Quantum-Safe SRv6 Fabric:
Intent-Based Routing for Mission-Critical Networks

Segment Routing with IPv6 across Post-Quantum Cryptography-ready MACsec and IPsec transports

August 2026

Introduction

Historically, highly engineered, mission-critical networks relied on legacy protocols like MPLS and RSVP-TE. While functional in their time, these architectures are fundamentally fragile. They suffer from constant signaling storms and demand massive core-router memory databases to maintain state, severely limiting network scalability, agility, and convergence speed.

Compounding these architectural bottlenecks is a rapidly evolving cyber-threat landscape. Sophisticated adversaries are actively employing "Harvest Now, Decrypt Later" (HNDL) strategies—intercepting and storing encrypted mission-critical data today with the intent to decrypt it tomorrow using advanced quantum computing. Consequently, deploying quantum-proof infrastructure is no longer a future consideration; it is an immediate need for digital resiliency and the defense of critical national infrastructure.

Today, mission-critical defense networks look like this: a layered, highly segmented architecture separating diverse physical transports (like IP, Carrier Ethernet, LEO Satellite, and 5G) from routing underlays and logical overlays, all governed by end-to-end orchestration.

Figure 1. Layered, segmented mission-critical network architecture

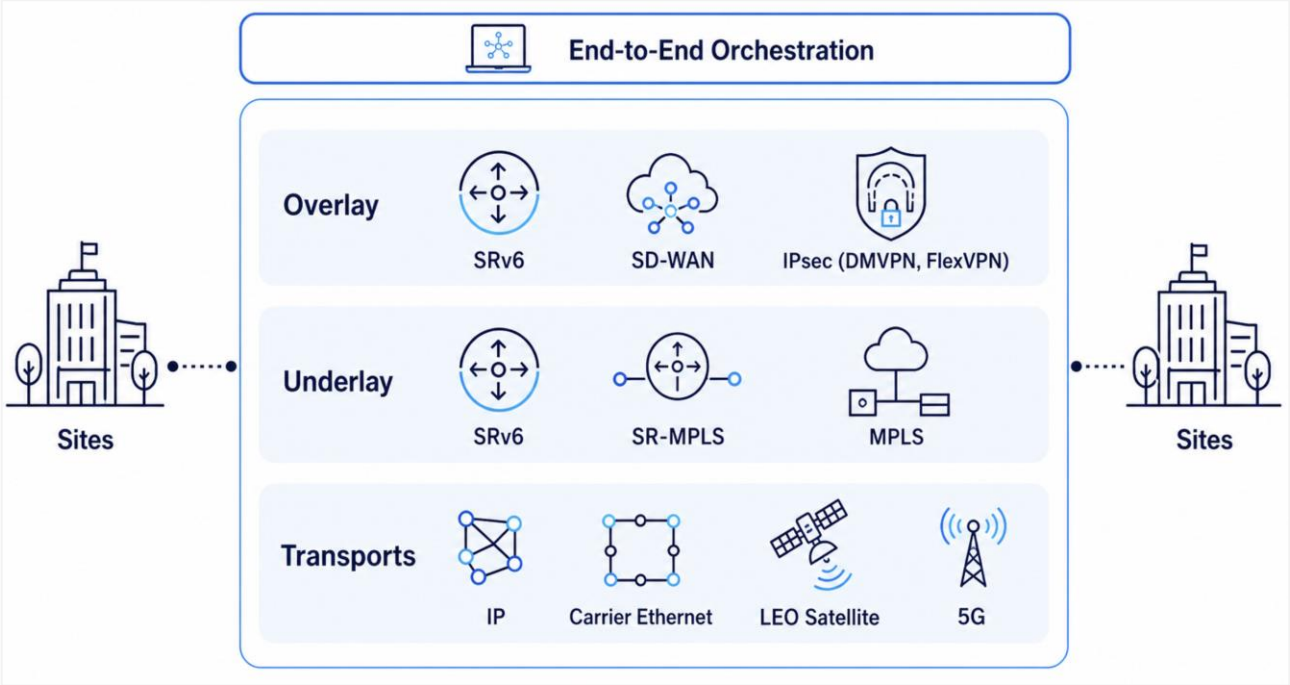

To secure logical data pathways across a shared core, this operational design strictly adheres to red/black isolation principles. It establishes a definitive cryptographic boundary between trusted secure enclave (the Red network) and the untrusted transport domain (the Black network). This deliberate separation guarantees that unencrypted information remains strictly confined within trusted edge security enclaves, while only fully secured, ciphertext traffic traverses the untrusted provider infrastructure.

Figure 2. Red/black isolation principles

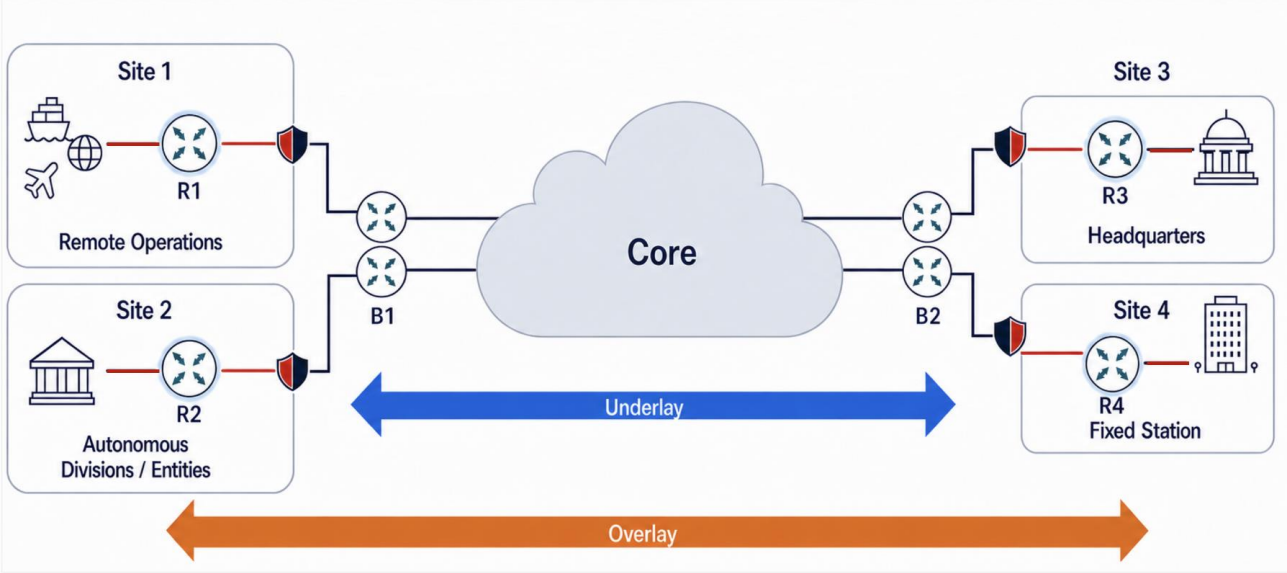

This Cisco Validated Design (CVD) serves as the architectural blueprint and technical guide for building a next-generation, zero-downtime, and quantum-safe network infrastructure. Engineered specifically for Cisco 8000 Series Secure Routers running the Cisco IOS XE operating system, this document bridges the gap between high-level security mandates and the exact physical, routing, and cryptographic configurations required to operationalize them.

The proposed architecture outlines a highly automated, programmable Segment Routing over IPv6 (SRv6) fabric. By embedding routing instructions directly into the IPv6 packet header, the network core becomes entirely stateless and infinitely scalable. To directly neutralize HNDL attacks, this design details how SRv6 intelligently forwards traffic across Post-Quantum Cryptography (PQC)-ready MACsec and IPsec paths, ensuring end-to-end data integrity and confidentiality.

Figure 3. Segment Routing over IPv6 (SRv6) Fabric

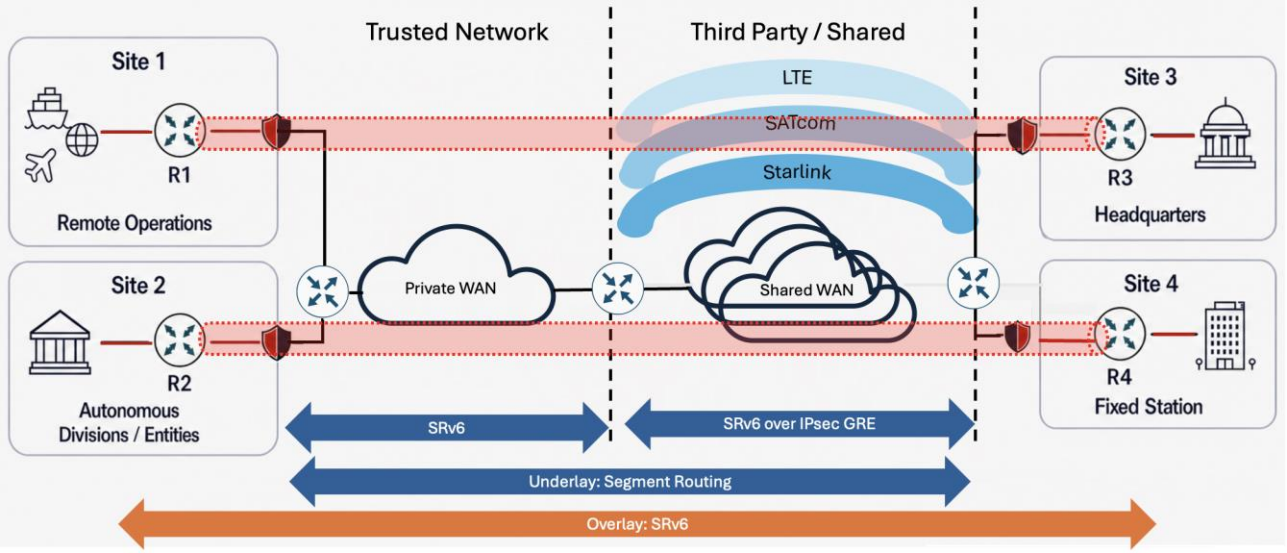

To achieve this level of digital resiliency and neutralize the HNDL threat, this architecture leverages the Cisco 8000 series secure routers as its foundational block.

Figure 4. Cisco 8000 series secure routers

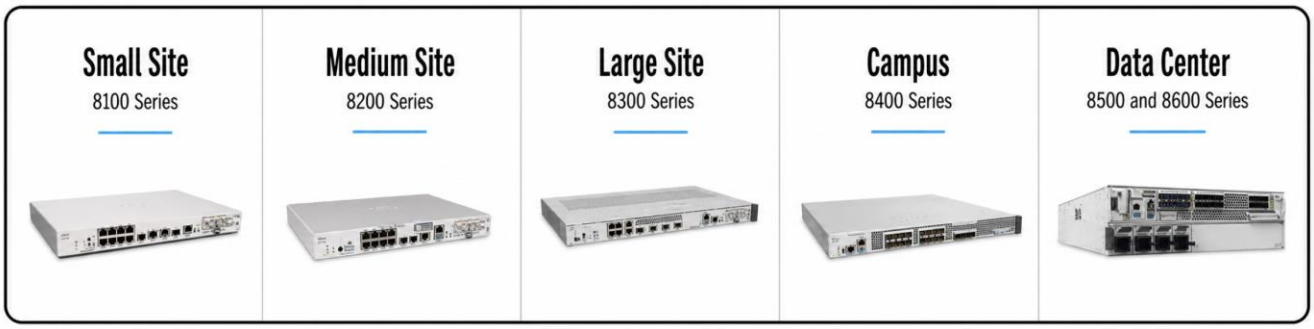

This platform transforms the legacy WAN into a highly automated, programmable Segment Routing over IPv6 (SRv6) fabric through the following mechanisms:

- Native SRv6 with uSID (F3216): Transport, traffic engineering, and overlay services are handled entirely by native IPv6 destination addresses. This eliminates core state bloat and protocol signaling storms, fundamentally simplifying network operations.
- Dynamic Performance Measurement (PM): Running two-way PM probes natively inside SR-TE policies at aggressive 3-second intervals allows routers to dynamically measure real-time latency, jitter, and packet loss. This telemetry mathematically steers critical traffic away from degraded links before packet loss impacts the mission.
- Topology Independent Loop-Free Alternate (TI-LFA): The local Point of Local Repair (PLR) natively pre-calculates the exact post-convergence path. During a hard fiber cut, TI-LFA encapsulates and steers traffic around the failure in under 50 milliseconds in hardware, guaranteeing deterministic survivability.
- Secure network processor: Powered by a custom secure network processor, the platform features dedicated, high-performance hardware crypto engines engineered specifically to handle the intensive mathematical calculations of hybrid PQC at line-rate. It also includes a built-in hardware trust anchor supporting secure boot and firmware integrity checks using quantum-safe ML-DSA (FIPS 204) signatures, mathematically ensuring the device has not been compromised at the physical layer.

By combining theoretical frameworks, strategic business drivers, design rules, and validated configuration syntaxes, this CVD provides everything required to deploy a highly scalable, automated, and quantum-ready network infrastructure.

What this document covers

To maintain a strict focus on scalable, natively routed SRv6 enterprise deployments across a quantum ready network the scope of the document is limited to following technologies and architectures, organized under the following sections:

- Why SRv6? – Analyzes the operational bifurcation between the Internal Service Provider (stateless transport, protocol minimization) and the WAN Consumer (logical service overlays, and local autonomy in DDIL environments).
- The SRv6 technology foundation – Establishes the mechanics of the Segment Routing Header (SRH) and the compressed uSID F3216 segment format (block, node, function). Maps Cisco IOS XE hardware end-behaviors, including uN (Node), uA (Adjacency), uDT4/uDT6 (VPN Decapsulation), and H.Encap.Red.
- SRv6 prerequisites and base configuration – Stages the IS-IS underlay routing, SRv6 base configuration and physical Ethernet interface requirements (isis network point-to-point) to support locator advertisements.

- L3VPN in SRv6 as the primary service – Details the VRF logical isolation boundaries, prefix-to-Service-SID mapping, and BGP Prefix-SID community advertisement using the optimized alloc-mode per-vrf design.
- Global BGP design and inter-domain scaling – Evaluates edge ASN routing loop-prevention mechanics (AS-Override vs. Allowas-in) and demonstrates the scaling benefit of configuring Next-Hop Unchanged on ASBR boundary routers to offload data-plane encapsulation.
- SRv6 traffic engineering – Configures dynamic constraint paths using static TE-metrics, Link Affinities, and dynamic delay. Documents SRv6 Performance Measurement (PM) active probing and maps ePBR classifications to parent PFPs for real-time traffic dispatching.
- Intent-based slicing using flexible-algorithm – Documents the network slicing and path isolation of Flex-Algo 128 (secured) and Flex-Algo 129 (low-latency) topologies. Establishes the design rule of binding separate locator blocks to distinct algorithms to guarantee physical isolation.
- Mission-critical resiliency using TI-LFA – Covers the pre-calculation of post-convergence repair lists, deployment of fast-detection BFD templates, and execution of hardware-level local rerouting under 50ms.
- SRv6 over secure transport – Builds quantum ready, EAP-TLS secured Layer 2 WAN MACsec sub-interfaces over provider private networks, and GRE over IPsec tunnels over untrusted internet. Includes a precise MTU guide to restrict tunnel fragmentation.
- SRv6 multicast design – Outlines the end-to-end configuration of BGP MVPN Ingress Replication, static Anycast RP mappings, MSDP source synchronization, and the dynamic allocation of the End.DTMC4 multicast decapsulation SID.
- Validated Case Study: "VeriVault" – A fictitious defense customer's design and deployment.

How to read this document

This document is specifically engineered for network architects, communications engineers, and senior network operators tasked with designing and maintaining a mission-critical infrastructure and is designed to be consumed in two ways:

- As a sequential blueprint: For teams deploying a greenfield SRv6 infrastructure or migrating from legacy MPLS, the document should be read cover-to-cover, following the strict engineering progression from base IS-IS enablement to advanced overlay traffic engineering.
- As a modular reference guide: For teams looking to implement specific operational features—such as enabling quantum-safe IPsec/ MACsec encryption or mitigating the "noisy neighbor" problem—individual sections are fully self-contained and can be referenced strictly for their design rules and exact IOS-XE command-line syntax.

Why SRv6?

This section outlines the strategic drivers for migrating legacy WAN environments to SRv6, demonstrating how this modern fabric addresses the distinct challenges of transport architects, application consumers, and tactical edge operators.

To truly understand the value of this modernization, we must look at it from the perspective of two personas: the network provider engineering the service, and the mission-critical end-user consuming it.

Persona 1: the Internal Service Provider (ISP)

The engineering team responsible for architecting and scaling the physical enterprise transport grid— terrestrial dark fiber, Metro-E connections, private provider dark-fiber E-Lines, and non-terrestrial LEO satellite links— now operates as an " Internal Service Provider" (SP).

By deploying a native, post-quantum secure Segment Routing over IPv6 (SRv6) fabric, the Internal SP abstracts the physical complexity of the Wide Area Network, delivering it as a highly scalable multi-tenant transport utility to the various mission enclaves.

This persona's responsibilities and WAN architectural execution are defined across the core operational pillars (note, if any of terms listed in the bullets below are hard to follow, read section 2 and return to this section):

1. Absolute ownership of E2E service " turn up"

As the Internal SP, this team owns and controls the entire " turn up" and lifecycle of overlay network services.

- Zero core state scaling: SRv6 unifies the underlay, overlay, and traffic engineering into a single IPv6 stack, deprecating complex legacy protocols (LDP, RSVP-TE, BGP-LU). By leveraging uSID F3216 encapsulation, routing instructions are encoded directly into the packet header, keeping core transit (P) routers completely stateless. Service activation is fully automated via On-Demand Next-Hop (ODN), dynamically instantiating SR-TE policies only when specific BGP Color communities are advertised.

2. Establishing and enforcing multi-service slas

The Internal SP is responsible for establishing and setting strict, end-to-end Service Level Agreements (SLAs) across the WAN, guaranteeing that high-priority mission payloads do not suffer from congestion or degradation.

- IGP-layer slicing (flex-algo) and active-link state probing: To support macro-level SLAs, the SP carves the physical network into mathematically isolated logical topologies using Flexible Algorithm (Flex-Algo). For granular application awareness, Per-Flow Policies (PFP) map packets to specific Per-Destination Policies (PDP) based on deep packet inspection, utilizing active 3-second PM probes to dynamically steer flows around congestion.

3. Cryptographic sovereignty

Operating in a high-consequence environment, the Internal SP controls, dictates, and enforces all security, encryption, and cryptographic compliance policies across the entire transit infrastructure.

- Post-quantum ready network: Private links utilize hardware-level WAN MACsec (802.1AE) negotiated via ML-KEM post-quantum EAP-TLS. Untrusted transits utilize GRE over IPsec tunnels secured with ML-KEM-1024 (Kyber). Strict Link Affinities are bound directly to SRv6 locators, mathematically guaranteeing that classified tenant VRFs never leak onto unencrypted paths.

Persona 2: WAN consumer and tactical conditions

While the Internal Service Provider (Persona 1) focuses on the underlying transport mechanics, network stability, and stateless core scaling, Persona 2 represents the mission-critical business units, tactical commanders, and enclave owners who rely on that infrastructure. To this persona, the underlying network parameters are merely a means to an end. Their primary mandate is to transmit highly sensitive, classified data—such as command-and-control (C2) telemetry and uncrewed aerial vehicle (UAV) video streams—safely and reliably across any available path.

While this encompasses general administrative overlay requirements for standard sites, its ultimate test lies at the tactical edge. In DDIL (Disconnected, Disrupted, Intermittent, and Low-bandwidth) environments, tactical operators, military communications engineers, and emergency response units operate under the harshest and most unpredictable physical conditions. Standard networking assumptions—such as stable terrestrial fiber and continuous reachability to a centralized operations center—are completely invalid.

In these environments, standard networking assumptions—such as stable terrestrial fiber, high bandwidth, and continuous reachability to a centralized operations center—are completely invalid. Instead, the network must operate under the constraints of a DDIL (Disconnected, Disrupted, Intermittent, and Low-bandwidth) Environment. Under these austere conditions, the communication "pipe" is structurally fragile, slow, and frequently severed by physical cuts, hostile jamming, geographic obstacles, or satellite orbital handoffs.

To survive and execute the mission, the tactical edge network requires a distributed routing architecture with the following "must-have" capabilities:

1. Absolute local autonomy (the controller-less mandate)

In a tactical DDIL space, relying on a centralized SDN controller (such as an SD-WAN manager or an external path computation cluster) to establish routes is a fatal design flaw. If the low-bandwidth satellite or radio link back to the central controller is disrupted, a controller-reliant edge router loses its ability to calculate new paths, dynamically steer around failures, or provision new secure enclaves, leaving the edge isolated and dark.

- Zero Centralized Dependencies: SRv6 shifts path-computation intelligence completely to the ingress edge nodes (head-ends). Edge routers natively run Constrained Shortest Path First (CSPF) calculations based on local IGP link-state updates, allowing the site to route traffic and enforce security enclaves even when completely disconnected from the global network.

2. Survivability over "spotty" multi-transport paths

Tactical deployments must aggregate whatever physical transmission medium is available in the field—combining private line-of-sight radios, public LTE/5G, commercial satellite (such as Starlink LEO), and high-latency military SATCOM.

- Transport-Agnostic Tunneling: SRv6 abstracts these disparate physical layers, running natively over WAN MACsec-secured point-to-point sub-interfaces on trusted links, or encapsulating inside Post-Quantum Cryptography (PQC)-secured GRE over IPsec tunnels across untrusted satellite and commercial internet connections.
- Active Link-State Probing: To handle highly degraded or "spotty" links, SRv6 abstracts these layers, utilizing active PM probing to track extreme latency or jitter spikes (e.g., satellite fade) and instantly steering critical telemetry to alternate paths without dropping active sessions.

3. Stateless core for narrow-bandwidth

In legacy tactical networks, signaling protocols like RSVP-TE were used to build traffic-engineered paths. However, RSVP-TE is exceptionally "chatty," requiring constant peer-to-peer signaling refresh messages

to maintain core tunnel states. Over low-bandwidth SATCOM or tactical radio links, this signaling overhead can easily consume a significant portion of the available bandwidth, leading to packet drops, protocol timeouts, and catastrophic network flaps.

- Source-Routed Data Plane: SRv6 uses uSID segments to embed forwarding instructions directly into the packet header, eliminating signaling storms and reserving narrow links strictly for mission payloads.

4. Instantaneous IGP-decoupled failover

When a physical link is cut or a transit node is destroyed in a tactical theater, waiting for global routing protocols to converge can take several seconds, resulting in massive data loss that can halt drone telemetry or command voice feeds.

- Pre-calculated hardware repair paths: By enabling Topology Independent Loop-Free Alternate (TI-LFA) natively under the IS-IS process, the local routing engine pre-calculates loop-free backup paths that mirror the post-convergence topology.
- Sub-50ms failover: The instant a physical link drop or liveness failure is detected by BFD, the local router immediately encapsulates the packet, inserting the pre-calculated SRv6 Segment SIDs to steer traffic around the failure. Because this recovery occurs locally in hardware in under 50 milliseconds, the failover is completely invisible to tactical operators, maintaining absolute service continuity during active combat or severe storms.

By understanding these personas, it becomes clear how a tactical site operating in a DDIL environment utilizes the SRv6 fabric to dynamically steer critical traffic across the fabric.

SRv6 technology foundation

This section outlines the native SRv6 data-plane mechanics, the Segment Routing Header (SRH) structure, and the micro-SID (uSID) network programming model as implemented on Cisco IOS-XE platforms—specifically targeting the Cisco 8000 Series Secure Routers.

Introduction to SRv6 network programming

The SRv6 architecture transitions the wide-area network from a basic destination-lookup topology into a programmable, distributed state machine. In this network programming model, standard 128-bit IPv6 addresses are no longer treated merely as topological interfaces; instead, they are instantiated as specialized network instructions (Segment Identifiers or SIDs) that execute deterministic functions on transit or egress routing nodes.

Table 1. SRv6 architecture

Architecture feature	SRv6 (Segment Routing with IPv6)
Underlying dataplane	Native IPv6 addresses (Segment Identifiers / SIDs)
Path encoding	IPv6 routing extension header (Segment Routing Header (SRH))
Scalability limit	Large IPv6 address space with highly flexible extensibility
Protocol footprint	Native IPv6 routing underlay; no MPLS-specific label protocols
Network integrity	Operates over standard IPv6 transit networks (stateless core forwarding)
Unified services	Unified transport and services (L3VPN/L2VPN) inside IPv6 headers

Segment Routing over IPv6 (SRv6) completely removes the need for complex MPLS overlays. It routes packets by appending a Segment Routing Header (SRH)—formally known as IPv6 Routing Extension Type 4—which contains a strict, ordered list of 128-bit network instructions known as Segment Identifiers (SIDs).

Segment Routing Header mechanism

In an SRv6 network, routing instructions are embedded directly into the IPv6 packet header itself. When a packet must traverse an explicit sequence of paths, link affinities, or service enclaves, the ingress source node appends a Segment Routing Header (SRH)—defined as IPv6 Routing Extension Type 4—directly behind the standard IPv6 network header.

The SRH acts as a native source-routing container carrying an ordered array of 128-bit Segment Identifiers (SIDs).

Figure 5. Segment Routing Header (SRH)

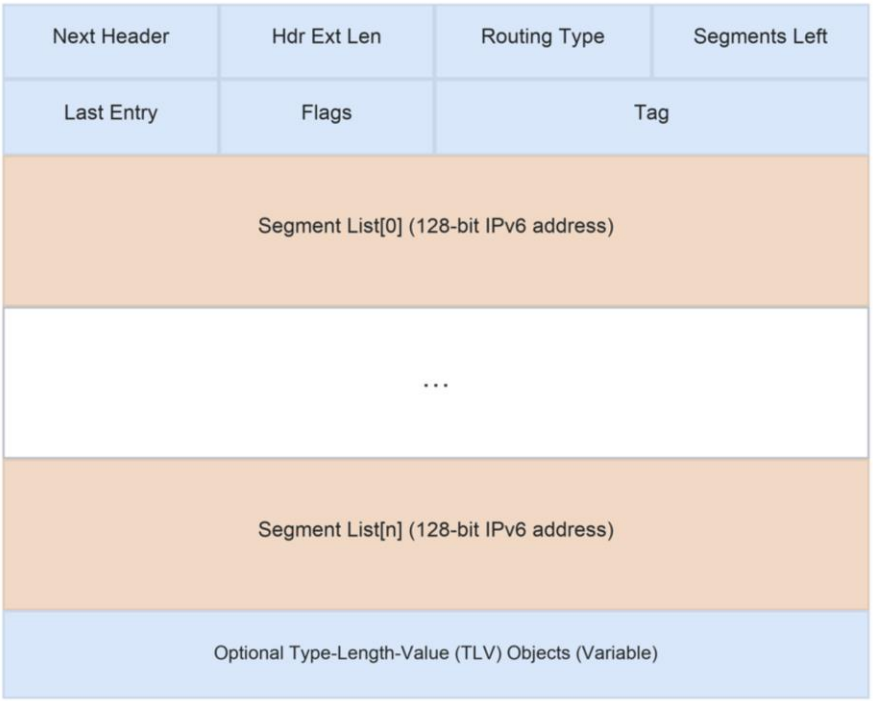

SRH field specifications:

Next header (8 bits): The "What's Inside" Label. This tells the router whether the payload is standard data (like TCP/UDP) or if it's carrying another encapsulated packet (like an inner IPv4 packet for an L3VPN).

Header extension length (8 bits): The "Header Size". It tells the router exactly how physically large this entire block of SRv6 instructions is. This ensures the router knows exactly where the routing instructions end and the actual user data begins.

Routing type (8 bits): The "ID Badge". Setting this to "4" is just the universal protocol code for "Hey, I am a Segment Routing header!"

Segments left (8 bits): The "Remaining Stops Counter". Think of this as an active countdown. If the packet must make 5 stops, this starts at 5. Every time the packet successfully reaches a waypoint and executes its instruction, this number ticks down by one. It tells the current router exactly which instruction in the list it is supposed to be executing right now.

Last entry (8 bits): The "Total Stops" Marker. This records how many total instructions were on the list when the journey originally started.

Flags (8 bits): The "Toggle Switches". These are binary on/off switches for extra features, like turning on OAM (troubleshooting/ping) or enforcing specific cryptographic rules.

Tag (16 bits): The "Sticky Note". A way to slap a label on the packet so that routers can group it with other similar packets and apply administrative policies to them all at once.

Segment list (array of 128-bit IPv6 addresses): The Driving Directions (in reverse) – A series of 128-bit IPv6 addresses representing the explicit path waypoints. The interesting part is that it is written backwards. The final destination is placed at the top of the list (Index 0) and the very first hop is at the

bottom of the list. As the "Segments Left" countdown ticks down, the router reads the list from the bottom to the top.

