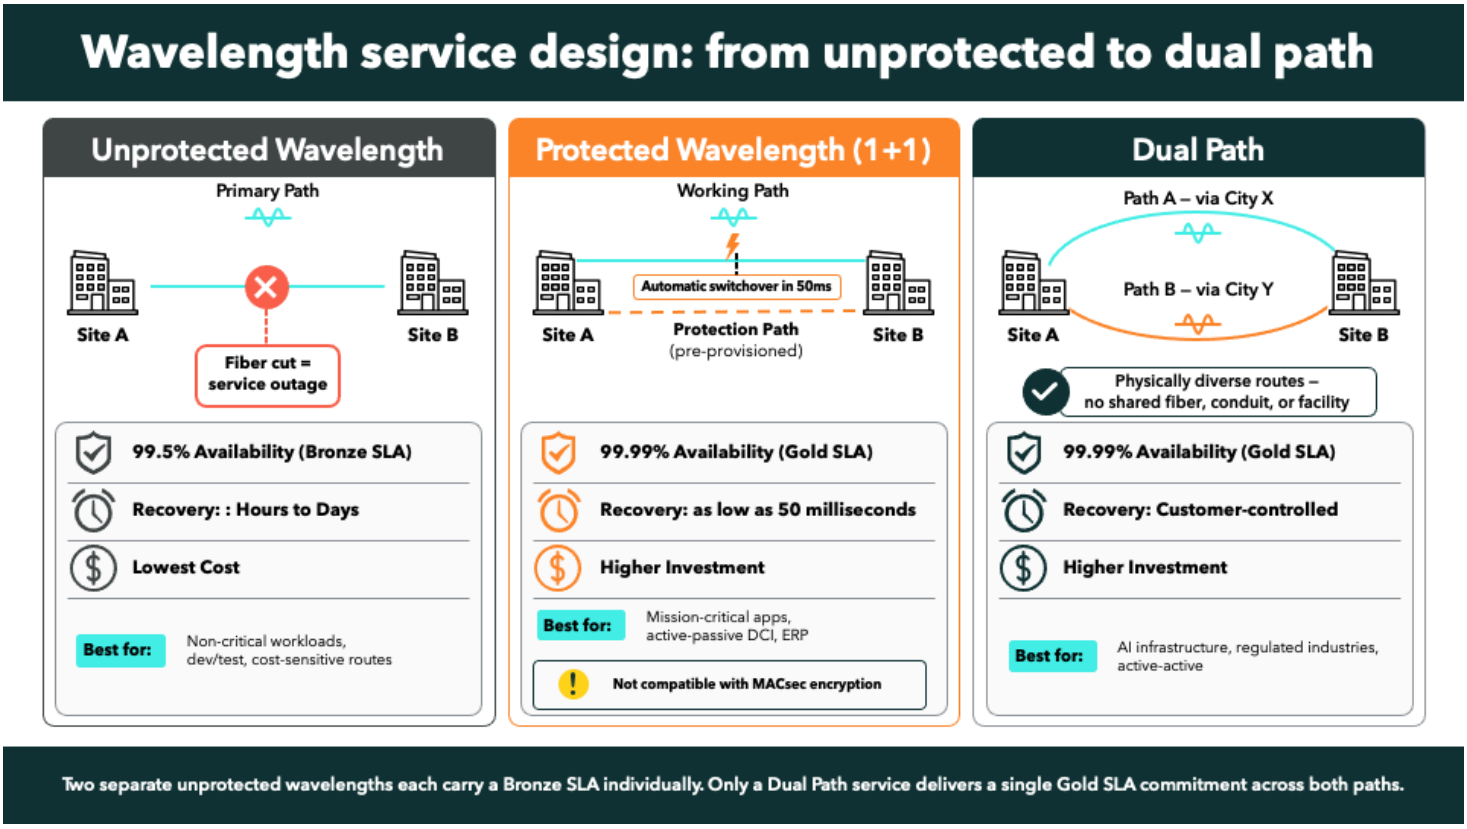
An **unprotected wavelength** provides a single, direct path between point A and point B. While it offers high capacity and low latency, a fiber cut or equipment failure on that specific route can interrupt the service until the path is repaired. It is appropriate for non-critical transfers or architectures where redundancy and recovery are handled elsewhere, such as in a higher-layer network, application, or second independently designed service.

A **protected wavelength** includes a working path, a protect path, and provider-managed optical protection switching within one logical wavelength service. Lumen's internal technical design guidance describes optical switching at as low as 50 milliseconds after a protection fault is detected. That timing describes the response inside the optical service; it is not a universal guarantee of end-to-end application recovery. Endpoint behavior, forwarding, capacity, application state, and timers all affect the customer's actual recovery experience.

© 2026 Lumen Technologies. All Rights Reserved.

A **diverse wavelength design** uses two or more independently routed wavelength services engineered not to intersect across defined infrastructure domains. The endpoints do not always have to be identical; for example, three data centers may be connected in a triangle where each leg is intended to be diverse from the others. The customer's Layer 2 or Layer 3 equipment—or an agreed operating process—detects failure and moves traffic. The customer therefore owns traffic movement, convergence, capacity planning, and application behavior.

Avoidance, or customer-specified routing, is related but distinct. It records a requirement to include or avoid a named city, corridor, facility, or route element. It is a valuable input to route design, but it does not independently prove that the final service is redundant, diverse, protected, or resilient.



Important compatibility consideration: MACsec may not maintain a link through an optical protection path change. Customers that require MACsec and optical protection should validate the combined behavior with their network design and the applicable product documentation.

