Cisco Agile Metro

This chapter covers the overview, benefits, architecture, technologies, and deployment models of Cisco Metro solution.

Traditional edge routing networks

Traditional edge routing network architecture

Traditional Edge routing networks are networks where the edge is a set of dedicated routers placed centrally in the network. Users usually have multiple Edge networks for services such as business VPN (L2 and L3), subscriber services, internet peering, data center interconnect (DCI), cloud gateways, and so on.

Figure 1: Traditional edge routing network architecture

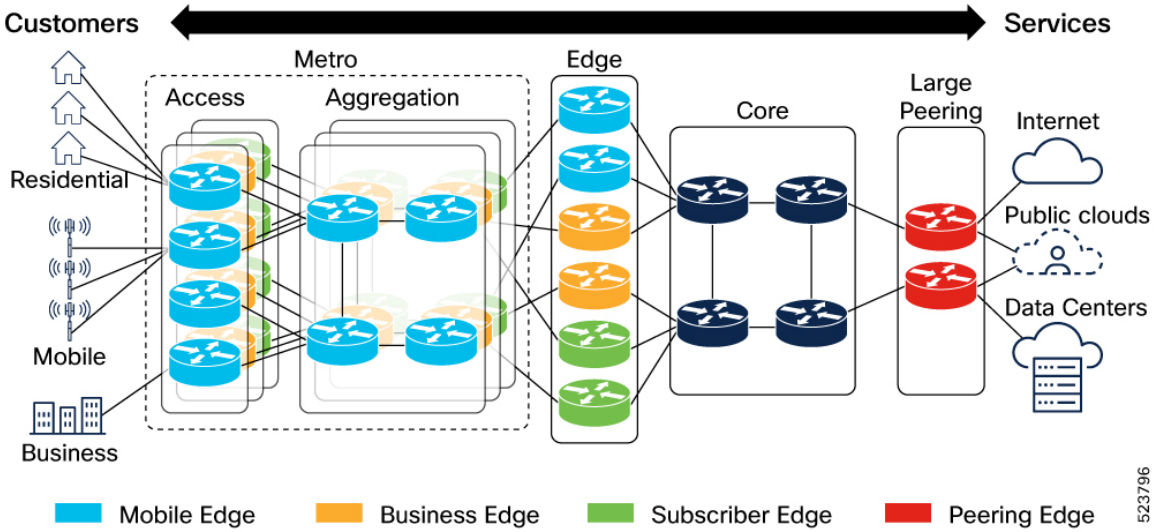

Challenges with traditional edge routing networks

Traditional edge routing networks pose challenges in these areas:

- Distributed services: The bandwidth profile and subscriber density for a distributed edge have changed significantly.
- Failure Impact: The large radius for failures demands more complex geographically redundant homing. It also necessitates the use of network hardware with redundant processors, fabric cards, line cards, power supplies, ports, and chassis. This increases the cost and complexity.
- Service optimization: Service providers might have different edge networks for various services. Optimization of these is difficult. The cost of building a solution that meets all the service requirements on a single platform is high. The most demanding services usually set the price for all others.
- Operational efficiency: Limits operators who want simplicity and efficiency.
- Platform lock-in: Large systems remain in place for many years thereby slowing down innovation.
- Upgrades: Upgrades are delayed as much as possible. Forklift upgrades have become a norm due to compatibility challenges between technology generations.

Business challenges

Internet traffic saw a compounded annual growth rate of 30% or higher over the last ten years. Primary challenges for modern networks include proliferation of devices, increase in end-user bandwidth speed, and the continued movement of applications to the cloud. Modern networks are also increasingly encountering challenges brought about by the era of Artificial Intelligence (AI). These include shifts in upstream and downstream traffic patterns, considerations for the optimal placement of AI inference processes, the need for dynamic, resilient, and secure connectivity, as well as ensuring service assurance and continuity.

Traditional edge routing designs face challenges by new traffic and business realities that include

- ever growing bandwidth with flat average revenue per user (ARPU)

- increasing cost and complexity of scaling centralized network architectures
- content and application evolution towards distributed network architectures, and
- arising connectivity matrix and traffic pattern for AI.

Solution

You must build networks to handle the advanced services and increased traffic associated with modern network services. Networks must evolve so the infrastructure layer can keep up with the service layer. The result of these shifts is driving traffic away from centralized delivery to a more distributed network.

Moving forward, we must consider new network architectures as design options that include these key aspects:

- Fully distributed and fixed routing systems to deliver services at any place in the network
- Fabric models delivering scale-out networks that can be built on-demand
- Flexible deployment options matching provider requirements
- Network simplification at all layers of the network

Cisco Agile Services Networking

Cisco Agile Services Networking is an architecture evolution of Cisco Converged SDN Transport (CSDN-T) that is focused on converging network infrastructure in multiple dimensions to change the way networks are built. The Metro solution considers edge as a set of functions which can be enabled anywhere in the network.

Cisco Agile Metro

Cisco Agile Metro is a dynamic and flexible edge solution that is part of Cisco Agile Services Networking. The solution introduces new Silicon One A100, K100, and P100-based fixed and centralized routers and line cards to deliver improved experiences for residential, business and mobile services with a network that is simpler and more cost-effective to build, operate, and scale from locations closer to end-users.

The Agile Metro architecture focuses on these key aspects:

- Enhanced scale and resiliency through distributed networking
- Simplified packet transport and overlay services
- Simplified and converged business, residential, and transport infrastructure
- Enhanced automation

Benefits of Agile Metro

These are the key benefits of Agile Metro:

- Technology benefits:
  - High-capacity edge silicon
  - Convergence of network service functions

- Flexible network design and systems to fit any size location in the network

- Business benefits:
  - Deliver services closer to users and applications
  - Cost savings
  - Sustainability benefits

- Operational benefits:
  - Improved services resilience
  - Network efficiency
  - Enhanced operations through network automation and orchestration

Agile Metro architecture

The Metro network evolution is driven by increasing bandwidth demands, resulting in network functions distributed in the network closer to the end user. This evolution is driving a consequent network architecture evolution. The classical split between access, pre-aggregation, aggregation, and edge leaves room for a more homogeneous network without distinct boundaries between the domains.

