White paper
Cisco public

Why Peering Networks Must Change

Create a localized peering strategy to meet new demands

It's time to reimagine peering

The capacity requirements of today's networks and the move to cloud require rethinking networks to make more efficient use of resources at all layers. The network architecture and technology needs to be reimagined with a focus on simplification and improving network efficiency. Internet peering is the exchange of traffic between two providers, and localized peering lowers network costs by reducing the distance and hops across the network. Service providers should consider a localized peering strategy that places peering or content provider cache nodes closer to traffic consumers.

Contents

The unquenchable demand for video shows no signs of slowing down, and the move to cloud has forced substantial changes in network environments. To deliver traffic to customers, more efficient peering will be required. Traffic continues to grow as more content sources are added and Internet connections speeds increase. To respond to these changes, service and content providers must reconsider the design of their peering networks. They need to be able to scale for a future that includes more connected devices with traffic sources and destinations that span the globe. Several key drivers are forcing the need for change: traffic growth, network simplification, and network efficiency.

Coping with traffic

Globally, Internet traffic is expected to grow 26 percent CAGR through 2022 and 82 percent of the traffic is expected to be video, which could add strain to your network if you must backhaul that traffic to aggregated peering points for off-loading.[1] The rise of cloud applications and infrastructure is also causing fundamental changes. By 2021 more than 95 percent of data center traffic will be cloud traffic. Annual global cloud IP traffic will reach 19.5 ZB (1.6 ZB per month) by the end of 2021. And by 2021, 94 percent of workloads and compute instances will be processed by cloud data centers with only 6 percent processed by traditional data centers.[2]

Simplifying complex networks

As requirements have changed, the networking stack has become ever more complex. To respond to specific problems, many protocols have been introduced at both the data and control plane levels. However, protocols have rarely been removed from the network. All of these protocols not only lead to complex network architectures, they also affect the efficiency of network operations. As networks scale to handle traffic growth, the level of network complexity must remain flat.

A prescriptive design using standard discrete components makes it easier for providers to scale from networks handling a small amount of traffic to tens of terabits per second (Tbps) without a complete network overhaul. Network fabrics with reduced control-plane elements and feature sets enhance stability and availability. Dedicating nodes to specific functions of the network also helps isolate the rest of the network from malicious behavior, defects, or instability.

Improving efficiency

To remain competitive, service providers also need to focus on improving the efficiency of their networks. Network efficiency not only refers to maximizing network resources but also to optimizing the environmental impact of the deployed network. Today, much of the Internet peering being performed occurs in third-party facilities where space, power, and cooling are at a premium. High-density, lower environmental footprint devices are critical to handling more traffic without exceeding the capabilities of a facility. In cases where multiple facilities must be connected, a simple and efficient way to extend networks must exist.

In many typical peering deployments, a traditional two-node setup is used where providers vertically upgrade nodes to support the higher capacity needs of the network. To support more connections, some service providers may employ technologies such as back-to-back or multichassis clusters while keeping what seems like the operational footprint low. However, failures and operational issues that occur in these types of systems are typically difficult to troubleshoot and repair. They also require lengthy planning and time frames for performing system upgrades.

A new peering strategy

To reduce network costs, service providers need a localized peering strategy that places peering or content provider cache nodes closer to traffic consumers. This approach reduces capacity on long-haul backbone networks carrying traffic from Internet exchange points (IXP) to end users. It also improves the quality of experience for users by reducing the latency to content sources. The same design also can be used for content provider networks that need to deploy a smaller footprint solution in a service provider location or third-party peering facility. The fabric architecture should be:

- Scalable so you can add more elements as traffic grows
- Resilient so traffic isn't impacted when a node or a link fails
- Efficient so you can increase bandwidth utilization by using equal-cost multipath routing (ECMP)
- Automated so you can upgrade software in a node without traffic loss