Figure 6. SRv6 packet details

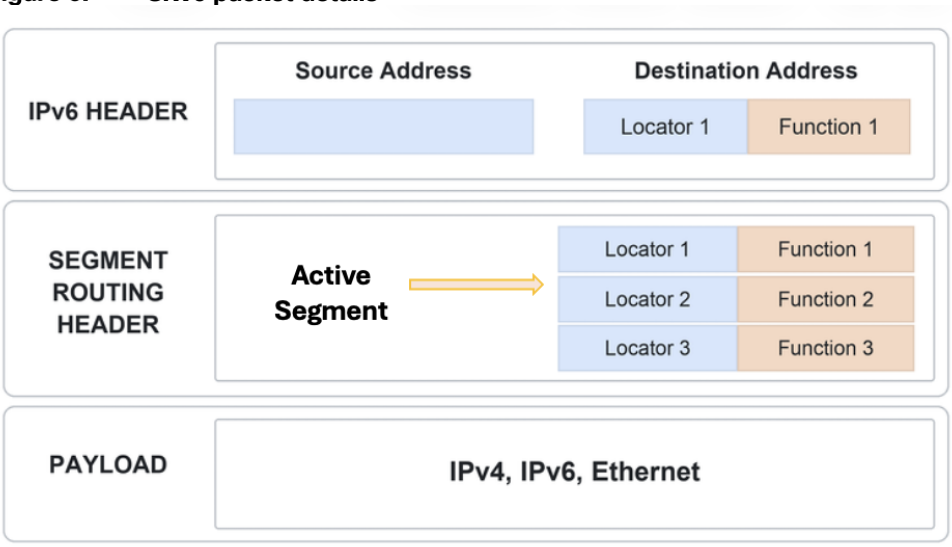

As the packet moves through the network, the router continuously places the "Active SID" into the standard IPv6 Destination Address (DA) field. Because of this, transit P-routers do not need to parse the entire SRH; they simply perform hardware-based IPv6 routing on the destination address. Intermediate SRv6-capable routers read the destination address, decrement the Segments Left counter, and copy the next ordered SID directly into the primary IPv6 DA field.

Segment Identifier structural architecture

In SRv6, a Segment Identifier (SID) is a 128-bit network instruction. While it looks exactly like a standard IPv6 address and is fully routable over an IPv6 network, it carries entirely different semantics. Instead of simply pointing to a destination device, a SID acts as an executable command.

A standard SID is traditionally divided into three primary structural components:

- Locator (block + node ID): The routing portion of the SID. It consists of a globally assigned SRv6 address space (the SID block) and a unique hardware identifier (the node ID). The locator allows the network to physically route the packet to the correct router.
- Function (FUNCT): The local instruction executed by the target node's network processor. Once the packet arrives, this section defines the exact action (end behavior) the router must execute—such as a VRF table lookup, payload decapsulation, or a physical interface cross-connect.
- Arguments and padding (ARG): Optional bits that provide additional parameters for the function to use, with remaining bits padded as zeros to complete the flexible 128-bit length.

Figure 7. Standard SID structure

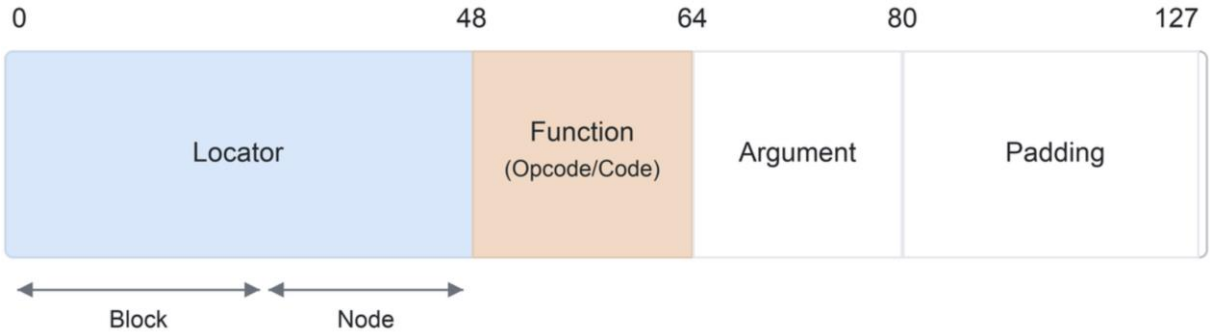

Highly scalable routing with uSID

While highly flexible, stacking multiple traditional 128-bit SIDs causes significant header bloat—adding up to 40 bytes of overhead per hop when steering traffic across complex, multi-hop paths.

To eliminate this overhead and scale the fabric, this architecture mandates the use of the uSID (Micro-SID) F3216 compressed format. The F3216 format partitions the standard 128-bit IPv6 container into a strict <32-bit Block> <16-bit Node ID> <16-bit Function> hierarchy:

- uSID Block (32 bits): A domain-wide prefix allocated to the entire network fabric (e.g., FCBB:DEAD::/32). Every participating node shares this identical block, ensuring cohesive routing and maximum compression.
- Node ID (16 bits): A unique identifier assigned to each specific router within the domain (e.g., 0009 or 1031) [3, 4]. It represents the topological locator of the device.
- Function/ arguments (16 bits): The explicit instruction the node must execute upon packet arrival. (e.g., local transit, table decapsulation, or adjacency cross-connect).

The true power of the F3216 format lies in its efficiency: by compressing the routing and function data, it allows the hardware to pack up to six distinct 16-bit Micro-SIDs into a single 128-bit IPv6 Destination Address container, drastically reducing bandwidth consumption across the core.

The "shift-left" operation

To achieve dynamic, wire-rate path routing without expanding packet sizes, transit routers perform a stateless, hardware-level "Shift-Left" operation directly on the IPv6 Destination Address (DA).

The step-by-step forwarding journey

When an encapsulated packet hits a transit router:

Step 1. The router identifies the incoming IPv6 Destination Address as matching its own local locator prefix (matching the Block and Node ID).

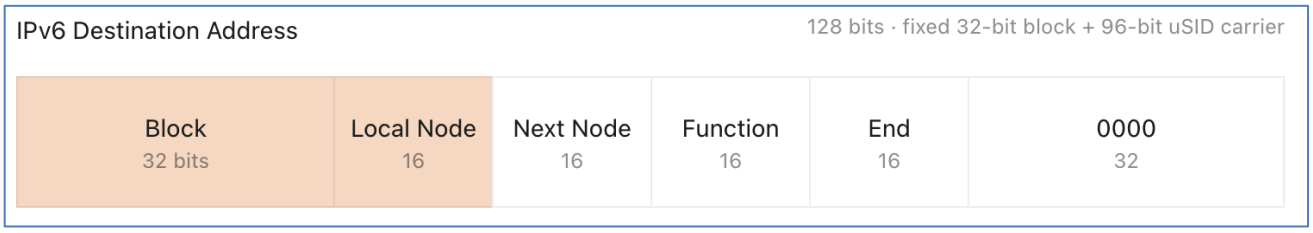

Step 2. The router's network processor removes its own active 16-bit Node ID from the address.

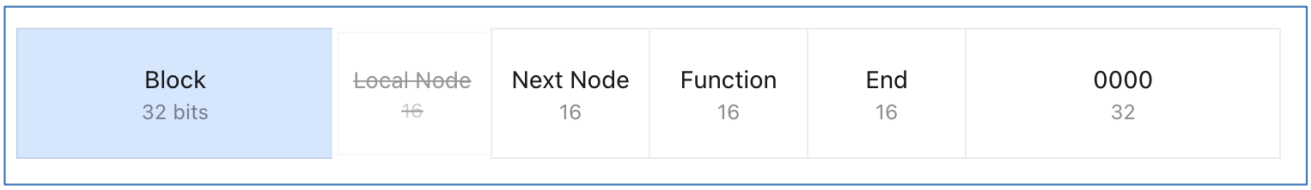

Step 3. The remaining bits of the IPv6 Destination Address are shifted exactly 16 bits to the left.

Step 4. The end of the address is padded with trailing zeros (0000).

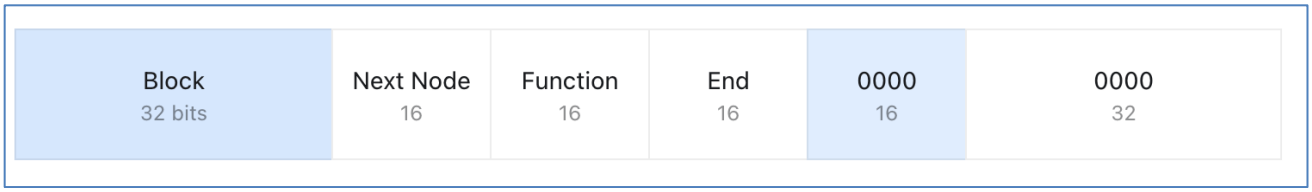

Step 5. This operation instantly promotes the next 16-bit Node/Function instruction into the active routing field of the IPv6 Destination Address. The transit router forwards the packet to the next hop using standard, stateless longest-prefix-match (LPM) lookups in hardware, completely bypassing the need to read or pop headers from an outer SRH payload.

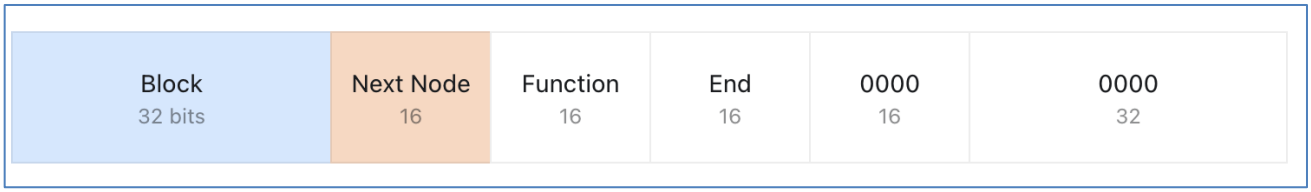

SRv6 node roles and headend behaviors

During its transit across the Wide Area Network, a secure packet interacts with three distinct node roles, each executing specific actions based on hardware capability:

Figure 8. SRv6 node roles and headend behaviors

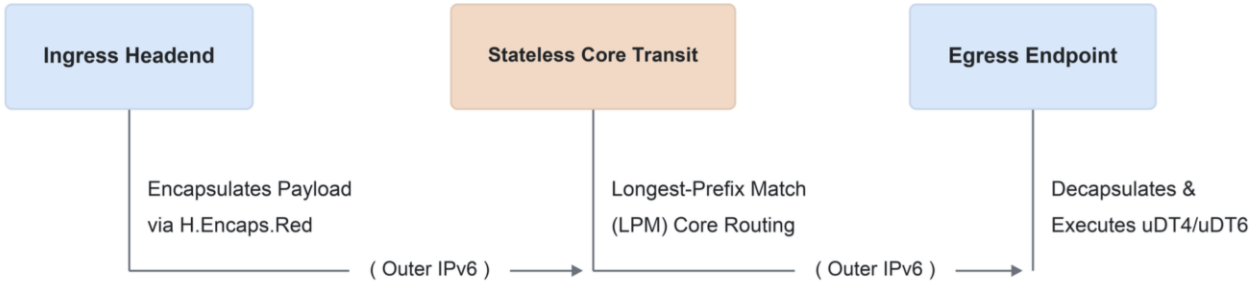

Ingress headend node

The entry point of the SRv6 fabric where customer payloads (IPv4, IPv6, or L2 frames) are ingested. The Ingress headend router classifies the incoming traffic (via ePBR and DSCP/NBAR maps), maps it to a Traffic Engineering policy, and enforces the routing path.

- H.Encaps.Red (headend encapsulation with reduced SRH): The mandatory encapsulation behavior supported on Cisco IOS-XE. The headend encapsulates the original packet into an outer IPv6 carrier packet. If the path contains multiple hops, the headend inserts an SRH but optimizes the header by omitting the first segment from the Segment List (since it is already placed directly into the outer IPv6 Destination Address), saving 16 bytes of overhead.
- Local candidate-path constraint: To ensure routing predictability, all traffic engineering policies calculate their candidate paths locally using Constrained Shortest Path First (CSPF) computed against the IGP link-state database.

Transit node

Intermediate core routers along the path. On a modernized network, transit P-nodes run extremely fast and remain completely stateless.

- The Longest-Prefix-Match (LPM) forwarding paradigm: Transit nodes do not inspect the inner payloads, do not maintain VPN/VRF routing tables, and do not parse the SRH. Because the active instruction is always shifted into the standard IPv6 Destination Address field, transit routers forward the frame purely as standard IPv6 packets using longest-prefix-match (LPM) lookups at line-rate.
- Core minimization: This design allows high-capacity core routers in the P-role, to be kept completely lean and free of VRF overlays, preventing control-plane signaling storms and core memory exhaustion.

Egress endpoint node

The destination router in the SRv6 domain where the outer transport segment is terminated. The endpoint router receives the packet, matches the destination address with its local locator, decapsulates the outer IPv6 header, and executes the functional instruction bound to the Service SID (such as dropping the payload into a tenant VRF).

SRv6 end behaviors and node capabilities

When the packet reaches the targeted router, it executes the embedded SID instruction, formally known as an End Behavior. Core capabilities in this design include:

- uN (Node SID): The standard endpoint behavior used to route to a specific router along the shortest IGP path.
- uA (Adjacency SID): An endpoint with a Layer-3 cross-connect. This forces the packet out of a highly specific local physical link, overriding the shortest path (critical for strict TE routing).
- H.Encap.Red (Headend Reduced Encapsulation): An optimization behavior used by the ingress edge router to reduce overhead when encapsulating packets into an SR-TE policy.
- uDT4/ uDT6: The core Service SIDs for L3VPNs. They instruct the egress router to decapsulate the outer IPv6 header and perform an explicit table lookup within a specific, isolated IPv4 (uDT4) or IPv6 (uDT6) Virtual Routing and Forwarding (VRF) instance.
- uDT46: A unified endpoint behavior that decapsulates and performs lookups for both IPv4 and IPv6 traffic inside the exact same dual-stack VRF.
- End.DTMC4: An endpoint behavior specifically used to decapsulate and route IPv4 Multicast traffic across the fabric.

SRv6 prerequisites and configuration

Before deploying advanced Traffic Engineering (TE) or L3VPN overlay services, the network underlay must be properly staged. In Cisco IOS-XE, SRv6 relies on the Interior Gateway Protocol (IGP)—specifically IS-IS—to distribute routing instructions (Locators and Micro-SIDs) across the fabric.

This section outlines the global IS-IS core underlay prerequisites, the compressed uSID F3216 locator allocation design, and the step-by-step base process configuration and verification procedures as implemented on Cisco IOS-XE platforms.

SRv6 prerequisites for IS-IS

To support the extensive topological telemetry required by SRv6 and SR-TE, the underlying IS-IS routing process must be strictly configured with specific extensions. Failure to implement these prerequisites will prevent the SRv6 control plane from establishing or cause SR-TE policies to fail.

- Global IS-IS parameters: Under the global `router isis process`, administrators must enable `metric-style wide` to support extended TLVs, along with `advertise link-attributes` and `distribute link-state`.
- IPv6 address family: Because SRv6 operates exclusively over IPv6, you must activate `the address-family ipv6` sub-mode within IS-IS and explicitly enable `multi-topology` to ensure the router maintains a dedicated, loop-free routing table for IPv6 traffic independent of IPv4.
- Crucial interface constraint (point-to-point): By default, Cisco IOS-XE treats standard Ethernet interfaces as broadcast network types. You must explicitly configure '`isis network point-to-point`' on all physical WAN interfaces participating in the SRv6 core. If the interfaces are left as broadcast, the SR-TE policies will silently fail to compute or steer traffic.

Example configuration reference

The following snippet demonstrates the applied prerequisites on a router's IS-IS process and a physical Ethernet interface.

```
router isis 1
net 49.0000.fc00.0009.00
 is-type level-2-only
 router-id Loopback0
 advertise link attributes
 metric-style wide
 distribute link-state
 passive-interface Loopback0
 !
 address-family ipv6
  multi-topology
  router-id Loopback0
 !
interface GigabitEthernet2
 no ip address
 negotiation auto
 ipv6 address 2001:DB8:ACAD:1::1/64
 ipv6 enable
 ipv6 router isis 1
```

```
 isis circuit-type level-2-only
 isis network point-to-point
```

Detailed explanation of the configurations:

- `router isis 1`: Initializes the IS-IS Interior Gateway Protocol (IGP) routing process under ID 1. This process serves as the foundational underlay control plane responsible for carrying the prefix locators and SIDs throughout the network.
- `net 49.0000.fc00.0009.00`: Configures the Network Entity Title (NET) for the routing node. Here, 49 represents the private authority/area identifier, 0000.fc00 defines the logical IS-IS area, 0009 specifies the unique 16-bit System ID of the physical node, and 00 is the N-selector identifying the router's local system itself.
- `is-type level-2-only`: Restricts the routing engine to Level-2 backbone adjacencies. This optimizes performance and scaling by preventing the router from expending CPU and memory overhead calculating separate Level-1 local area topologies.
- `router-id Loopback0`: Establishes Loopback0 as the unique Traffic Engineering Router ID. This is published to the network database to serve as a fixed endpoint identifier for path-computation algorithms.
- `advertise link attributes`: Instructs IS-IS to advertise stethoscopic Traffic Engineering (TE) link properties—such as dynamic link delay, jitter, packet loss, and administrative link affinities—into the IGP link-state database. This is a mandatory requirement for dynamic, application-aware SR-TE path selection.
- `metric-style wide`: Enables wide metrics (24-bit link metrics and 32-bit path metrics). Narrow metrics (the default) cannot carry the extended Type-Length-Value (TLV) sub-structures required to advertise SRv6 locators, Micro-SIDs (uSIDs), and TE attributes across the underlay.
- `distribute link-state`: Configures the IS-IS process to export its Link-State Database (LSDB) topology. This allows the local router and centralized controllers to build an accurate, real-time map of the physical topology.
- `passive-interface Loopback0`: Configures the interface as passive, ensuring that the Loopback prefix is advertised throughout the network while suppressing the generation of IS-IS Hello routing packets out of the logical port.
- `address-family ipv6`: Enters the IPv6 address family configuration sub-mode. Because SRv6 operates natively in the IPv6 data plane, all segment routing, locator bindings, and fast-reroute policies are configured under this specific address family.
- `multi-topology`: Activates Multi-Topology Routing (MTR) under the IPv6 address family. This forces IS-IS to calculate entirely separate, independent Shortest Path First (SPF) routing tables for IPv4 and IPv6, preventing routing "black holes" in dual-stack environments where a physical link is up but lacks IPv6 forwarding capabilities.
- `isis network point-to-point` (on interface GigabitEthernet2): Explicitly overrides the default broadcast network type, forcing the physical Ethernet link to operate as a direct point-to-point connection. By default, Cisco IOS XE treats Ethernet interfaces as broadcast, which triggers a Designated Intermediate System (DIS) election and creates virtual pseudo nodes. Pseudo nodes obscure the exact physical interface topology, which silently breaks the calculation of dynamic Adjacency SIDs (uAs) and dynamic SR-TE candidate paths. Enforcing a point-to-point network type is an absolute, non-negotiable architectural rule to guarantee successful SR-TE policy calculation and forwarding.

SRv6 IPv6 locator allocation and design

Before applying the SRv6 base configuration to the routing nodes, architects must establish a strict, scalable IPv6 addressing schema for the SRv6 Locators. In an SRv6 network, the locator acts as the summary routing aggregate for all local Segment Identifiers (SIDs)—such as node, adjacency, and service behaviors—instantiated on a specific router. Because this locator block is natively advertised into the IGP (IS-IS) as a standard IPv6 route, the transit P-routers in the core can forward encapsulated traffic using highly efficient, line-rate longest-prefix-match (LPM) lookups in hardware, completely remaining stateless and oblivious to individual VPN overlay routing tables.

To prevent header bloat and packet-size overhead across complex, multi-hop enterprise and defense topologies, this design mandates the use of the highly compressed uSID F3216 (Micro-SID) format. Under the F3216 structure, the IPv6 locator is meticulously structured into a specific <32-bit Block><16-bit Node> hierarchy:

- The 32-bit uSID block (the global domain)

  The first 32 bits of the IPv6 address represent the global SRv6 domain block. All routers operating within the same routing domain must share this exact same 32-bit prefix.

  Example: FCBB:DEAD::/32

- The 16-bit node ID (the specific router)

  The next 16 bits uniquely identify the specific Provider Edge (PE) or Customer Premises Equipment (CPE) router within that domain. When combined with the 32-bit global block, it provides a unique /48 IPv6 prefix strictly allocated to that single hardware node.

  Example: Assigning Node ID 0001 creates the unique router locator FCBB:DEAD:0001::/48.

Reserving locator blocks for network slicing

While a single /32 global block is sufficient for basic, single-topology connectivity, mission-critical networks require Network Slicing to isolate different types of traffic across the network (such as separating secure military data from unsecure bulk administrative traffic).

To implement strict, intent-based slicing, architects must reserve a separate 32-bit SRv6 Locator block for each network slice. Under this model, the IGP computes entirely independent Shortest Path First (SPF) topologies for each slice based on custom metrics or link-affinity constraints:

- Base topology block (default / algorithm 0): FCBB:DEAA::/32
  - Used for standard enterprise routing, management, and best-effort transport.
- Reserved secure topology block (secure slice / flex-algo 128): FCBB:DEAF::/32
  - Reserved exclusively for highly classified enclaves. The IGP computes a topology restricted strictly to physical links secured by WAN MACsec or post-quantum IPsec GRE tunnels.
- Reserved low-latency topology block (low-latency slice / flex-algo 129): FCBB:DEAB::/32
  - Reserved for real-time voice, video, and critical Command-and-Control (C2) telemetry. The IGP computes paths dynamically optimized for the absolute lowest latency based on active Performance Measurement (PM) link probes.

(Note: The exact control-plane configurations and dynamic data-plane steering mechanisms used to map traffic into these isolated topologies via Flexible Algorithms will be covered in depth in section 7).

SRv6 base process initialization

Once the physical interface types are locked down and the /48 aggregate locators are meticulously planned, network architects must initialize the global Segment Routing over IPv6 (SRv6) process and tie it directly to the underlay IGP routing domain. On Cisco IOS XE secure platforms, this initialization is achieved through a coordinated three-step operational sequence:

- Establishing a stable loopback anchor: Instantiating a logical Loopback interface with a /128 IPv6 address to act as the global encapsulation source.
- Configuring the global SRv6 process: Binding the encapsulation engine to this Loopback IP and carving out the active /48 uSID aggregate locator block utilizing the mandatory usid-f3216 compression format.
- IGP adjacency binding: Injecting and advertising the locator prefix natively inside the address-family ipv6 submode of the active IS-IS process.

Verified baseline configuration blueprint

Below is the validated baseline configuration for an SRv6 Provider Edge (PE) or Customer Premises Equipment (CPE) router:

```
! 1. Establish the foundational Loopback for SRv6 Source Encapsulation
interface Loopback0
 ip address 10.10.1.1 255.255.255.255
 ipv6 address FC00::1/128
 ipv6 enable
 isis circuit-type level-2-only
! 2. Configure Physical WAN Interfaces (Applying the Point-to-Point Prerequisite)
interface GigabitEthernet2
 no ip address
 negotiation auto
 ipv6 address 2001:DB8:ACAD:1::1/64
 ipv6 enable
 ipv6 router isis 1
 isis circuit-type level-2-only
 isis network point-to-point
! 3. Instantiate the Global SRv6 Process and uSID Locator
segment-routing srv6
 encapsulation
  ! Define the source address used when this router encapsulates traffic
  source-address FC00::1
  traffic-class propagate
 !
 locators
  locator DEFAULT
   ! Allocate a /48 prefix block specifically to this router
   prefix FCBB:DEAD:1::/48
   ! Enforce the highly compressed F3216 Micro-SID format mandated by the design
   format usid-f3216
!
! 4. Bind SRv6 to the IS-IS Underlay
```

```
router isis 1
 net 49.0000.fc00.0001.00
 is-type level-2-only
 router-id Loopback0
 metric-style wide
 advertise link attributes
 distribute link-state
 passive-interface Loopback0
 !
 address-family ipv6
  multi-topology
  router-id Loopback0
  ! Enable SRv6 natively within the IGP and attach the locator
  segment-routing srv6
   locator DEFAULT
   level-2
```

Architectural and configuration deep-dive

1. Defining the encapsulation source address

SRv6 requires a stable, highly available IPv6 address to serve as the Source IP (SA) stamped into the outer IPv6 carrier header when user payloads enter an SRv6 tunnel at the ingress node.

- The loopback0 anchor: In this design, the dedicated /128 address FC00::1 is bound to Loopback0.
- Encapsulation enforcement: This address is explicitly referenced under the global segment-routing srv6 > encapsulation > source-address FC00::1 block. If this binding is missing or configured with an unreachable IP, the router's network processing unit will fail to generate outer encapsulation headers, blackholing all ingress customer VRF flows.
- SLA flag propagation: The traffic-class propagate command ensures that Type of Service (ToS) or Differentiated Services Code Point (DSCP) markings from inner payload headers are automatically copied to the outer IPv6 tunnel header, preserving class-of-service priority mappings across transit core links.

2. Configuring the SRv6 locator and usid structure

The locator block is the core IPv6 prefix assigned to an individual node. It acts as a standard IPv6 route aggregate, summarizing all functional instructions (SIDs) executing locally on the routing engine.

- uSID F3216 format mandate: To circumvent the severe packet-size overhead associated with traditional 128-bit Segment Routing headers, this configuration enforces the format usid-f3216 compression hierarchy. This strictly divides the address block into a 32-bit global domain block (FCBB:DEAD) and a 16-bit node identifier (0001), providing a neat /48 routing aggregate (FCBB:DEAD:1::/48) specifically assigned to this single hardware node.
- Stateless scaling mechanics: By aggregating SIDs under a single /48 locator, transit routers (P-nodes) do not need to maintain state for VPN tables or individual VRF routes. Core transit nodes route traffic purely on this aggregate /48 locator prefix using standard, line-rate Longest Prefix Match (LPM) routing tables.

3. Advertising the locator via IS-IS

For the SRv6 fabric to forward packets end-to-end, intermediate transit nodes must dynamically learn how to reach each edge router's locator block.

- Dynamic IGP redistribution: Under the router isis 1 -> address-family ipv6 submode, the locator is dynamically bound using the segment-routing srv6 -> locator DEFAULT directive.
- Link-state convergence: Once applied, the IS-IS process automatically generates custom Level-2 sub-TLVs to advertise the locator prefix and its operational parameters to all neighboring nodes in the IGP domain. Upon receiving these Link State Packets (LSPs), transit routers install the /48 prefix into their global IPv6 Routing Information Base (RIB), enabling seamless end-to-end underlay reachability.

Verify base SRv6 configuration

To guarantee that the physical underlay transport has converged and the local hardware is prepared to receive overlay services, network operators must execute a standardized sequence of verification procedures. Some commands to use to verify the configuration using `show segment-routing srv6 locator`, `show segment-routing srv6 sid`, `show isis srv6 locator`, `show ipv6 route`.

L3VPN in SRv6

To enforce logical multi-tenancy and secure application containment across a high-consequence network, legacy virtualization methods—which rely on complex, hop-by-hop VRF-lite stitching or resource-heavy MPLS label stack allocations—must be retired. BGP-based SRv6 L3VPN overlay services provide the architectural replacement. By leveraging Multiprotocol BGP (MP-BGP) to dynamically map tenant VRF instances directly to functional, compressed uSID Service SIDs, SRv6 establishes a unified, transport-agnostic overlay. This overlay natively encapsulates both legacy IPv4 and modern IPv6 payloads into a single IPv6 core, maintaining flawless cryptographic and logical separation between different routing domains without requiring transit P-routers to maintain any service state.

This section outlines the design, topology, and operational mechanics of Layer 3 Virtual Private Networks over SRv6 (L3VPN-SRv6) on Cisco IOS XE platforms. It details the step-by-step control plane and data plane operations, examines the BGP locator binding and inheritance hierarchy, and provides verified baseline configurations and verification schemas to successfully establish multi-tenant isolated enclaves across the secure WAN.

Design of L3VPN in SRv6 (L3VPN-SRv6)

Modern enterprise and defense environments require strict Layer 3 isolation across the Wide Area Network. For example, in national security enclaves, architects must completely isolate critical physical infrastructure, such as drone command-and-control systems, from legacy IPv4 third-party vendor maintenance tunnels or general office networks.

Historically, this required running complex, hop-by-hop VRF-lite configurations (which do not scale and are operationally intensive) or deploying distinct MPLS overlays (which introduce protocol bloat like LDP/RSVP-TE). SRv6 solves this by utilizing Multiprotocol BGP (MP-BGP) to provide native, highly scalable L3VPN Overlay Services. The SRv6 fabric acts as a universal, programmable transport, natively encapsulating customer VRF traffic into standard IPv6 packets, maintaining flawless logical separation between different routing domains without touching the intermediate core P-routers.

L3VPN over SRv6 architectural flow

Rather than allocating traditional MPLS labels, the SRv6 control plane replaces the standard MPLS VPN label with a 128-bit functional SRv6 Service SID (such as a uDT4, uDT6, or uDT46 SID). This SID is advertised in the MP-BGP control plane within the BGP Prefix-SID attribute.