Figure 2: Cisco Agile Metro architecture

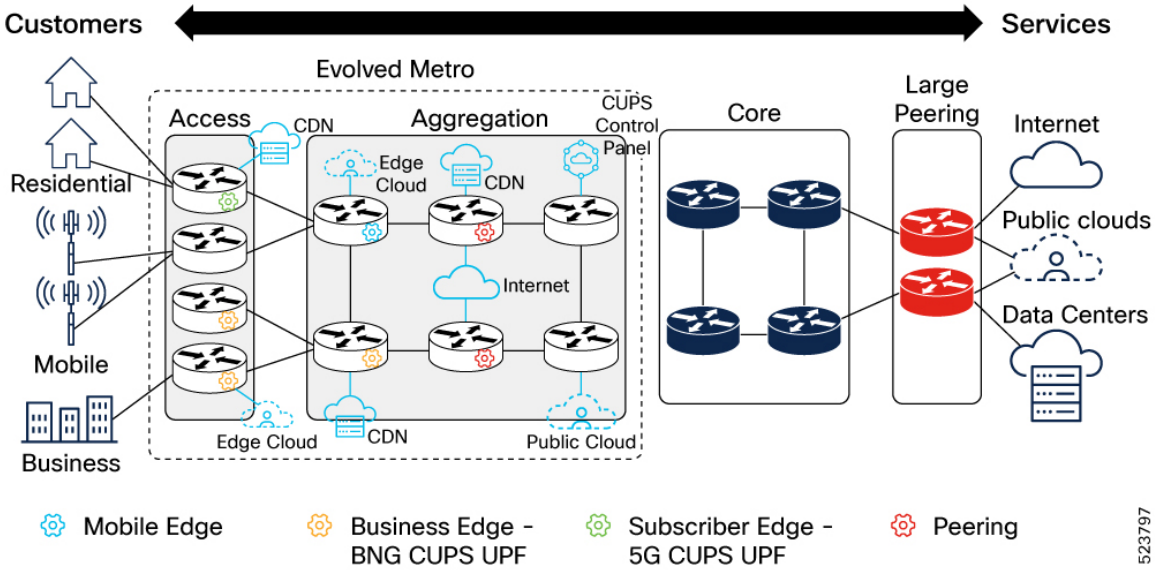

How the Agile Metro architecture differs from the traditional edge routing networks

The Agile Metro architecture differs from the traditional edge routing networks in terms of these key aspects:

- Automation-first architecture that is focused on simplified deployment and operation
- Edge fabric for service and bandwidth scale, which is the concept of taking a traditional modular chassis multi-service edge and disaggregating it into a multi-device fabric with greater scale and flexibility
- Edge service placement flexibility using common silicon and feature sets at any place in the network
- A consolidated architecture to handle trends in metro deployments
- A vehicle to introduce new service generating elements such as Edge Service Gateway and inline services

Key pillars of Agile Metro architecture

These are the key pillars of Agile Metro architecture:

- Wide range of supported interfaces:
  - 1/10/25/50/100/400GE and beyond on unified family of Metro devices
  - Any speed user–network interface (UNI) with any service
  - High speed network-to-network interfaces (NNI) and Routed Optical Networking
- Simplified connectivity model and protocols:
  - Segment Routing IPv6 (SRv6) and SR-MPLS underlay networks; SRv6-TE and SR-MPLS TE for advanced Traffic Engineered use cases
  - Secured infrastructure using Trusted Cisco platforms and advanced distributed DDoS protection
  - Co-existence with legacy underlay and overlay technologies
- Business, residential, and mobile subscriber services:
  - EVPN and L3VPN in services layer
  - Private Line Emulation (PLE) for bit-transparent transport of Ethernet and non-Ethernet (OTN, SONET, Fiber Channel)
  - Next-generation subscriber edge using control plane and user plane separation (CUPS)
  - Converged business and subscriber access using Cisco Routed PON
- High performance end-to-end timing and synchronization
- Automation across all components in the architecture covering provisioning, monitoring, and service assurance

High-level use cases of Agile Metro

The Agile Metro architecture covers these high-level use cases:

- Next-generation residential subscriber networks deployments
- Enterprise business services

- Mobile network IP transport
- Centralized and edge data center connectivity including networks that are built to support artificial intelligence
- Internet peering, content delivery, and cloud connectivity

Agile Metro technologies

This topic covers the various technologies used in Cisco Agile Metro architecture.

Network technologies and protocols

The table gives a comparison of the common network technologies and protocols that are used in legacy networks vs. the Agile Metro.

Table 1: Common network technologies and protocols used in legacy networks vs. the Agile Metro

Network technology and protocol	Legacy network		Agile Metro
xVPN Services	LDP	BGP	BGP for all L2VPN, and L3VPN
IP Network Scaling	BGP-LU		Segment Routing
Traffic Engineering Fast Reroute	RSVP-TE		
MPLS Overlay Protocol	RSVP-TE	LDP	
IPv6 Transport Overlay	None		
IP to DWDM Transition	Transponder or Muxponder		Routed Optical Networking
	Grey Router Interface		
Private Line Services	Dedicated OTN	Dedicated Ethernet over DWDM	Private Line Emulation
Subscriber BNG	Physical Integrated BNG		Cisco CUPS using Cloud Native BNG
PON Access	Dedicated PON Equipment		Cisco Routed Passive Optical Networking

Metro Edge Fabric

This section details the new disaggregated Metro Edge Fabric, including its components and distributed control plane.

The Metro Edge Fabric is a component of Agile Metro architecture that is designed to provide scalable edge services termination. The Metro Edge Fabric is designed to enhance network efficiency and scalability by separating network functions into distinct physical layers. Cisco Fabric-based Edge solution is a composition of multiple routers in a leaf-spine architecture to accommodate required functionality and scale that cannot be met in a standalone multi-service edge (MSE) model.

