

Lumen® Multi-Cloud Gateway and Ethernet Fabric Connect

CIO FAQs

How does Multi-Cloud Gateway support our multi-cloud strategy and network investments for the future?

Multi-Cloud Gateway provides a distributed Layer 3 routing service that connects multiple public clouds, private clouds, and enterprise sites over Lumen's private backbone. It helps eliminate the need for multiple direct connects or complex VPN meshes, minimizing operational overhead. The platform supports major cloud service providers (CSP) like AWS, Azure, Google, and Oracle, with three service instance tiers (10/50/ Unlimited¹ Gbps) and diverse paths to deliver scalable, highly available routing for customer premises-to-cloud, cloud-to-cloud, and customer premises-to-data center connectivity. It is API-first and is integrated with the Lumen portal for near real-time provisioning and automation. Roadmap features, such as IP VPN integration and SaaS/private exchange connectivity, are expected to enable scalability and flexibility as your cloud footprint grows without requiring major network redesigns. Multi-Cloud Gateway also supports qualified off-net locations, so sites can join the same private routing fabric even where Lumen does not own the fiber – extending reach to over 10 million locations without changing the ordering experience.



What are the business continuity and disaster recovery advantages of using Multi-Cloud Gateway for interconnecting data centers and cloud regions?

Multi-Cloud Gateway supports resilient interconnection between data centers and cloud regions by leveraging private, high-performance paths rather than the public internet. This helps reduce exposure to internet outages and improve overall reliability. It also supports flexible topologies, including mesh-like connectivity between sites, enabling rapid failover and minimizing downtime during disruptions. For example, financial institutions can replicate data between regions in milliseconds, to help support compliance and uptime during failover scenarios. This approach simplifies disaster recovery planning and helps accelerate recovery times.

How does Multi-Cloud Gateway help us reduce operational complexity and network costs across our global footprint?

By consolidating CSP ports and routing policies into one gateway, enterprises avoid duplicate circuits and optimize egress fees through aggregation by leveraging Lumen's private connectivity instead of sending traffic over the public internet. Self-service provisioning and automation cut manual configuration time from days to minutes. Procurement benefits from predictable pricing models (hourly or monthly) to enabling fewer vendor contracts. This can help reduce both CapEx and OpEx while enhancing agility.

Is Multi-Cloud Gateway a physical device?

No. Multi-Cloud Gateway is a virtual service built on Lumen® Connectivity FabricSM and PE routers.

What security and compliance controls are built into Multi-Cloud Gateway, and how do they help us meet regulatory requirements)?

Multi-Cloud Gateway enforces private connectivity, MD5-authenticated BGP sessions, and prefix filtering. These controls minimize exposure to the public internet and help maintain data integrity during transit, which supports compliance frameworks.

How does Multi-Cloud Gateway enable seamless migration of legacy applications to the cloud without disrupting user experience?

Multi-Cloud Gateway treats cloud VPCs as part of the corporate WAN, so applications moved to the cloud remain accessible via private paths without changing user workflows. This minimizes disruption during migration and preserves performance for latency-sensitive workloads. By routing application traffic over private connectivity rather than the public internet, organizations can migrate confidently without disrupting the end-user experience..

What is the customer experience for onboarding, provisioning, and managing Multi-Cloud Gateway—can we self-serve, and what support is available?

Multi-Cloud Gateway is a DIY service. Create instances, attach interfaces, and configure routing policies through the Lumen Connect® portal or APIs. The platform offers near real-time provisioning, monitoring, and scaling, providing a single pane of glass for visibility and control. Support includes enterprise-grade service SLAs and access to Lumen's NOC and customer success teams for troubleshooting and guidance.

How does Lumen's ownership of the underlying network infrastructure translate into performance, reliability, and accountability for our enterprise?