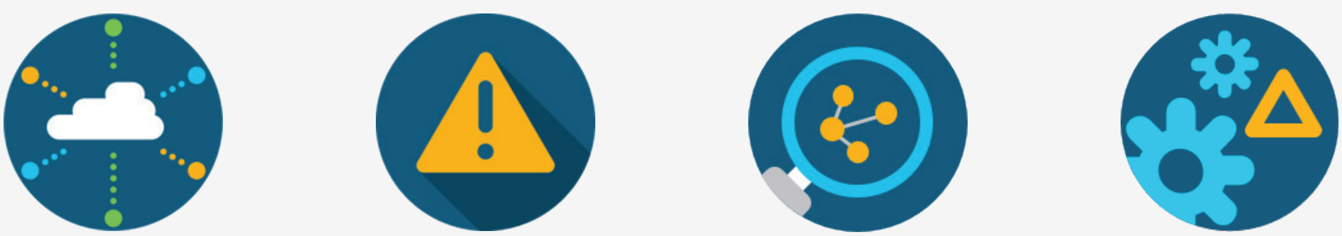

Scalable **Resilient** **Efficient** **Automated**

The Cisco Peering Fabric solution

We refer to this design as the Cisco Peering Fabric solution, and it introduces two options for fabric topology and peer termination. The first option is similar to traditional peering deployments. It collapses the peer termination and core connectivity network functions into a single physical device using the device's internal fabric to connect each function. The second option uses a fabric that isolates the network functions into separate physical layers, which are connected using an external fabric that runs over standard Ethernet.

The design minimizes the loss of peering capacity, which is important for both ingress-heavy service providers and egress-heavy content providers. The loss of local peering capacity means that traffic must ingress or egress a suboptimal network port. Making a conscious design decision to spread peer connections, even to the same peer, across multiple edge nodes helps increase resiliency and limit traffic-affecting network events.

The design of the Cisco Peering Fabric solution makes it possible to:

- Extract operational data. You can embrace Model-Driven Telemetry to collect high-fidelity data at very high-frequency and enrich information with BGP Monitoring Protocol and Netflow/IPFIX.
- Automate operations. You can on-board new device with zero-touch provisioning (ZTP) and install Cisco Network Services Orchestrator (NSO) peering function packs to manage peers and common peering configuration.
- Optimize the network. You can use segment routing traffic engineering capabilities to signal ingress and egress traffic optimization.
- Make your fabric trustworthy. You can protect against distributed denial of service (DDoS) attacks and detect hijacking of border gateway protocol (BGP) routes.

Solution components

The Cisco Peering Fabric solution introduces a simplified control-plane built upon IPv4/IPv6 with segment routing. In the collapsed design, each peering node is connected to External Border Gateway Protocol (EBGP) peers and upstream to the core.

In the distributed design, network functions are separated. Peer termination happens on peering fabric leaf (PFL) nodes. Peering fabric spine (PFS) aggregation nodes are responsible for core connectivity and perform more advanced label edge router (LER) functions. The PFS routers use ECMP to balance traffic between PFL routers and are responsible for forwarding within the fabric and to the rest of the provider network. Each PFS acts as an LER, incorporated into the control-plane of the core network.

The Cisco Peering Fabric solution uses telemetry to enable an unprecedented level of insight into network and device behavior. Common tasks are automated such as peer interface configuration, peer BGP configuration, and adding physical interfaces to an existing peer bundle. It also supports ZTP operation for automated device provisioning.

The Cisco NCS5500 platform is ideal for edge peer termination, given its high-density, large RIB and FIB scale, buffering capability, and IOS-XR software feature set. The Cisco Peering Fabric solution also includes advanced security capabilities using BGP Flowspec and QoS Policy Propagation using BGP or QPPB.