Figure 3: Metro Edge Fabric in Metro architecture

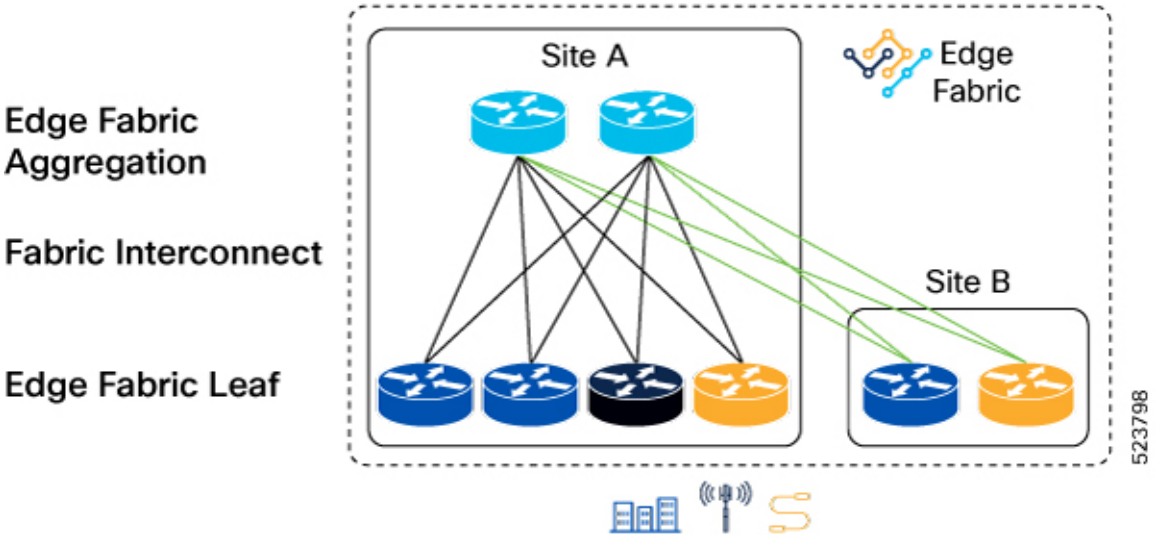

Edge Fabric leaf

Leaf nodes are the routers that are used for network service termination use cases. You can split the specific services across a set of leaf devices based on the design and network services. The leaf nodes may be collapsed into a universal leaf for all functions or split between different network or even VPN service type. All Cisco IOS XR platforms can be used as a leaf in the deployment depending on the feature requirements and feature scale.

Edge Fabric spine

The Edge Fabric aggregation routers or spines are the nodes that provides underlay connectivity to all leaf nodes that include service termination nodes, core networking connecting nodes, edge DC connecting nodes, and so on. These spine nodes act as L3/SR-MPLS switch that carry overlay services across leaf nodes. Spine nodes have advanced policy-based traffic management functionalities to support end-to-end QoS for selective overlay services.

Fabric interconnect

The fabric interconnects are the links connecting leaf nodes to spine nodes. Each leaf node must be connected to every spine node to provide maximum resiliency and load balancing across the fabric. It is recommended to standardize local interconnects to one type—copper (CU) or active optical cables (AOC) being the most cost-effective method. Interconnects may also utilize WAN connectivity in the case of remote leaf devices. Longer distances can be covered using Routed Optical Networking components such as ZR/ZR+, DP04QSDD-ER1, and QSFP-DD 100G ZR coherent optics.

Fabric control plane

The Fabric uses standard routing protocols; it does not use proprietary communication between the elements. This allows providers to easily insert any type of node, including third-party node, into the fabric.

Routed Optical Networking

Routed Optical Networking is a Cisco architecture for simplified IP and optical convergence. It is a component of the Metro architecture, utilized as an interconnect technology for WAN links. The high-bandwidth networks require agile, flexible connectivity taking full advantage of the high bandwidth capabilities of next-generation NPUs. Routed Optical Networking must be utilized in places where longer distance interconnects are required between core and edge elements.

For details, see Cisco Routed Optical Networking Solution Guide.

Subscriber services

CUPS-based subscriber management system

The figure depicts the high-level architecture of Control Plane and User Plane Separation (CUPS)-based subscriber management system in the network.