The following trace details the step-by-step control-plane and data-plane operations,

Figure 9. Step-by-step control-plane and data-plane operations

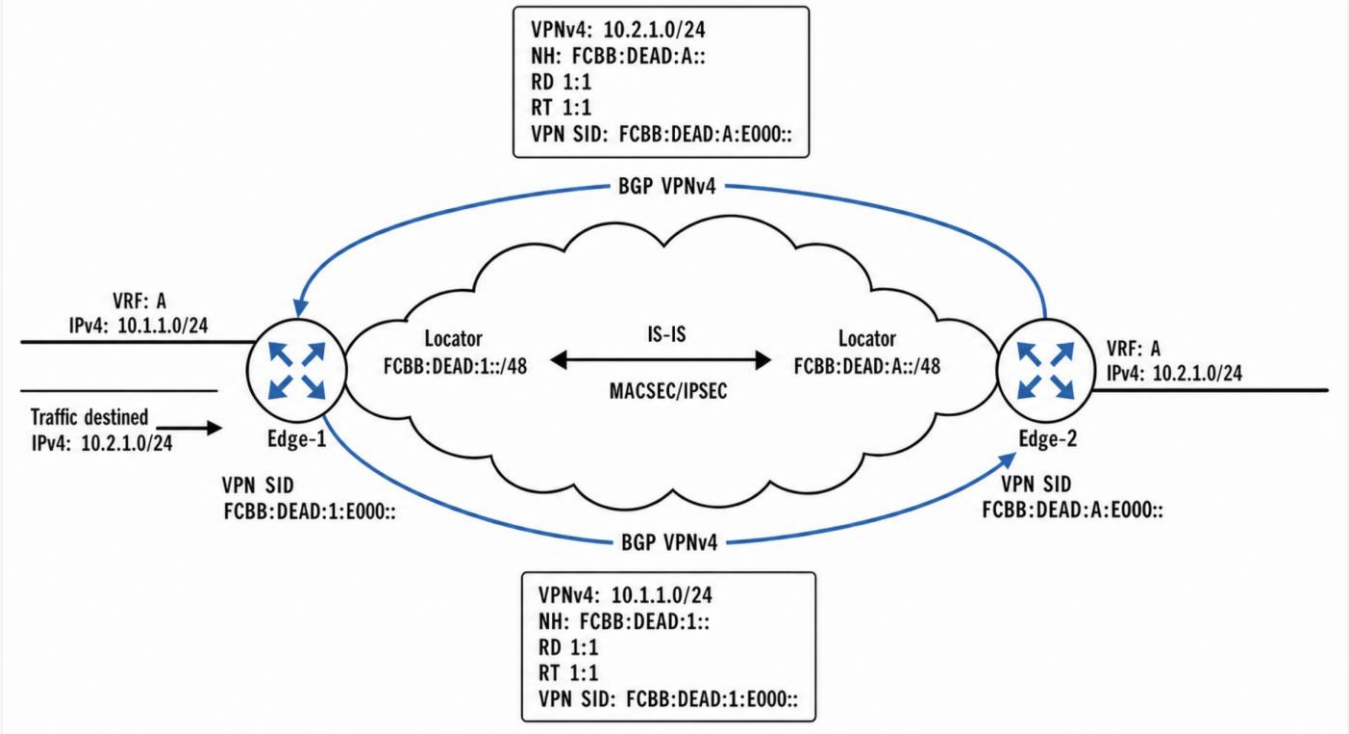

Control plane operations:

- Service SID allocation: The egress router, Edge-2, utilizing locator FCBB:DEAD:A::/48 dynamically requests a Service SID from the local SID Manager. Because the platform is strictly configured for high-scale forwarding, BGP allocates a single, unified SID for the entire VRF using alloc-mode per-vrf. For IPv4 unicast routing inside VRF: A, it instantiates the uDT4 Service SID FCBB:DEAD:A:E000::.
- MP-BGP advertisement: Edge-2 originates an MP-BGP VPNv4 routing update for prefix 10.2.1.0/24. The standard MPLS label block is replaced by the 128-bit SRv6 Service SID (configured as uDT4: FCBB:DEAD:A:E000::), appended inside the BGP Prefix-SID attribute. The update also explicitly carries the Route Distinguisher (RD 1:1), Route Target (RT 1:1), and sets the next-hop address to Edge-2's global IPv6 locator (FCBB:DEAD:A::).
- Ingress import: The ingress router (Edge-1) receives the BGP update, updates its localized routing table for the matching VRF A, and associates the tenant prefix directly with the remote Service SID.

Data plane operations:

- Packet reception: A native, unencapsulated IPv4 packet destined for 10.2.1.10 enters the ingress interface on Edge-1, which is assigned to the secure VRF A.
- Encapsulation: Edge-1 performs a route lookup within the specific VRF A table, identifying the remote prefix bound to the SRv6 Service SID. Edge-1 encapsulates the inner IPv4 packet inside an outer IPv6 transport header, setting the IPv6 Destination Address directly to the remote Service SID (FCBB:DEAD:A:E000::).
- Core forwarding: The intermediate nodes route the packet across the IS-IS underlay—which in this architecture is secured via MACsec or IPsec—using standard IPv6 prefix lookup mechanics based strictly on the Locator portion of the address (FCBB:DEAD:A::/48). Core P-routers remain entirely oblivious to the inner payload or the VRF binding.
- Egress processing: The packet arrives at Edge-2. Because the Destination Address matches its localized locator space, Edge-2 reads the Function bits (E000). This matches the specific uDT4 opcode, instructing

Edge-2 to strip away the outer IPv6 header and forward the raw inner IPv4 packet into the isolated tenant VRF A table toward its final 10.2.1.0/24 destination.

Critical design rule: per-vrf allocation mode

When designing and deploying the MP-BGP overlay, network architects must manage how the router generates these Service SIDs. Currently, Cisco IOS XE supports alloc-mode per-vrf for SRv6 L3VPN overlay services. Traditional per-prefix SID allocation (where every individual prefix inside a VRF is assigned a unique SID) is entirely unsupported.

Under the mandatory alloc-mode per-vrf model, BGP allocates a single, unified Service SID for the entire VRF instance (such as one uDT4 SID for all IPv4 routes and one uDT6 SID for all IPv6 routes inside the VRF).

- The scalability benefit: Instead of consuming thousands of hardware entries in the router's local SID table (which would quickly exhaust the TCAM space on edge routers), a single SID aggregates the entire tenant space.
- Decapsulation logic: When an encapsulated packet hits the egress PE, the outer header is stripped, and the router is forced to perform a standard IP lookup in the localized VRF table to find the exit interface. This guarantees absolute tenant isolation at massive scale on enterprise edge platform.

BGP locator binding and inheritance

For the MP-BGP process to generate Service SIDs and attach them to VPN routes, it must be bound to an underlying SRv6 locator block. Cisco IOS XE utilizes a strict, hierarchical Locator Inheritance model to determine how VPN traffic is encapsulated, operating from the most generic scope down to the most specific:

- Global level (most generic): Configured directly under the base segment-routing srv6 process. If a locator is configured at this level, BGP automatically applies it to all VRFs and address families across the router.
- VPN address-family level: Configured under the BGP address-family vpnv4 or address-family vpnv6 sub-modes. A locator applied here dictates the SID generation for all VRFs participating in that specific address family, overriding the Global default.
- VRF address-family level (most specific): Configured directly under the individual BGP address-family ipv4 vrf <vrf-name> or address-family ipv6 vrf <vrf-name> sub-modes.

Hierarchical order of inheritance:

This hierarchical design provides exceptional flexibility for mission-critical enclaves. While standard enterprise traffic can inherit the global default locator, an architect can explicitly bind a highly classified tenant VRF directly to a dedicated, secure locator block (such as a locator mapped to Flex-Algo 128 that is mathematically constrained only to MACsec-encrypted physical links). This VRF-level configuration completely overrides the default inheritance chain, ensuring that classified data plane traffic is structurally and topologically isolated across the fabric.

BGP L3VPN configuration blueprint

The following configuration represents the validated Cisco IOS XE blueprint used to establish L3VPN services inside VRF:A on EDGE1. The configuration demonstrates the instantiation of the tenant VRF, the binding of the TACTICAL_CORE IS-IS process, and the MP-BGP configuration enforcing alloc-mode per-vrf.

```
! 1. Instantiate the Tenant Isolation VRF instance
vrf definition A
 rd 1:1
```

```
 address-family ipv4
  route-target export 1:1
  route-target import 1:1
 exit-address-family
!
! 2. Configure global SRv6 process and locator
segment-routing srv6
 encapsulation
  source-address FC00::1
  traffic-class propagate
 locators
  locator DEFAULT
   prefix FCBB:DEAD:1::/48
   format usid-f3216
!
! 3. Bind the VRF to the core Interior Gateway Protocol (IS-IS)
router isis 1
 net 49.0000.fc00.0001.00
 is-type level-2-only
 router-id Loopback0
 advertise link attributes
 metric-style wide
 passive-interface Loopback0
 !
 address-family ipv6
  multi-topology
  router-id Loopback0
 exit-address-family
!
! 4. Configure MP-BGP to dynamic map VRF services to functional SIDs
router bgp 65001
!
 segment-routing srv6
  locator MAIN-LOCATOR
 exit-srv6
 !
 bgp router-id interface Loopback0
 bgp log-neighbor-changes
 bgp graceful-restart
 no bgp default ipv4-unicast
 neighbor FC00::3 remote-sa 65001
 neighbor FC00::3 update-source Loopback0
 !
 address-family vpnv4
  ! Enforce the Per-VRF allocation mode mandated by IOS-XE
  segment-routing srv6
   locator MAIN_LOCATOR
```

```
   alloc-mode per-vrf
  exit-srv6
  !
  neighbor FC00::3 activate
  neighbor FC00::3 send-community extended
  neighbor FC00::3 next-hop-self
 exit-address-family
 !
address-family link-state link-state
  neighbor FC00::3 activate
 exit-address-family
 !
 address-family ipv4 vrf A
  network 10.1.1.0 mask 255.255.255.0
 exit-address-family
 !
```

Verify base underlay and SRv6 process

To verify that the overlay control plane has successfully established and mapped tenant routing domains to their corresponding SRv6 service SIDs, network operators should execute diagnostic commands directly in the Cisco IOS XE CLI. Active L3VPN connectivity can be verified using show segment-routing srv6 sid to confirm that the local SID manager has successfully generated the uDT4 and uDT6 decapsulation behaviors and bound them to the target tenant VRFs. Additionally, operators can validate route advertisement and neighbor next-hop mappings by auditing the BGP database with `show ip bgp vpnv4 rd <RD> <prefix>`, examining prefix-to-SID resolution using `show ip route vrf <vrf> <prefix>`, and inspecting CEF-level hardware lookup properties using `show ip cef vrf <vrf> <prefix> internal` or `show ipv6 cef <prefix>`. Sourced pings across isolated VRFs should be executed to confirm 100% end-to-end data-plane reachability over the secure, stateless transport core.

Global BGP design and scaling

Once the foundational SRv6 L3VPN fabric is established, defense and enterprise architects must ensure the BGP control plane can scale across hundreds of isolated sites and multiple Autonomous System (AS) boundaries without degrading the hardware data plane.

This section outlines the global MP-BGP control-plane design, Autonomous System Number (ASN) management strategies, and inter-domain scaling models required to support hundreds of secure multi-tenant enclaves over an end-to-end SRv6 fabric.

ASN design and scalability

When deploying an automated SRv6 L3VPN fabric across hundreds of regional site locations, tactical enclaves, or Forward Operating Bases (FOBs), ASN management is a primary scalability concern. Network architects are faced with two main design paths: allocating a unique Autonomous System Number to every individual site router, or standardizing on a single, shared private ASN (e.g., AS 65001) across all remote sites.

Unique ASN per site (the explicit model)

- Mechanics: Each remote site is assigned its own unique private ASN.
- Benefits: Prevents BGP loop-prevention conflicts natively, as the AS-Path attribute naturally traces unique hop histories.
- Drawbacks: Severely complicates Zero-Touch Provisioning (ZTP) and automated deployment templates. Each site configuration must be customized with unique ASN parameters, drastically increasing operational overhead and the risk of configuration drift.

Shared ASN across sites (the standardized model)

- Mechanics: All edge routers share a single, standardized private ASN.
- Benefits: Heavily simplifies ZTP and template standardization. Edge templates remain identical, allowing rapid deployment without customizing BGP peer processes per site.
- Drawbacks: Triggers eBGP's native AS-PATH loop prevention check , which must be systematically resolved to allow inter-branch communication.

Shared ASN loop prevention

The major challenge of standardizing on a single, shared private ASN across multiple edge sites lies in BGP's default loop-prevention behavior. By design, when an eBGP router receives a prefix advertisement, it inspects the AS-PATH attribute. If the router detects its own local ASN in the path, it assumes a routing loop has occurred and silently discards the update.

In a traditional MPLS or IP network, a dropped BGP route simply results in a prefix reachability failure. However, in an SRv6-enabled L3VPN fabric, the consequences are far more catastrophic:

$$Dropped BGP Route \Longrightarrow Missing Service SID Block \Longrightarrow Data - Plane\ Blackhole$$

Figure 10. Shared ASN loop prevention and the srv6 sid drop impact

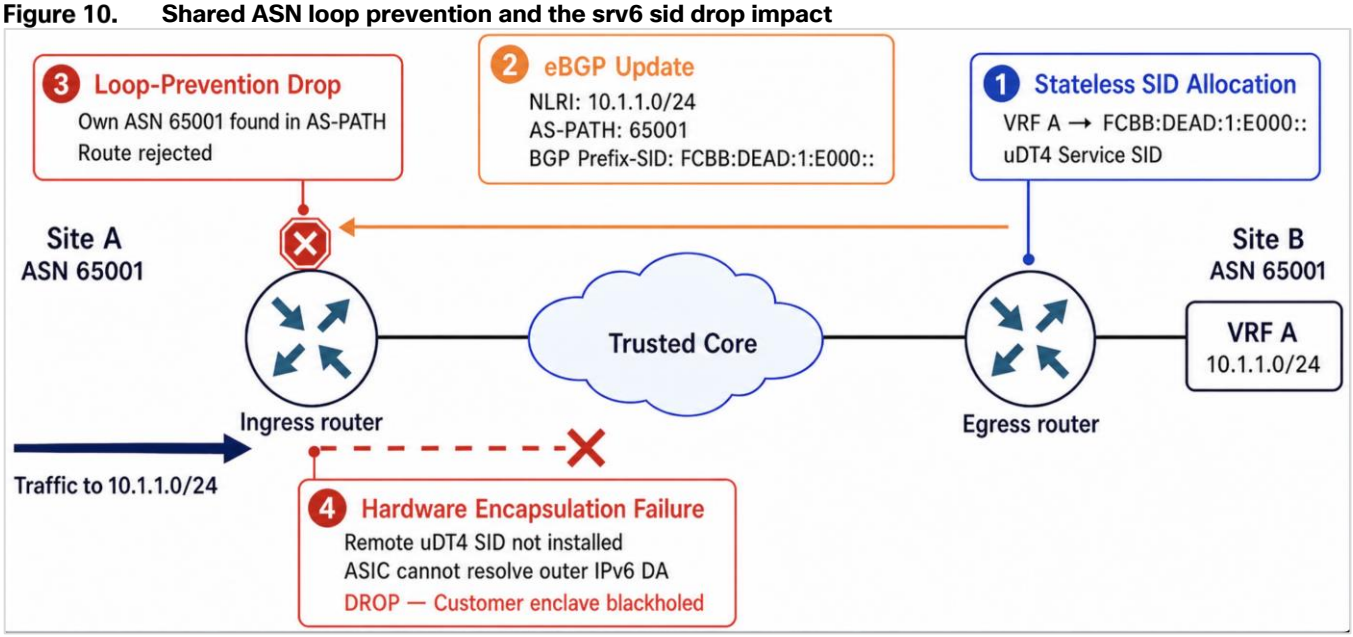

- Stateless SID allocation: The egress router allocates a unique 128-bit SRv6 Service SID (configured as a micro-SID aggregate, e.g., FCBB:DEAD:1:E000:: representing a uDT4 behavior) directly to a tenant VRF.
- SID propagation via BGP: The egress router advertises the customer prefix (e.g., 10.1.1.0/24) to the rest of the network via eBGP. Rather than appending a standard MPLS label, BGP attaches the 128-bit Service SID inside the BGP Prefix-SID Attribute of the BGP update.
- The loop-prevention drop: When a remote edge router sharing the same ASN receives this BGP update, it detects its own ASN in the path and drops the route.
- Hardware encapsulation failure: Because the BGP route is rejected, the ingress router never learns the remote Service SID (uDT4). When customer traffic destined for 10.1.1.0/24 arrives at the ingress PE, the hardware cannot resolve the outer IPv6 destination address (which must be stamped with the remote Service SID). The ingress router is forced to drop the traffic at the ASIC layer, blackholing the customer enclave.

Control-plane loop resolution

To prevent the data-plane encapsulation failures caused by shared ASN loop prevention, network designers must choose between two primary control-plane override models: AS-Override (a PE-centric solution) or Allowas-in (a CE-centric solution).

Resolution a: using as-override

Figure 11. Resolution a: using as-override

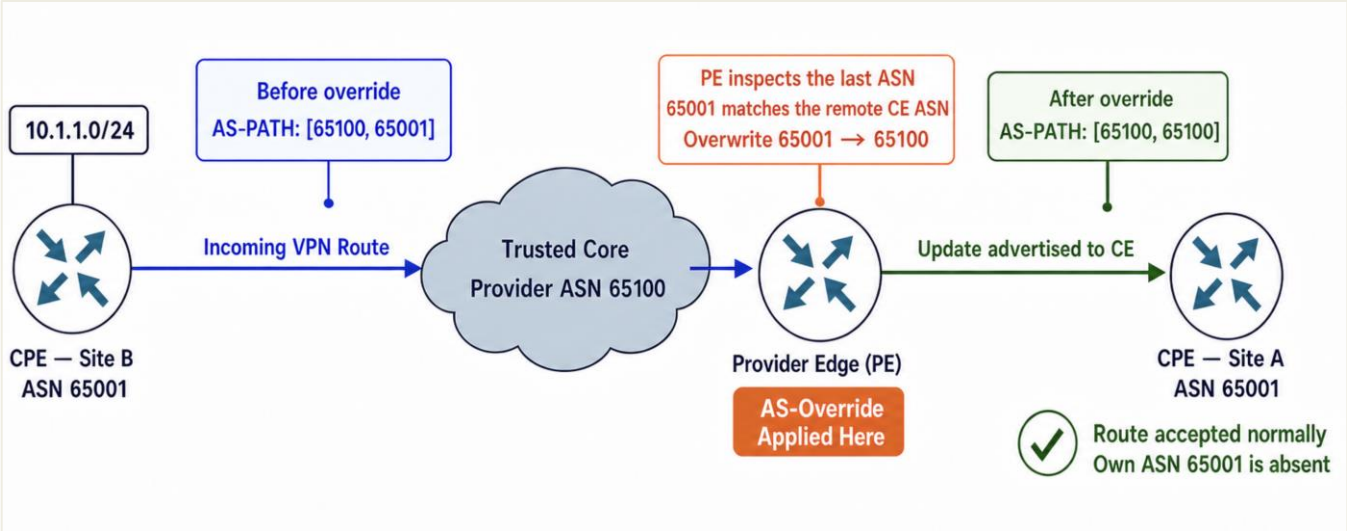

- **Concept:** Applied directly on the core Provider Edge (PE) routers.
- **Operation:** When advertising a prefix to a Customer Premises Edge (CPE) router, the PE inspects the AS-PATH. If the last ASN in the path matches the remote CPE's ASN, the PE overwrites it with its own ASN before sending the update.
- **Key advantage:** This requires zero custom loop-prevention configuration on the branch devices, enabling perfectly homogeneous "dumb-edge" templates that accelerate zero-touch provisioning at massive scale.

Resolution b: using allowas-in

Figure 12. Resolution b: using allowas-in

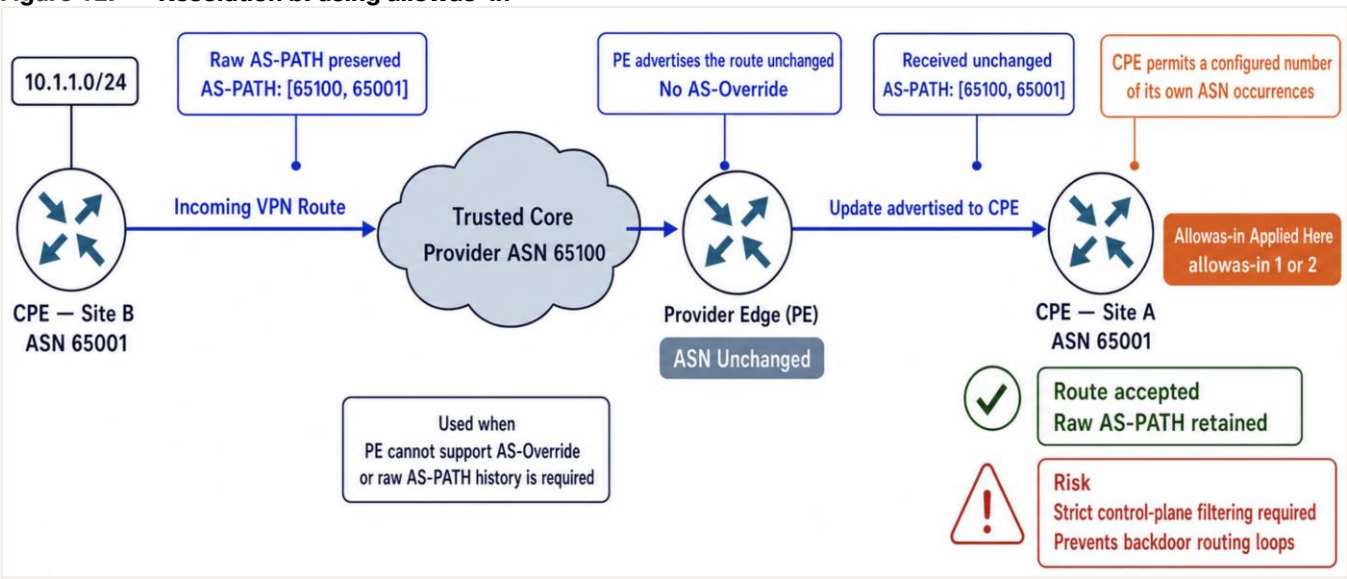

- **Concept:** Applied directly on the Customer Premises Edge (CPE) routers.
- **Operation:** The CPE is explicitly configured to override its default loop-prevention behavior, allowing a specified number of occurrences (typically allowas-in 1 or allowas-in 2) of its own ASN in the received AS-PATH.

- **Why it is used:** Necessary when the service provider cannot or will not support AS-Override on the PE routers, or when maintaining the raw, unmodified AS-PATH history is required for path-steering analytics.
- **Risk:** Demands rigorous control-plane filtering on the CPE to prevent backdoor routing loops.

Scaling the inter-domain data plane

For global networks spanning multiple administrative domains or separate Autonomous Systems (e.g., AS 64001 and AS 64002), standard eBGP behavior mandates that when BGP routes are advertised across an AS boundary, the exiting ASBR automatically rewrites the BGP Next-Hop attribute to its own interface IP address (next-hop-self). In legacy networks, this Next-Hop rewrite was a routine control-plane operation; however, in a Segment Routing over IPv6 (SRv6) architecture, this default behavior creates a severe Data-Plane Bottleneck.

The next-hop changed bottleneck

Figure 13. The next-hop changed bottleneck (default behavior: next-hop-self)

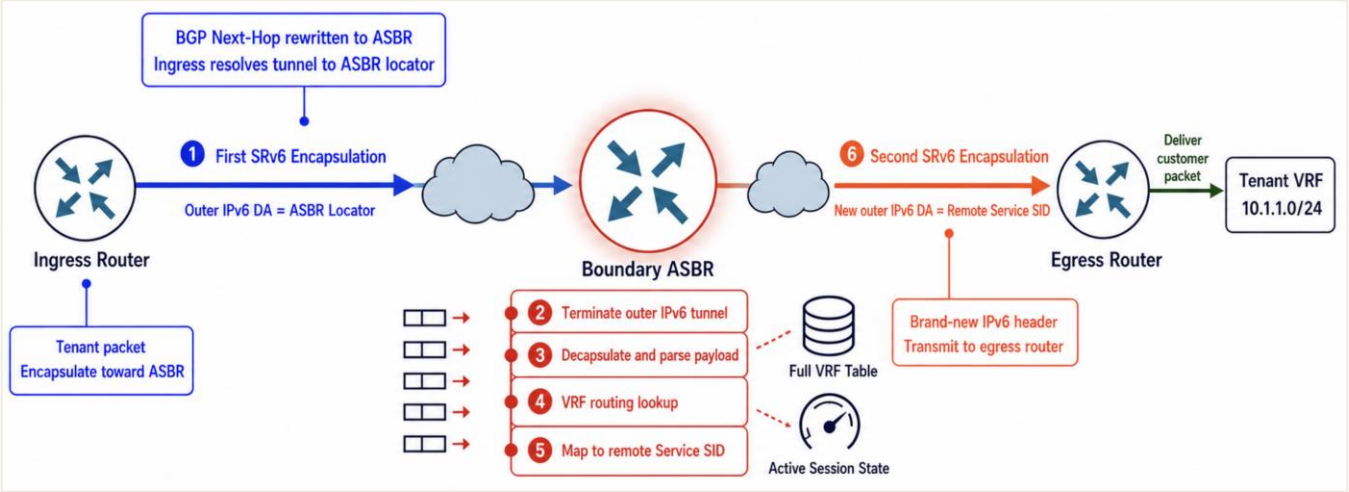

Because the BGP Next-Hop directly dictates the target IPv6 Locator of the egress router, rewriting the next-hop to the ASBR forces the ingress router to build its SRv6 tunnel terminating directly on the ASBR. The data plane behaves as follows:

- **Stateful re-encapsulation:** The ingress router encapsulates tenant traffic using the boundary ASBR's locator prefix as the destination IPv6 carrier address.
- **Hardware decapsulation at the border:** When the encapsulated packet hits the ASBR, the ASBR is forced to terminate the outer IPv6 tunnel. It must decapsulate the packet, parse the payload, and perform a hardware-intensive IP routing lookup in its localized VRF table to find the next-hop toward the downstream domain.
- **Data-plane re-encapsulation:** The ASBR performs a second lookup, maps the packet to the remote domain's Service SID, encapsulates it in a brand-new outer IPv6 header, and transmits it toward the egress router.
- **The scalability toll:** This model destroys the stateless core paradigm. The boundary ASBRs must maintain full VRF tables, active session states, and perform double-lookup encapsulation for every single packet, turning the ASBR into an expensive, resource-saturated choke point that severely restricts throughput .

The next-hop unchanged solution

Figure 14. The next-hop unchanged solution (neighbor x.x.x.x next-hop-unchanged)

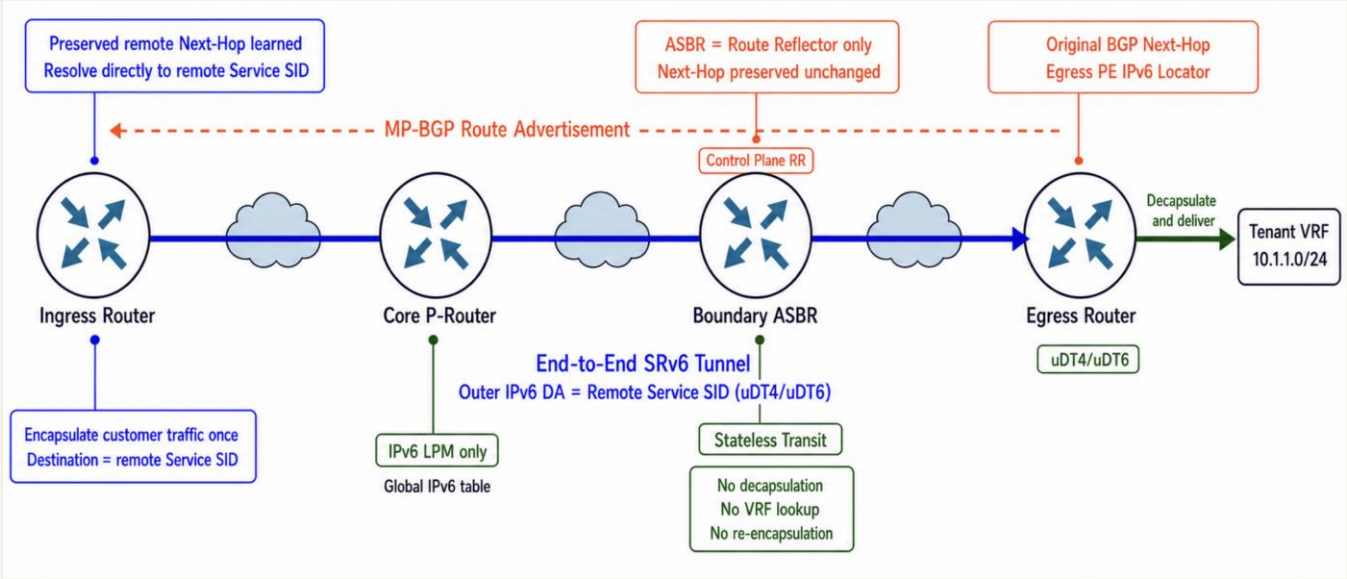

To preserve complete statelessness and maximize forwarding throughput, multi-domain networks utilize the Next-Hop Unchanged BGP design rule. The ASBRs are explicitly configured to act strictly as control-plane Route Reflectors (RRs), propagating the MP-BGP updates across AS boundaries while preserving the original Next-Hop attribute (the egress PE's native IPv6 locator) untouched.