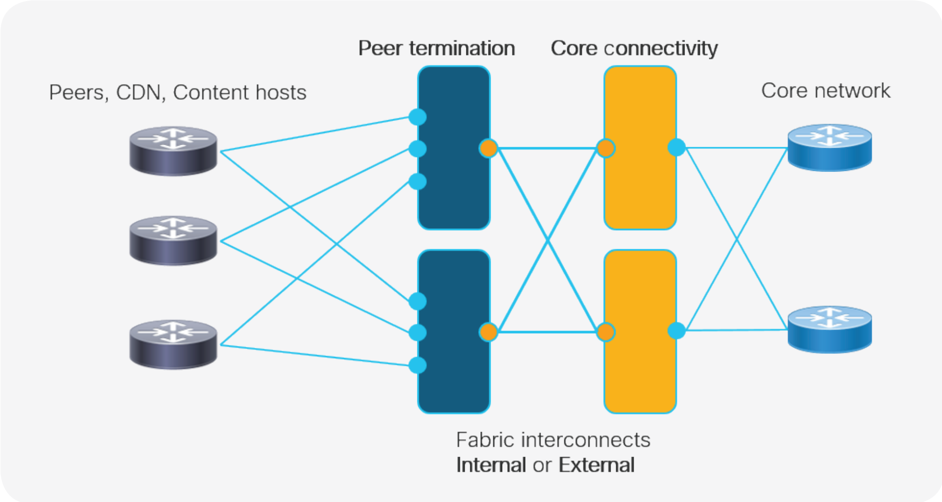

Figure 1. The Cisco Peering Fabric solution design

Use cases

The design of the Cisco Peering Fabric solution is suitable for multiple use cases.

Traditional IXP peering to peering fabric migration

A traditional service provider IXP design typically uses one or two large modular systems that terminate all peering connections. Because providers often have space and power constraints, they use a collapsed design where the minimal set of peering nodes not only terminates peer connections but also provides services and core connectivity to the location.

The Cisco Peering Fabric solution uses high-density, low-footprint hardware that requires less space than older generation modular systems. Many older systems provide densities at approximately 4x100GE per rack unit, while the Cisco solution PFL nodes start at 24x100GE or 36x100GE per 1RU with high forwarding information base (FIB) capability. Due to the superior space efficiency, there's no longer a limitation of using just a pair of nodes for these functions. In either a collapsed function or distributed function design, peers can be distributed across a number of devices to increase resiliency and lessen collateral impact when failures occur. Figure 2 shows a fully distributed fabric, with peers distributed across three PFL nodes, each with full connectivity to upstream PFS nodes.