© 2026 Lumen Technologies. All Rights Reserved.

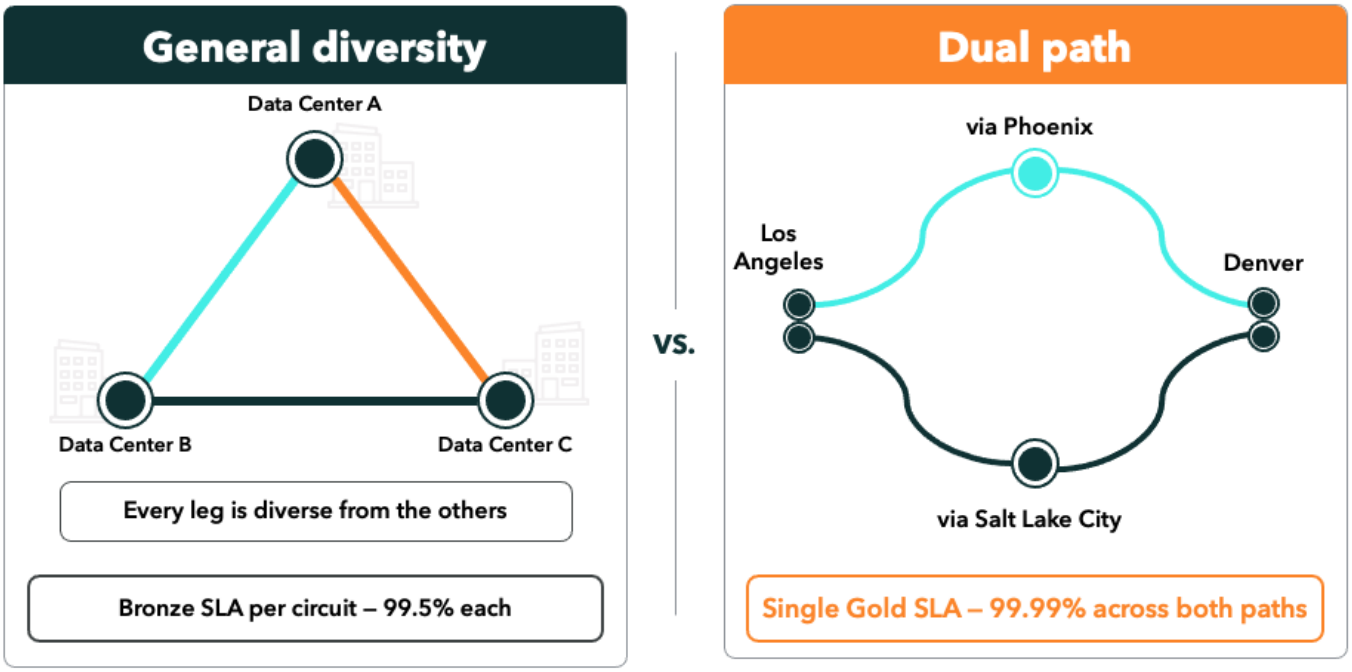
dual path: Two active circuits with customer-controlled traffic movement

Diversity and dual path are related but distinct. General diversity means two or more paths are engineered not to intersect across defined infrastructure domains, and the circuits do not need to share the same endpoints. dual path is a specific case: **Two active, full-bandwidth, diversely routed unprotected wavelengths between the same A and Z endpoints, linked as a service**.

Dual path gives the customer two separately usable circuits. Both circuits can be available to carry traffic, but the transport service does not automatically load balance or fail over customer traffic. The customer's routing, switching, load-balancing, or operating process determines how traffic is distributed and how it moves if one circuit fails.

The circuits terminate on different cards within the same chassis. This provides card diversity, but it does not establish chassis or shelf separation. Room separation, building entrance separation, gateway separation, and every other relevant infrastructure domain must be evaluated and documented separately.

Current product boundaries are equally important. A dual path offering does not include a latency guarantee or Lumen optical encryption as part of the offering. These capabilities should not be assumed because the design provides two diversely routed wavelengths between the same endpoints. Current availability commitments, outage-credit thresholds, and ordering conditions should be confirmed through the applicable approved service documentation.