- **End-to-end tunneling:** The ingress router receives the route with the preserved remote Next-Hop. It encapsulates customer traffic directly to the remote egress PE's Service SID (uDT4/uDT6).
- **Pure IPv6 transit:** Because the Next-Hop is preserved as the remote PE's locator, the intermediate ASBRs and core P-routers do not participate in L3VPN decapsulation or VRF routing. They forward the packet using simple, standard longest-prefix-match (LPM) lookups against their global IPv6 tables.
- **ASBR offloading:** The boundary nodes are kept entirely oblivious to individual customer prefixes, VRFs, or service labels, allowing the border routers to scale infinitely without resource saturation.

Tech tip: While BGP next-hop manipulation is frequently analyzed through the lens of basic routing reachability and ASBR hardware offloading, its impact on Segment Routing Traffic Engineering (SR-TE) is far-reaching. When a VPN prefix is advertised with a specific BGP Color Extended Community, the ingress router matches this color against a local template and instantiates a dynamic SR-TE policy.

Crucially, the destination endpoint of this dynamic SR-TE policy is the BGP Next-Hop of the received prefix. With ‘next-hop self’ as the next-hop is rewritten to the ASBR, the ingress PE's Constrained Shortest Path First (CSPF) algorithm can only compute a path up to the ASBR. However, applying ‘next-hop-unchanged’ on ASBRs restores full traffic engineering integrity across administrative boundaries, Because the BGP Next-Hop is preserved as the remote egress PE's Loopback0 address, the ingress PE's ODN engine instantiates a seamless, end-to-end SR-TE policy targeting the true final endpoint.

SRv6 traffic engineering

Standard shortest-path routing protocols are fundamentally inadequate for high-consequence mission-critical enclaves. Relying strictly on default interface costs forces high-priority telemetry, real-time command-and-control (C2), and voice/video traffic onto the same physical paths as bulk, unclassified backups. This exposure can lead to congestion, packet drops, and severe latency spikes.

Segment Routing over IPv6 Traffic Engineering (SRv6-TE) completely decouples the service intent from the physical network topology, allowing network administrators to enforce strict Service Level Agreements (SLAs), topological path-pinning, and line-rate encryption constraints at the ingress edge of the network. Crucially, because SRv6 uses a source-routing model, the packet carries its own path waypoints directly in the outer IPv6 destination address and Segment Routing Header (SRH). The intermediate core transit routers (P-nodes) remain completely stateless and oblivious to individual traffic-engineering tunnels, resolving the massive control-plane signaling storms and core memory bloat that crippled legacy RSVP-TE architectures.

This section outlines the advanced Segment Routing Traffic Engineering (SR-TE) toolkit as implemented on Cisco IOS XE platforms—specifically targeting the Cisco 8000 Series Secure Routers and Catalyst edge platforms. It details the technical mechanics, dynamic path-calculation parameters, and configuration templates required to deploy a highly secure, application-aware, and self-healing transport fabric.

Per-Destination (PDP) vs. Per-Flow (PFP)

Cisco IOS XE secure platforms support two distinct Traffic Engineering policy models depending on the required level of granularity:

Per-Destination Policy (PDP) – "macro-steering"

A Per-Destination Policy (PDP) is a simpler destination-based routing model. It steers all traffic destined for a specific endpoint (such as an egress Hub or remote branch loopback) over a single, specific candidate path.

Figure 15. Per-Destination Policy (PDP) – "macro-steering"

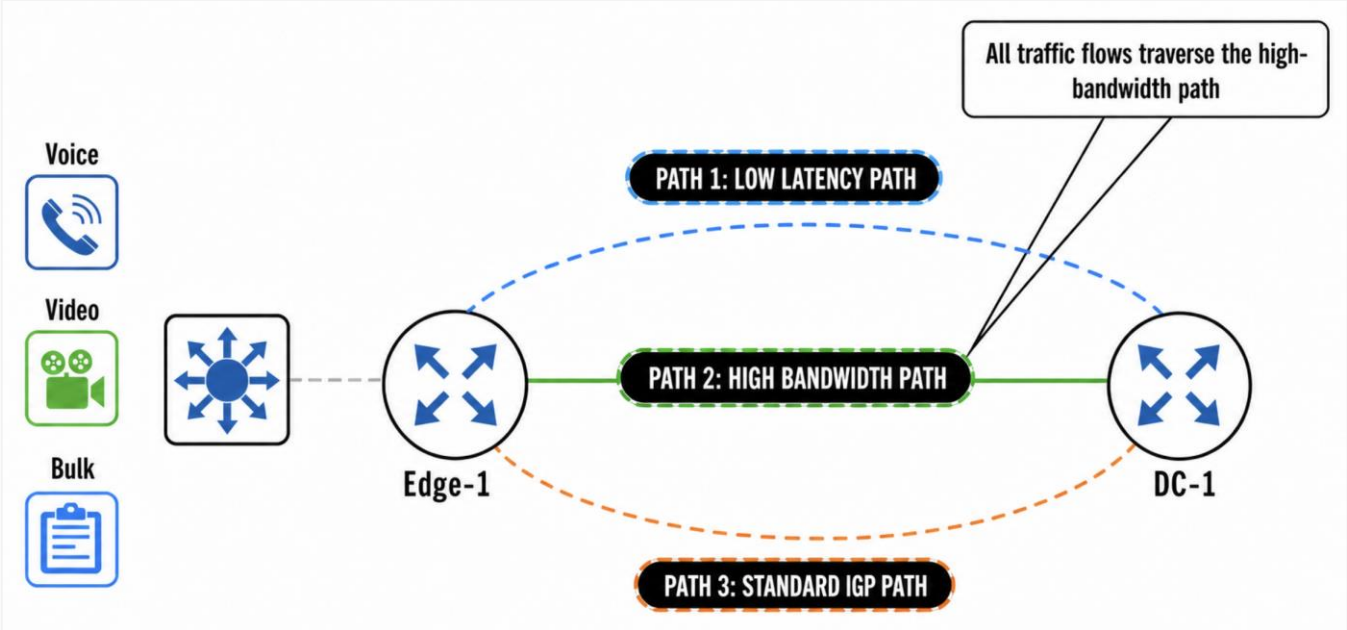

- **Deciding factor:** The path selection is based purely on the destination IP address of the incoming packet.

- **Behavior:** Whether the packet contains real-time critical Voice payload or a background unclassified bulk file transfer, it will take the exact same traffic-engineered path to that destination, assuming they are bound to the same color prefix.

Per-Flow Policy (PFP) – "micro-steering"

A Per-Flow Policy (PFP) provides a highly granular, application-aware routing model. It allows a single destination to be mapped to multiple completely different physical paths across the fabric based on the packet's internal attributes (such as QoS DSCP markings or NBAR application classifications).

Figure 16. Per-Flow Policy (PFP) – "micro-steering"

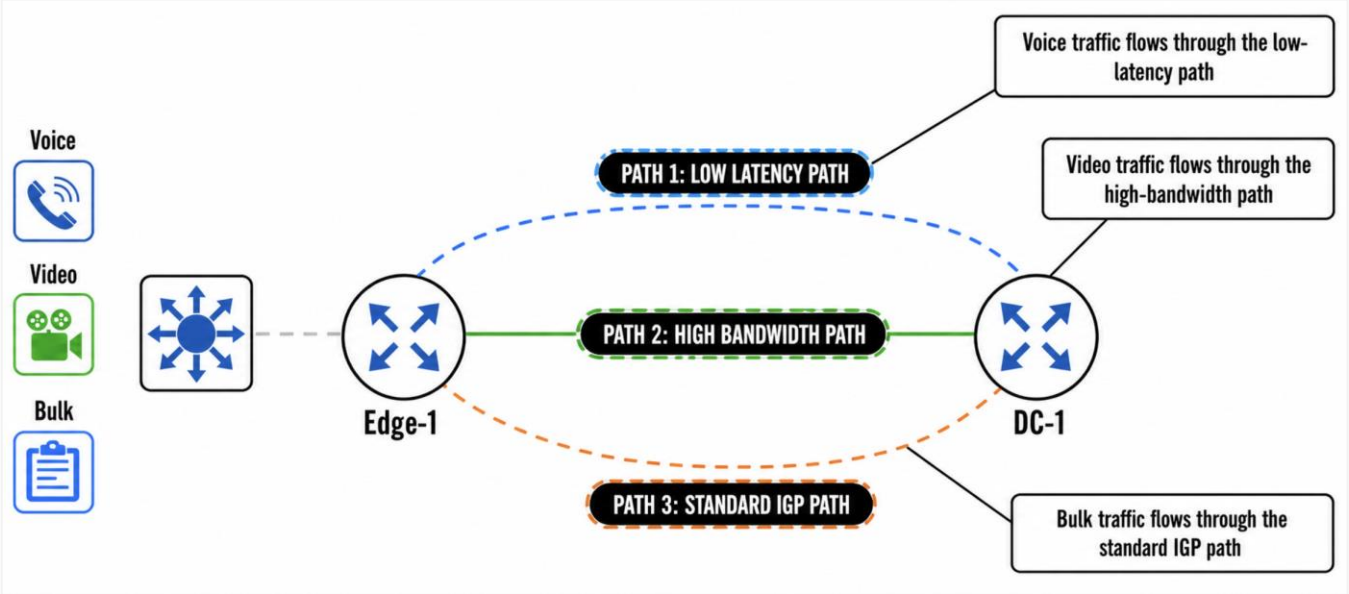

The Parent-Child Relationship: In a PFP architecture, a strict division of labor is enforced:

- **The Parent PFP (the dispatcher):** The Per-Flow Policy acts as an edge classification dispatcher. It does not actually compute paths across the physical topology. Its sole responsibility is to intercept incoming packets, match their QoS or NBAR classification, assign them an internal Forwarding Class (0 through 7), and map those forwarding classes to specific child "Color" tags.
- **The Child PDP (the path builder):** The child Per-Destination Policies represent the actual physical paths. Each child policy corresponds to a specific Color and endpoint. This is where the physical traffic engineering, path computation (CSPF), metric constraints (delay, TE, IGP), and link affinities are configured.

Ingress classification, ePBR, and NBAR pipeline

At the ingress edge of the SRv6 WAN, Cisco IOS XE utilizes Endpoint Policy-Based Routing (ePBR) and Network-Based Application Recognition (NBAR) to intercept and classify LAN traffic. This process marks the transition from standard destination lookup to intent-based path dispatching.

Figure 17. Ingress classification, ePBR, and NBAR pipeline

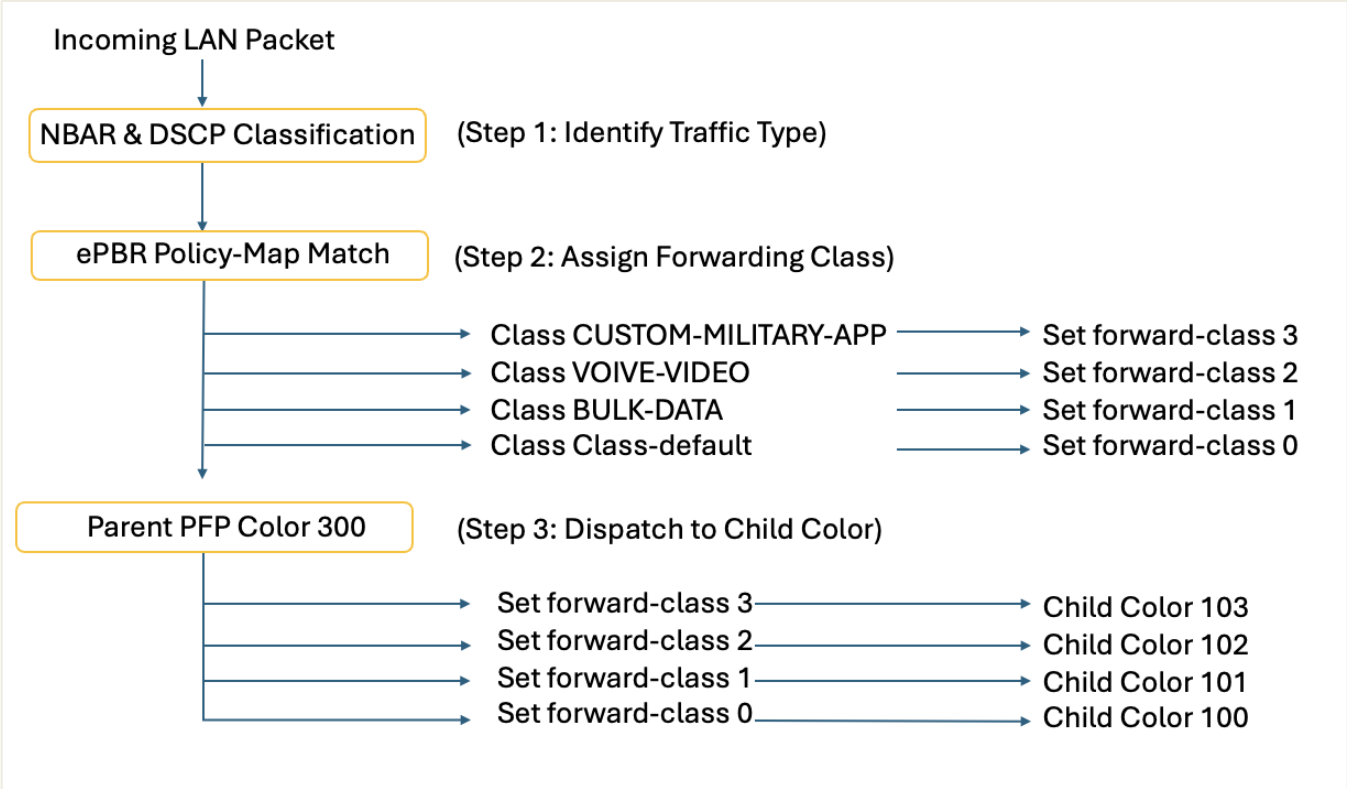

The execution sequence:

- **Application identification:** The ingress physical interface or VLAN subinterface actively inspects incoming payloads. High-priority traffic is identified using either Layer 3 DSCP markings or Layer 7 deep packet inspection via NBAR.
- **ePBR mapping:** An ePBR policy-map intercepts the classified flows and maps them to internal forwarding classes. For example, classified military C2 traffic is marked as forward-class 3, voice/video as forward-class 2, and bulk data as forward-class 1.
- **Parent-policy handoff:** The ingress PE hands these forwarding classes to the parent PFP (Color 300). The parent policy maps forward-class 1 to child color 101 (Best Effort) and forward-class 2 to child Color 102 (Low-Latency).

Path selection mechanics: metrics and affinities

Once a packet is steered into a child Color (PDP), the ingress router must compute the mathematically optimal path to the remote endpoint using the Constrained Shortest Path First (CSPF) algorithm. In this design, path selection is driven by three foundational constraint parameters:

TE-metric

- **Definition:** The Traffic Engineering Metric is an administrative cost manually assigned to physical interfaces.
- **Behavior:** It allows architects to define a "policy cost" completely independent of the standard IGP (IS-IS) cost. Unless otherwise specified, the TE-metric acts as the default parameter for all constraint calculations, enabling absolute, deterministic control over which links are preferred.

Link affinity

- **Definition:** Affinities act as topological gatekeepers across the WAN. Operators map human-readable names to specific bit-positions (0 through 31) representing physical interface characteristics.
- **Usage:** By applying constraints like include-all, include-any, or exclude-any to a child policy, administrators can force traffic only over approved links. For instance, the affinity name SECURED (bit-position 0) is mapped to MACsec-encrypted fiber, while UNSECURED (bit-position 1) is mapped to standard internet-facing links.

Dynamic delay

- **Definition:** Rather than relying on static, rigid costs, dynamic delay-based steering allows the network to automatically route around congestion, link degradation, and micro-bursts.
- **Integration:** By enabling Performance Measurement (PM), the router actively monitors link latency and advertises real-time millisecond delay values throughout the IS-IS link-state database. The CSPF engine can then dynamically calculate the absolute lowest-latency path to a remote destination in real-time.

  The next section further explains how this parameter works.

Dynamic delay and Performance Measurement

Mission-critical networks must dynamically adapt to real-world degradation. Relying on basic link state (up/down) is no longer sufficient; a link may be physically "up" but suffering from severe congestion or micro-bursts that destroy voice and video quality.

To solve this, Cisco IOS-XE utilizes SRv6 Performance Measurement (PM). PM transforms the router into an active telemetry sensor, allowing the network to steer traffic based on real-time mathematical latency rather than static costs.

How latency is calculated

Figure 18. How Latency is Calculated: The PM Probing Mechanism

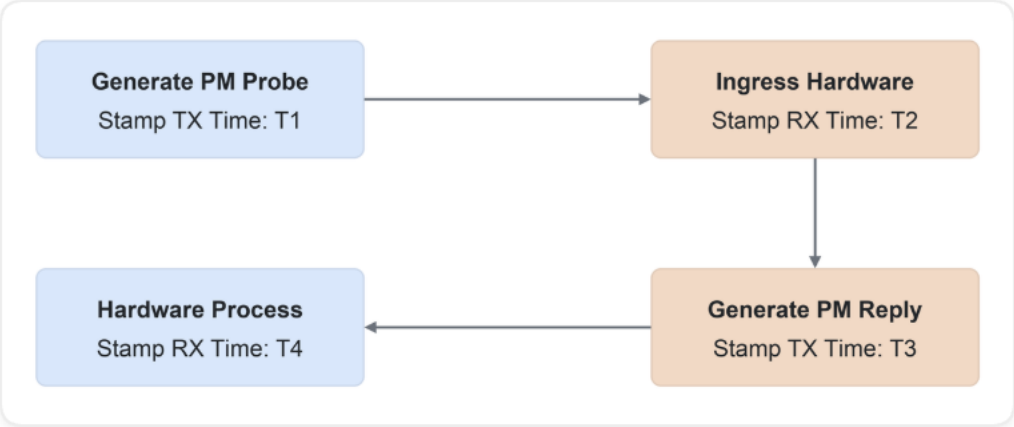

Instead of guessing the health of a path, the SRv6 headend router continuously sends active PM probe packets across the segment lists of candidate paths. By default, these probes are fired at aggressive 3-second intervals.

The latency calculation relies on a precise, four-way timestamping mechanism between the Querier (the ingress router) and the Responder (the egress router or next-hop):

- T1: The exact transmit (TX) timestamp when the Querier sends the probe.

- T2: The exact receive (RX) timestamp when the Responder gets the probe.
- T3: The transmit (TX) timestamp when the Responder sends the reply.
- T4: The receive (RX) timestamp when the Querier receives the reply.

Using these hardware-level timestamps, the router calculates the exact latency. The network can be configured to measure Two-Way Mode (Round-Trip Delay) using the formula (T4 - T1) - (T3 - T2), which automatically deducts the processing time spent on the Responder. Alternatively, it can calculate a highly accurate One-Way Delay by halving the round-trip result.

IGP advertisement and dynamic path selection

Once the router calculates the real-time delay, it does not keep this information to itself. The router injects this delay metric into the underlay routing protocol (IS-IS). IS-IS then floods this delay as an extended link attribute throughout the entire topology.

When an SR-TE Per-Destination Policy (PDP) is configured with metric type delay, the router completely ignores standard shortest-path IGP costs. Instead, it runs its path-computation algorithms strictly against the real-time delay metrics advertised by IS-IS, dynamically calculating and forcing traffic over the absolute lowest-latency path available across the continent.

Liveness detection and auto-reoptimization

Beyond just measuring latency, PM acts as a rapid-failover trigger through Liveness Detection.

If the PM probes stop returning (by default, after 3 consecutive missed probes), the router immediately flags the path state as "Down". If the policy is configured with invalidation-action down, the router will instantly tear down the impaired hardware forwarding path and automatically reoptimize the traffic to a healthy alternate candidate path.

Validation in action: defeating the "noisy neighbor"

The true power of PM delay metrics is against an artificial "noisy neighbor" impairment.

For instance, if a primary provider transit link (P1) was choked using an 8 Mbps output shaper to simulate severe network saturation. Without PM, standard routing protocols would have continued blindly sending traffic into the choke point because the physical link never dropped.

However, with PM enabled, the routers immediately detected the congestion:

- The active probes registered a massive latency spike across the tunnel.
- IS-IS instantly updated the SR-TE topology with the new degraded delay values.
- The SR-TE policies utilizing metric type delay instantly rejected the congested route, re-calculating and seamlessly steering Voice (DSCP EF) and Business Critical (DSCP AF31) traffic onto a completely clean alternate path.
- The validation recorded zero packet loss for the protected applications during the congestion event.

Tech tip: When defining the dynamic path for a PDP, the routing algorithm can optimize against only a single metric. You can configure the path computation to use metric type te OR metric type delay, but you cannot optimize on both simultaneously. However, you can seamlessly combine your chosen metric with Link-Affinity constraints to enforce uncompromising SLAs. For example, a mission-critical Voice policy can be configured to optimize for the lowest real-time latency (metric type delay) AND strictly traverse only encrypted links (affinity include-all name secure). Conversely, a bulk data policy can optimize based on static business cost (metric type te) while excluding expensive satellite links (affinity exclude-any name satcom).

Scaling with automated steering

Manually configuring static traffic-engineered tunnels and segment lists between hundreds of secure enclaves is operationally unsustainable and destroys the stateless simplicity of SRv6. To solve this, this architecture enforces the use of On-Demand Next-Hop (ODN) and Automated Steering (AS).

Rather than pre-provisioning end-to-end paths, the ingress router dynamically instantiates the appropriate traffic-engineered SR-TE policy only when active customer service flows require it.

Here's an example topology on SRv6 L3VPN architecture and the operational sequence of how the fabric dynamically builds traffic-engineered paths without manual tunnel configuration:

Figure 19. SRv6 L3VPN architecture and the operational sequence

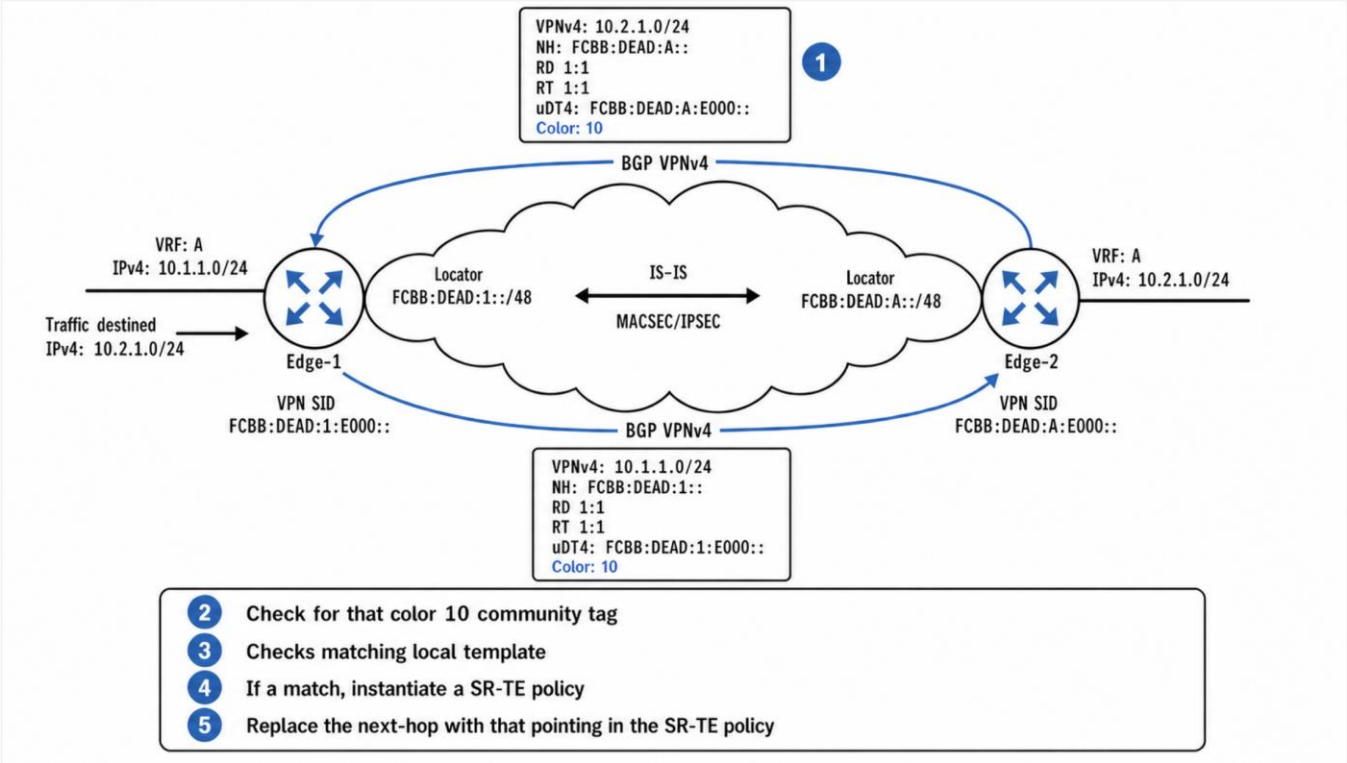

Step 1. BGP advertisement and tagging: The egress router (Edge-2) advertises a tenant prefix (e.g., VPNv4: 10.2.1.0/24) across the control plane. Alongside the standard route targets and its SRv6 Service SID (uDT4: FCBB:DEAD:A:E000::), it attaches a specific SLA hint—the BGP Color Extended Community (Color: 10).

Step 2. Color identification: When the ingress router (Edge-1) receives the BGP prefix update, it intercepts and checks for that specific Color 10 community tag.

Step 3. Template matching: Edge-1 checks its local configuration database for a matching on-demand color 10 SR-TE policy template that defines the physical routing constraints (such as low-latency or MACsec-only link affinities) required for that SLA.

Step 4. Policy instantiation (ODN): If a matching template is found, the ingress router dynamically instantiates an SR-TE policy tailored specifically to those constraints, bound directly to the egress destination locator (FCBB:DEAD:A::).

Step 5. Next-hop replacement (automated steering): The routing engine automatically overrides the standard BGP next-hop table entry for the 10.2.1.0/24 prefix, replacing it with a pointer directing traffic into the newly instantiated SR-TE policy. As data-plane traffic arrives, the ASIC encapsulates the

payload using the explicit path-constrained header, completely automating the traffic engineering lifecycle.

Tech tip: The Routing ID and SLOC Prerequisite: For the local CSPF algorithm to compute valid traffic-engineered paths, a strict separation of duties must be maintained in the IS-IS database:

1) Control-plane path computation source: The router's Loopback0 address must be explicitly advertised within the IS-IS process and configured as the ipv6 router-id. This provides the local CSPF graph with a fixed, reachable starting node.

2) Data-Plane Encapsulation Source: The aggregate SLOC (Service Locator) address is utilized strictly as the outer IPv6 source address for hardware encapsulation.

3) The Golden Rule: Both the Loopback0 control-plane router ID and the SLOC data-plane aggregate locator must be completely reachable and fully converged within the IS-IS database. If either prefix is missing or unreachable, dynamic ODN path calculations will silently fail, leaving candidate paths in an inactive state.

Bringing it all together

The following blueprints represent the Cisco IOS XE configurations required to establish parent Per-Flow Policies, Child PDP constraints, ePBR classification, and active Performance Measurement probing.