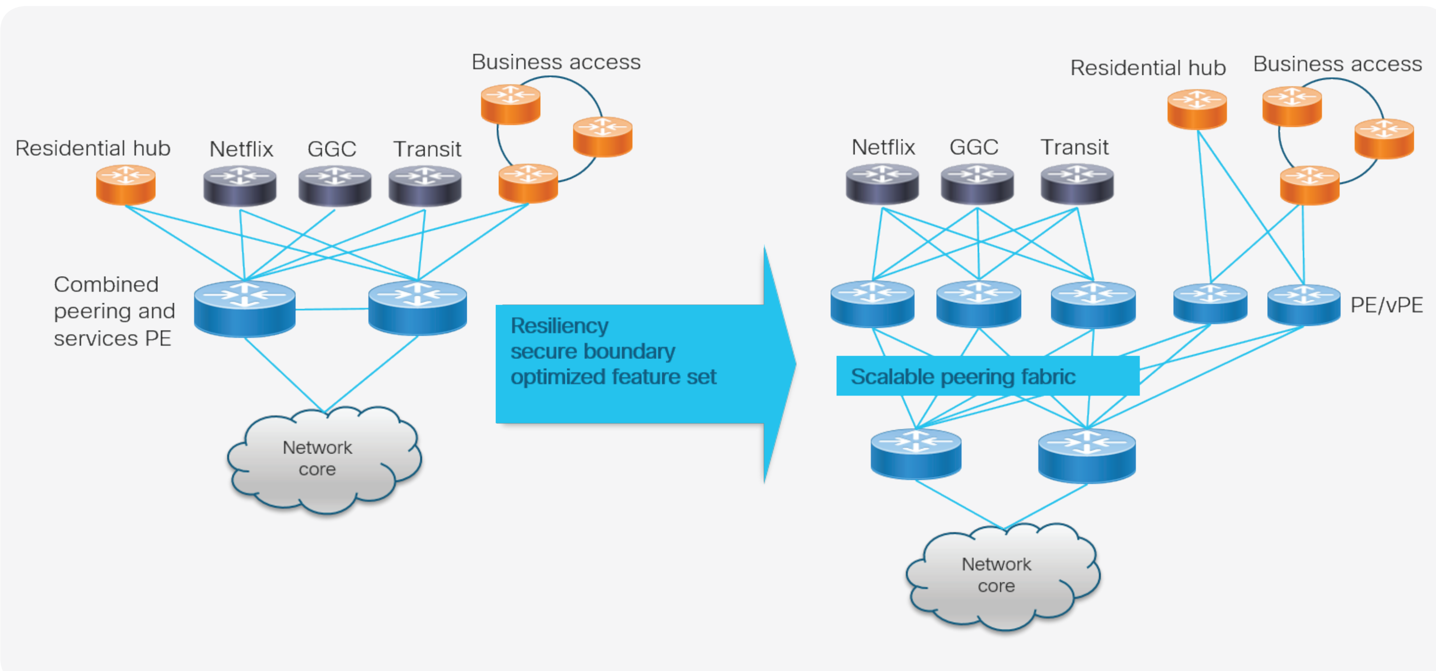

Figure 2. Fully distributed fabric

Peering fabric extension

Sometimes peering facilities within close geographic proximity may need to integrate into a single fabric. This situation can occur if there are multiple third-party facilities in a close geographic area that each have unique peers, which need to be connected. Multiple independent peering facilities within a small geographic area also may exist, yet you don't want to install a complete peering fabric into them. In those cases, you can connect remote PFL nodes to a larger peering fabric using optical transport or longer range gray optics.