Dual Path: two diverse routes, one SLA-backed service

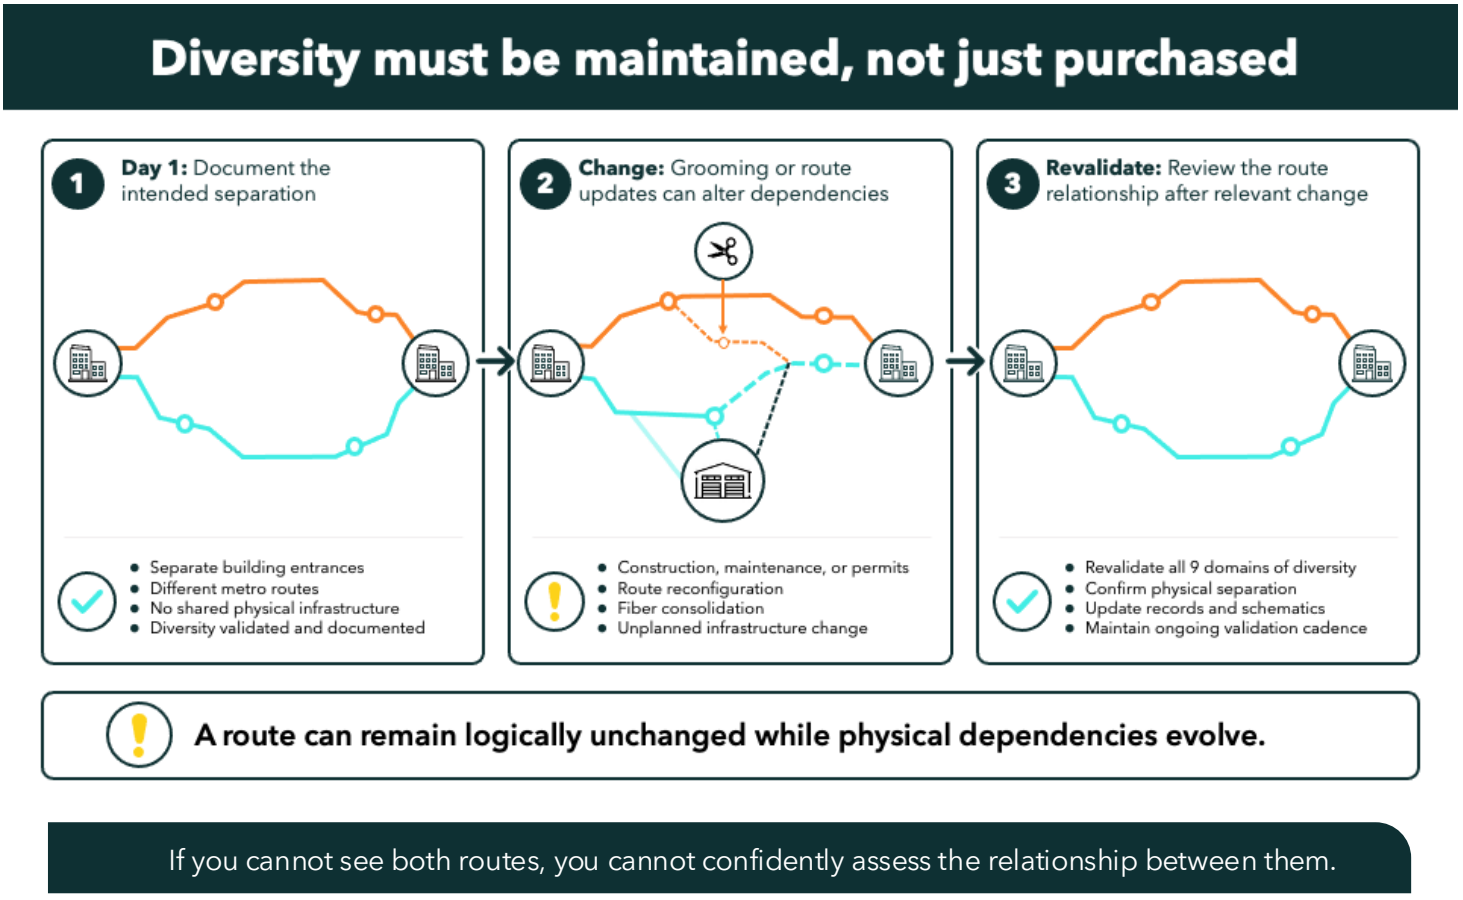
© 2026 Lumen Technologies. All Rights Reserved.

The multi-vendor diversity trap and operational drift

When building a network, enterprises have historically purchased services from different providers to reduce risk. The assumption is that different vendors necessarily mean different physical infrastructure. That assumption is not sufficient. **Carrier diversity is commercial and operational separation between provider relationships; it is not proof that the underlying physical routes, conduit, facilities, rooms, or restoration dependencies are separate**.

Infrastructure evolves. Services can be groomed, restoration paths can change, capacity can move, and optical platforms can be replaced. A circuit may remain logically the same from the customer's perspective while its underlying physical route or equipment relationship changes. If a replacement route later shares a conduit, metro segment, gateway, room, or another dependency with a previously diverse circuit, some of the original separation can be reduced.

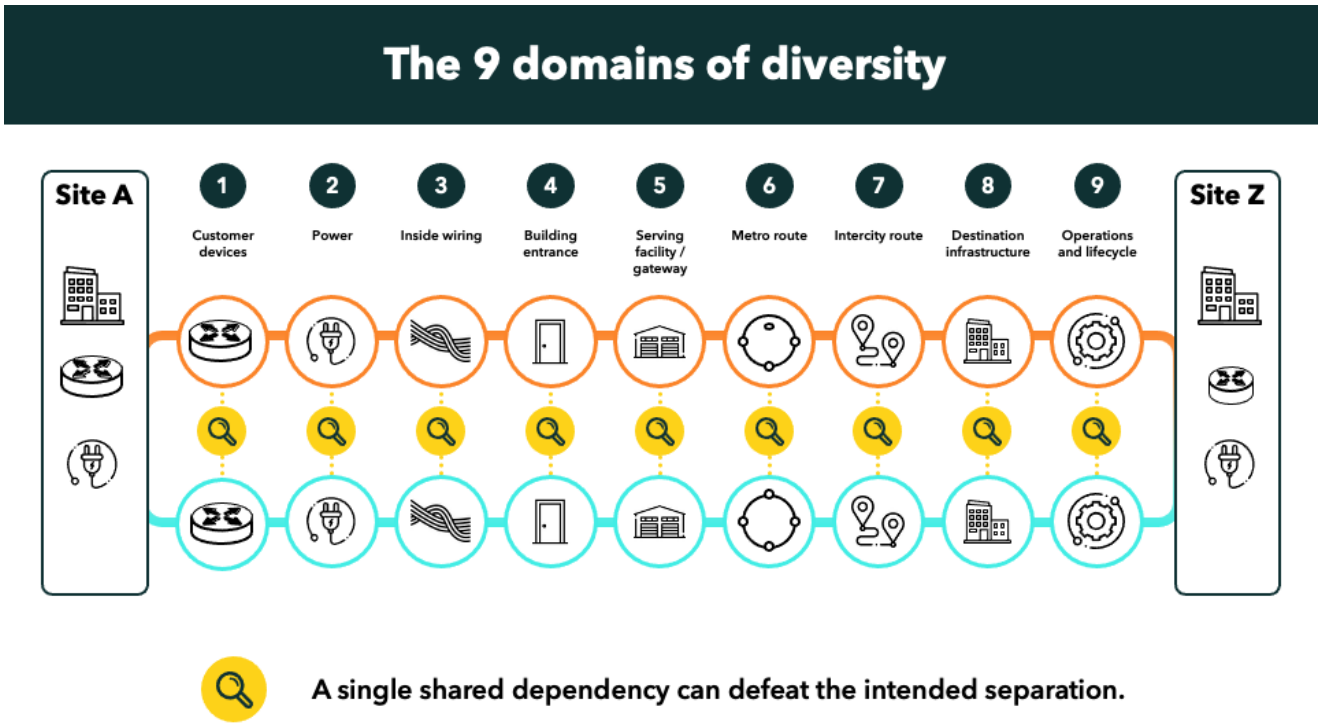
This is known as **operational drift**. It is why path validation should not end at procurement. An integrated provider can improve visibility, documentation, and lifecycle governance, but complete resiliency still requires a clear design record, review of relevant changes, and revalidation after the change.