Step 1. Define the underlay affinity bitmaps

```
router isis SRTE-OVLY
 affinity-map SECURED bit-position 0
 affinity-map UNSECURED bit-position 1
 !
 address-family ipv6
  router-id Loopback0
  segment-routing srv6
   locator SLOC
!
```

Step 2. Traffic classification (NBAR & DSCP)

```
1a. Define the Custom Military Application via NBAR DPI
ip nbar custom SECURE-C2 transport tcp id 1
 port 444
 direction destination

! 1b. Match the NBAR Custom App
class-map match-all CUSTOM-MILITARY-APP
 match protocol SECURE-C2

! 1c. Match standard DSCP tags
class-map match-all VOICE-VIDEO
 match dscp ef
class-map match-all BULK-DATA
 match dscp af11
```

Step 3. Forwarding class assignment via ePBR

```
policy-map type epbr SRv6_PFP
 class CUSTOM-MILITARY-APP
  set forward-class 3
 class VOICE-VIDEO
  set forward-class 2
 class BULK-DATA
  set forward-class 1
 class class-default
  set forward-class 0
```

Step 4. Define the Per-Destination Policies (PDP) constraints

```
segment-routing traffic-eng
 ! Military App Path (Color 103)
 on-demand color 103
  authorize
  candidate-paths
   preference 100
   constraints
    segments
     dataplane srv6
    affinity
     include-all
      name secure

 ! Voice/Video Path (Color 102)
 on-demand color 102
  authorize
  candidate-paths
   preference 100
   constraints
    segments
     dataplane srv6
    affinity
     include-all
      name voice-video
    dynamic
     metric
      type delay
```

Step 5. Create the Parent Per-Flow Policy (PFP)

```
segment-routing traffic-eng
 ! Parent PFP
 on-demand color 1000
  authorize
  candidate-paths
```

```
   preference 100
   per-flow
    forward-class 0 color 100
    forward-class 1 color 101
    forward-class 2 color 102
    forward-class 3 color 103
```

Step 6. Apply the policy and enable NBAR

```
interface GigabitEthernet0/0/0
 ! Enable NBAR Deep Packet Inspection on the interface
 ip nbar protocol-discovery
 ! Apply the ePBR Classification Policy
 service-policy type epbr input SRv6_PFP
```

Operational verification commands

Network operators can leverage the `show segment-routing traffic-eng policy name * detail` command to audit the overall health, administrative/operational status, and candidate path binding SID allocations for active policies. To verify the active Performance Measurement probing engine, issue the `show performance-measurement summary` and `show performance-measurement sr-policy` commands, which display query transmission counts, hardware-based round-trip latency statistics, and state tracking. Furthermore, operators can use `show performance-measurement history interfaces adv` to trace the dynamic advertisement of delay values into the IS-IS link-state database . Finally, running `show segment-routing traffic-eng topology ipv6` provides an exhaustive trace of the CSPF topological graph, showing available nodes, adjacency SIDs, static TE-metrics, and active affinity-map bit positions.

Topological isolation using flexible-algorithm

While Segment Routing Traffic Engineering (SR-TE) provides highly granular, flow-based path control at the edge, building a secure, multi-tenant network infrastructure requires a macro-level mechanism to isolate entire enclaves and secure routing planes. Flexible Algorithm (Flex-Algo) provides the definitive architectural framework for native, topology-level network slicing. Natively integrated into the Interior Gateway Protocol (IGP) layer, Flex-Algo enables the underlay to dynamically compute entirely distinct, mathematically isolated routing topologies over a single physical infrastructure. By allowing operators to define custom algorithms optimized for specific constraints—such as restricting transport to strict hardware-level encryption or minimum latency—Flex-Algo eliminates accidental data leakage and ensures that critical enclaves operate as if the rest of the network simply does not exist.

This section outlines the native Flex-Algo data-plane mechanics, the step-by-step Cisco IOS XE configuration templates, and the strict design rules governing locator isolation and application-aware steering.

Foundational principles of flex-algo slicing

In a native SRv6 underlay, the Interior Gateway Protocol (IGP)—typically IS-IS—is responsible for calculating the shortest path between nodes using standard interface metrics. By default, all routing nodes participate in Algorithm 0, which runs the standard Shortest Path First (SPF) algorithm against the default IGP link metrics to construct a single, best-effort forwarding table.

Flexible Algorithm completely redefines this model by allowing operators to define custom path-computation algorithms (numbered 128 through 255) natively within the IGP. Rather than computing a single shortest-path topology, the IGP process runs entirely independent SPF calculations for each configured algorithm. Each customized algorithm computes a unique, mathematically isolated routing topology over the same physical hardware based on user-defined constraints:

- **Metric minimization:** Optimizes the topology against specific metrics, such as minimizing dynamically measured link delay or administrative Traffic Engineering (TE) costs, rather than standard IGP interface costs.
- **Resource exclusion:** Excludes specific links from the path-computation graph based on administrative link characteristics (Link Affinities).

Because these calculations are executed natively inside the distributed IGP, Flex-Algo provides an exceptionally scalable traffic engineering solution. Core transit nodes (P-routers) do not require manual tunnel configurations, active state tables, or centralized controller coordination. The network automatically and autonomously converges on the optimized topologies, dynamically routing traffic around failed links or congested nodes while strictly honoring the custom algorithmic constraints.

Design topologies: secure vs. low delay slices

To enforce strict isolation and meet distinct Service Level Agreements (SLAs) across a shared physical underlay, this validated design carves out two specialized, custom logical topologies alongside the best-effort Algorithm 0 core:

Figure 20. Designing the topologies: secure vs. low delay slices

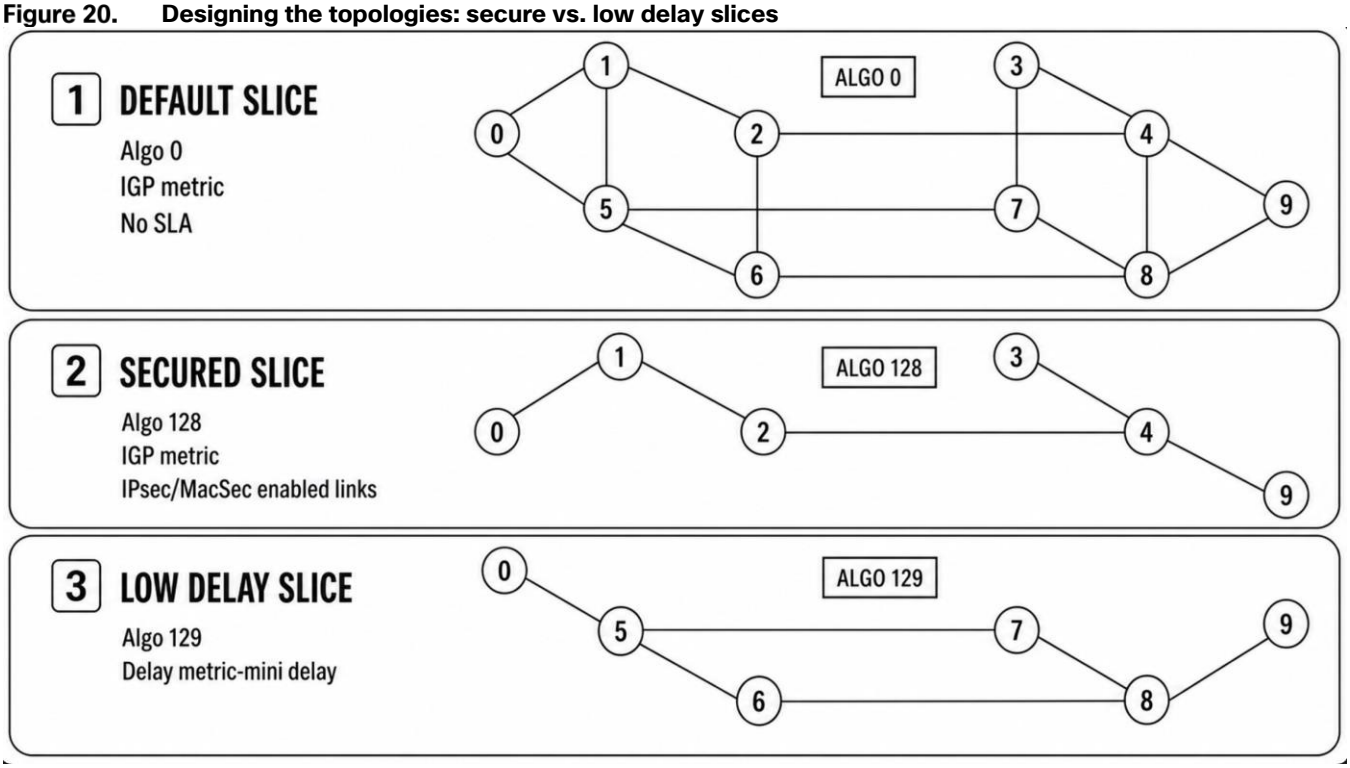

- Algorithm 128: the secure slice

  In high-consequence enclaves, architects must guarantee that classified tenant traffic is strictly restricted to secure physical paths. Algorithm 128 is designed as a secure slice that mathematically restricts routing to links protected by hardware-level encryption.

  ◦ **Algorithmic constraint:** The IGP calculates paths utilizing the manual Traffic-Engineering (TE) metric as its base optimization parameter to maintain administrative cost control.

  ◦ **Resource enforcement:** It explicitly restricts path calculation utilizing link affinities. Any physical interface participating in Algorithm 128 must be tagged with a specific administrative group (e.g., SECURE_FIBER).

  ◦ **Security guarantee:** The SPF algorithm completely prunes any unencrypted or untrusted links from the topology graph prior to path calculation. To the secure tenant traffic routed within Algorithm 128, unsecure links simply do not exist, rendering accidental data leakage or routing spillover physically impossible.

- Algorithm 129: the low-delay slice

  For real-time enclaves transporting low-latency voice, video, or critical telemetry, paths optimized for physical distance often fail during network congestion. Algorithm 129 is designed as a low-delay slice that dynamically adapts to physical latency.

  ◦ **Algorithmic constraint:** Instead of static costs, Algorithm 129 optimizes path calculations strictly against the dynamic link delay metric.

  ◦ **Dynamic latency tracking:** The routers leverage active Performance Measurement (PM) delay-measurement probes, firing active telemetry queries across physical links at aggressive 3-second intervals to calculate real-time latency.

  ◦ **Automated convergence:** These dynamically measured delay values are flooded natively via IS-IS as extended link attributes. The IGP then calculates the absolute lowest-latency path across the

continent, automatically re-routing traffic around links experiencing congestion, queueing delay, or fiber degradation.

Operational separation: dedicated locators per slice

For topological isolation to be successfully executed in the data plane, a strict architectural mandate must be enforced: A separate, dedicated SRv6 Locator block must be assigned to each Flexible Algorithm.

If a network platform utilizes multiple customized algorithms, each Provider Edge (PE) or Customer Premises Equipment (CPE) router must be provisioned with distinct, non-overlapping aggregate locator blocks mapped directly to those algorithms:

- Default topology (Algo 0): Mapped to FCBB:DEAD:E::/48
- Secure topology (Algo 128): Mapped to FCBB:DEAF:E::/48
- Latency topology (Algo 129): Mapped to FCBB:DEAB:E::/48

Stateless data-plane forwarding mechanics

When a tenant VRF or secure subnet is onboarded, its service routes are advertised across MP-BGP with a next-hop pointing exclusively to the locator block of its designated flex-algo (e.g., binding the secure VRF to the Algo 128 locator space FCBB:DEAF).

Figure 21. Data-plane forwarding with flex-algo

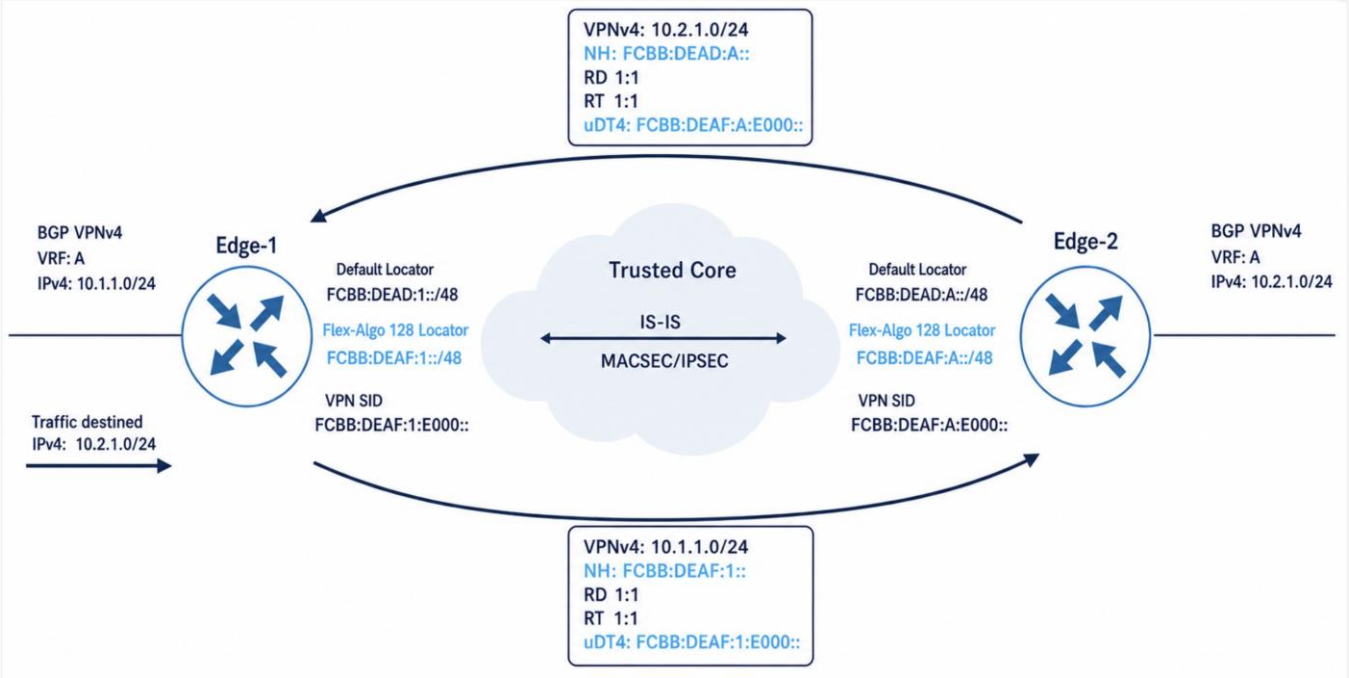

Notice how the prefix blocks have transitioned from DEAD to DEAF.

An important control-plane distinction is that MP-BGP does not carry an explicit "Flex-Algo: 128" attribute in its route updates to signal the topology. Instead, the association is handled implicitly through locator reachability:

- **The IGP foundation:** The underlay IGP (IS-IS) independently calculates the custom topology for Algorithm 128 and floods the dedicated locator prefixes (e.g., FCBB:DEAF:A::/48) throughout the network, explicitly tying them to Algorithm 128.

- **BGP inheritance:** On the egress PE, BGP inherits the locator assigned to that VRF address family (e.g., binding the VRF to the SECURE-128 locator)

When the ingress edge router encapsulates a tenant packet, it stamps the remote Service SID—which inherits the Flex-Algo 128 prefix—directly into the outer IPv6 Destination Address field. As the packet moves across the transit core:

- **Stateless lookup:** Core transit P-routers perform standard IPv6 longest-prefix-match (LPM) lookups against the destination address.
- **Algorithmic enforcement:** The hardware matches the packet's destination prefix (FCBB:DEAF::/48) to its local routing table computed strictly for Algorithm 128.
- **Strict Path isolation:** The ASIC forwards the packet exclusively over links participating in the Secure Slice, honoring the encryption constraints without intermediate core nodes maintaining any VPN state, label bindings, or per-flow context.

Configuration blueprint

The following validated configuration blueprint outlines the sequential commands required to instantiate global link affinities, define custom Flexible Algorithms, bind dedicated uSID locators, and advertise them natively under the IS-IS process on Cisco IOS XE platforms:

```
! Define global affinity mappings for algorithmic selection
router isis TACTICAL_CORE
 affinity-map SECURE_FIBER bit-position 1
 !
 ! Provision the custom Flex-Algo execution definitions
 flex-algo 128
  ! Base optimization parameters on custom administrative TE costs
  metric-type te
  ! Limit path computation exclusively to secure links
  advertise-definition
  affinity include-any SECURE_FIBER
 !
 flex-algo 129
  ! Base path optimization entirely on real-time link delay performance
  metric-type delay
  advertise-definition
 !
 ! Bind individual router locators to explicit algorithm topologies
 segment-routing srv6
  locators
   locator SECURE_SLICE_LOCATOR
    prefix fcbb:dead:1128::/48
    ! Explicitly tie this locator space to the Flex-Algo 128 calculation engine
    algorithm 128
   !
   locator LATENCY_SLICE_LOCATOR
    prefix fcbb:dead:1129::/48
    ! Explicitly tie this locator space to the Flex-Algo 129 calculation engine
    algorithm 129
```

Design rules: prefix-based slicing vs. flow-based steering

When planning a scalable WAN architecture, network designers must maintain a clear operational division between Topology-Based Slicing (Flex-Algo) and Application-Aware Flow Steering (PFP/PDP), as they utilize completely different layers of the Cisco IOS XE control plane to solve distinct routing challenges:

Table 2. Design rules: prefix-based slicing vs. flow-based steering

Design variable	Topology-based slicing (flex-algo)	Application-aware flow steering (PFP/PDP)
Operational layer	Natively calculated at the IGP (IS-IS) routing table level.	Enforced at the Edge Policy Layer using Endpoint Policy-Based Routing (ePBR).
Granularity	Prefix-Level (" All-or-Nothing"): Applies macro-routing decisions to entire subnets, prefixes, or VRFs.	Flow-Level (Application-Specific): Applies micro-routing decisions based on L4-L7 packet headers, NBAR signatures, or DSCP QoS markings.
Packet inspection	Ignored. Forwarding is executed strictly on the outer IPv6 destination address aggregate.	Active. Edge interfaces inspect inner payload parameters to map packets to specified forwarding classes (0-7).
Data plane path	Natively calculated shortest path within the algorithm's topology; no explicit segment lists required.	Explicitly directed over dynamic SR-TE policies utilizing pre-calculated segment lists.

Architectural guidelines:

- Deploy flex-algo (prefix-based slicing): When you need to enforce absolute isolation for entire enclaves (such as isolating a classified VRF to encrypted links). If a secure VRF is bound to the Flex-Algo 128 locator, all traffic originating from that VRF is mathematically constrained to the secure slice. No deep packet inspection is required, ensuring maximum edge throughput and security.
- Deploy PFP/PDP (flow-based steering): When you must differentiate traffic types destined for the exact same endpoint (such as ensuring Webex media takes a low-latency path while bulk HTTP traffic takes a best-effort path to the same remote branch). In this model, the parent Per-Flow Policy acts as an edge dispatcher, inspecting packet markings to steer individual flows into distinct Per-Destination child policies.

Tech tip: Order of operations, if both SRv6 Traffic Engineering and a flex-algo network slice is carved impacting packet flow: Cisco IOS XE evaluates flow-based SR-TE policies (such as a Per-Flow Policy matching voice) first at the ingress edge, steering classified applications into dedicated tunnels and completely overriding the standard RIB. If no active SR-TE policy matches, the routing engine falls back to destination-prefix lookups in the routing table, where flexible algorithm (flex-algo) topology constraints are evaluated second.

Verification of sliced environments

Specifically, administrators can run the `show segment-routing srv6 locator` command to confirm that the custom SECURE-128 and TRUSTED-130 locators are active and mapped to their respective algorithmic engines. The distributed topology calculations can be verified using the `show isis ipv6 flex 128` database audit, which details the algorithm's definition source, priority, and link-affinity metric parameters, alongside `show isis ipv6 topo flex 128` to inspect the mathematically resolved shortest-path graph. Finally, the service-plane mapping must be verified using the `show ip bgp vpnv4 vrf [vrf-name]` command, checking that tenant BGP VPNv4 routing updates are successfully resolved in the

forwarding table with an outer SRv6 Service SID (`srv6 out-sid`) originating strictly from the authorized Flex-Algo locator prefix block.

Mission-critical resiliency using TI-LFA

To build a modern, high-consequence network infrastructure, legacy Wide Area Network (WAN) designs—which rely on complex, stateful link-protection mechanisms like RSVP-TE Fast Reroute (FRR)—must be retired. Segment Routing over IPv6 (SRv6) provides the architectural replacement. By consolidating sub-50ms physical path resiliency, multi-topology flexible algorithm protection, and active performance telemetry directly into the link-state underlay, SRv6 drastically simplifies the recovery stack, making the core network entirely stateless and infinitely scalable. As mission-critical enclaves scale across heterogeneous transport mediums, physical link cuts and node failures must be resolved instantaneously at the hardware level without waiting for global routing convergence or relying on centralized SDN controller intervention.

This section outlines the high availability and resiliency architecture as implemented on Cisco IOS XE platforms—specifically targeting the Cisco 8000 Series Secure Routers platforms. It details the mathematical foundations of Topology Independent Loop-Free Alternate (TI-LFA) fast reroute, the strict order of operations during Flexible Algorithm integration, the active Performance Measurement (PM) telemetry engine for congestion mitigation, and the command-line configurations and diagnostics required to validate a self-healing fabric.

Sub-50ms fast reroute with TI-LFA

Protecting traffic against fiber cuts required complex RSVP-TE Fast Reroute (FRR) tunnels that crippled the core with state bloat. Standard fast-reroute protocols suffer from micro-loops and cannot guarantee complete protection coverage across complex topologies. Topology Independent Loop-Free Alternate (TI-LFA) modernizes this by providing a highly predictable, automated FRR mechanism that guarantees 100% sub-50ms protection coverage across any network topology.

Topology Independent Loop-Free Alternate (TI-LFA) modernizes link and node protection by providing a highly predictable, automated Fast Reroute mechanism natively within Segment Routing over IPv6 (SRv6). Instead of establishing complex signaled tunnels, TI-LFA leverages the stateless, source-routing nature of SRv6. The local router adjacent to the failure—known as the Point of Local Repair (PLR)—proactively pre-computes the optimal backup path for every single destination prefix in the routing table. Upon detecting a hard physical cut (typically triggered via Bidirectional Forwarding Detection, BFD, or loss of signal), the PLR instantly encapsulates in-flight traffic with an explicit SRv6 Segment List, steering packets around the failed link in under 50 milliseconds. Because the backup path is hardcoded into the outer IPv6 destination headers of the redirected packets, intermediate transit nodes forward them using standard longest-prefix-match (LPM) lookups, completely bypassing the micro-loops that plague standard fast-reroute protocols.

Mathematics of TI-LFA failure mitigation

To understand the mathematical superiority of TI-LFA over legacy Local LFA, we must examine the mechanics of network convergence and the structural causes of " micro-loops" . When a physical link fails, standard IGP routing protocols (such as IS-IS) require several seconds to flood Link State Packets (LSPs) across the network, run the Shortest Path First (SPF) algorithm, and update local Routing Information Base (RIB) and Forwarding Information Base (FIB) tables.

During this transient convergence window, routers possess inconsistent views of the topology. For example, in a topology where traffic flows from Node A to Node C via Node B, a physical failure on link B-C will cause Node B to immediately detect the link drop. If Node B simply deflects the packet to an alternate neighbor, Node E, a micro-loop is born. Since Node E has not yet converged, its local FIB still states that the optimal path to Node C is back through Node B. The packet bounces back and forth between Node B and Node E until the IGP converges, saturating the link and dropping critical real-time telemetry.

Figure 22. TI-LFA mechanics

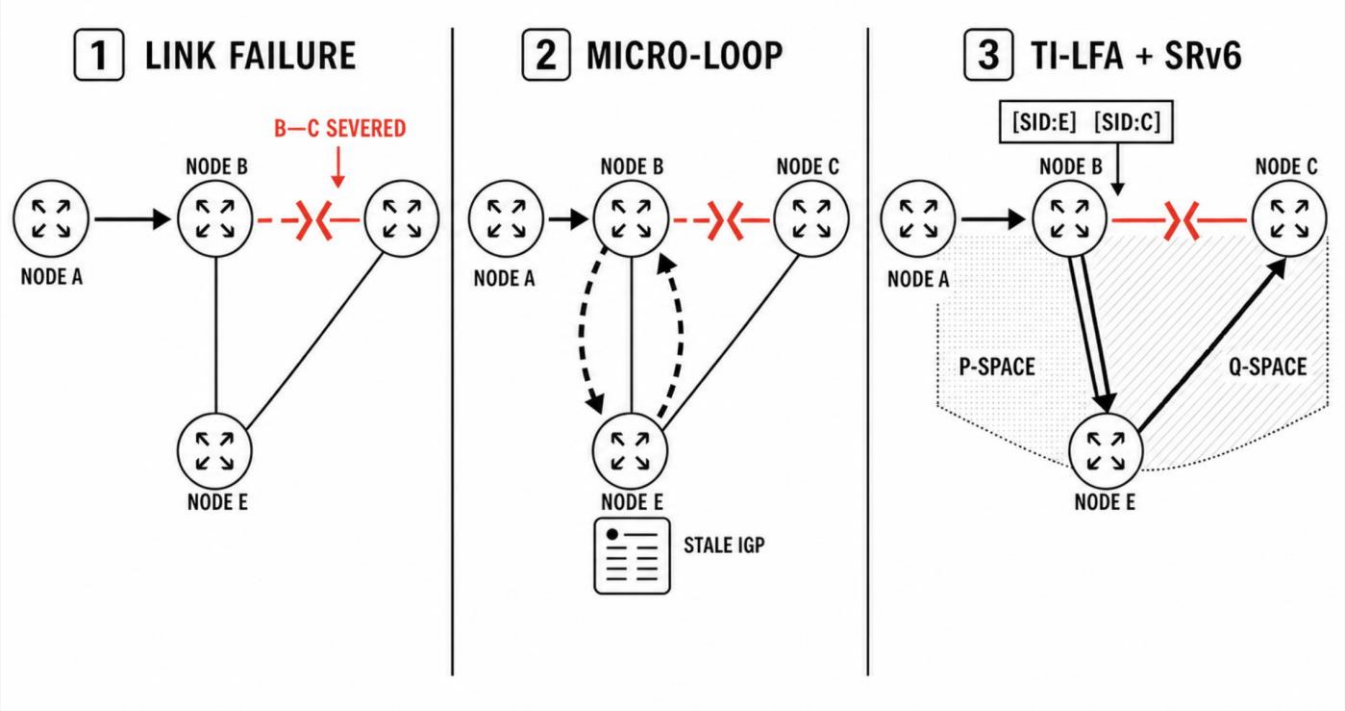

The TI-LFA Solution: TI-LFA solves this by using the link-state database of the IGP to pre-compute an explicit backup path for every destination prefix before a failure occurs. This backup path is mathematically split into two non-looping regions:

- **P-Space:** The set of nodes reachable from the calculating router without passing through the failed primary link.
- **Q-Space:** The set of nodes that can reach the destination without passing through the failed primary link.

Where P-Space and Q-Space do not naturally intersect, TI-LFA dynamically calculates the exact segment identifier or sequence of segments required to bridge the gap across the core network. This explicit backup instruction is programmed directly into the router's hardware forwarding tables.

Back to our topology, when the link breaks, Node B encapsulates the packet with an SRv6 Extension Header containing the exact SIDs for Node E and Node C. Node B forces the packet to Node E. Because of the strict SRv6 instructions, Node E ignores its own outdated routing table and forwards the packet strictly to Node C.

Flex-algo and SLA constraint-aware resiliency

In mission-critical defense networks, a primary architectural risk is cryptographic data leakage during failover events. If a defense agency routes highly classified data inside a dedicated "Secure Slice" (e.g., Flex-Algo 128) constrained strictly to MACsec-encrypted physical fiber, a physical link cut must not force that traffic to failover onto an unencrypted public transit link.

Unlike legacy fast-reroute protocols which calculate backup paths blindly based on default physical shortest-path metrics, TI-LFA is completely Flexible Algorithm and Constraint Aware. When a link fails inside a customized logical slice, the PLR recalculates the post-convergence backup route strictly using the link metrics, administrative groups (affinities), and resource exclusion rules defined by that specific Flexible Algorithm.