Figure 4: CUPS architecture

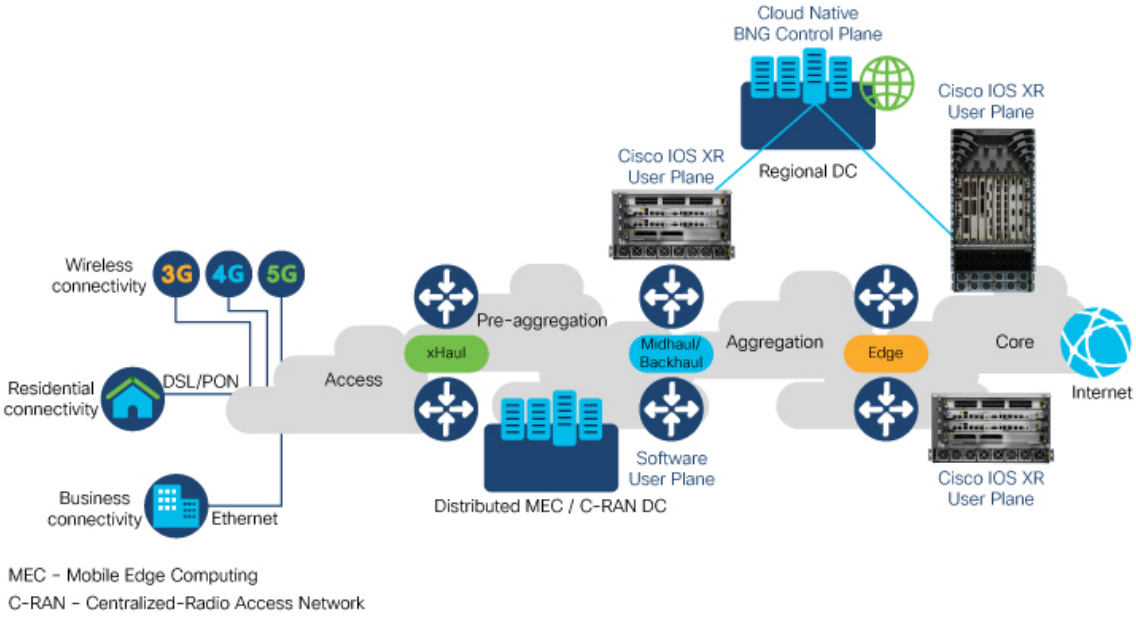

In the CUPS architecture, the subscriber control plane (CP) and user plane (UP) are separated. The CP is responsible for functions such as AAA of subscriber sessions, IP address management, forwarding policies and session specific protocols such as PPPoE, or IPoE. In addition to this, CP assumes the responsibility to onboard, monitor and manage each UP node through the standard interface. One CP instance can manage several UP instances distributed throughout the metro network.

In Metro Release 1.0, the UP is based on Cisco ASR 9000 Series Routers. The onboarding and provisioning of the BNG ecosystem is done using the Cisco NSO-based solution.

For details, see Cloud Native BNG Control Plane Configuration Guide, and Cloud Native BNG Control Plane User Guide.

Cisco Routed Passive Optical Networking

Cisco Routed PON is a solution created by combining the rich access routing portfolio from Cisco and a pluggable PON solution. The core of the solution is a smart SFP acting as a single port OLT, the PON controller software managing the OLT SFPs on a device, and the PON manager software managing the overall PON network deployment. The PON SFP can be plugged in to selective set of Cisco IOS XR routers (Cisco NCS 540 Series routers) on which the PON controller software is also installed. Cisco Routed PON manager is a centralized PON manager that manages several PON controllers. The connectivity between Cisco Routed PON manager and the PON controller is secured using TLS.

Figure 5: Cisco Routed PON deployment

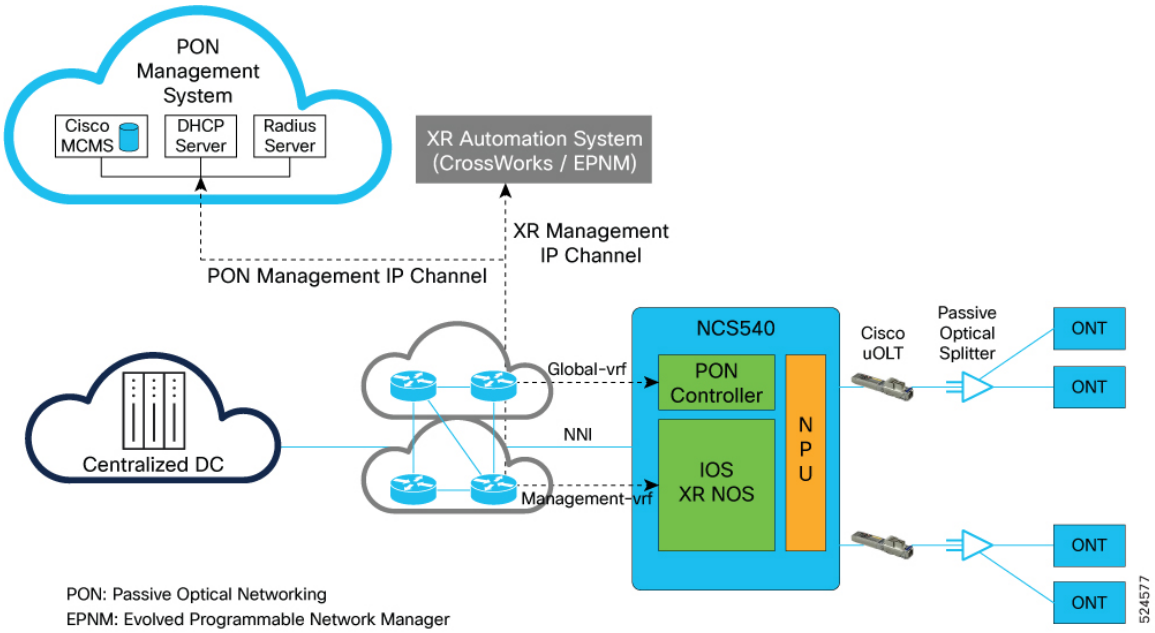

For details on Routed PON solution components, see the Cisco Routed PON Deployment section.

Agile Metro deployment models

The Agile Metro architecture has the flexibility to support a mix of deployments models based on factors such as service scale, service traffic bandwidth, interface types, and other network constraints.

Figure 6: Cisco Agile Metro deployment models

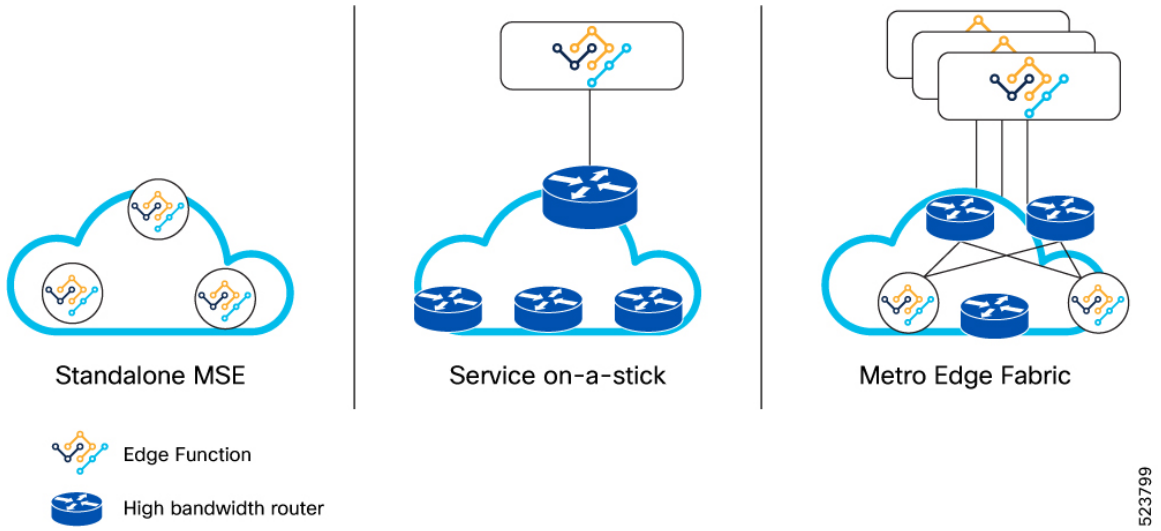

Agile Metro supports these deployment models:

- Standalone MSE: Traditional multi-service Edge (MSE) deployment model for a single Metro edge router where all functions and scale are handled at a single box level. The model is best suited for smaller edge sites which do not need the higher service scale of the Metro Edge Fabric distributed solution. Agile Metro enables distribution of devices with MSE functions deeper into the network, increasing overall network efficiency.
- Service on-a-stick: Out-of-line network functions model where the edge services are taken out from the main forwarding path and hosted elsewhere, thereby decoupling the transport and service architectures. The model is ideal for deployments with a high volume of low bandwidth services compared to the volume of transit traffic through the aggregation router.
- Metro Edge Fabric: Scale out Edge services infrastructure model which is a composition of multiple routers in a leaf-spine architecture. The leaf nodes themselves perform the Edge functions of a traditional MSE router but using a distributed horizontally scaled approach as opposed to a vertically scaled approach. There is flexibility in the nodes deployed as a fabric spine, with the primary attribute being they do not terminate many services and are primarily used as transit. The spine nodes could be new nodes in large aggregation sites or re-purpose core nodes already deployed. The key point of the fabric model is horizontal service edge scaling and simplified management using automation.

All models support the *Edge as-a-function* concept where a service Edge function can be deployed anywhere in the network.

For details on the deployment options for these models, see the Metro Edge Fabric deployment options, on page 13 section.

Network infrastructure deployment

Apart from the traditional network deployments, the Agile Metro architecture covers these new network infrastructure deployments:

- Metro network infrastructure: Modern underlay infrastructure including routing control plane, overlay service signaling, and timing using Precision Time Protocol (PTP), and Synchronous Ethernet (SyncE)
- Edge as-a-function: Edge features at any place in the network (PIN)
- Service on-a-stick: Out-of-line network functions
- Metro Edge Fabric: Scale out Edge services infrastructure

Metro network infrastructure

These are the Metro network infrastructure requirements for all devices participating in the Agile Metro .

- **Common packet transport**: The underlay control and forwarding plane technologies used to support all network services. The common packet transport updates the underlay design outlined in Cisco Converged SDN Transport design. The underlay network is built upon proven technology such as IS-IS for v4 or v6 IGP and Segment Routing using the SRv6 or SR-MPLS data plane.
- **Base infrastructure**: Additional protocols and technologies outside the scope of routing and forwarding required in the infrastructure to support network services. The solution architecture utilizes PTP G.8275.1 with SyncE for timing distribution, supporting G.8275.2 interoperability wherever required for endpoints requiring G.8275.2.
- **Base device automation**: Device support for APIs and protocols required to properly automate the network. The solution architecture leverages gNMI-based telemetry using OpenConfig and Cisco native models to monitor all aspects of the network. The interface is used by automation tools such as Crosswork Network Controller and Provider Connectivity Assurance to provide end-to-end infrastructure assurance.
- **Base operations**: Additional operational tooling and features that are required for operators to properly monitor and troubleshoot the network infrastructure and the services using the network.
- **Base security**: Base criteria for securing the network infrastructure. This does not cover service security.
- **Base QoS**: Base set of CoS and QoS feature support to ensure that network service SLAs are met across the network.

Edge as-a-function

The edge as-a-function deployment model removes the traditional location of edge services and allows edge to be served at any point in the network. That is, any router at any place in the network can act as an MSE device that is limited only by the scale required. This is becoming critical and common as higher touch edge services must be distributed to scale. Also, centralizing the edge services is no longer feasible due to service bandwidth growth.

The figure shows the evolution of fabric-based deployment from traditional Edge deployment.

Figure 7: Edge network evolution

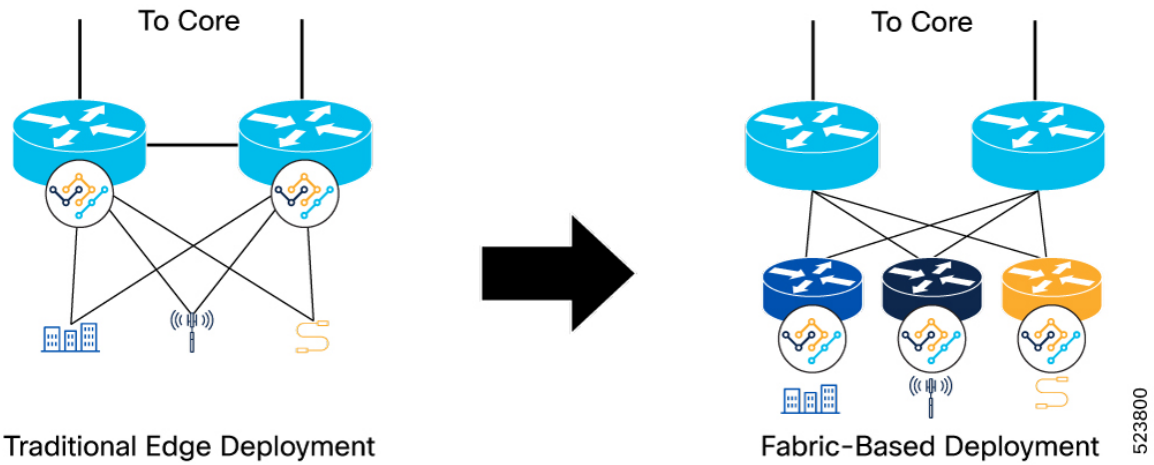

The table lists the characteristics of traditional edge deployment and Fabric-based edge deployment models.