© 2026 Lumen Technologies. All Rights Reserved.

End-to-End diversity depends on the complete path

A line on a map can appear separate while the underlying design still converges. End-to-end diversity should be evaluated across customer handoff, equipment, power, room, building entrance, metro access, gateway or wire center, intercity corridor, destination-side infrastructure, and operations.



Domain	What to Validate
Customer handoff, devices, and power	Separate routers, switches, line cards, optics, PDUs, UPS systems, and generators as required by the recovery objective.
Card, chassis/shelf, and room	Dual path uses different cards within the same chassis. Evaluate chassis, shelf, and room separation as distinct domains.
Inside wiring	Separate patch panels, risers, and cross-connect pathways where needed.
Building entrance and lateral	Determine whether separate entrance conduits or laterals are used so a single excavation event does not defeat upstream diversity.
Serving facility, gateway, and wire center	Confirm whether paths converge at the same central office, optical node, gateway, or serving network building.
Metro route	Review conduit, cable, ring segments, bridge, tunnel, railroad crossings, and intermediate facilities.

© 2026 Lumen Technologies. All Rights Reserved.

Domain	What to Validate
Intercity route	Evaluate shared fiber, conduit, regeneration facilities, and long-haul corridor dependencies.
Destination-side infrastructure	Repeat the same entrance, room, equipment, power, and route analysis at the Z location.
Operations and lifecycle	Identify maintenance, configuration, restoration, grooming, notification, and revalidation processes that could create common impact.

Key Takeaway: Real resiliency requires more than multiple carrier logos. It requires confirmed separation across the domains that matter, a clear failure-response model, sufficient surviving-path capacity, and a process for maintaining those conditions over time.

© 2026 Lumen Technologies. All Rights Reserved.

Chapter 6: Capacity, latency, and traffic engineering

Capacity planning and AI workload behavior

Start with measured peak and average throughput by workload. Then account for growth, protocol overhead, replication bursts, backup windows, recovery operations, and concurrent application events.

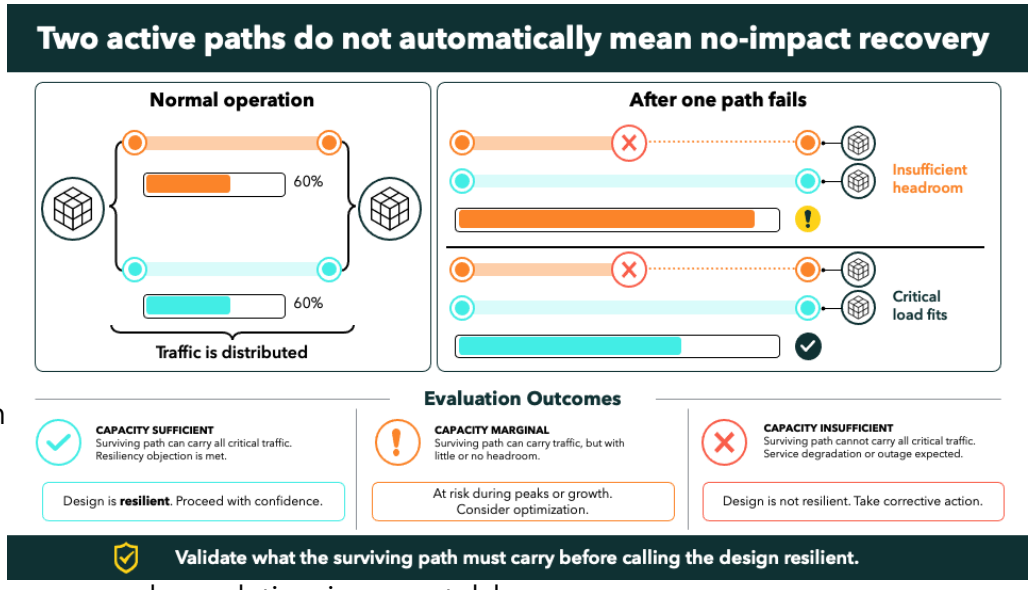
This is particularly critical for AI workloads. AI model training and inference drive massive, sustained, and often asymmetric data flows between compute clusters and storage repositories. Unlike traditional enterprise applications, where traffic peaks and valleys are somewhat predictable, AI training workloads can saturate network links for hours or days. The location where data lives versus where models are trained determines these large flows.