This strict order of operations is executed natively in the hardware ASIC at the moment of failure:

Figure 23. Order of operations executed in ASIC at failure

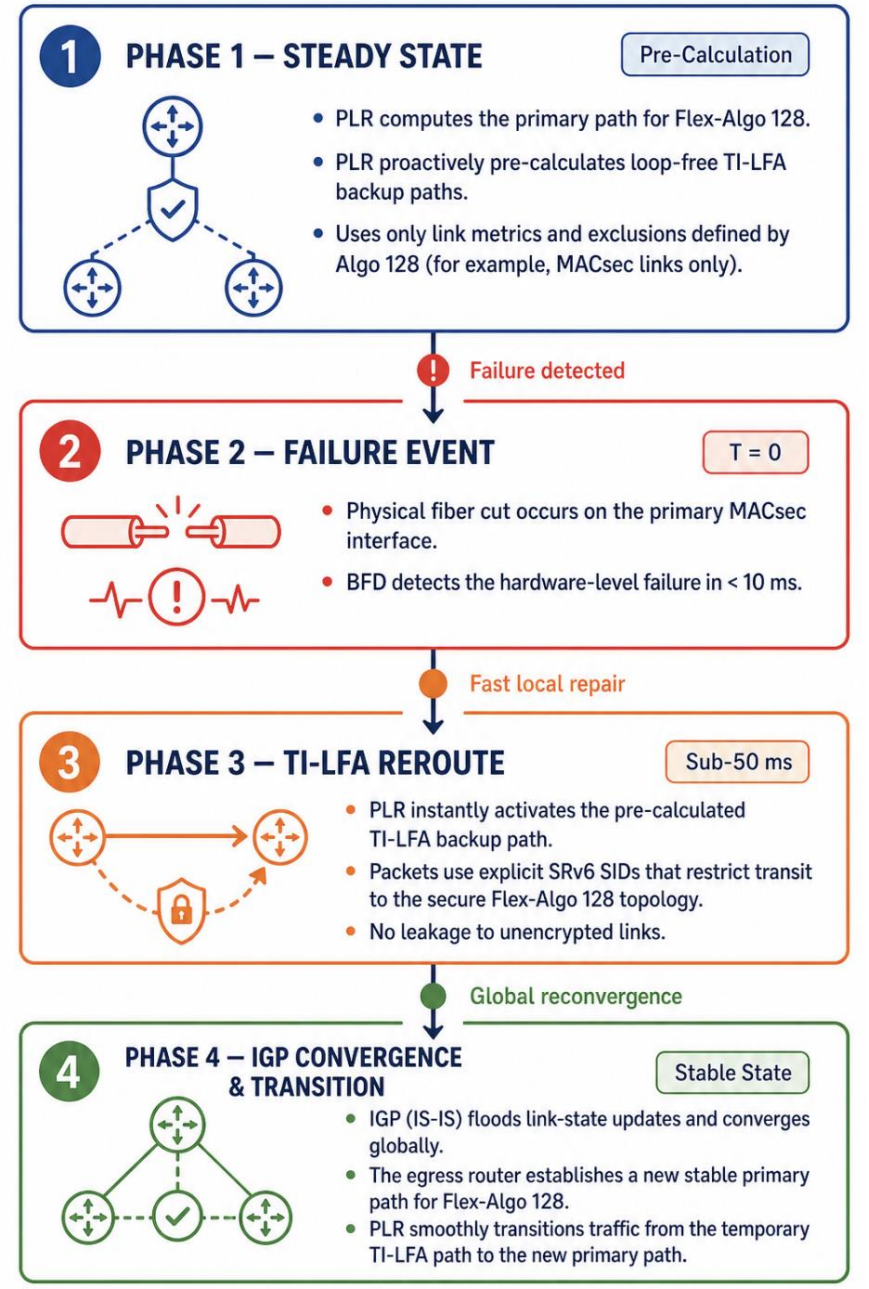

Advanced resiliency tiebreakers

When multiple repair paths exist around a failed link, Cisco IOS XE uses a highly optimized set of mathematical heuristics—known as Fast Reroute Tiebreakers—to prune the candidate list and select the single most resilient path. These tiebreakers can be configured globally under the routing protocol address-family or applied directly to specific physical interfaces to override default.

Cisco IOS XE supports the following core tiebreakers for SRv6 TI-LFA:

- Lowest-backup-path-metric (enabled by default):
  Prefers the repair path that minimizes the total accumulated link-state metric to the destination, ensuring that the backup path closely mirrors the future post-convergence shortest path.

- Linecard-disjoint (enabled by default):
  Instructs the router to select a backup path that exits the chassis via a completely different physical linecard than the failed primary port. This protects against total linecard ASIC failures or power disruptions.
- Node-protecting (user-prioritized):
  Computes a backup path that completely bypasses the downstream neighbor router (the next-hop node) rather than just the adjacent link. If the primary next-hop router crashes entirely, the PLR successfully steers traffic around the dead node directly to the subsequent downstream hop.
- SRLG-disjoint (user-prioritized):
  Shared Risk Link Groups (SRLGs) identify links that share common physical vulnerabilities, such as sharing the same fiber conduit, utility pole, or geographic trench. When enabled, this tiebreaker ensures that the selected backup path uses interfaces that do not share any SRLG IDs with the primary failed path, protecting against multi-link cuts caused by a single physical impact.

If both node and SRLG protection modes are configured, Cisco IOS XE executes a combined Node-SRLG protection algorithm. This removes the failed next-hop node and all physical interfaces sharing the same SRLG group from the link-state topology database before calculating the post-convergence shortest path tree (SPT).

Architectural decoupling: underlay resiliency vs. overlay SLAs

A critical design rule for high-consequence network engineering is the absolute control-plane decoupling of physical infrastructure protection from application-level service steering. In the Cisco IOS XE architecture, TI-LFA and SR-TE Performance Measurement (PM) operate at completely different layers of the routing stack and solve distinct operational failures:

Table 3. Architectural decoupling: underlay resiliency vs. overlay SLAs

Operational dimension	Fabric resiliency (TI-LFA)	SLA enforcement (SR-TE delay metric)
Control plane layer	Native IGP layer (IS-IS CORE)	Policy Layer (Segment Routing Traffic Engineering)
Target failure type	Hard failure: physical fiber cuts, link-down events, next-hop node crashes	Soft failure: circuit congestion, latency spikes, " noisy neighbor" packet queuing
Trigger mechanism	Hardware loss-of-signal, Bidirectional Forwarding Detection (BFD)	Two-way active delay measurement probes
Failure Detection Speed	Instantaneous: < 10 milliseconds	Telemetry interval: probe cycles at 3-second intervals; path re-optimization within 1 second
Path computation	Autonomous local computation (PLR) using pre-calculated backup SIDs	Ingress head-end CSPF calculations based on flooded link attributes
Policy scale impact	Zero manual configuration; calculated natively in the background by IS-IS	Dynamic On-Demand Next-Hop (ODN) instantiation triggered by BGP Color communities

This separation ensures that under extreme stress, TI-LFA acts as the rapid-failover shield that keeps the physical transport alive. It automatically repairs the physical underlay path before the overlay SR-TE policy engine even registers that a physical link went down. While TI-LFA guarantees that the network

infrastructure survives, SR-TE ensures that the applications and data traffic traversing that infrastructure continue to meet their strict constraints.

Base configurations and verification commands

Base configuration blueprint

To successfully deploy TI-LFA within a validated Segment Routing over IPv6 (SRv6) framework, fast-reroute parameters must be explicitly configured under the IPv6 address family of the IS-IS routing process. The following configuration blueprint represents the baseline applied globally to the ISIS routing domain:

```
router isis 1
 net 49.0099.0099.0000.1041.00
 is-type level-2-only
 router-id Loopback0
 metric-style wide
 distribute link-state instance-id 99
 !
 address-family ipv6
  multi-topology
  router-id Loopback0
  segment-routing srv6
   locator SLOC
  ! Enable local loop-free alternate calculation
  fast-reroute per-prefix level-2 all
  ! Enable Topology-Independent LFA (TI-LFA) fast-reroute
  fast-reroute ti-lfa level-2
  ! Configure advanced tiebreakers (Linecard and Node protection prioritized)
  fast-reroute tie-breaker node-protecting index 100
  fast-reroute tie-breaker linecard-disjoint index 200
 exit-address-family
!
```

Verification details

To verify and validate local underlay resiliency, network operators can utilize a targeted suite of Cisco IOS XE CLI diagnostics: auditing global fast-reroute metrics with `show isis ipv6 fast-reroute summary`, verifying active backup tunnel paths with `show isis fast-reroute ti-lfa tunnel`, and auditing CEF hardware-routing tables with `show ipv6 cef`. Furthermore, executing `show ipv6 cef [prefix] internal` allows engineers to trace the output chain and confirm that the hardware ASIC is dynamically pre-programmed with the remote repair Segment List (such as the pre-calculated P-node locator FCBB:DEAD:6::), mathematically proving that the router is armed to execute a sub-50ms link-failover the instant a physical link failure is registered.

SRv6 over secure transport

As mission-critical networks migrate to next-generation Segment Routing over IPv6 (SRv6) architectures, achieving absolute zero-trust data-plane confidentiality across untrusted transits is an uncompromising mandate. Traditionally, WAN encryption was treated as a static, hop-by-hop overlay that operated completely divorced from underlay routing semantics. In a modern, enterprise or defense fabric, security must be unified natively with the routing plane. Standard cryptographic protocols, however, are facing an existential threat from the rapid development of quantum computing.

While the SRv6 routing core provides stateless, transport virtualization, it contains no native cryptographic protection. This section outlines the design rules, cryptographic frameworks, and Cisco IOS XE configurations required to run a high-performance SRv6 fabric over Post-Quantum Cryptography (PQC) ready transport layers. By combining port-level WAN MACsec for private, high-speed fiber transits with post-quantum protected IPsec GRE tunnels for shared, untrusted middle-mile transits, defense and national security enclaves can secure today's communications against tomorrow's computational breakthroughs.

Post-quantum imperative and HNDL threat

Earlier WAN security architectures relied on classical public-key cryptography—specifically RSA, Diffie-Hellman (DH), and Elliptic Curve Diffie-Hellman (ECDH)—to perform identity verification and securely negotiate session keys during the IPsec and MACsec handshake phases. While these mathematical algorithms are computationally secure against today's most powerful classical supercomputers, they are fundamentally vulnerable to a Cryptographically Relevant Quantum Computer (CRQC) running Shor's algorithm.

To evaluate this vulnerability, network security parameters must be divided into two distinct components:

- The lock (asymmetric key exchange): The public-key exchange used to verify identity and securely negotiate symmetric encryption keys. This component relies on the mathematical difficulty of integer factorization and discrete logarithms, which Shor's algorithm can solve in minutes, rendering classical handshakes entirely vulnerable.
- The contents (symmetric data payload): The actual user packets encrypted using a symmetric cipher. Symmetric encryption utilizing AES-256-GCM is considered mathematically safe and quantum-resistant, as the only known quantum threat is Grover's algorithm, which merely reduces the effective key strength from 256 bits to 128 bits—still well beyond the reach of any foreseeable computational attack.

Adversaries are actively exploiting this asymmetric vulnerability today through "Harvest Now, Decrypt Later" (HNDL) attacks. Hostile actors intercept and archive petabytes of encrypted high-value enterprise and defense data traversing public transits. While they cannot decrypt the AES-GCM-256 payloads today, they store the captured handshakes and packets with the explicit intent to run them through a CRQC once quantum hardware matures.

If data transmitted across a WAN has an operational shelf-life of 10 or more years, it is already compromised if sent over classical IPsec VPNs today. To defeat the HNDL threat, organizations must immediately adopt Post-Quantum Cryptography (PQC).

Native post-quantum algorithms

To defend against Cryptographically Relevant Quantum Computers (CRQCs) and align with the National Security Agency's Commercial National Security Algorithm Suite 2.0 (CNSA 2.0) mandates (which require full PQC compliance by 2030), Cisco IOS XE implements the standardized post-quantum cryptographic primitives finalized by NIST.

1. ML-KEM (Module-Lattice-Based Key-Encapsulation Mechanism - FIPS 203)

Formally known as CRYSTALS-Kyber, ML-KEM is the foundation for establishing shared secret keys over unsecure channels. In Cisco IOS XE 26.1, ML-KEM-1024 (the highest security category, equivalent to AES-256 strength) is utilized across three primary enclaves:

- **IPsec/IKEv2 (RFC 9370):** Negotiates VPN session keys using a Post-Quantum Hybrid Key Exchange. This hybrid model combines a classical Diffie-Hellman exchange (e.g., DH Group 19 or 21) with an ML-KEM-1024 exchange in a single handshake. A security association is established only if both exchanges succeed, ensuring immediate protection against both classical and future quantum attacks while mitigating early-adopter software risks.
- **WAN MACsec:** Secures high-speed Layer 2 point-to-point links using EAP-TLS with ML-KEM via TLS 1.3.
- **Management Plane (PQ SSH):** Protects administrative routing access by enforcing ML-KEM key exchange algorithms natively inside the SSH server and client daemons.

2. ML-DSA (Module-Lattice-Based Digital Signature Algorithm - FIPS 204)

Formally known as CRYSTALS-Dilithium, ML-DSA replaces classical RSA and ECDSA signatures to protect identity and authentication. It will be utilized natively in Cisco IOS XE 26.2 to verify the identity of the router itself during tunnel authentication and secure-boot microcode verification, preventing attackers from forging identity to hijack the routing control plane.

3. Transitional protection: Post-Quantum Pre-Shared Keys (PPK)

For legacy, multi-vendor, or transitional enclaves where native ML-KEM is not yet supported, Cisco platforms support the Post-Quantum Pre-Shared Key (PPK) method (complying with RFC 8784).

- **The PPK mechanism:** A high-entropy, quantum-safe symmetric key is configured out-of-band and mixed directly into IKEv2's session key derivation. Even if an adversary intercepts the key exchange and subsequently cracks the classical ECDH handshake using a quantum computer, the session key remains completely secure due to the high-entropy PPK.
- **Cisco SKIP (Secure Key Integration Protocol):** To eliminate the massive operational overhead and security risks of manually configuring static PPKs across hundreds of branches, Cisco developed SKIP. SKIP allows edge routers to dynamically request and fetch high-entropy, quantum-safe keys from an independent, trusted third-party Key Provider or Quantum Key Distribution (QKD) system on the fly. This maintains complete cryptographic agility at scale.

SRv6 over private WAN: WAN MACsec integration

High-speed layer 2 WAN MACsec and carrier CLNS drops

For private dark-fiber paths or provider-managed Ethernet Private Line (E-Line) services, WAN MACsec (IEEE 802.1AE) provides standards-based, Layer 2 encryption. Unlike Layer 3 IPsec, which incurs packet-encapsulation overhead, MACsec encrypts the entire SRv6 IPv6 frame directly on the physical interface in hardware, maximizing throughput while minimizing latency.

In this design, point-to-point MACsec sessions are established directly between the edge routers and the datacenter hubs. Device authenticity is verified using certificate-based mutual authentication via EAP-TLS (utilizing TLS 1.3 with ML-KEM key exchanges), enabling automated, quantum-safe MACsec key derivation.

Figure 24. High-speed layer 2 WAN MACsec and carrier CLNS drops

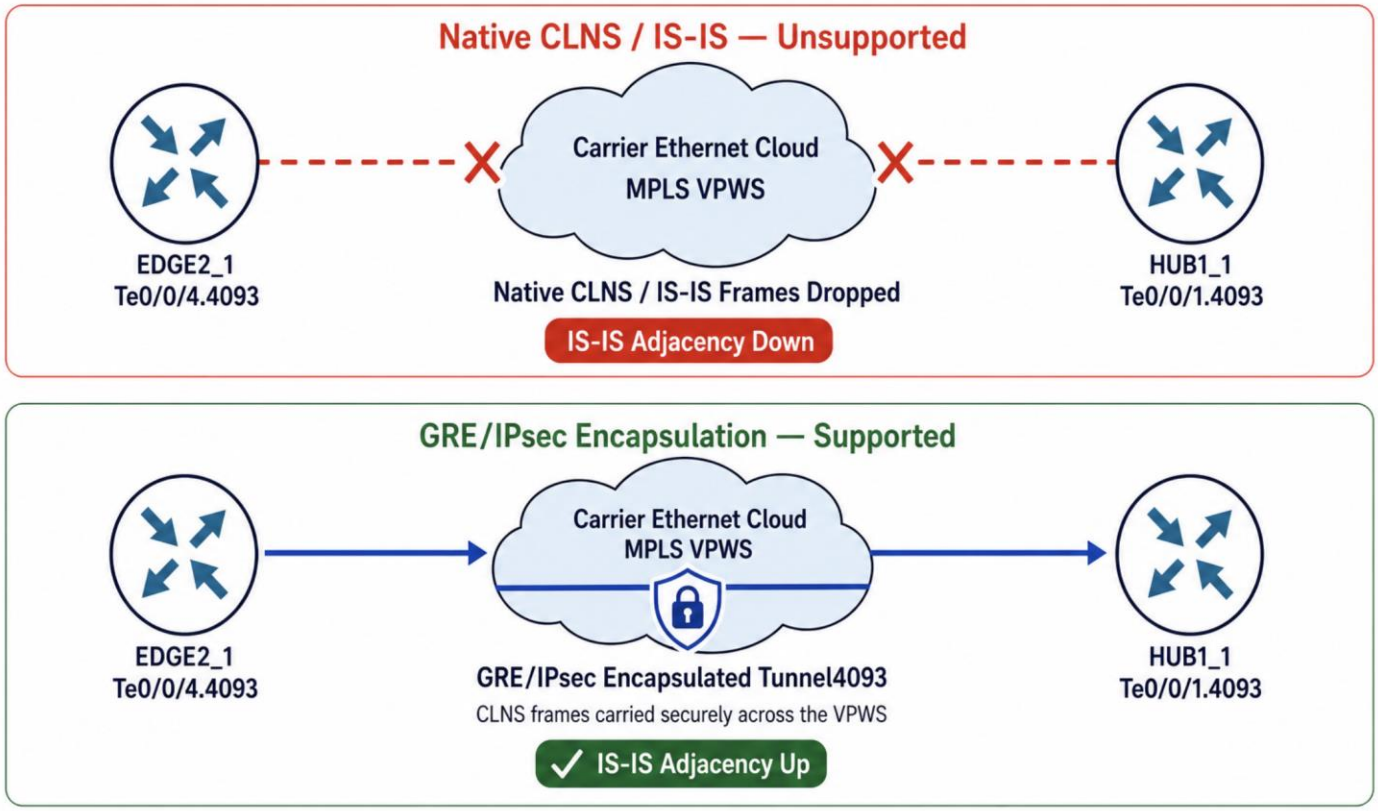

Carrier CLNS drop constraint and GRE fallback

While native IS-IS routing directly over MACsec is ideal for dark-fiber connections, validation testing reveals a critical constraint when traversing commercial Carrier Ethernet services (such as MPLS VPWS). Many provider clouds do not transparently forward Layer 2 Connectionless Network Service (CLNS) frames, silently dropping the IS-IS Hello packets, and preventing routing adjacencies from forming.

To overcome this carrier drop constraint, the design implements GRE over WAN MACsec:

- **Physical encryption:** Port-level WAN MACsec secures the raw, point-to-point physical interface or dot1q sub-interface.
- **Logical encapsulation:** A point-to-point IPv6 GRE tunnel is built over the MACsec-protected link.
- **Routing under MACsec:** The IS-IS routing process is enabled directly on the GRE tunnel interface. The GRE header safely encapsulates the CLNS frames as unicast IPv6 packets, allowing routing updates and locator reachability to pass flawlessly through the provider cloud while remaining fully encrypted on the wire.

Shared WAN: GRE over IPsec (post-quantum)

When routing SRv6 traffic over shared, untrusted public transits (such as commercial ISPs or LTE/SATCOM links), the architecture leverages IPv6 GRE tunnels protected by IPsec. To defend these public paths against "Harvest Now, Decrypt Later" (HNDL) retrospective decryption, the underlay IPsec tunnel is secured with native ML-KEM.

- **The cryptographic negotiation:** The edge routers negotiate the IKEv2 security association utilizing an IKEv2 proposal configured with `pqc mlkem1024` optional or required. This triggers the RFC 9370 hybrid key exchange, establishing the IKEv2 SA with a quantum-safe symmetric key.

- **The protected tunnel:** An IPv6 GRE tunnel is bound directly to the IPsec profile (tunnel protection ipsec profile). The tunnel interface is assigned a native IPv6 address and placed directly into the unified IS-IS process.
- **Data plane forwarding flow:**
    - Standard VRF tenant traffic enters the ingress router.
    - The PE performs a BGP VPN lookup, mapping the VRF destination to the egress PE's aggregate SLOC locator.
    - The ingress router encapsulates the packet inside an outer IPv6 header (using the remote Service SID as the Destination Address).
    - The packet is routed into the GRE tunnel interface, which appends a GRE header.
    - The IPsec engine heavily encrypts the GRE frame using AES-GCM-256, protecting it from retrospective decryption before transmitting it over the untrusted network.

Secure the management plane: post-quantum SSH

While MACsec and IPsec secure the active data plane and routing telemetry, architects must ensure that the Management Plane remains equally impenetrable to quantum cryptanalysis. Standard Secure Shell (SSH) sessions rely on classical Key Exchange (KEX) algorithms (such as Diffie-Hellman Group 14 or ECDH Curve 25519) to securely negotiate session keys.

If an adversary intercepts SSH administrative sessions today, a future quantum computer running Shor's algorithm can easily decrypt the session. This exposes plaintext administrator credentials, configuration templates, and internal cryptographic parameters, resulting in complete network compromise.

To neutralize this threat, Cisco IOS XE implements Post-Quantum ML-KEM algorithms directly inside the SSH Key Exchange (KEX) engine. This allows security administrators to globally restrict SSH negotiations to quantum-safe algorithms, ensuring that both incoming connections (the SSH Server) and outgoing hops (the SSH Client) are mathematically immune to quantum decryption:

```
ip ssh server algorithm kex mlkem1024nistp384-sha384 mlkem768nistp256-sha256
ip ssh client algorithm kex mlkem1024nistp384-sha384 mlkem768nistp256-sha256

show ip ssh
SSH Enabled - version 2.0
Authentication timeout: 120 secs; Authentication retries: 3
Minimum expected Diffie Hellman key size : 2048 bits
IOS Keys generated:
  SSH-RSA-2048 : 2048 bits
KEX Algorithms:mlkem1024nistp384-sha384,mlkem768nistp256-sha256,mlkem768x25519-sha256
```

MTU, TCP MSS, and IKEv2 fragmentation management

Running stateless SRv6 overlays inside secure underlay transports introduces multiple layers of encapsulation overhead. Each packet must accommodate the original customer payload, the 128-bit SRv6 uSID carrier IPv6 header, the GRE tunnel header, and the IPsec ESP tunnel or transport headers.

Explicit MTU configurations

To prevent performance-degrading IP-layer fragmentation and packet drops, operators must apply strict, explicit MTU boundaries on all secure interfaces:

- **Unsecured GRE/IPsec tunnels:** Must be tuned to ip mtu 1400 and ipv6 mtu 1400.
- **WAN MACsec GRE tunnels:** Because Layer 2 MACsec introduces its own 802.1AE header overhead, these GRE tunnels must be dropped to a restrictive ip mtu 1300, ipv6 mtu 1300, and clns mtu 1300.
- **IS-IS routing protocol LSP MTU:** The routing process must have its Link State Packet (LSP) size explicitly matched using lsp-mtu 1300 to prevent large routing updates from fragmenting and flapping the IGP adjacencies.
- **TCP MSS clamping:** Restricts the maximum segment size of TCP packets by applying ip tcp adjust-mss 1360 on all LAN-facing interfaces to prevent user data from fragmenting inside the WAN IPsec tunnels.

Handshake failures and IKEv2 fragmentation with ML-KEM-1024

A major operational hurdle when transitioning to post-quantum cryptography is IKEv2 Handshake Failure due to packet fragmentation.

Classical elliptic curve keys are tiny (e.g., ECDH Curve25519 is only 32 bytes). In contrast, lattice-based public keys are massive; the ML-KEM-1024 public key is 1,568 bytes, and the ciphertext is 1,568 bytes.

When negotiating a quantum-safe IPsec SA, the carrying IKEv2 messages (such as IKE_SA_INIT or IKE_INTERMEDIATE carrying the Key Encapsulation Method payload) easily exceed the standard 1,500-byte IP MTU limit.

The fragmentation bottleneck

Without proper handling, the router is forced to fragment these oversized IKEv2 packets at the IP Layer. However, most enterprise firewalls, transit provider ACLs, and shared WAN gateways block IP fragments by default as a security best practice to prevent fragmentation-based Denial of Service (DoS) attacks. When these fragments are dropped, the IKEv2 negotiation hangs indefinitely, and the IPsec tunnel fails to establish.

The RFC 7383 solution: IKEv2 fragmentation

To prevent these drops, operators MUST configure native IKEv2 Fragmentation (RFC 7383) globally on all secure edge routers:

```
crypto ikev2 fragmentation mtu 1400
```

Enabling this command instructs the IKEv2 process to perform fragmentation at the Application layer rather than the IP layer. The router splits the massive ML-KEM key exchange payload into multiple distinct, fully formed IKEv2 messages, each encased in its own independent UDP/IP packet structure that easily slides under the 1,400-byte threshold. Because these packets are standard, unfragmented IP packets, firewalls transit them without drops.

During active validation on Cisco 8000 Series Secure Routers, executing show `crypto ikev2 sa` detailed verifies that application-layer fragmentation is active and executing successfully:

```
Router# show crypto ikev2 sa detailed
  Tunnel-id Local            Remote           Status
  1         172.16.11.2/500  172.16.22.2/500  READY
  ...
  PQC Key Exchange: ML-KEM-1024
  Quantum-safe Encryption using PQC: ML-KEM-1024
  IETF Std Fragmentation MTU in use: 1372 bytes
```

This output confirms that the security engine has successfully negotiated the ML-KEM-1024 post-quantum key exchange, using a standardized IETF fragmentation MTU of 1,372 bytes to prevent packet drops across the WAN.

Step-by-step configuration templates

Below is a configuration for a Cisco 8000 Series Secure Router, combining SSH hardening, a quantum-safe IPsec GRE tunnel, WAN MACsec sub-interfaces, and underlay IS-IS SRv6:

```
! 1. Hardening SSH with NIST-Approved lattice-based KEX
ip ssh server algorithm kex mlkem1024nistp384-sha384 mlkem768nistp256-sha256
ip ssh client algorithm kex mlkem1024nistp384-sha384 mlkem768nistp256-sha256

! 2. Define the Quantum-Safe IKEv2 Proposal (RFC 9370 Multiple Key Exchanges)
crypto ikev2 proposal U-IKEV2-PROP
 encryption aes-gcm-256
 prf sha256
 group 19
 pqc mlkem1024 optional
!
crypto ikev2 policy U-IKEV2-POLICY
 proposal U-IKEV2-PROP
!
crypto ikev2 fragmentation mtu 1400

! 3. Configure the IPsec Profile executing transport-mode encapsulation
crypto ipsec transform-set U-TS esp-gcm 256
 mode transport
!
crypto ipsec profile U-H12-MLKEM-IPSEC
 set transform-set U-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile U-IKEV2-PROFILE

! 4. Configure Layer 2 WAN MACsec on the tagged provider subinterface
key chain KC-WAN-MACSEC macsec
 key 01
  cryptographic-algorithm aes-256-cmac
  key-string 3637fc51ed2f71b5a354e1aca9a91adb0f5a6744f1cda2cc7e18417cbd025e51
  lifetime local 00:00:00 Jan 1 2026 infinite
!
mka policy POL-WAN-MACSEC
 key-server priority 10
 macsec-cipher-suite gcm-aes-256
!
interface TenGigabitEthernet0/0/4.4093
 description UC4_WAN_MACSEC_PILOT_TO_HUB1_1_VLAN4093
 encapsulation dot1Q 4093
 ip address 172.31.93.1 255.255.255.252
```

```
 ip mtu 1468
 ipv6 address 2001:DB8:4093:1::1/64
 eapol destination-address d862.caba.a401
 eapol eth-type 876F
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec

! 5. Configure the Point-to-Point GRE Tunnel running IS-IS
interface Tunnel4093
 description UC4_GRE_ISIS_RECHECK_OVER_WAN_MACSEC
 no ip address
 ip mtu 1300
 ipv6 address 2001:DB8:5093:1::1/64
 ipv6 mtu 1300
 ipv6 router isis SRTE-OVLY
 tunnel source 172.31.93.1
 tunnel destination 172.31.93.2
 clns mtu 1300
 isis circuit-type level-2-only
 isis network point-to-point
 isis affinity flex-algo
  name SECURED

! 6. Configure the IS-IS Underlay process for stateless uSID Locator reachability
router isis SRTE-OVLY
 net 49.0099.0099.0000.1041.00
 is-type level-2-only
 metric-style wide
 lsp-mtu 1300
 distribute link-state instance-id 99
 affinity-map SECURED bit-position 0
 affinity-map UNSECURED bit-position 1
 !
 address-family ipv6
  router-id Loopback0
  segment-routing srv6
   locator SLOC
  fast-reroute per-prefix level-2 all
  fast-reroute ti-lfa level-2
```

Phased post-quantum migration framework

Transitioning hundreds of production sites to a post-quantum posture can be executed as a phased migration to ensure zero downtime and prevent configuration mismatches.

```
Phase 1: Readiness Audit  ===> Phase 2: In-Place Hybrid  ===> Phase 3: Enforced PQC
```

```
 * Identify Legacy Assets       * Deploy Cisco Secure Router * Set 'pqc required'
 * Procure Cisco Secure Routers * Set 'pqc optional'         * Decommission Catalyst routers
 * Audit with Cisco IQ          * Coexist Legacy/PQC         * Validate performance
```

For more details on Quantum migration methods, refer to the Quantum-Ready Migration Guide.

SRv6 multicast design

In enterprise and defense networks, multicast delivery is an operational prerequisite for real-time video surveillance, sensor telemetry distribution, tactical command-and-control, and rapid one-to-many software image deployment. However, transporting multicast streams across a Wide Area Network (WAN) has historically been an architectural compromise. Traditional multicast models force the core transit routers to run stateful protocol stacks—such as Protocol Independent Multicast Sparse Mode (PIM-SM) or Multipoint LDP (mLDP)—to build and maintain active multicast distribution trees. This stateful coupling violates the primary design goal of an SRv6 fabric: keeping the transit underlay entirely stateless, quiet, and unburdened by active tunnel signaling.