Lumen owns and operates its fiber backbone and 125+ cloud on-ramps, delivering low latency and high reliability without third-party dependencies. This ownership allows Lumen to monitor and optimize performance at every hop, allowing for accountability and exceptional service quality compared to overlay solutions. However, Multi-Cloud Gateway is a 100% DIY service, giving you full control over your connectivity.

How does Multi-Cloud Gateway scale to support AI workloads, high-throughput applications, and future bandwidth needs?

Multi-Cloud Gateway offers an unlimited bandwidth tierⁱ, with Ethernet services supporting up to 100G today and roadmap-ready for 400G. This enables high-throughput data movement for AI training, analytics, and other data-intensive workloads. The architecture is designed for scalability without requiring hardware changes.


What is the roadmap for integrating MPLS/IP VPN, internet egress, and SaaS/exchange connectivity into Multi-Cloud Gateway?

Future capabilities are expected to include IP VPN integration to extend MPLS networks into cloud environments, centralized internet egress for secure breakout, and service cloud access for private connectivity to SaaS and cloud exchanges. Partner/extranet connectivity is also planned, enabling secure B2B integration.

Cloud Architect FAQs

How does Multi-Cloud Gateway simplify the design and management of multi-cloud and hybrid cloud topologies?

Multi-Cloud Gateway centralizes routing and policy enforcement, avoiding the need for complex site-to-site tunnels and minimizing latency for inter-cloud traffic. It supports scalable hub-and-spoke designs, where the hub is the Lumen Connectivity Fabric rather than a physical location. This architecture makes it easy to add new clouds or sites without redesigning the network.



What is Ethernet Fabric Connect and how does it work with Multi-Cloud Gateway?

Lumen® Ethernet Fabric Connect is Lumen's portfolio of Layer-2 Ethernet service connection types (e.g., point-to-point, multipoint-to-multipoint, and cloud connectivity variants). These connections can be used standalone (without Multi-Cloud Gateway) to link ports/locations or connect ports to cloud on-ramps, or they can be used with Multi-Cloud Gateway as the Layer-2 underlay for routed multi-cloud architectures. In the Multi-Cloud Gateway workflow, customers create gateway interfaces and then attach Ethernet Fabric Connects to those interfaces to make connectivity operational.

What routing protocols and policy controls (e.g., BGP, prefix lists, route filtering) are supported for granular traffic engineering?

Multi-Cloud Gateway supports BGP (IPv4), static routes, prefix lists, and route filtering for granular control. These features allow architects to enforce routing policies and optimize traffic paths. Support for IPv6 is on the roadmap.

How do routing tables work in Multi-Cloud Gateway?

Routing tables in Multi-Cloud Gateway define how traffic is allowed to flow between connected interfaces, sites, and cloud environments. They act as a centralized policy layer that determines which routes are learned, preferred, or advertised using static routes or BGP policies. By separating routing intent from physical connectivity, routing tables allow customers to scale multi-cloud and hybrid environments without re-architecting their network.

How are connections to cloud on-ramps provisioned- can we automate via API, and what clouds are natively supported?

Connections to cloud on-ramps are provisioned via hosted cloud connections to AWS, Azure, Google, and Oracle. All functionality is API-first (OpenAPI 3.0), enabling automation and easy integration into workflows. Provisioning is near real-time through the portal or API, allowing architects to spin up connections in minutes. Dedicated private on-ramps are expected to be introduced in future phases.

What are the options for bandwidth tiers, redundancy, and high availability across regions and clouds?

Multi-Cloud Gateways are provisioned in instance bandwidth tiers of 10 Gbps, 50 Gbps, and Unlimitedⁱ, with total throughput based on the aggregate bandwidth of all attached Ethernet Fabric Connects. For redundancy and high availability, you can deploy multiple connections across diverse paths and regions to maintain service continuity. This design minimizes single points of failure and supports resilient interconnects between clouds and data centers.

Can Multi-Cloud Gateway connect sites at off-net locations where Lumen does not own the fiber?