Furthermore, a recovery scenario can produce substantially more traffic than normal operations because data may need to be resynchronized after an outage. A narrow recovery window may justify more capacity than normal operations require—especially for AI model training, where an interrupted training run can be costly to restart.

High bandwidth alone does not guarantee fast transfer. TCP windowing, storage performance, encryption overhead, packet loss, host limits, and application concurrency can constrain throughput. End-to-end testing should validate the business workload, not only a synthetic line-rate test.

Capacity headroom Is part of resiliency

Two active paths do not automatically mean the design can lose one path without performance impact. If both paths normally run near capacity, a failure can overload the surviving circuit. The design should establish whether one path can carry all critical traffic, whether traffic is prioritized during a failure, and whether temporary degradation is acceptable.



© 2026 Lumen Technologies. All Rights Reserved.

Designing resilience around business impact

The correct architecture begins with the business impact of an interruption. Not every workload needs the same network design. A practical approach connects recovery requirements to network, endpoint, and application behavior.

Design Question	Why It Matters
What is the maximum tolerable interruption?	Define the recovery time objective for the business service, not only the circuit.
How much data loss is acceptable?	Replication design and transport capacity affect the recovery point objective.
Which failures must the architecture survive?	Examples include fiber cuts, card failures, serving-node events, building events, metro disruptions, and regional incidents.
Where should failover occur?	Decide whether the optical layer, customer routing, storage, application layer, or a coordinated combination moves traffic.
Can the surviving resource carry the critical load?	Validate full-load, priority-traffic, and degradation assumptions before the design is considered resilient.

Key Takeaway: Size the connection from measured traffic, AI workload behavior, growth, recovery windows, and the surviving-path requirement—not only from interface speed. A 100G interface does not necessarily mean 100G of sustained traffic, but the cost of interrupting a critical workload can justify dedicated capacity and meaningful headroom.

© 2026 Lumen Technologies. All Rights Reserved.

Chapter 7: Comparing connectivity options

Choosing the right optical solution requires balancing performance, operational complexity, cost, security requirements, and the customer's ability to manage failure behavior. The most effective approach is not to select a design from a generic product matrix. It is to map the customer's specific requirement to the architecture conversation needed to validate it.

Wavelength design decision framework

Customer Requirement to Validate	Design Conversation It Triggers
Lowest-cost transport for a non-critical workload	Determine whether an unprotected service is acceptable and what recovery exists elsewhere in the customer architecture.
Fast, provider-managed response within one logical service	Assess protected wavelength behavior, protection domains, and the application's compatibility with optical switching.
Two usable circuits for customer-directed failover or load sharing	Assess diverse wavelengths or dual path; validate routing, headroom, convergence, and application behavior.
Risk separation across a city, corridor, facility, entrance, or gateway	Document avoidance requirements and diversity domains; confirm what remains shared or needs site-specific validation.
Encryption or strict compliance requirements	Identify the encryption layer, performance needs, and whether the selected service architecture supports that option.
Low-latency application behavior	Assess physical route, distance, regeneration, and the required latency commitment. Do not presume a latency guarantee from dual path.