This section deconstructs the validated design for native, stateless multicast delivery over an IPv6 SRv6 underlay. By combining BGP Multicast VPN (BGP MVPN) control-plane signaling with Segment Routing Ingress Replication (IR), the network limits multicast state strictly to the VRF-aware edge routers. The transit P-routers remain entirely oblivious to multicast groups, forwarding replication packets as standard, unicast IPv6 longest-prefix-match (LPM) lookups.

Ingress replication architecture vs. stateful core multicast

Traditional service provider networks require the core (P) routers to participate directly in multicast path signaling. In a standard PIM-SM or mLDP deployment, every transit node along the path must maintain active state tables—tracking (*, G) and (S, G) pairs—and dynamically build physical branch trees. If a link flaps or a receiver leaves, core routers must execute complex tree pruning and rebuilding operations, creating control-plane signaling storms that saturate low-bandwidth or degraded WAN links.

SRv6 Multicast Ingress Replication (IR) completely decouples the transit underlay from multicast awareness. Under this model, the underlay core runs standard unicast IPv6 routing. When a source transmits a multicast frame into an VRF-enabled LAN interface, the ingress PE performs the replication on the local forwarding ASIC. For each remote egress PE that has actively signaled receiver interest for that group, the ingress PE creates a dedicated, unicast-encapsulated IPv6 packet. The destination address of each replica is set to the specific egress PE's local multicast service SID (End.DTMC4).

The operational trade-offs

- **Core simplification:** Core transit routers are entirely offloaded from multicast routing tables, PIM join/prune tracking, and mLDP state maintenance. They route each replica based strictly on standard, line-rate longest-prefix-match lookups against their global IPv6 tables.
- **Edge replication load:** The replication burden is shifted entirely to the ingress PE [86]. Replicating a high-bandwidth video stream N times across the first common WAN link segment consumes N times the unicast bandwidth.
- **Capacity planning:** Because of this "N-copy" characteristic, architects must strictly validate the maximum number of egress branch leaves (N), the aggregate multicast stream bitrate, and the physical packet-replication scale of the ingress PE’s forwarding ASIC before production rollout.

BGP MVPN overlay signaling and leaf membership

To prevent the ingress PE from blindly flooding copies of every multicast stream to every site in the WAN, the design utilizes Partitioned Ingress Replication. Under this architecture, the overlay control plane is entirely receiver-driven, ensuring that packets are only replicated to egress PEs that host active, interested receivers.

Figure 25. BGP MVPN overlay signaling and leaf membership

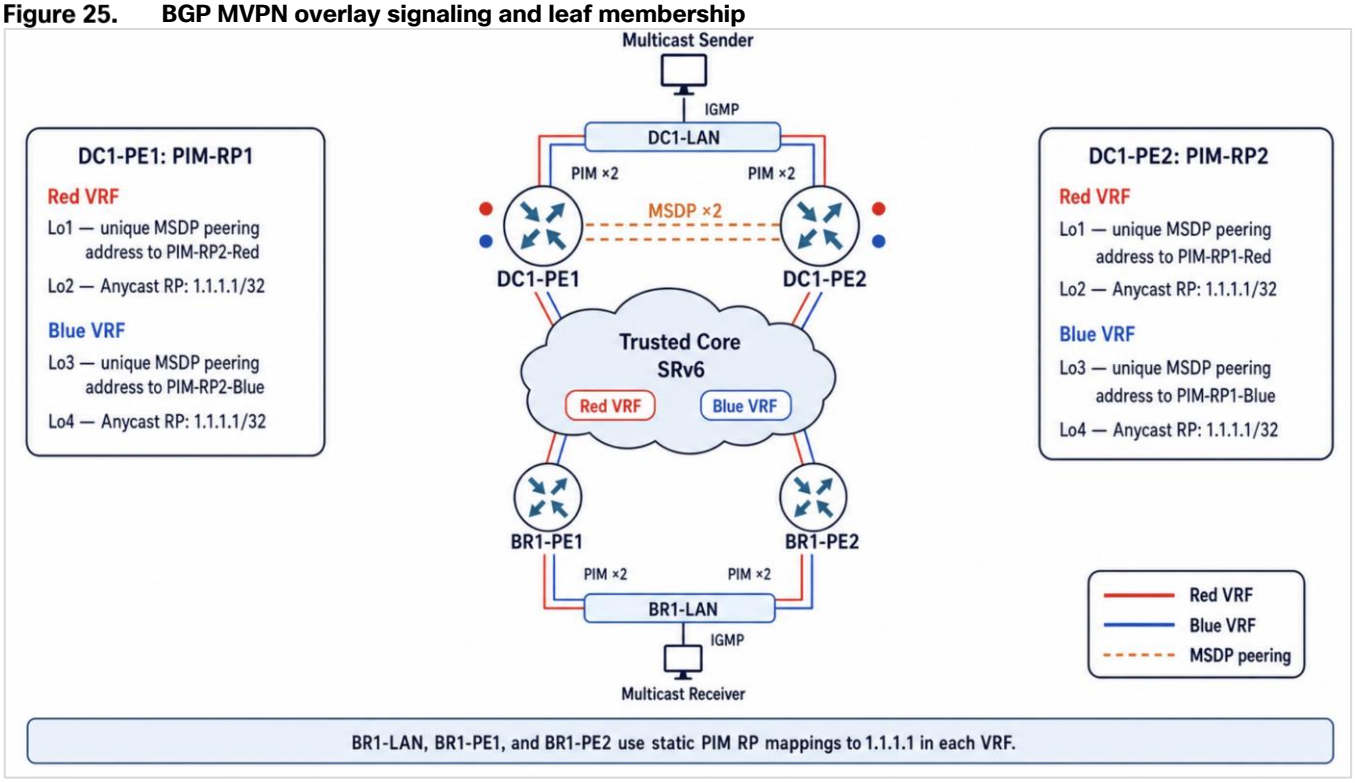

In the figure above, every PIM router in RED and BLUE is configured with the appropriate group range and the shared 1.1.1.1 RP address. The address is hosted concurrently by DC1-PE1 and DC1-PE2 in each VRF, so unicast routing directs PIM Registers and Joins to the closest reachable datacenter-hub PE. Auto-RP candidate, mapping-agent, and listener commands are not part of this baseline.

Functional Roles

- **DC1-PE1 and DC1-PE2:** Act as ingress PEs for a datacenter source, as egress PEs for datacenter receivers, and as the redundant Anycast RP pair for RED and BLUE. Each router owns the shared RP /32 and a unique per-VRF MSDP /32.
- **BR1-PE1 and BR1-PE2:** Act as egress PEs when branch receivers join a group, and as ingress PEs if a branch source transmits. They learn the RP mapping and advertise receiver interest through BGP MVPN signaling.
- **SRv6 core:** Provides IPv6 reachability to PE locators and service SIDs. It forwards unicast SRv6 replicas and does not maintain PIM (*,G) or (S,G) state.
- **Anycast RP and MSDP:** The common RP address directs PIM Registers and Joins to the closest live RP. The unique MSDP addresses form the peer relationship that shares Source-Active information between DC1-PE1 and DC1-PE2 within each VRF.

The control-plane workflow operates in five distinct phases:

- **Receiver joint discovery (phase 1):** A host on the branch LAN transmits an IGMPv2/v3 membership report (e.g., for Group 239.10.10.10 in VRF RED). The branch PE (e.g., BR1-PE1) processes this report, builds local receiver state, and sends a standard PIM Join toward the VRF's Rendezvous Point (RP).
- **MVPN leaf signaling (phase 2):** The branch PE translates the customer's multicast interest into a BGP Multicast VPN (BGP MVPN) Route. It advertises a Type 4 Leaf Auto-Discovery Route across the MP-BGP overlay to the Route Reflector. This route carries:
    - The target Multicast Source and Group (S,G) or (*,G).

- The BGP Route Target (RT) to ensure correct VRF importing [92].
- The branch PE's dynamically allocated End.DTMC4 Service SID inside the BGP Prefix-SID attribute [90, 91].

- **Source registration (phase 3):** At the Datacenter, the active source begins transmitting its multicast stream. The source-facing Ingress PE intercepts the stream, and its local PIM Designated Router (DR) encapsulates the payload inside a PIM Register message, sending it directly to the nearest reachable instance of the Anycast RP.
- **Anycast RP source synchronization (phase 4):** Since the Anycast RP address is concurrently hosted by both DC1-PE1 and DC1-PE2, the PIM Register may land on DC1-PE1 while the branch PIM Join landed on DC1-PE2. To bridge this gap, an active Multicast Source Discovery Protocol (MSDP) peering session running over unique, per-VRF loopback interfaces synchronizes Source-Active (SA) information between the RPs.
- **SRv6 data plane replication (phase 5):** The ingress PE ingests the BGP MVPN Type 4 Leaf routes. It parses the leaf list, identifying only those egress PEs with registered receivers. It encapsulates the incoming payload into N distinct IPv6 envelopes, sets the destination IPv6 address of each packet to the respective branch's local End.DTMC4 SID, and forwards them into the stateless core.

The End.DTMC4 service SID

End.DTMC4 is the SRv6 endpoint behavior that binds an arriving SRv6 packet to an IPv4 multicast table. It is the multicast counterpart to the unicast VPN decapsulation behaviors described earlier in the guide. The critical distinction is that the final lookup is a multicast lookup in the service VRF, not an IPv4 unicast lookup.

The operator enables SRv6 multicast on the VRF and configures BGP to use the SRv6 locator with per-VRF allocation. BGP MVPN then requests the multicast service SID from the SRv6 SID manager. The resulting End.DTMC4 SID is dynamically allocated on each participating PE and advertised in the MVPN control plane. It should be verified, recorded, and monitored, but it should not be copied from one router to another as a hard-coded address.

> **Operational rule:** A different egress PE can allocate a different End.DTMC4 SID for the same VRF. The ingress PE uses the SID advertised by the specific egress PE. The locator and function allocation policy must therefore be consistent and reachable, but the final SID value is node-local.

Example addressing and group plan

The following values make the configuration readable and deterministic. They are documentation examples and must be replaced with the validated addressing plan before deployment.

Purpose	RED VRF	BLUE VRF	Design rule
Group range	239.10.0.0/16	239.20.0.0/16	Do not overlap production group ownership.
Anycast RP	1.1.1.1/32	1.1.1.1/32	Reuse is safe only because the VRFs are isolated.
DC1-PE1 MSDP	192.0.2.11/32	198.51.100.11/32	Unique and routable inside its VRF.
DC1-PE2 MSDP	192.0.2.12/32	198.51.100.12/32	Must differ from the Anycast address.
Route target	65000:10	65000:20	Common across PEs in the same VPN.

Use a unique RD per PE and VRF. The examples use RED RDs 65000:101, 65000:102, 65000:111, and 65000:112 for DC1-PE1, DC1-PE2, BR1-PE1, and BR1-PE2 respectively. BLUE uses 65000:201, 65000:202, 65000:211, and 65000:212. The route target remains common within a VPN so that all participating PEs import the same service membership.

Deployment prerequisites

- The IPv6 IS-IS underlay must provide stable reachability to every PE SRv6 locator and BGP update source.
- The platform and Cisco IOS XE image must support MVPN ingress replication over SRv6 and the End.DTMC4 behavior. Confirm hardware and software support before applying the configuration.
- The existing MP-BGP design must carry VPNv4 and IPv4 MVPN address families through the route reflectors. Route reflectors must preserve the required extended communities.
- Each VRF must have a unique RD on every PE and consistent import/export route targets across the VPN.
- Unicast reachability to the multicast source, Anycast RP /32, and MSDP peer /32 must be present inside the correct VRF. Strict RPF is enabled by default for this SRv6 MVPN model, so an inconsistent unicast path can suppress multicast forwarding even when the BGP MVPN session is healthy.
- PIM-SM must be enabled only on the required customer-facing and RP loopback interfaces. IGMP must be enabled on receiver-facing LANs; IGMPv3 is recommended when hosts and applications support it.
- RED and BLUE configuration must remain symmetrical but independent. Never leak RP, MSDP, multicast source, or receiver routes between the VRFs unless an explicitly designed multicast extranet is required.

Configuration workflow

The examples below show the required configuration layers in dependency order. They use autonomous system 65000, route reflector 2001:DB8:0:FF::10, and SRv6 locator SLOC. Adapt interface names and addressing to the validated platform configuration.

1. Enable SRv6 multicast ingress replication in each VRF

Configure the VRF on every participating PE. The example below represents RED on DC1-PE1. Apply the same structure to BLUE with its own RD and route target, and to the remaining PEs with their node-specific RDs.

```
vrf definition RED
 rd 65000:101
 address-family ipv4
  route-target export 65000:10
  route-target import 65000:10
  srv6-mcast ingress-replication partitioned data
 exit-address-family
!
ip multicast-routing vrf RED distributed
!
vrf definition BLUE
 rd 65000:201
 address-family ipv4
  route-target export 65000:20
  route-target import 65000:20
  srv6-mcast ingress-replication partitioned data
 exit-address-family
!
ip multicast-routing vrf BLUE distributed
```

The partitioned keyword restricts the replication list to egress PEs that signal interest. The optional data keyword enables the data MDT profile. Maintain identical service intent across all PEs; a partial deployment produces asymmetric control-plane state and can leave receivers without a valid SRv6 leaf.

2. Enable BGP MVPN and per-VRF SRv6 allocation

The MVPN address family distributes multicast VPN routes, while the locator and per-VRF allocation statements allow BGP to allocate and advertise the SRv6 service SID. The following additive example assumes that the unicast VPNv4 neighbor and base BGP policy already exist.

```
router bgp 65000
 neighbor 2001:DB8:0:FF::10 remote-as 65000
 neighbor 2001:DB8:0:FF::10 update-source Loopback0
 !
 segment-routing srv6
  locator SLOC
 exit-srv6
 !
 address-family vpnv4
  neighbor 2001:DB8:0:FF::10 activate
  neighbor 2001:DB8:0:FF::10 send-community both
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
 exit-address-family
 !
 address-family ipv4 vrf RED
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
 exit-address-family
 !
 address-family ipv4 vrf BLUE
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
 exit-address-family
 !
 address-family ipv4 mvpn
  neighbor 2001:DB8:0:FF::10 activate
  neighbor 2001:DB8:0:FF::10 send-community both
 exit-address-family
```

Apply the IPv4 MVPN address family to the route reflectors and all PEs that participate in multicast VPN service. The route reflector does not allocate an End.DTMC4 SID unless it is also a service PE, but it must reflect the MVPN NLRI and associated tunnel and SRv6 attributes without stripping extended communities.

3. Enable PIM-SM and IGMP at the LAN edge

Configure multicast only inside the service VRF. The source-facing and receiver-facing interfaces use PIM-SM. Receiver LANs use IGMP to learn host interest. The sample below illustrates RED sub-interfaces; BLUE follows the same pattern with its own VLAN and subnet.

```
 ! DC1 source-facing LAN
```

```
interface TenGigabitEthernet0/0/0.110
 description DC1-LAN - RED multicast source segment
 encapsulation dot1Q 110
 vrf forwarding RED
 ip address 10.10.10.1 255.255.255.0
 ip pim sparse-mode
 ip igmp version 3
 no shutdown
!
! BR1 receiver-facing LAN
interface TenGigabitEthernet0/0/0.120
 description BR1-LAN - RED multicast receiver segment
 encapsulation dot1Q 120
 vrf forwarding RED
 ip address 10.10.20.1 255.255.255.0
 ip pim sparse-mode
 ip igmp version 3
 no shutdown
```

When both PEs attach to the same LAN, PIM Designated Router election and the first-hop redundancy design must be deterministic. Set an explicit PIM DR priority only when the intended active forwarding role is understood. Confirm that the unicast gateway, PIM DR, and multicast RPF path do not produce an unintended asymmetric topology.

4. Build the Per-VRF anycast RP pair

On both DC1 PEs, create the same Anycast RP /32 in each VRF. Create a second, unique loopback in each VRF for MSDP peering. The same 1.1.1.1 address can be reused in RED and BLUE because each VRF has an independent routing table.

```
! DC1-PE1 - RED
interface Loopback100
 description RED Anycast RP
 vrf forwarding RED
 ip address 1.1.1.1 255.255.255.255
 ip pim sparse-mode
!
interface Loopback101
 description RED MSDP peering address - DC1-PE1
 vrf forwarding RED
 ip address 192.0.2.11 255.255.255.255
 ip pim sparse-mode
!
! DC1-PE1 - BLUE
interface Loopback200
 description BLUE Anycast RP
 vrf forwarding BLUE
 ip address 1.1.1.1 255.255.255.255
 ip pim sparse-mode
!
interface Loopback201
 description BLUE MSDP peering address - DC1-PE1
 vrf forwarding BLUE
 ip address 198.51.100.11 255.255.255.255
 ip pim sparse-mode
```

On DC1-PE2, use the same Loopback100 and Loopback200 Anycast addresses, but configure 192.0.2.12/32 and 198.51.100.12/32 as the unique MSDP peer addresses. Ensure all four loopbacks are advertised into the correct VRF unicast control plane. The unique peer address, not the shared Anycast address, must be used for MSDP and for any manually selected routing protocol router ID.

5. Configure MSDP source synchronization

MSDP allows the two Anycast RPs to exchange Source-Active state. Build one independent MSDP relationship per VRF. DC1-PE1 uses the following configuration; reverse the peer addresses on DC1-PE2.

```
! DC1-PE1
ip msdp vrf RED peer 192.0.2.12 connect-source Loopback101
ip msdp vrf RED originator-id Loopback101
!
ip msdp vrf BLUE peer 198.51.100.12 connect-source Loopback201
ip msdp vrf BLUE originator-id Loopback201
! DC1-PE2
ip msdp vrf RED peer 192.0.2.11 connect-source Loopback101
ip msdp vrf RED originator-id Loopback101
!
ip msdp vrf BLUE peer 198.51.100.11 connect-source Loopback201
ip msdp vrf BLUE originator-id Loopback201
```

The MSDP TCP session must follow a valid unicast path inside the same VRF. If the RPs cannot reach one another through their unique loopbacks, Anycast routing can still direct local joins to an RP, but sources learned by the other RP may remain invisible. The failure can therefore look like a data-plane problem even though the root cause is the missing per-VRF MSDP session.

6. Configure static RP mappings for each VRF

Configure the same explicit group-to-RP mapping on every PIM router that participates in the RED or BLUE multicast domain. This includes DC1-PE1, DC1-PE2, BR1-PE1, BR1-PE2, and any customer multicast router between the PE and attached hosts. The common 1.1.1.1 address resolves to the local or closest reachable Anycast-RP instance at the DC hub, while the access lists keep group ownership distinct by VRF.

```
ip access-list standard RED-MCAST-GROUPS
 permit 239.10.0.0 0.0.255.255
!
ip access-list standard BLUE-MCAST-GROUPS
 permit 239.20.0.0 0.0.255.255
!
ip pim vrf RED rp-address 1.1.1.1 RED-MCAST-GROUPS
ip pim vrf BLUE rp-address 1.1.1.1 BLUE-MCAST-GROUPS
```

The mapping is identical on both datacenter hub PEs and on both branch PEs. RED resolves 239.10.0.0/16 to 1.1.1.1 inside the RED VRF; BLUE resolves 239.20.0.0/16 to the same numerical address inside the BLUE VRF. VRF separation makes that address reuse safe. The unique MSDP loopbacks remain different on DC1-PE1 and DC1-PE2 and must never be replaced with the shared Anycast address.

Static RP baseline: Do not configure Auto-RP candidate, mapping-agent, or listener functions for these group ranges. A future migration to dynamic RP discovery must be treated as a coordinated end-to-end change and validated on every multicast router.

Verification commands

Validating an end-to-end SRv6 Multicast fabric requires auditing the underlay locators, the BGP MVPN control plane, and the hardware-programmed CEF chains. Use the following commands to confirm operational health:

Auditing Dynamic End.DTMC4 SIDs

Run show segment-routing srv6 sid on the egress branch PE to verify that the router has dynamically allocated a multicast Service SID inside the locator block.

Verify BGP MVPN signaling and ingress replication trees

To confirm that the Route Reflector has propagated receiver-driven joins and constructed correct partitioned replication branches, execute:

- show bgp ipv4 mvpn summary – Verifies established peerings for the MVPN address family.
- show mvpn ipv4 vrf RED auto-discovery detail – Confirms that the PE has discovered remote multicast hubs and leaf nodes.
- show mvpn ipv4 vrf RED leaf-information detail – Traces which remote egress PEs have requested replication, along with their advertised End.DTMC4 SIDs.

Audit datacenter anycast RP and MSDP state

Confirm that active source registrations are successfully synchronizing between the hub Anycast RP pair:

- show ip pim vrf RED rp mapping – Audits the active standard group-to-RP mappings.
- show ip msdp vrf RED summary – Verifies that the MSDP TCP session is established through the unique Loopback peering IPs.
- show ip msdp vrf RED sa-cache – Audits synchronized active multicast sources.

Verify MFIB and ingress replication hardware programming

Verify that the incoming multicast packets are actively mapped to ingress replication Srvif tunnels and programmed directly into the router's hardware ASIC:

- show ip mroute vrf RED 239.10.10.10 verbose – Verifies that the incoming LAN interface is correct and that the Outgoing Interface List (OIL) lists the virtual Segment Routing interface (Srvif0).
- show ip mfib vrf RED 239.10.10.10 – Verifies hardware forwarding state.
- show mvpn replication lsm-id [id] – Displays the exact replication branches and lists each remote egress PE's destination End.DTMC4 SID, proving stateless data-plane delivery across your WAN core.

Operational design considerations

Replication scale: Ingress replication exchanges core multicast state for edge replication load. Validate the maximum number of egress leaves, aggregate multicast throughput, and hardware replication capacity on the selected platform.

Dual-homed LAN behavior: Two PEs on a common LAN can both observe IGMP and PIM traffic. The first-hop redundancy, PIM DR election, and unicast RPF design must identify the active forwarder and prevent duplicate delivery.

RPF determinism: MVPN signaling does not override an invalid customer unicast topology. Source prefixes, RP addresses, and MSDP peer addresses must resolve through the expected VRF interfaces.

Static RP consistency: The group ACL and 1.1.1.1 RP mapping must be identical on every PIM router in a VRF. A missing or mismatched mapping can create site-specific failures even when the Anycast RP and MSDP pair are healthy.

Anycast address selection: The RP address must be a /32 and must not be used as the MSDP or BGP router ID. The unique MSDP loopback is the correct source for MSDP and a safer explicit router ID.

Change control: Changing a route target, RD, locator, group ACL, or RP discovery mechanism can alter both signaling and data-plane state. Treat multicast service changes as end-to-end changes and verify all four PEs after every maintenance event.

Feature support: Cisco IOS XE feature availability varies by release, platform, forwarding ASIC, and license. Validate the exact configuration syntax and scale in the target software image before declaring the section configuration production-ready.

Design summary

The SRv6 multicast design extends the CVD without compromising its central architectural principle: the transport core remains a simple IPv6/SRv6 forwarding fabric, while service intelligence is implemented at the edge. BGP MVPN communicates receiver interest, ingress replication creates a unicast SRv6 copy for each interested egress PE, and End.DTMC4 returns the packet to the correct IPv4 multicast VRF after decapsulation.

At the customer edge, PIM-SM and IGMP preserve familiar multicast operations. Explicit static mappings bind each group range to the 1.1.1.1 Anycast RP, while per-VRF MSDP sessions synchronize active sources between the two datacenter-hub PEs. RED and BLUE can reuse the same Anycast address because every control-plane and forwarding operation remains scoped to its VRF.

The result is a multicast service that is isolated, redundant, receiver-driven, and operationally aligned with the broader mission-critical SRv6 fabric. Its acceptance criteria are measurable: correct RP mapping, established MSDP peers, valid BGP MVPN leaf state, a dynamically allocated End.DTMC4 SID, deterministic RPF, hardware MFIB programming, and successful failure tests without cross-VRF leakage or duplicate delivery.

Validated case study — the VeriVault

Modern military operations and intelligence-gathering systems depend on the absolute sovereignty, integrity, and continuous availability of tactical telemetry and classified sensor data. When forward-deployed tactical teams collect high-consequence intelligence in contested enclaves, a network compromise or a disruption in communication can immediately jeopardize lives and compromise national security. For defense and intelligence operators, legacy WAN infrastructures have reached a breaking point: they are crippled by signaling protocol bloat, bound to vulnerable centralized control planes, and structurally unequipped to defend against next-generation quantum-decryption threats.

This section deconstructs the validated WAN modernization of VeriVault Defense, a high-consequence government contractor providing secure data collection, tactical telemetry relay, and intelligence-processing services in support of national security operations. It details how VeriVault transitioned from a fragile, controller-reliant legacy network to an autonomous, segment-routed, and quantum-safe transport fabric utilizing Segment Routing over IPv6 (SRv6) on Cisco secure routing platforms.

VeriVault mission and the quantum threat

Criticality of enclave sovereignty

As a primary government contractor supporting defense and intelligence enclaves, VeriVault Defense operates at the absolute frontier of national security. Their primary mission involves the real-time ingest, processing, and distribution of highly classified blueprints, tactical telemetry feeds, and life-safety command-and-control operations. Forward-deployed teams gather Top Secret sensor data, drone surveillance video, and raw tactical telemetry in contested field environments, transmitting this sensitive payloads back to centralized analytical facilities for correlation and assessment.

In this environment, logical and physical security are not merely operational preferences; they are strict national security mandates. The exposure of a single packet could leak tactical plans, compromise active military enclaves, or betray forward intelligence sources. Every segment of the WAN must enforce absolute, uncompromising logical isolation between the Critical Enclave (Top Secret telemetry, command traffic) and the Non-Critical Enclave (administrative and general internet traffic).

Looming crisis: harvest now, decrypt later

While logical segmentation isolates enclaves inside virtual routing tables, the physical transit paths must traverse third-party commercial transport networks, fiber links, and public internet transits. Traditionally, these transports have been secured using classical asymmetric (public-key) cryptography—such as RSA, Diffie-Hellman (DH), and Elliptic Curve Diffie-Hellman (ECDH)—to securely negotiate encryption keys during IPsec and MACsec handshakes.

These mathematical primitives are secure against today's classical supercomputers, but they are completely vulnerable to a Cryptographically Relevant Quantum Computer (CRQC) running Shor's Algorithm. Adversaries and hostile nation-states are actively executing "Harvest Now, Decrypt Later" (HNDL) attacks today. They quietly intercept and archive petabytes of encrypted, highly classified defense data crossing shared WAN transits.

While they cannot read this data today, it acts as a digital time capsule. The instant a functional CRQC is developed (projected within the 2030–2035 horizon), adversaries will run Shor's algorithm to retroactively crack the historical public-key handshakes, obtain the master session keys, and decrypt decades of hoarded state secrets.

Post-quantum compliance mandate

Because data sent across the WAN today is effectively vulnerable to future retro-decryption, the White House (via EO 14412 / OMB M-26-15) and the National Security Agency (via CNSA 2.0) have accelerated post-quantum transition timelines. CNSA 2.0 mandates that all "Traditional Network Equipment" (including routers, switches, and firewalls) must support and enforce NIST Level 5 Post-Quantum Cryptography (PQC) for key establishment and digital signatures by 2027–2030.

To secure their data sovereignty and maintain federal compliance, VeriVault's WAN modernization required immediate, hardware-anchored integration of ML-KEM (FIPS 203) and ML-DSA (FIPS 204) lattice-based algorithms to mathematically immune their cryptographic handshakes from quantum decryption.

Case for SRv6 at the edge

Fragility of legacy transport patchworks

Prior to modernization, VeriVault operated a fragmented, highly complex hybrid WAN infrastructure. This legacy network consisted of two completely disjointed architectures:

- The private overlay: An MPLS-TE VPN running over private Metro-E links, utilizing stateful RSVP-TE signaling to enforce traffic engineering paths.
- The public overlay: A classical SD-WAN overlay running over public internet and cellular LTE links, managed by a centralized, software-defined controller cluster.

This model introduced severe operational, scalability, and security challenges:

- Extreme control-plane state bloat: Under legacy RSVP-TE, every intermediate transit (P) router in the core had to maintain an active state entry and a label-database allocation for every single traffic-engineered tunnel across the WAN.
- The "signaling storm" hazard: On low-bandwidth, disrupted, or high-latency tactical transport paths (such as satellite links or radio connections), the constant keepalive chatter required to synchronize these state tables would frequently saturate the links. A minor packet-loss event would trigger massive "signaling storms," leading to cascading protocol timeouts, routing flaps, and total network instability.
- Inefficient overlay-underlay splicing: The SD-WAN overlay remained completely blind to the physical underlay's pathing, routing metrics, and link degradation, resulting in the "noisy neighbor" problem where bulk background data transfers would saturate a link and drop critical voice or SCADA telemetry packets.