Table 2: Traditional edge deployment vs. Fabric-based edge deployment

Traditional Edge deployment	Fabric-based Edge deployment
Large modular chassis	Service edge lead based on the function
Multi-service termination	Prime device for service edge
L2 or L3 aggregation that is common between end sites and edge PEs	Edge leaf can be in the same site or deeper in the network

Service on-a-stick

The service on-a-stick deployment is one where service functions are not performed on devices inline between the edge and core of the network. Instead, they are performed on devices out-of-line.

It is inefficient, and cost and power prohibitive to build higher touch network services into the distributed routing hardware when it may not be applicable to all deployments. You can do service on-a-stick deployment either as part of an edge fabric or a more centralized architecture where ingress and egress from the network function are connected to upstream routers. Service scale and bandwidth require a more distributed approach, while compute intensive functions may require them to be located at a centralized data center. The Metro architecture supports both these deployment options.

Figure 8: Service on-a-stick deployment model

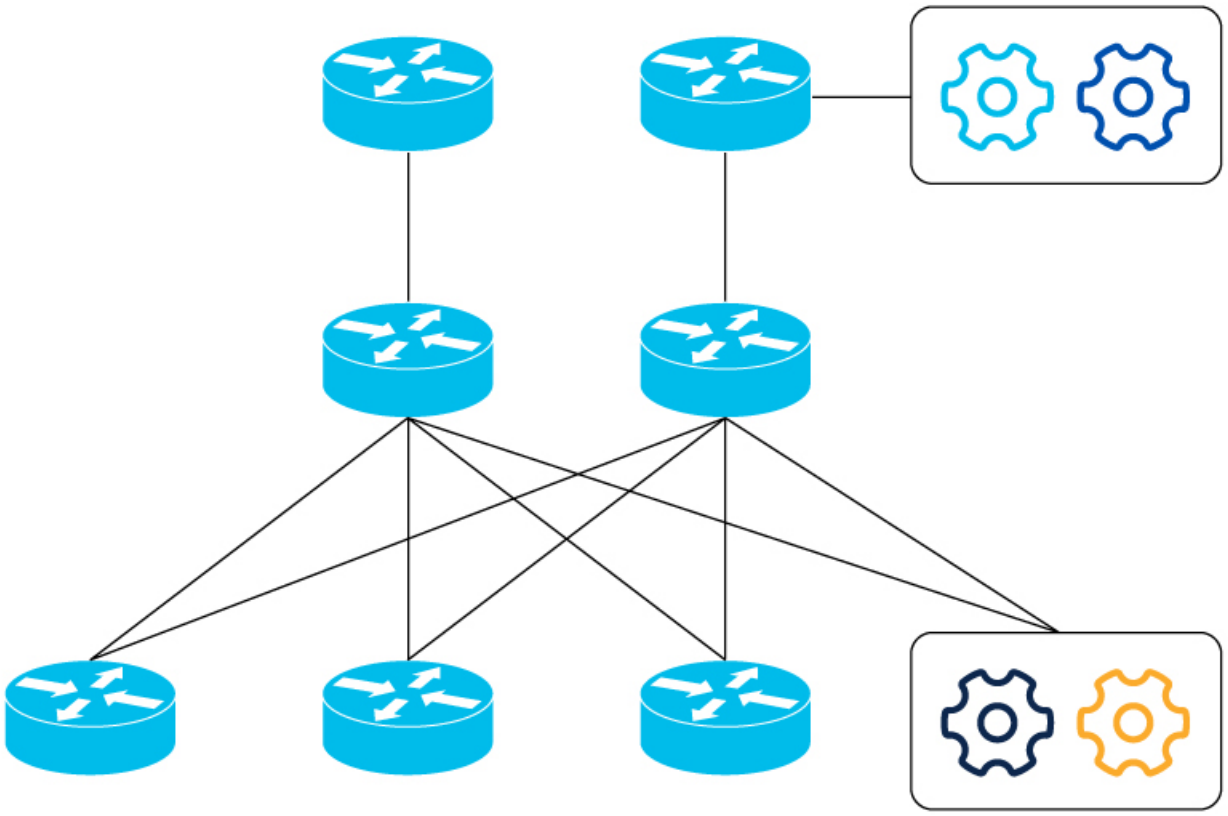

Service on-a-stick deployment model is suitable for these functions and services:

- High-touch functions that are applicable only to a small subset of traffic

  For example, a NAT appliance capable of handling exception traffic.

- High-touch services that require high performance and scale, such as firewall deployments, where firewalls are placed *on a stick* to perform security functions

In Metro Release 1.0, the service on-a-stick deployment is applicable to BNG for subscriber services.

Metro Edge Fabric

Creating a network capable of serving a variety of use cases on a common fabric is best served by decomposing a modular chassis into a set of leaf nodes with a specific service role. The disaggregated Edge Fabric functions as a leaf-spine topology with edge service facing ports located on leaf devices connected through fabric ports to an upstream aggregation node. For more details on the Metro Edge Fabric components, see the Metro Edge Fabric, on page 6 section.

Metro Edge Fabric deployment options

Metro Edge Fabric can have different deployment options based on various factors such as space, power, geographic diversity, or aggregate bandwidth required for specific services. Certain service types such as business edge services lend themselves to being distributed all the way to the access portion of the network.

Other types of services such as value-added inline services which are high compute and lower bandwidth maybe be better suited in a centralized deployment model.