Yes. Qualified off-net ports connect to Multi-Cloud Gateway through Ethernet Fabric Connect, delivering parity between on-net and off-net access options. Multi-Cloud Gateway is no longer limited by where Lumen owns the fiber, so distributed sites can be brought into the same private routing fabric wherever those sites are located – extending Multi-Cloud Gateway availability to over 10 million locations. Architecturally, nothing else changes: an off-net port attaches to a Multi-Cloud Gateway interface through Ethernet Fabric Connect in the same way as an on-net port, and the same routing, policy, and bandwidth tier options apply.

What does off-net reach mean for designing large, distributed enterprise architectures?

Enterprises with many branch or remote locations no longer need to design around Lumen's on-net footprint. Sites that were previously unqualifiable can now be included in the same Layer 3 routing fabric, which means a single hub-and-spoke or any-to-any design can span the entire estate rather than only the on-net portion of it. Off-net locations do require individual qualification, so confirm each site during the design phase before committing to bandwidth, term, or delivery timing.

How does Multi-Cloud Gateway handle NAT, static routes, and advanced routing scenarios (e.g., hub-and-spoke or any-to-any)?

NAT is on the roadmap; the current release focuses on Layer 3/BGP routing and policy control. Static routes are supported, and the architecture allows flexible routing topologies such as hub-and-spoke or any-to-any connectivity between sites and clouds.

What visibility is available today to understand my Multi-Cloud Gateway connections and how the network is evolving over time?

Today, customers can view their services and connectivity context through the Lumen Connect® map dashboard, which provides a visual representation of how sites, clouds, and gateway connections are organized. Additional capabilities—such as deeper route views, topology awareness, and expanded operational telemetry—are planned for future releases.

What is the process for connecting new sites, clouds, or data centers to Multi-Cloud Gateway?

Provision Ethernet Fabric Connect to Multi-Cloud Gateway interfaces and configure IP/BGP. All of this can be performed in Lumen Connect® or via API! This process is fully automated and can be completed in minutes. This applies to both on-net and qualified off-net locations; once an off-net Ethernet Fabric Connect or Ethernet On-Demand port is activated, it appears in the Multi-Cloud Gateway port list alongside on-net ports.

How does Multi-Cloud Gateway support integration with existing Layer 2/3 architectures and private WANs?

Multi-Cloud Gateway acts as a Layer 3 hub for multi-cloud routing and integrates with Ethernet environments today. Future IP VPN integration will allow MPLS networks to extend into the cloud via Multi-Cloud Gateway. For example, your public cloud environment, which may be hosting multiple VPCs, can appear as a single site on your private WAN – like a private data center. This enables seamless routing between on-premises and cloud environments.

What is the expected latency and throughput between major CSPs and enterprise sites?

Low latency via Lumen backbone. Throughput scales with selected bandwidth tier.

How does connectivity scale in Multi-Cloud Gateway – is it one connection per interface?

Yes. In Multi-Cloud Gateway, each interface represents a single logical connection to a cloud, site, or service. Customers attach one connection per interface and can scale by adding additional interfaces as needed. This model keeps routing and policy simple and predictable, while allowing customers to expand connectivity incrementally without redesigning their architecture.

Network Engineer FAQs

What are the supported interface types, speeds, and physical/logical connection options for Multi-Cloud Gateway?

Multi-Cloud Gateway provides Layer 3 IP-based interfaces that support speeds from 1 Gbps up to 100 Gbps. Logical connections can attach to Lumen Ethernet Fabric Connects—which link enterprise sites, data centers, and cloud on-ramps—all connecting within the Lumen Connectivity Fabric. Both on-net and qualified off-net ports are supported, connecting through Ethernet Fabric Connect with full parity between the two access types.



What role do ports play in the Multi-Cloud Gateway solution?

Ports are the physical connection to the Lumen Connectivity Fabric that are used to connect your network to the Multi-Cloud Gateway.