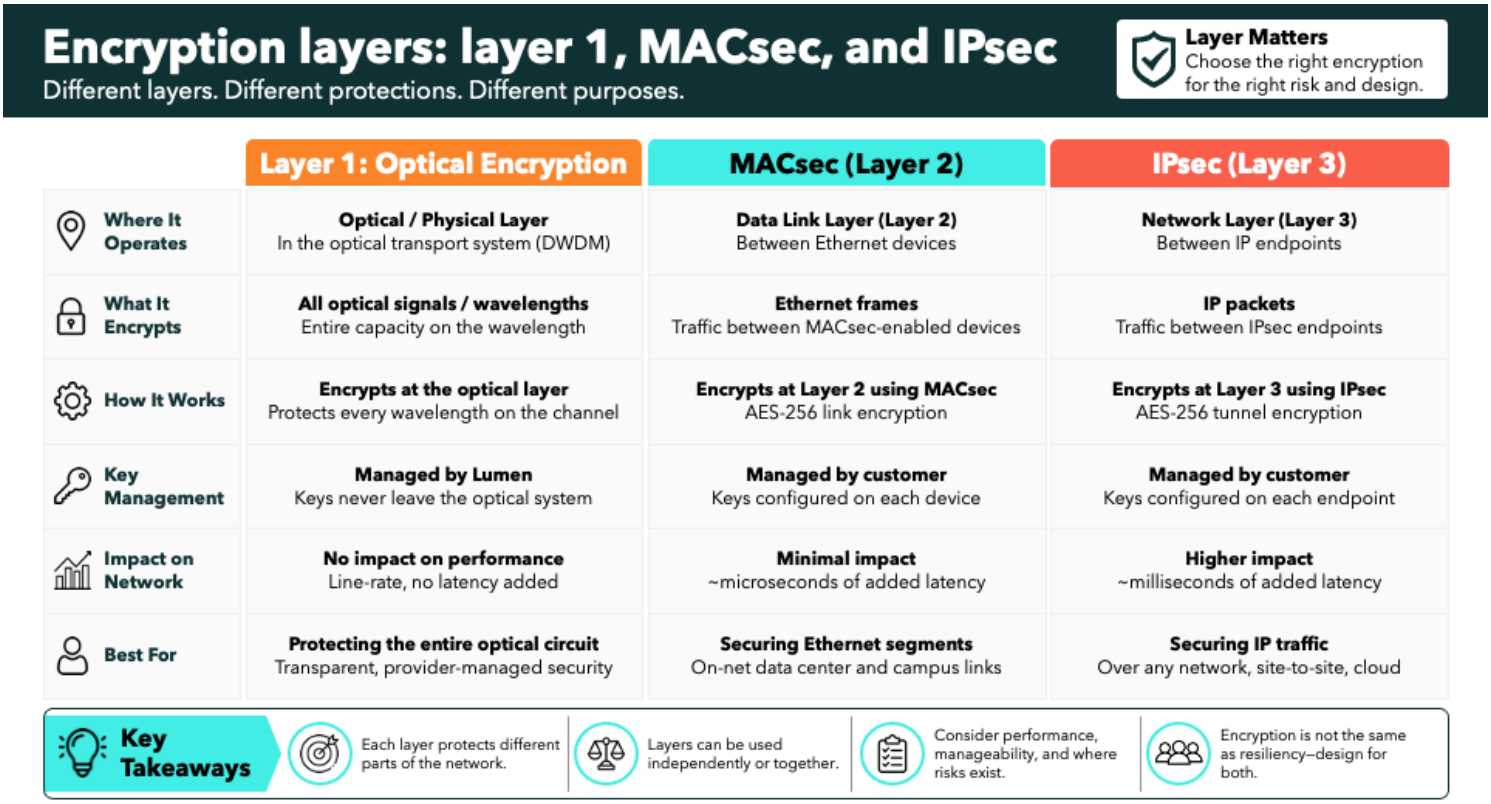
Security: Optical layer encryption

As data travels outside the secure perimeter of a data center, it can become vulnerable to interception. Optical encryption protects in-flight data at the transport layer for supported services, subject to the selected implementation, route, and product availability.

Unlike IPsec at Layer 3 or MACsec at Layer 2, Layer 1 encryption can be performed by DWDM hardware at wire speed. The appropriate encryption method depends on security policy, operational model, endpoint capabilities, latency tolerance, and the chosen resiliency architecture. Lumen offers optical encryption in standalone and transponder-integrated deployment models for applicable offerings.

© 2026 Lumen Technologies. All Rights Reserved.

A private Layer 1 service should not automatically be interpreted as an encrypted service. Equally, an encrypted architecture should not be assumed to be compatible with every failover design. Layer 1 optical encryption is **not** part of the dual path offering; MACsec and optical protection switching require design validation.



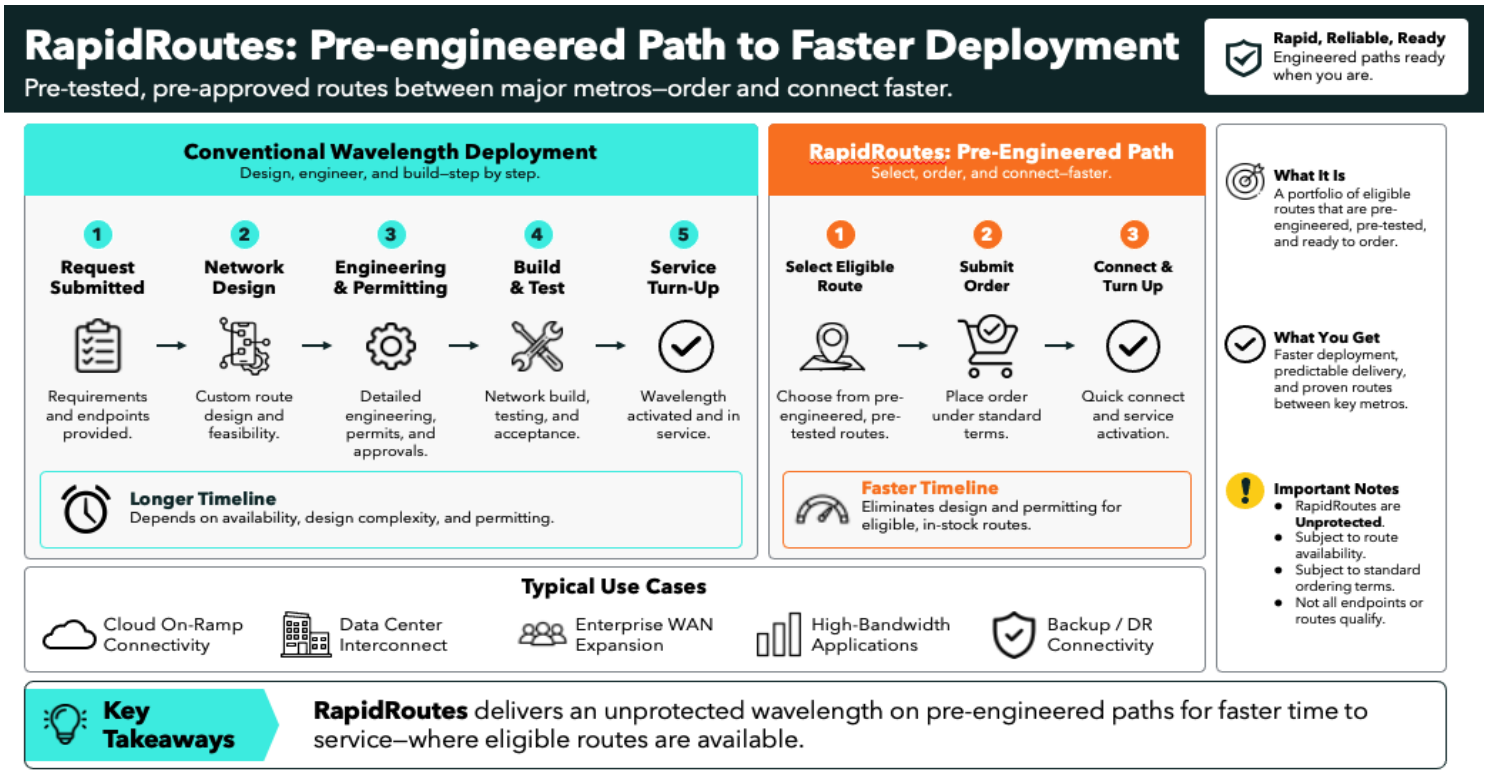
Key Takeaway: The design recommendation follows the customer's failure, capacity, security, and operating requirements—not a one-size-fits-all label. State the desired security outcome first, then validate the encryption layer and its interaction with the selected resiliency architecture.