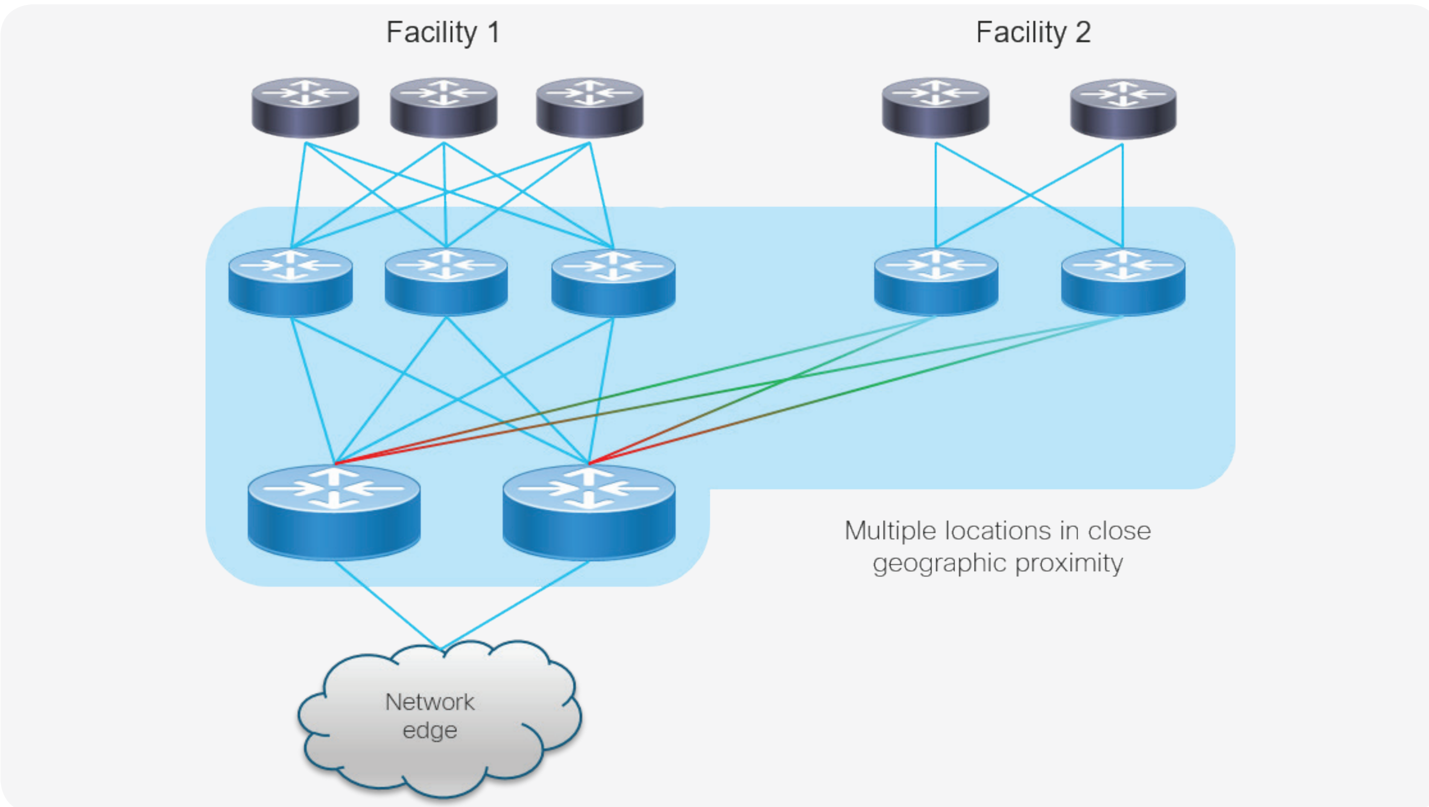

Figure 3. Peering fabric extension

Localized metro peering and content delivery

To drive greater network efficiency, content sources should be placed as close to the end destination as possible. Traditional wireline and wireless service providers have heavy inbound traffic from content providers delivering over the top (OTT) video. Providers may also have their own IP video services going to on-net and off-net destinations using a service provider content delivery network (CDN). Peering and internal CDN equipment can be placed within a localized peer or content delivery center, connected using a common peering fabric. In these cases, the PFS nodes connect directly to the metro core to enable delivery across the region or metro.