How does an off-net port get attached to Multi-Cloud Gateway?

After an Ethernet Fabric Connect (EFC) or Ethernet On-Demand (EoD) off-net port is activated, it becomes available in the port drop-down list within the Ethernet Fabric Connect Service Detail section of Lumen Connect®. Customers select from a complete list of activated on-net and off-net ports supported by their chosen BAN. There is no change to the ordering experience in the portal.

Are there limitations to be aware of when using off-net access with Multi-Cloud Gateway?

Off-net sites require individual qualification – confirm each location before committing to bandwidth, term, or delivery timing. Second, Internet On-Demand (IoD) requires its own port and cannot share a port with Ethernet Fabric Connect, Ethernet on Demand, or Multi-Cloud Gateway. If an off-net port already carries IoD, that service must be disconnected before Multi-Cloud Gateway can be added, or the customer needs a separate port. EFC and EoD can coexist on the same port.

What are Fabric Ports and when will they be available?

Fabric Ports will provide high-capacity, dedicated connectivity into the Lumen Connectivity Fabric, enabling direct, private on-ramps to CSPs and other services. Fabric Ports will support API-based provisioning for automation to help improve flexibility and performance for multi-cloud deployments. Fabric Ports are now available.

How do ports fit into the overall Multi-Cloud Gateway architecture?

Ports assign IP subnets for routing, connect to PE/INT and PE/METV for WAN and metro services, anchor Ethernet Fabric Connect for dynamic CSP links, and support HCI VLAN segmentation for multi-cloud environments.

What types of Ethernet Fabric Connect can be used with Multi-Cloud Gateway (and can they also be used without it)?

Ethernet Fabric Connect supports multiple Layer-2 service types—including virtual and transparent Ethernet constructs (e.g., EVPL/EPL and related EVC forms) and cloud connectivity options. These can be used without Multi-Cloud Gateway for port-to-port or port-to-cloud connectivity or used with Multi-Cloud Gateway by attaching the selected Ethernet Fabric Connect type to a Multi-Cloud Gateway interface.

What are the configuration steps for BGP peering, prefix filtering, and route policy enforcement?

Configuration involves assigning IP addresses to the connection and configuring both the local peer (Multi-Cloud Gateway) and the remote peer (cloud or router). Steps include setting the ASN for each peer, enabling MD5 authentication for secure BGP sessions, applying prefix lists and route filters to control inbound and outbound advertisements, and defining routing policies for traffic prioritization and failover. All steps can be executed via the Lumen Connect® portal or automated through APIs.

Will customers be able to run ping or trace tests from the portal?

Not initially. Diagnostics or troubleshooting tools are planned for future releases to enhance troubleshooting capabilities.

What security features are available at the network layer (e.g., segmentation, MD5 authentication)?

Multi-Cloud Gateway provides segmentation of public and private traffic, MD5 authentication for BGP sessions, and prefix filtering to prevent route leaks. These features help protect routing integrity and provide secure connectivity between sites and clouds.

What is the process for provisioning, scaling, and modifying connections via portal or API?

Provisioning steps include creating a Multi-Cloud Gateway, adding interfaces and attaching Ethernet Fabric Connects, configuring IP addressing and routing protocols. Scaling is achieved by adjusting bandwidth tiers or adding new interfaces, all through self-service portal or API.


How does Multi-Cloud Gateway support integration with existing Ethernet, MPLS, and SD-WAN environments?

Multi-Cloud Gateway complements existing architectures by acting as a Layer 3 hub for multi-cloud routing. Future roadmap includes IP VPN integration for MPLS networks, allowing seamless extension of private WANs into cloud environments.

IT Manager FAQs

How does Multi-Cloud Gateway streamline the process of connecting new branch offices or remote sites to cloud workloads?