© 2026 Lumen Technologies. All Rights Reserved.

Chapter 8: Accelerated deployment with pre-engineered routes

The physical relationship between the customer locations and the provider's network dictates the solution type. To support accelerate deployment and optimize costs, providers can offer predefined, high-capacity routes.

For example, Lumen RapidRoutes℠ provides pre-engineered 100G and 400G unprotected wavelengths across established corridors, subject to route availability and applicable ordering terms. Traditional wavelength provisioning can take months, while RapidRoutes uses pre-engineered paths to support accelerated deployment on eligible routes. Current delivery commitments, location eligibility, and ordering conditions should be confirmed for the specific request.



RapidRoutes solves a time-to-service challenge. It does not remove the need to design for capacity, failure behavior, security, or route separation. Customers with resiliency requirements should identify whether they require protection, diverse circuits, avoidance constraints, or customer-managed recovery in addition to the accelerated unprotected transport path.

Continuous Visibility: Lumen Connect and the LASER Platform

Procuring a resilient route is only the first step; maintaining visibility into that route is designed to protect the investment over time.

© 2026 Lumen Technologies. All Rights Reserved.

Lumen's support and engineering teams use **LASER**—an AI-based internal platform with topology-based map visualization, proactive fiber-cut prediction, protect-path switching controls, and automated workflows. LASER is an internal Lumen tool; its capabilities support Lumen's ability to monitor, validate, and act on network events on the customer's behalf.

Customers interact with approved service capabilities through the **Lumen Connect portal**, where they can view service status, verify available service information, open and track trouble tickets, and perform selected operational actions. The exact customer-facing functions available for a service should be confirmed through current portal and product documentation.

For a diversity-dependent design, the customer and service team should agree on the route evidence available at order and activation, the notification and escalation process for changes that could affect the route relationship, the way paired services and exceptions are documented, and the process for revalidating diversity after relevant change.

Key Takeaway: Designing a simpler optical paths can enable improved performance and reduced operational complexity. Pre-engineered routes can accelerate deployment, but the intended resiliency design still requires explicit validation, documentation, and lifecycle governance.

© 2026 Lumen Technologies. All Rights Reserved.

Chapter 9: Getting to service — what to expect and how to prepare

Ordering a wavelength service is not like ordering internet access. Because a wavelength is a dedicated, infrastructure-level connection—engineered to a specific physical path, speed, and resiliency model—the quality of the outcome depends heavily on the quality of the information exchanged before the order is placed.

This chapter is written for you, the customer. Its purpose is to help you understand what questions will be asked, what information you will need to provide, and why each step directly is designed to protect your investment and your ability to meet recovery and performance objectives.

What your Lumen team will need to understand

Before a wavelength can be designed, the Lumen solution team needs to understand your environment—not just the technical specifications, but the business context behind them.

About Your Business and Applications

- Which applications and data flows will use this connection, and what happens to the business if it is unavailable?
- What are the Recovery Time Objective and Recovery Point Objective for each workload?
- Is the traffic production, backup, replication, migration, AI training, inference, or a combination? Each has different capacity and resiliency implications.
- How much growth do you expect over the next one, three, and five years?
- Which specific failures must the design withstand: fiber cut, equipment event, building event, facility event, metro disruption, or regional incident?
- Can the surviving path carry all critical traffic, and which workloads take priority if a failure occurs?
- Does the application move traffic automatically, or must a network or optical-layer mechanism respond?

About Your Sites and Technical Environment

- What are the exact addresses, suites, and meet-me-room details for both ends of the connection?
- What client interface, protocol, and speed does your equipment require?
- Do you need third-party cross-connects or building extensions at any sites?
- Who are the local site contacts for each location?
- Do you need optical protection, physical route diversity, or both—and across which domains?

© 2026 Lumen Technologies. All Rights Reserved.

- Are there city, corridor, facility, building-entrance, room, or gateway elements that must be included or avoided?
- Is encryption required, and at which layer—optical, MACsec, or IPsec?
- Are there latency commitments or application thresholds that the design must meet?

What happens before your service is ordered

Once requirements are understood, the service team can confirm the exact service addresses and demarcation points and validate route availability for the selected speed and architecture. If route diversity is part of the design, you may be asked to review proposed route evidence, such as a KMZ map, before the order is placed. This is an important step: reviewing the route before ordering is the clearest way to confirm that the physical path meets the stated diversity requirements rather than discovering a gap after the service is live.

The design record should also state which dependency domains are intended to be separate, which are known to be shared, and which require site-specific validation. It should document any avoidance constraints, exceptions, acceptance criteria, test requirements, operational contacts, and the applicable service-level language.