Figure 9: Metro Edge Fabric deployment options

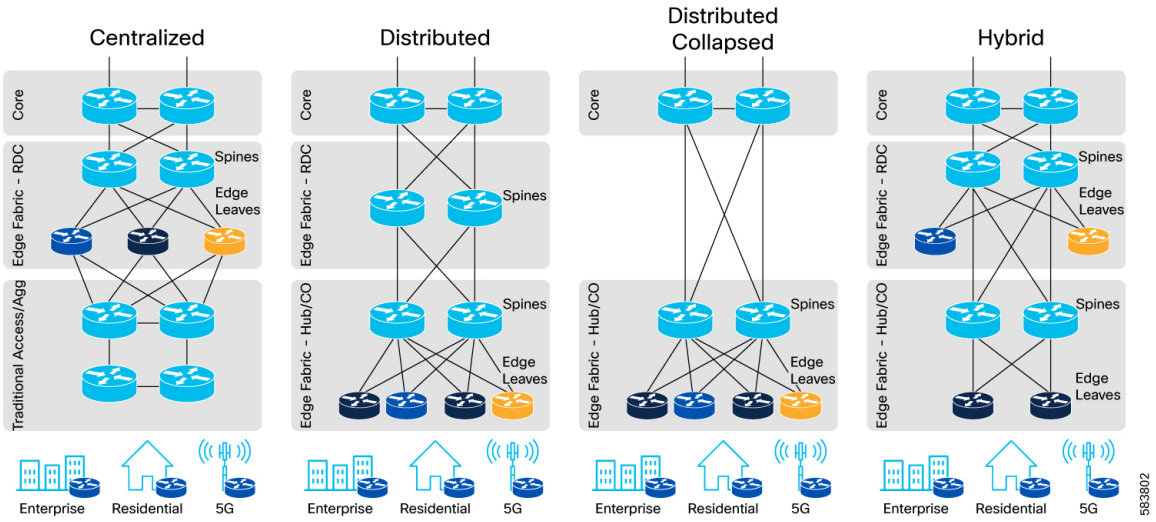

These are the main deployment options for Metro Edge Fabric:

- **Centralized**: This option follows the centralized MSE approach by housing the edge functions within a regional or core data center but distributing the edge services across a fabric instead of a traditional large modular MSE router. This case is ideal for types of services which are high compute and lower bandwidth such as value-added inline services.
- **Distributed**: This option fully distributes services to what would traditionally be the access or aggregation level of the network. This case is ideal for most service termination types such as business edge services (L2VPN, L3VPN, mobile backhaul, and dedicated internet access (DIA)) and subscriber edge services using a CUPS user plane.

  Distributed Collapsed: In this option, the spine nodes in the access or aggregation directly connect to the core routers without connecting to the intermediate spine nodes.
- **Hybrid**: This option distributes some edge services to further into the network, but also has some service termination at a centralized data center following the service on-a-stick network deployment.

Network infrastructure operations

The network infrastructure operations use cases of Agile Metro include network commissioning and infrastructure assurance operations.

Network commissioning

Zero Touch Provisioning

Zero Touch Provisioning (ZTP) is the method of provisioning network devices without manual intervention. ZTP helps to seamlessly onboard new devices across the network within a short span of time. ZTP also reduces the manual tasks required to scale network capacity.

Network infrastructure assurance

Infrastructure assurance refers to monitoring of the health of the base infrastructure network.

These are the factors contributing to the assurance of the network:

- Base Cisco IOS XR software features: SR-PM to measure latency and loss across logical links

  Starting with CNC 6.0 and the TSDN 6.0 there is an example function pack to configure and deploy performance measurement profiles across the network. This can be used to maintain consistent performance measurement definitions across the network.

- Provider Connectivity Assurance (PCA) container-based agents to measure network performance (latency, loss, and jitter) between edge devices to millisecond precision
- PCA compute SFP+ for higher accuracy to microsecond resolution when needed
- CNC to monitor and visualize the performance (bandwidth and latency) of the network
- PCA to aggregate infrastructure assurance data

These are the various network infrastructure assurance use cases of Agile Metro:

- Segment Routing performance monitoring (SR-PM)
- SRv6 integrated performance measurement (SR-IPM)
- SRv6 path tracing
- SRv6 traffic accounting
- Cisco Provider Connectivity Assurance

Segment Routing performance monitoring

Segment Routing performance monitoring (SR-PM) is enabled to perform link monitoring for the end-to-end metro network to support both infrastructure monitoring and end-to-end TE path monitoring. Latency measurement is performed on all links as per the hardware capability.

SRv6 integrated performance measurement

SRv6 integrated performance measurement (SR-IPM) supports loss, latency plus jitter, and liveness detection by using high frequency TWAMP packets between SRv6 endpoints. This feature is available on Cisco ASR 9000 Series Routers, and Cisco 8000 Series fixed routers and modular routers with Q200-based Silicon One ASICs.

SRv6 path tracing

SRv6 path tracing uses a small IPv6 header of 3 bytes that is added to the transit packets to record the outgoing interface, time, and load at each hop. At the head-end node these probe packets are sent out to each ECMP path across the network. This is done to see a true view of the network even when ECMP is utilized at each hop along the path between the source and the destination nodes. The data can then be used by various tools to analyze the health of the network.

SRv6 traffic accounting

SRv6 traffic accounting is applicable to infrastructure monitoring and long-term network capacity planning. SRv6 traffic account provides statistics for each egress locator per egress interface. That is, if ECMP is available on the head-end node, the statistics per Locator and OIF is recorded and made available through both CLI show commands and model-driven telemetry. The SRv6 traffic accounting feature is available on Cisco ASR 9000 series routers, and Cisco 8000 Series fixed routers and modular routers with Q200-based Silicon One ASICs.

Cisco Provider Connectivity Assurance

Agile Metro leverages the existing Cisco Provider Connectivity Assurance (PCA) components for automation workflows.

Cisco Provider Connectivity Assurance use cases

These are the Cisco PCA use cases of Agile Metro:

- Monitor loss, latency, and jitter between infrastructure service endpoints (PE to PE).
- Monitor loss, latency, and jitter.