Multi-Cloud Gateway avoids the need for site-specific VPNs by allowing branches to connect through Lumen Ethernet Fabric Connect into the gateway. Each branch attaches via a single interface. Once connected, branches gain private access to multiple clouds without additional hardware or complex configurations. Because Multi-Cloud Gateway also supports qualified off-net access, branches can be connected regardless of whether Lumen owns the fiber at that location, bringing distributed estates that were previously unqualifiable into scope.



Does off-net access change how Multi-Cloud Gateway is priced?

Multi-Cloud Gateway pricing is the same for off-net and on-net. Off-net access is delivered over Ethernet Fabric Connect, so off-net EFC pricing applies to that port – build it into the quote up front.

What is the operational model for managing, and scaling Multi-Cloud Gateway instances and interfaces?

Management is centralized through the Lumen Connect® portal and APIs, providing near real-time provisioning and scaling, unified visibility of all connections, and automated workflows for adding sites or clouds.

How does Multi-Cloud Gateway reduce the need for site-to-site VPNs and simplify firewall management?

By routing traffic through a centralized gateway, Multi-Cloud Gateway removes the complexity of maintaining multiple VPN tunnels and firewall rules across sites. Policy enforcement is centralized, helping to reduce operational overhead and improve security consistency.

What are the options for centralized policy enforcement and visibility across all connected sites and clouds?

Multi-Cloud Gateway provides global VRF and BGP policy control, enabling administrators to apply routing and policies across all connected endpoints.

How quickly can we provision or modify connectivity for new projects or changing business needs?

Provisioning and scaling can be completed in minutes via portal or API. This agility supports dynamic workloads, seasonal demand, and rapid cloud adoption without waiting for traditional circuit provisioning. Customers can adjust bandwidth tiers or add new connections on demand, with hourly billing options for short-term projects.

How does Ethernet Fabric Connect help scale connectivity as clouds or sites are added (with or without Multi-Cloud Gateway)?

With Multi-Cloud Gateway, scaling follows a one-to-many model: a single Multi-Cloud Gateway can support multiple interfaces, and each interface attaches to a single Ethernet Fabric Connect. As new clouds, regions, or sites are introduced, customers add additional interfaces and corresponding Fabric Connections, allowing the gateway to centrally route across an expanding set of connections without redesigning the underlying network. Without Multi-Cloud Gateway, Ethernet Fabric Connect scales by provisioning additional Layer-2 Fabric Connections as new sites, ports, or cloud endpoints are added.

What support and SLAs are provided for troubleshooting and incident response?

Lumen offers enterprise-grade SLAs for service availability and performance, backed by 24/7 NOC support and dedicated customer success teams. Incident response includes proactive monitoring and escalation procedure.

How does Multi-Cloud Gateway help us consolidate network vendors and simplify billing?

Multi-Cloud Gateway consolidates multi-cloud connectivity under one provider and one bill for network services to help reduce vendor management complexity. It does not include CSP usage charges. Pricing is transparent, with options for hourly or monthly billing to align with cloud consumption models.

What training or documentation is available for IT staff to manage Multi-Cloud Gateway day-to-day?

Comprehensive resources include self-service/portal guides and API documentation, onboarding support, and access to Lumen's technical account managers for ongoing assistance.

What are the key operational benefits compared to traditional cloud connectivity solutions?

Benefits include fast provisioning and scaling (benefits of as-a-Service), low latency (routing at cloud on-ramps) and enhanced performance compared to VPN tunnels, centralized management and visibility, and minimized complexity and cost compared to building a full mesh network with lots of point-to-point circuits

¹ Unlimited bandwidth subject to the capabilities and availability of the underlying Lumen network infrastructure

Why Lumen?

Lumen is unleashing the world's digital potential. We ignite business growth by connecting people, data, and applications – quickly, securely, and effortlessly. As the trusted network for AI, Lumen uses the scale of our network to help companies realize AI's full potential. From metro connectivity to long-haul data transport to our edge cloud, security, managed service, and digital platform capabilities, we meet our customers' needs today and as they build for tomorrow.