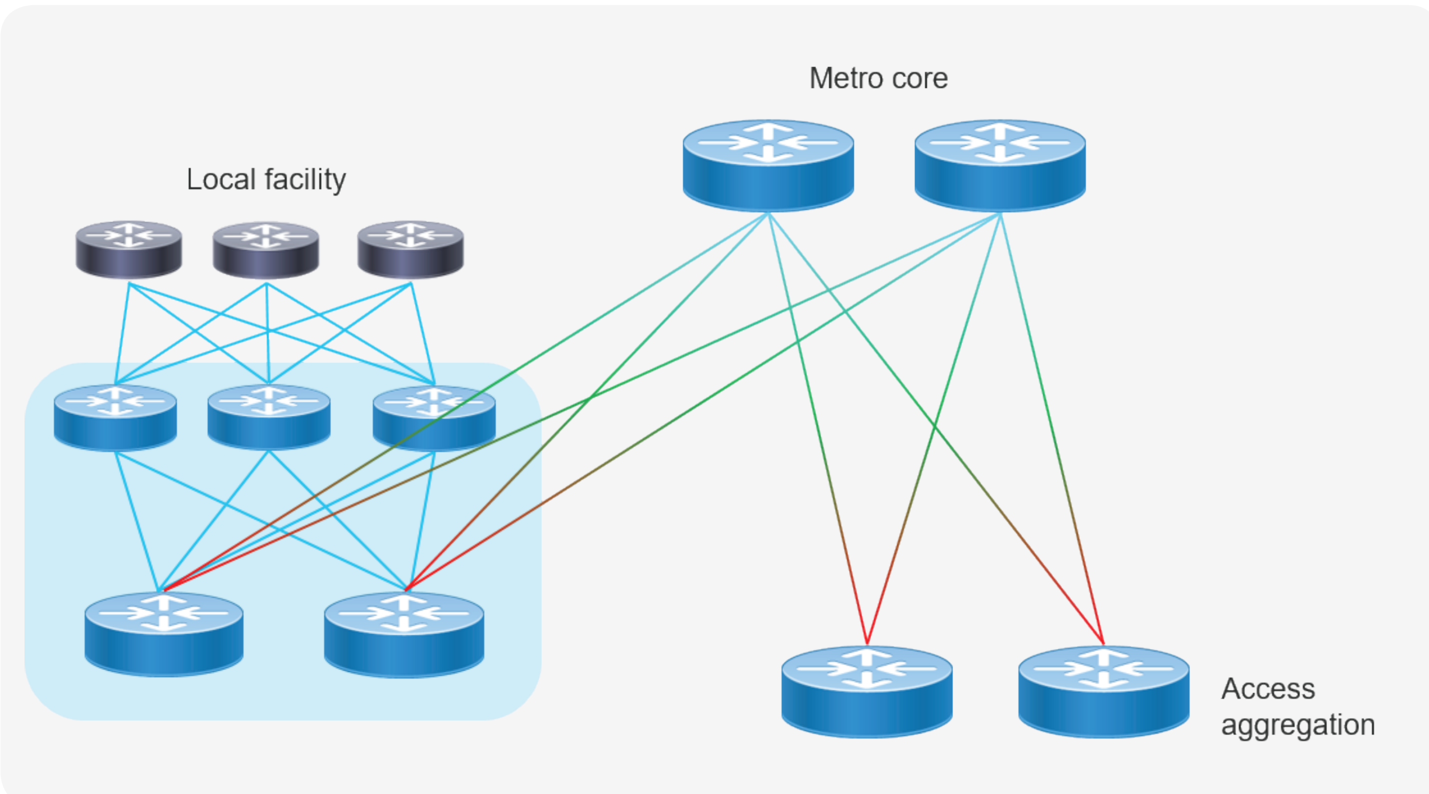

Figure 4. Localized metro peering and content delivery

Express peering fabric

An evolution to localized metro peering is to interconnect the PFS peering nodes directly or a metro-wide peering core. The main driver for direct interconnection is minimizing the number of router and transport network interfaces traffic must pass through.