Cisco Provider Connectivity Assurance deployment

These are the Cisco PCA deployment use cases of Agile Metro:

- PCA hardware-based agents at each endpoint- both inline and out-of-line
- PCA hardware-based agents at one endpoint with Cisco device acting as a reflector
- Use of PCA to monitor the overall health of the metro network including any fabric elements

Network infrastructure security

The network infrastructure security use cases of Agile Metro include base device security, and edge protect use cases.

Base device security

The base architecture provides a level of device security to protect against common device-level attacks.

Edge Protect DDoS

Cisco Edge Protect DDoS is a solution that is used to add another layer of protection against DDoS attacks to the network infrastructure.

In the initial release of Metro, these use cases are covered by Edge Protect—Mobile Edge Protect on Cisco NCS 540 Series Routers, and Edge Protect for peering use cases on the Cisco NCS 5500 and NCS 5700 Series Routers.

Edge Protect is a small, lightweight, and containerized DDoS detection engine, running inside Cisco IOS XR operating system on selected Cisco routers. As a feature of the router itself, the solution provides out-of-path, full detection, and granular mitigation capabilities against volumetric DDoS attacks. This allows the router to become the first line of defense against DDoS attacks where attack traffic can be blocked in-place and no longer needs to be backhauled to dedicated scrubbing infrastructure. This gives you the opportunity to extend

DDoS protection to the edge of the network which would otherwise be cost-prohibitive and negatively impacting the application latency and the quality of experience.

You can apply this solution to the existing install base and to new router deployments without requiring any additional hardware.

The figure shows the peering and internet custom use case for Edge Protect.

Figure 10: Peering and internet custom use case for Edge Protect

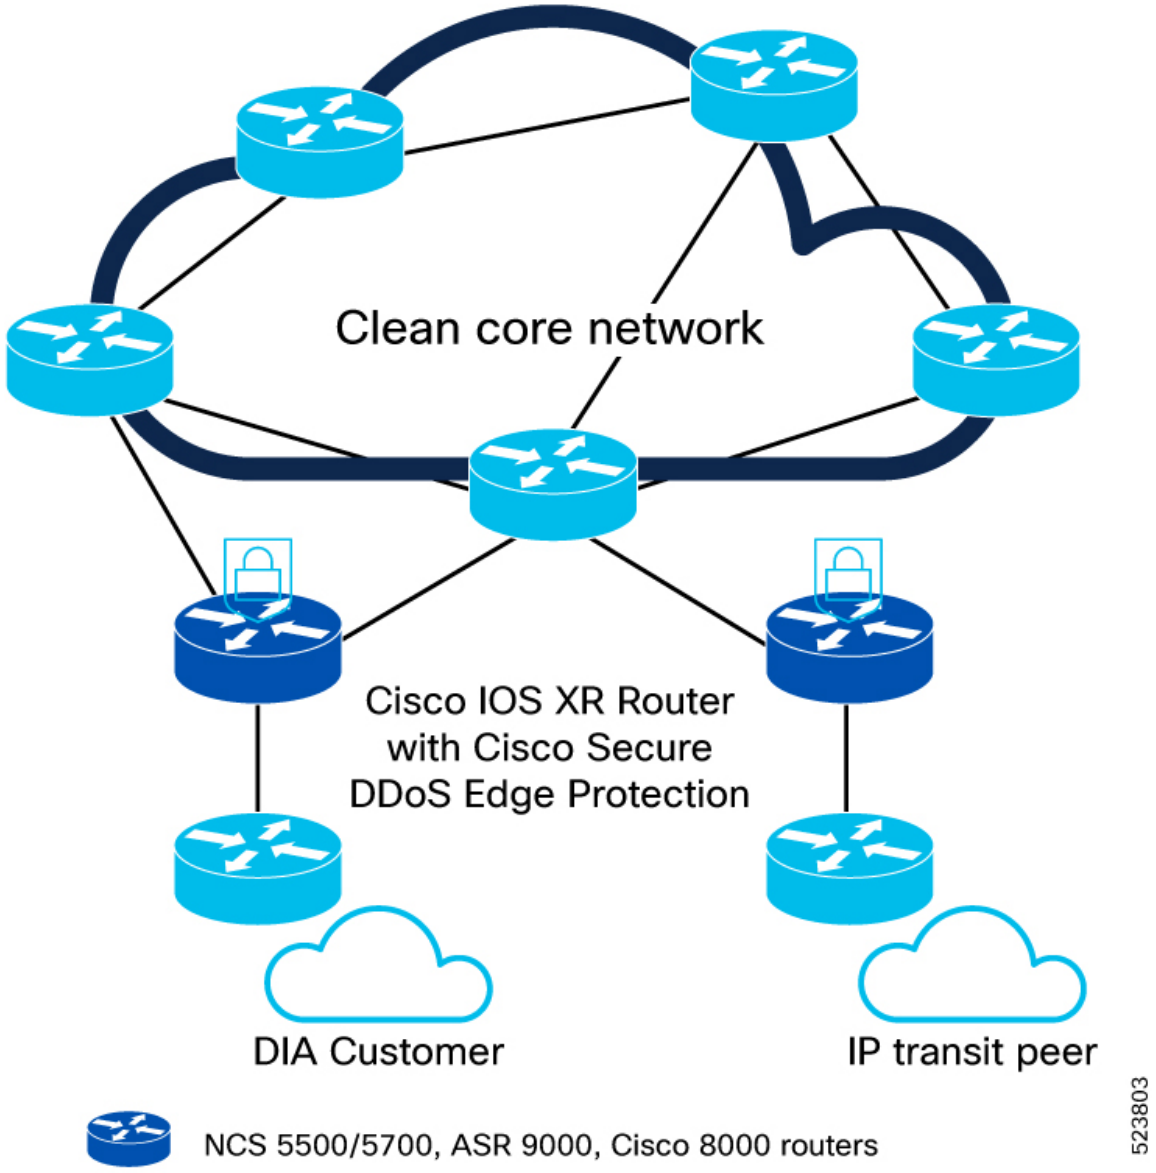

The figure shows the mobile network protection use case for Edge Protect.

Figure 11: Mobile network protection use case for Edge Protect