Keeping your design valid over time

A wavelength design that is correct today may not remain correct after a network change, capacity move, site move, platform replacement, or contract renewal. The underlying service can be groomed from one route, optical platform, or facility relationship to another while remaining logically unchanged to the customer. For a diversity-dependent design, this makes lifecycle governance essential.

After activation, the most important operational habits are:

- Confirming Lumen Connect portal access for approved service visibility, support interactions, and ticketing.
- Baselining utilization, latency, and error performance so degradation is detectable early.
- Exercising customer-managed failover and recovery procedures on an appropriate schedule.
- Reviewing the expected behavior of the selected protection or customer-managed recovery design.
- Revalidating diversity after grooming, capacity moves, route changes, site moves, facility changes, platform changes, or other events that may affect a shared dependency.
- Reviewing permanent or temporary replacement routes when the original separation is material to the business requirement.
- Documenting paired-service relationships, exceptions, revalidation outcomes, and current escalation contacts

© 2026 Lumen Technologies. All Rights Reserved.

Key Takeaway: This chapter exists to help you arrive at the design conversation prepared. The questions your Lumen team will ask are not administrative hurdles; they are the inputs that determine whether the service is engineered, documented, tested, and maintained to survive the failures that matter most to your business.

Citations

1. ARK Investment Management LLC. Big Ideas 2023: Innovation Trends with Exponential Growth Potential. ARK Invest, 31 Jan. 2023.

2. Mathematically, transferring a 1-petabyte dataset (10^{15} bytes) at a theoretical 100% line rate requires approximately 9.26 days over a 10 Gbps link, whereas a 400 Gbps network reduces the theoretical transfer time to roughly 5.56 hours. In practical application, network architects must factor in an additional 10% to 20% overhead for protocol layers and TCP window constraints.

© 2026 Lumen Technologies. All Rights Reserved.

Glossary

Term	Definition
Attenuation	The reduction in signal strength as light travels through optical fiber.
Avoidance / Customer-Specified Routing	A requirement to include or avoid a named city, corridor, facility, room, gateway, or other route element. It is a route constraint, not a resiliency classification.
CDC (Colorless, Directionless, Contentionless)	Advanced capabilities of modern ROADM systems that increase flexibility in routing optical signals without hardware constraints at fixed paths.
Carrier Diversity	Commercial and operational separation between provider relationships. It does not by itself prove physical separation of underlying routes or infrastructure.
CLLI (Common Language Location Identifier)	A standardized code used to identify physical locations in telecommunications networks and help determine equipment presence at a site.
DCI (Data Center Interconnect)	Technology and architecture used to link two or more data centers over short, medium, or long distances for data exchange, replication, and workload mobility.
Diversity	Two or more services engineered to avoid intersection across defined infrastructure domains. The design must identify what is separated, shared, or subject to site-specific validation.
Dual path	Two active, full-bandwidth, diversely routed unprotected wavelengths between the same two endpoints, linked as a service. Customer equipment or operations control traffic movement. The service does not include a latency guarantee or Lumen optical encryption.
DWDM (Dense Wavelength Division Multiplexing)	Technology that combines multiple optical signals onto a single fiber by using different wavelengths, or colors, of laser light.
East-west Traffic	Data traffic moving laterally within or between data centers, such as server-to-server traffic, storage replication, and AI model training.
EDFA (Erbium-Doped Fiber Amplifier)	An optical amplifier that boosts signals without converting them to electrical form, commonly used for long-haul transmission.
Grooming	Planned or operational movement of a live circuit, wavelength, or service group to another network resource, route, or platform. Grooming can change physical dependencies while the customer-facing service remains logically unchanged.
KMZ	A file format used to display geographic data in mapping software. In optical networking, it can be used to review the physical fiber route of a wavelength circuit.
MTU (Maximum Transmission Unit)	The size of the largest protocol data unit that can be communicated in a single network-layer transaction. Lumen Wavelength Solutions® support a 9600-byte MTU, subject to service specifications.

© 2026 Lumen Technologies. All Rights Reserved.

Term	Definition
North-south Traffic	Data traffic moving between an external user or client and a data center or cloud application.
Network Terminating Equipment (NTE) / Demark	Where a service provider's network infrastructure ends and the customer's private network begins. NTE is the physical hardware device provided and owned by the Internet Service Provider (ISP) or telecom company.
Operational Drift	Reduction in the originally intended route or infrastructure separation over time due to grooming, restoration changes, capacity shifts, platform changes, or ungoverned operational practices.
Protected Wavelength	One logical wavelength service with working and protect paths and provider-managed optical protection switching.
Proactive Notifications	An optional service capability that can alert a customer when Lumen detects a potential issue and opens a ticket, subject to current product availability and service terms.
Redundancy	Duplicate capacity, component, circuit, or path. Redundancy supports resiliency only when dependencies are sufficiently separated and the failure response works as intended.
Reliability	How consistently a service performs without interruption over time.
Resiliency	The ability of a service or system to continue operating, limit disruption, or recover within the business requirement after a failure.
RPO (Recovery Point Objective)	The maximum targeted period in which data might be lost from an IT service due to a major incident.
RTO (Recovery Time Objective)	The maximum tolerable length of time that a system, network, or application can be unavailable after a failure or disaster.
TCP Windowing	A TCP mechanism that controls the amount of data that can be sent before an acknowledgment is required. Over long-distance, high-latency links, improper window sizing can constrain throughput regardless of raw wavelength capacity.
Unprotected Wavelength	One wavelength service without service-level protection switching.

© 2026 Lumen Technologies. All Rights Reserved.