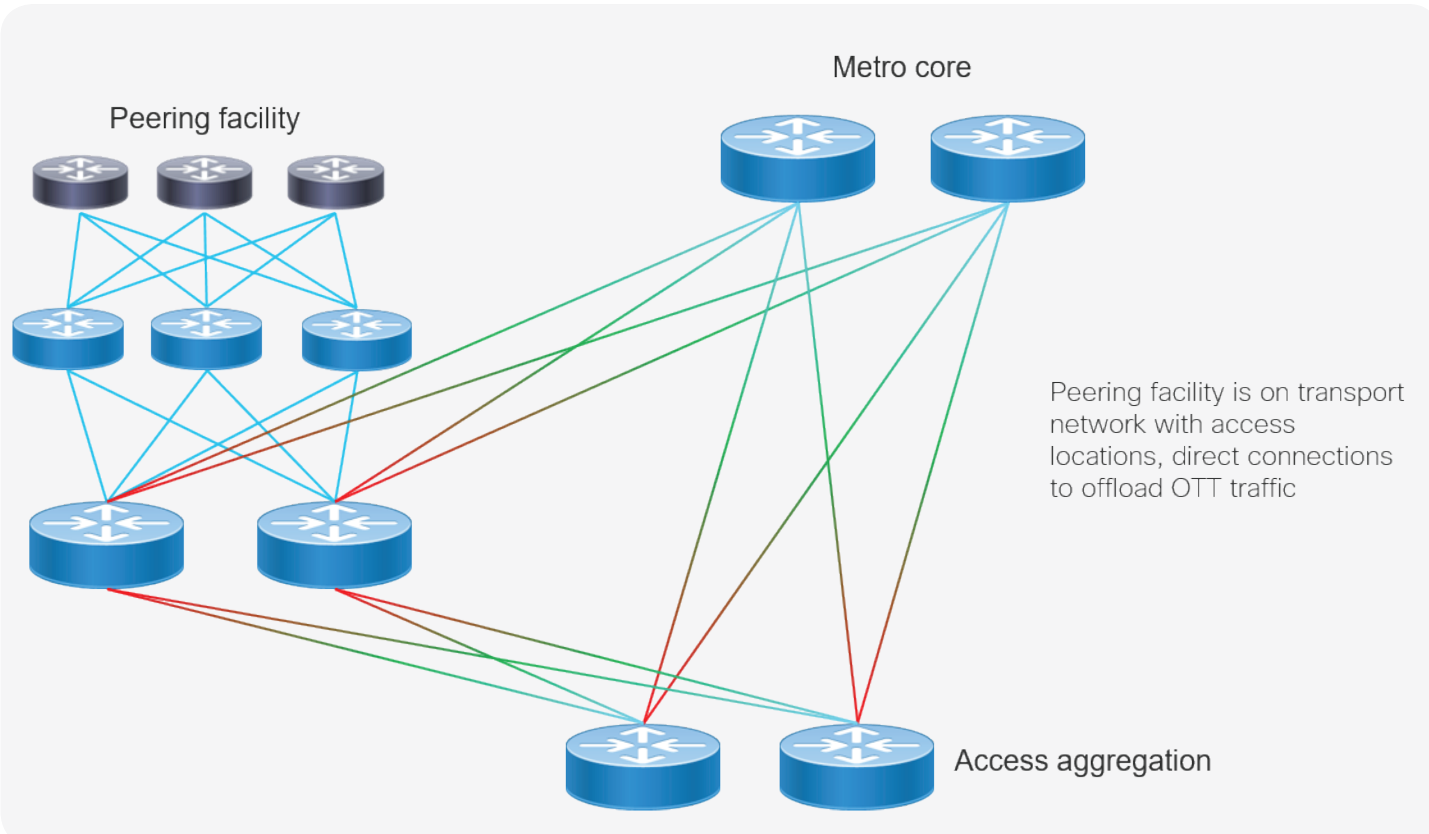

Figure 5. Express peering fabric

Data center edge peering

To serve traffic as close to consumer endpoints as possible, a provider may construct a peering edge attached to an edge or central data center. As gateway functions in the network become virtualized for applications such as virtual provider edge (vPE), virtual customer premises equipment (vCPE), and mobile 5G, the need to attach Internet peering to the service provider data center becomes more important.

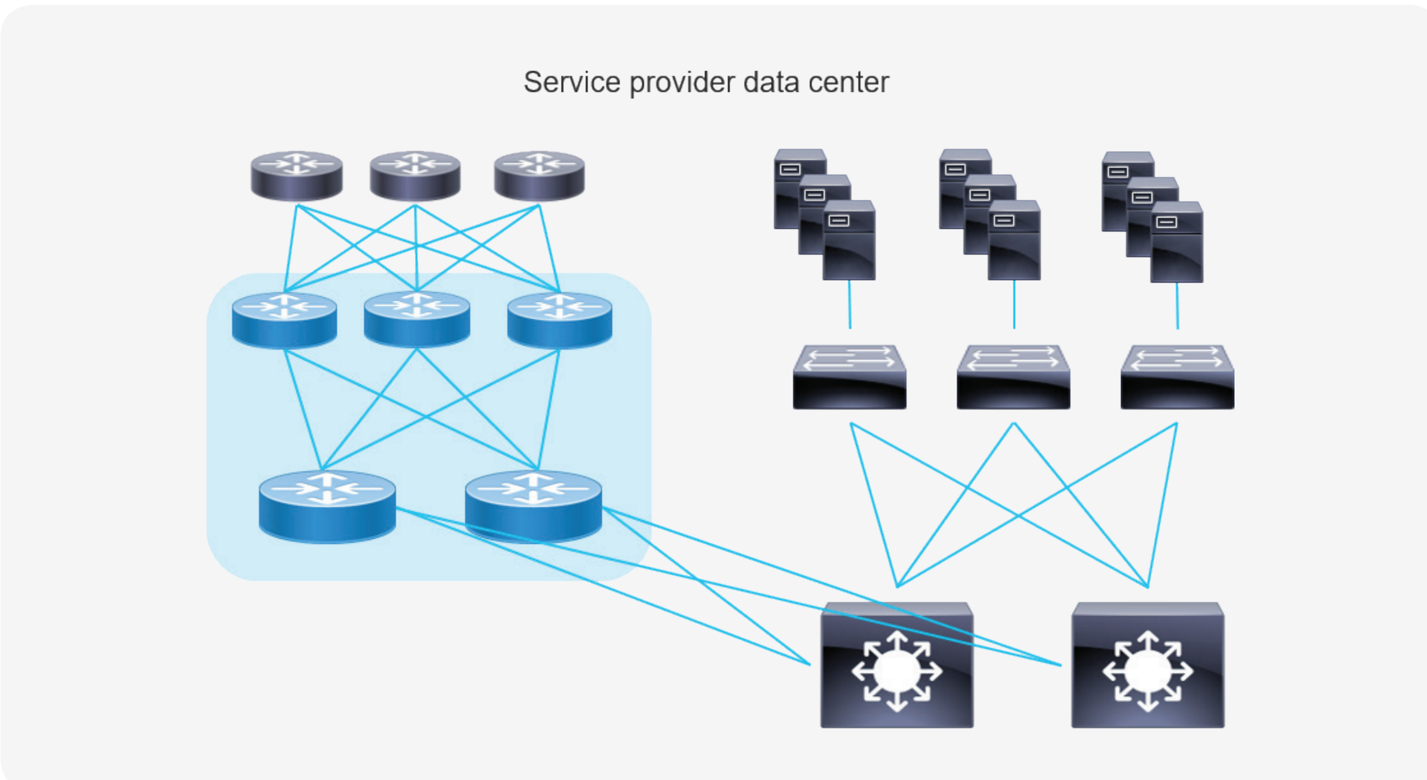

Figure 6. Data center edge peering

Video and Cloud

Video

It would take an individual more than 5 million years to watch the amount of video that will cross global IP networks each month in 2021. Every second, a million minutes of video content will cross the network by 2021.

Globally, Internet traffic is expected to grow 26 percent CAGR through 2022 and 82 percent of the traffic is expected to be video.[3]

Cloud

Global cloud IP traffic will more than triple (3.3-fold) over the next 5 years. Overall, cloud IP traffic will grow at a CAGR of 27 percent from 2016 to 2021.

Global cloud IP traffic will account for 95 percent of total data center traffic by 2021.[4]

Adapting to new demands

The rise in video and cloud traffic is placing new demands on the network. Today, peering needs to move closer to the edges of the network to serve wireline broadband subscribers along with high-bandwidth 5G mobile users. Reducing the distance and network hops is a key priority for service providers in lowering network costs. The right design can help drive efficiency into service provider networks, and the Cisco Peering Fabric solution offers the scalability, flexibility, and security you need.

Learn more about peering fabric

If you are facing increasing demands on your network and considering new approaches to the design, you can find more detailed information on peering fabric at: https://xrdocs.io/design/

1. Cisco Visual Networking Index: Forecast and Methodology, 2017-2022.2018.
2. Cisco Global Cloud Index: Forecast and Methodology, 2016-2021. 2018.
3. Cisco Visual Networking Index: Forecast and Methodology, 2017-2022. 2018..
4. Cisco Global Cloud Index: Forecast and Methodology, 2016-2021. 2018.