Architectural vulnerability of centralized controllers

More critically, VeriVault's legacy network was fundamentally unsuited for Disrupted, Degraded, Intermittently-connected, and Low-bandwidth (DDIL) tactical environments. In a controller-reliant SD-WAN design, path steering, tunnel orchestration, and dynamic policy execution are dictated by a centralized controller cluster.

If a link is severed by a physical fiber cut, jammed by electronic warfare, or severely degraded by storm conditions, a branch site loses connectivity back to the central controller. The moment this control link is broken, the site loses its ability to dynamically recalculate paths, provision backup tunnels, or steer around network failures—introducing a single point of failure that is completely unacceptable for mission-critical defense operations.

Modernization directive: why SRv6?

To address these critical vulnerabilities, they considered the adoption of Segment Routing over IPv6 (SRv6) as the consolidating architectural foundation. SRv6 unifies transport, traffic engineering, and VPN

overlay services into a single, native IPv6 data plane, eliminating legacy signaling protocol bloat (LDP, RSVP-TE, and BGP-labeled unicast).

Figure 26. Why SRv6?

Architectural dimension	Legacy RSVP-TE and SD-WAN patchwork	Native SRv6 secure fabric
Control plane state	Massive and stateful: transit core nodes must maintain active state entries and label allocations for every active tunnel.	Zero and stateless: core routers maintain zero path or customer VPN state; they forward packets based strictly on standard IPv6 LPM.
Signaling overhead	High: constant chatty peer-to-peer keepalives and refresh messages risk saturating tactical links.	None: path directives are self-contained within the IPv6 packet header itself.
Path computation	Centralized: vulnerable to complete policy failure if the controller connection is lost in DDIL conditions.	Autonomous Edge: locally computed via CSPF against the local IGP link-state database.
Protocol complexity	High Bloat: Multiple stacked protocols running simultaneously (IP, LDP, RSVP-TE, BGP, SD-WAN IPsec).	Unified: Single-layer native IPv6 handles transport, traffic engineering, and enclave isolation.
Control plane state	Massive and stateful: transit core nodes must maintain active state entries and label allocations for every active tunnel.	Zero and stateless: core routers maintain zero path or customer VPN state; they forward packets based strictly on standard IPv6 LPM.

By collapsing the stack, SRv6 shifts path computation and service execution directly to the ingress edge nodes (head-ends). Path computation is performed locally in hardware using Constrained Shortest Path First (CSPF) against the local Interior Gateway Protocol (IS-IS) link-state database. Even if a site is completely cut off from central orchestration, the local edge routers continue to make autonomous, real-time routing decisions, guaranteeing Autonomous Edge Resilience in DDIL environments.

Site profiles and functional hierarchy

Rather than deploying bespoke, complex configurations at every location, VeriVault standardized on three distinct, modular site profiles executing the same unified SRv6 fabric model.

Figure 27. Site profiles

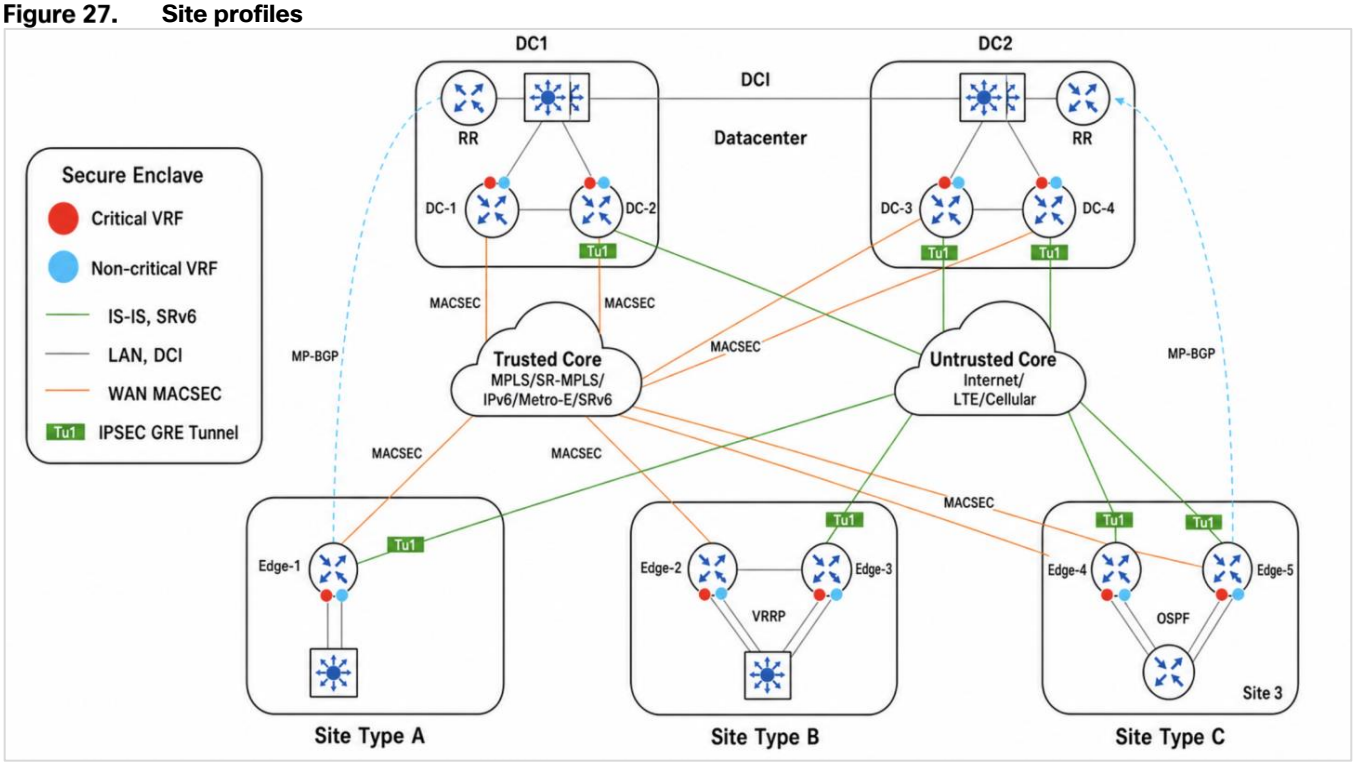

1. Field sites (site type a)

- Operational role: Lean, forward-deployed locations staffed by small collection teams. These sites are primarily responsible for the rapid ingest, aggregation, and transmission of raw sensor telemetry, drone video streams, and tactical field data.
- Hardware platform: Deployed using a Cisco Secure Router, C8161-G2 Secure Router running Cisco IOS XE 26.1. This compact platform was selected to minimize the physical, power, and thermal footprint in austere tactical conditions.
- Handoff constraints: Provisioned with a Layer 2 Metro-E E-Line subinterface (VLAN 4094) for the private core, and a routed Gigabit Ethernet interface for the public internet transport.
- Platform exception model: The C8161-G2 platform does not natively support hardware-level WAN MACsec (802.1AE). To maintain the strict security architecture without a costly hardware swap, the validated design implements a secured transport exception: it establishes a GRE over IPsec tunnel directly over the private E-Line to cryptographically protect the packets before handing them off to the provider cloud.

2. Secure processing centers (site type B and site type C)

- Operational role: Site Type B and Site Type C are regional analytical facilities staffed by intelligence analysts who sort, correlate, process, and act on the telemetry collected from Field Sites. These centers demand high throughput, physical and hardware-level redundancy, and uncompromising logical isolation.
- Hardware platforms: Deployed using Cisco Secure Routers, C8375-G2 and C8475-G2 Secure Edge Routers running IOS XE 26.1 in each if these site types.
- Dual-transport execution: Both edge nodes are physically connected to both the Trusted Core (private Metro-E E-Line) and the Untrusted Core (shared routed internet).
- Site type B has a cross-connect link which is used to redirect traffic between the two routers, similar to SD-WAN TLOC-extension.

- LAN redundancy variants:
    - L2 LAN variant (split-router spoke): Designed for sites where the transport enclaves are terminated on two separate physical platforms (C8375-E-G2). The routers share a common LAN segment. Virtual Router Redundancy Protocol (VRRP) is configured across both nodes to provide a single, highly available default gateway for LAN hosts, with EDGE2_1 designated as Master for the secure Critical VRF and EDGE2_2 as Master for the unsecure Non-Critical VRF.
    - L3 LAN variant: Designed for larger, pre-existing routed LAN infrastructures and sites that host the C8475-E Routers. Both edge routers operate as fully active/active transport gateways, using dynamic routing protocols to distribute and balance tenant traffic symmetrically across both private MACsec links and IPsec public fallback paths.

3. Datacenter hubs (DC1 and DC2 profiles)

- Operational role: Centralized high-performance core aggregation hubs responsible for terminating all site-originated secure tunnels, enforcing global routing intent, and hosting the BGP control-plane.
- Hardware platforms: Deployed as redundant pairs consisting of Cisco Secure Router 8570-G2 and 8550-G2 platforms running the IOS XE 26.1 engineering image with hardware-accelerated cryptographic offload engines.
- The Route Reflection (RR) anchor: To scale the Multi-VRF BGP control plane without creating an unmanageable full mesh of peerings, a dedicated Cisco ASR 1001-HX platform is deployed on the Data Center Interconnect (DCI) to act purely as a control-plane Route Reflector. This RR distributes service routes and attached SRv6 Service SIDs across all nine WAN endpoints without participating in transit packet encryption.

VeriVault design and architecture

The VeriVault architecture achieves logical, physical, and cryptographic security by layering several advanced IOS XE segment-routing capabilities into a single, cohesive operating model.

Figure 28. Layered SRv6 and quantum-safe capabilities

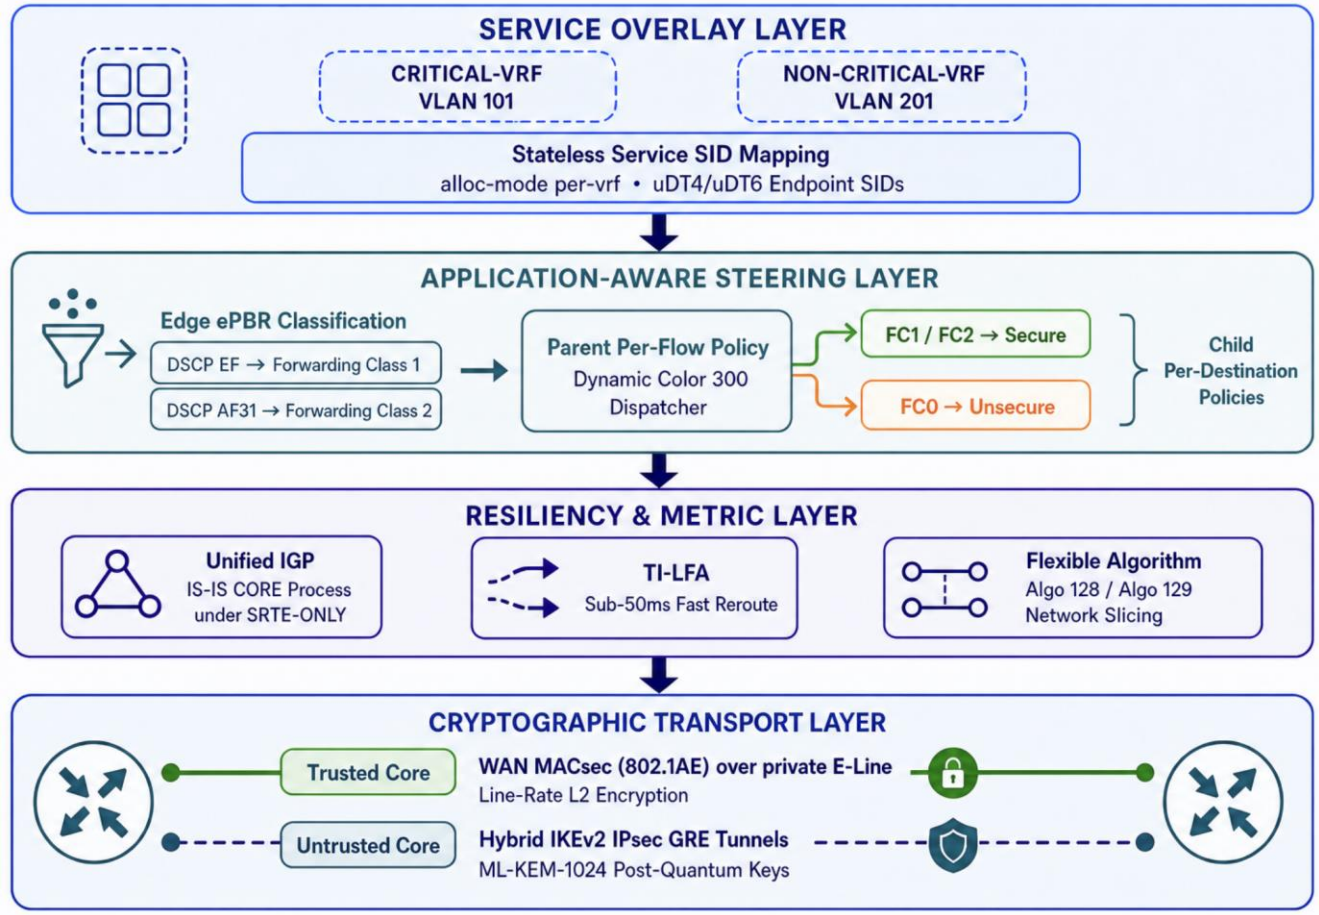

1. Enclave isolation and stateless scaling (Service SIDs)

To maintain absolute separation between classified and unclassified traffic at the edge, VeriVault instantiates two dedicated Virtual Routing and Forwarding (VRF) tables on all routers: CRITICAL_VRF (for secure, classified telemetry mapped to VLAN 101) and NON-CRITICAL_VRF (for administrative and internet traffic mapped to VLAN 201).

In traditional MPLS, separating VRFs required allocating independent labels per prefix, rapidly exhausting core-router TCAM memory and creating massive routing-table overhead. SRv6 solves this by enforcing the Cisco IOS XE Critical Design Rule: Service SIDs are allocated strictly using alloc-mode per-vrf.

Rather than generating SIDs for individual subnets, BGP allocates a single, highly aggregated functional Service SID for the entire VRF instance (such as End.DT4 for IPv4 VRF lookup). Egress routers advertise these VRF-specific Service SIDs via MP-BGP. Transit routers perform standard IPv6 longest-prefix-match routing purely on the aggregate Locator /48 block. They remain entirely oblivious to individual customer VPN prefixes, ensuring near-infinite control-plane and data-plane scalability.

2. Dual-transport slicing (link affinities)

VeriVault's physical transports are carved into distinct cryptographic planes:

- The trusted core: A private Metro-E VPWS E-Line backbone secured via hardware-level, line-rate WAN MACsec (802.1AE) encryption. Point-to-point MACsec Key Agreement (MKA) sessions are established directly between edge platforms and DC hubs, utilizing certificate-based mutual authentication via EAP-TLS (TLS 1.3 with ML-KEM).

- The untrusted core: Shared public internet and cellular links secured end-to-end via GRE over IPsec tunnels. IPsec tunnels are negotiated using RFC 9370 (IKEv2 Multiple Key Exchanges) to establish hybrid ML-KEM-1024 (Kyber) and Diffie-Hellman Group 19 keys.

To prevent these disparate underlay technologies from being treated as equal paths by the IGP, VeriVault engineers configure Link Affinities within the unified IS-IS underlay process (SRTE-OVLY). Using the affinity-map command, private interfaces are tagged with SECURED (bit-position 0) and public internet interfaces are tagged with UNSECURED (bit-position 1) [59, 202]. This topological awareness allows the Segment Routing path computation engine to mathematically exclude insecure paths during route calculation.

2. Application-aware steering hierarchy (PFP and PDP)

To defeat the " noisy neighbor" problem and guarantee strict SLAs for critical applications, VeriVault deploys an intent-based, layered traffic-engineering hierarchy.

- The ingress classifier: Ingress LAN interfaces are bound to an Endpoint Policy-Based Routing (ePBR) policy-map. Using Deep Packet Inspection via Network-Based Application Recognition (NBAR) and DSCP analysis, the ePBR policy intercepts incoming traffic and stamps it with an internal router Forwarding Class (FC) metadata tag.
    - Voice and SCADA Telemetry (DSCP EF) is mapped to Forwarding Class 1.
    - Business-Critical Data (DSCP AF31) is mapped to Forwarding Class 2.
    - Best Effort/General Administrative traffic falls into the default Forwarding Class 0.
- The parent policy (Per-Flow Policy - PFP): Service routes are advertised across the WAN carrying a BGP Color Extended Community of 300, which dynamically instantiates a parent Per-Flow Policy (PFP). The PFP does not calculate physical paths itself; instead, it acts as an intelligent dispatcher. It reads the internal Forwarding Class metadata stamped by the ePBR policy and redirects each flow to a specific child Per-Destination Policy (PDP):
    - FC 1 and FC 2 (Telemetry/Voice): Dispatched to a Child PDP restricted strictly to the Trusted Core (using link affinities, e.g., include-all SECURED and optimized for minimum real-time latency via metric type delay).
    - FC 0 (Administrative/Best-Effort): Dispatched to a Child PDP optimized for static business cost and routed over the Untrusted Core (UNSECURED public internet).

3. Network slicing (flexible algorithm)

For macro-level slicing, VeriVault deploys flexible algorithm (flex-algo). While PFPs inspect individual packets to split flows, flex-algo operates natively at the IGP layer (IS-IS) to calculate entirely distinct, mathematically isolated shortest-path routing topologies over the same physical hardware.

- Algorithm 128 (the secure slice): Computes a logical topology that strictly includes links explicitly flagged with the SECURED affinity (private WAN MACsec).
- Dedicated Locators per Slice: In accordance with strict SRv6 design rules, each flex-algo is assigned a separate, dedicated locator block (e.g., FCBB:DEAD for Algo 0, FCBB:DEAF for Algo 128, and FCBB:DEAB for Algo 130). Traffic destined for the Critical VRF is encapsulated using the Algo 128 locator, forcing intermediate transit routers to forward the packets exclusively along the isolated, encrypted slice.

4. Sub-50ms local resiliency (TI-LFA)

To guarantee zero-downtime survivability during hard physical link failures (such as fiber cuts), Topology Independent Loop-Free Alternate (TI-LFA) is enabled natively within the IS-IS process.

- Pre-calculated backup paths: The local routing engine proactively pre-calculates backup repair paths that perfectly mirror the post-convergence topology of the network. If a physical fiber cut occurs, the adjacent Point of Local Repair (PLR) immediately activates the backup path, encapsulating transit packets in an outer IPv6 header carrying the pre-calculated uSID Segment List. This transitions traffic locally in hardware in under 50 milliseconds, completely bypassing the micro-loops and cascading flaps that plague legacy protocols.
- Symmetry and outbound tuning: At dual-homed medium and large sites, outbound BGP local preference is adjusted (depressing the local preference on the unsecured router EDGE2_2 to 50, while keeping the secured router EDGE2_1 at 100). This configuration guarantees that return data plane traffic flows symmetrically back across the Trusted Core MACsec underlay under normal conditions, while still permitting a clean, automated fallback to the post-quantum IPsec tunnel if the primary link fails.

5. Guaranteed path symmetry

Because the secure enclave at Large sites is split across two physically separate platforms (secured EDGE3_1 and unsecured EDGE3_2), return traffic would normally be subject to asymmetric routing across the untrusted core, presenting a severe security risk. To guarantee absolute path symmetry, BGP outbound Local Preference Tuning is configured on the unsecured node (EDGE3_2). By depressing the local preference of advertisements exiting EDGE2_2 toward the Route Reflectors to 50 (relative to the default preference of 100 on the secured EDGE2_1), return data-plane traffic is attracted strictly back across the Trusted Core via EDGE3_1's native WAN MACsec underlay under normal conditions. If the primary MACsec underlay fails, BGP seamlessly falls back to routing return traffic through the post-quantum secured IPsec tunnel terminating on EDGE3_2.

Similar configurations and path symmetry are maintained and configured for the medium-sized split-transport sites (Site Type B) as well. Refer to the appendices for the exact device configurations of EDGE2_1 (Trusted Core termination) and EDGE2_2 (Untrusted Core termination).

Crucially, this site design incorporates a physical and logical inter-router cross-connect link directly connecting the two edge routers. This cross-connect link is vital for maintaining traffic symmetry, enforcing cryptographic path alignment, and guaranteeing deterministic local resiliency:

- Symmetric flow alignment: Under normal operations, outbound LAN traffic is directed to either EDGE2_1 or EDGE2_2 based on application classification (e.g., ePBR forwarding classes). The cross-connect link ensures that if traffic returns on an asymmetric WAN path, it can be bridged locally between the two routers to reach the correct LAN gateway without having to hairpin back through the provider network, ensuring consistent latency and session state preservation.
- Dynamic TI-LFA local repair: If the primary WAN interface on EDGE2_1 suffers a physical link failure, TI-LFA immediately activates. The pre-computed backup path (repair list) leverages the cross-connect link to instantly steer in-flight packets over to EDGE2_2 in under 50ms, allowing traffic to failover seamlessly to the IPsec GRE underlay.
- Flex-Algo and SR-TE path continuity: To prevent traffic from traversing unverified paths during a failover, the cross-connect link must be actively configured as a member of both Flex-Algo 128 (Secure) and Flex-Algo 129 (Low-Latency) topologies. This ensures that when a policy-based SR-TE or Flex-Algo path is calculated, the routing engine recognizes the cross-connect as a valid, mathematically isolated transit link. Consequently, constrained traffic can cross from one edge platform to another while strictly preserving its original slicing constraints (e.g., remaining entirely within the PQC-encrypted or low-delay topology).

6. Tactical edge double encryption and MTU/MSS optimization

A critical operational challenge validated in VeriVault's tactical edge design is the handling of Double Encryption. In many military and national security environments, defense enclaves operate high-

assurance Type 1 Inline Network Encryptors (INEs, such as KG-175 TACLANEs) directly within their secure local enclaves before the traffic ever reaches the WAN CPE router.

When these encrypted classified packets are handed off to a secure Cisco edge router, the router encapsulates them into SRv6 L3VPN service overlays and subsequently encrypts them a second time using Layer 2 WAN MACsec (on trusted links) or Layer 3 IPsec GRE (on public links) to secure the transport fabric against HNDL harvesting.

This double-encryption architecture adds significant header overhead, which must be carefully managed to prevent performance-degrading packet fragmentation and packet drops.

The double-encryption overhead stack

When a classified IP packet leaves a tactical LAN client and traverses the secure WAN, it is subjected to multiple encapsulation layers:

- Original IP payload (classified): Standard customer IP packet.
- Type 1 encryption layer (classified IPsec ESP): Encapsulated by the site's local inline encryptor. Adds a new IPv4/IPv6 header, an ESP header, an Initialization Vector (IV), padding, and an ESP Trailer/ICV (Integrity Check Value). This adds approximately 50 to 60 bytes of cryptographic overhead.
- SRv6 service encapsulation overlay: Encapsulated by the Cisco Ingress Edge router. The inner packet is wrapped in an outer IPv6 transport header with an active Destination Address representing the egress router's VRF Service SID (uDT4 / uDT6 or End.DTMC4). Using the F3216 compressed Micro-SID (uSID) format with Reduced Headend Encapsulation (H.Encap.Red), this outer IPv6 header adds exactly 40 bytes. If an explicit SR-TE Segment Routing Header (SRH) is inserted, it adds an additional 8 bytes per transit segment.
- GRE tunnel transport layer (shared WAN fallback): The SRv6 packet is wrapped inside an outer GRE header to tunnel IS-IS routing adjacencies and locator reachability across public or third-party networks. This adds an extra 24 bytes (IPv4 GRE) or 44 bytes (IPv6 GRE).
- Transport IPsec layer (shared WAN fallback): The GRE tunnel is protected by a Post-Quantum IPsec Profile (ML-KEM-1024). Running in IPsec Transport Mode adds an ESP header, an IV, padding, and an ESP trailer/integrity MAC, adding approximately 56 to 70 bytes.
- WAN MACsec encryption layer (private WAN): Applied at the egress physical or sub-interface. Adds a 16-byte SecTAG header and a 16-byte Integrity Check Value (ICV), introducing 32 bytes of Layer 2 overhead.

The overhead impact and fragmentation risks

In a standard WAN environment, the maximum transitable physical packet size is bounded by the standard Ethernet Maximum Transmission Unit (MTU) of 1500 bytes.

If a classified tactical payload of 1420 bytes is first encrypted, the resulting packet grows to ~1480 bytes. When this packet hits the secure WAN router and is encapsulated with SRv6 uSIDs (+40 bytes) and wrapped inside a shared WAN GRE over IPsec tunnel (+80 bytes), the cumulative packet size swells to 1,600 bytes, far exceeding the 1500-byte physical WAN MTU.

In an IPv6 underlay (which native SRv6 requires), intermediate transit routers do not perform fragmentation. If a packet exceeds a link's MTU, the transit router drops the packet and transmits an ICMPv6 Packet Too Big (PTB) message back to the sender. If these ICMPv6 messages are blocked or filtered by security firewalls along the path, the packet is silently dropped, creating a permanent MTU Black Hole that causes active TCP connections to stall and hang indefinitely.

For IPv4 payloads (or when running over IPv4 public underlays), any packets exceeding the MTU that do not have the Don't Fragment (DF) bit set are fragmented by the router. This forces the terminating or egress router to perform reassembly in software, placing a massive computational tax on the router's Central Processing Unit (QFP engine) and causing a severe degradation in overall network throughput and latency.

Mitigation strategies

To guarantee that packet fragmentation and MTU black holes are eliminated across VeriVault's double-encrypted transport fabric, the following four-tier engineering design must be globally enforced:

1. Rigid interface MTU clamping

To prevent the physical underlay from dropping highly encapsulated packets, strict, non-default MTU limits must be configured on all transport interfaces:

- **Unsecured GRE/IPsec tunnels:** Must be locked to ip mtu 1400 and ipv6 mtu 1400 to safely absorb standard IPsec and GRE tunnel encapsulation overhead.
- **WAN MACsec GRE tunnels (exception / exception models):** Must be tuned to a conservative ip mtu 1300, ipv6 mtu 1300, and clns mtu 1300. Dropping the CLNS MTU and the IS-IS Link State Packet (LSP) size (lsp-mtu 1300) guarantees that vital routing updates and locator advertisements pass flawlessly across the highly encapsulated provider E-Line without fragmentation.
- **Physical WAN MACsec interfaces:** The physical parent and tagged sub-interfaces terminating native MACsec must be configured with an IP MTU of 1468 (ip mtu 1468 / ipv6 mtu 1468 / clns mtu 1400) to safely absorb MACsec Layer 2 frame expansion while maintaining wire-speed line-rate forwarding.

2. Aggressive TCP Maximum Segment Size (MSS) clamping

While MTU boundaries protect the router interfaces, the most effective way to prevent fragmentation is to force TCP endpoints to send smaller segments in the first place. This is achieved by enabling TCP MSS Clamping on all LAN-facing service interfaces:

- Operators must apply ip tcp adjust-mss 1220 (for IPv4 enclaves) and ipv6 tcp adjust-mss 1200 (for IPv6 enclaves) on all LAN-facing customer VRF access interfaces (such as GigabitEthernet0/1/2.101 or TwoGigabitEthernet0/0/0.101).
- **How it works:** When a client initiates a TCP connection, the router actively intercepts the TCP SYN packet during the three-way handshake. If the client's advertised MSS exceeds the configured clamp value, the router dynamically rewrites the TCP MSS field down to 1220 (or 1200 for IPv6) before forwarding the packet. This mathematically guarantees that the subsequent TCP data payloads, even when wrapped in the cumulative Type 1 ESP header (+60 bytes), the SRv6 service header (+40 bytes), and the WAN IPsec GRE tunnel header (+80 bytes), will never exceed the 1500-byte WAN MTU limit, preventing any downstream fragmentation.

3. Tunnel Path MTU Discovery (PMTUD)

To dynamically adapt to restricted path MTUs across public or provider transport clouds, Tunnel Path MTU Discovery must be enabled on all virtual tunnel interfaces:

- Apply tunnel path-mtu-discovery directly under the tunnel configuration block.
- This instructs the router to set the Don't Fragment (DF) bit in the outer IP header and dynamically adjust its encapsulation packet sizes based on ICMP/ICMPv6 feedback from the network transit nodes. (Ensure firewall policies globally permit ICMP Type 3 Code 4 "Fragmentation Needed" and ICMPv6 Type 2 "Packet Too Big" messages to allow PMTUD to function).

4. Mandating IKEv2 fragmentation

Because VeriVault utilizes post-quantum ML-KEM-1024 algorithms for secure key exchange, the public keys exchanged during the initial IKEv2 handshake are significantly larger than classical Diffie-Hellman keys (the ML-KEM-1024 public key payload alone is 1,568 bytes).

- These massive key exchange packets will naturally be fragmented at the IP layer by the router. However, many intermediate firewalls and commercial ISPs block IP fragments by default, preventing the IKEv2 session from ever completing.
- To resolve this, operators must configure IKEv2 Fragmentation globally on all secure edge routers using the command: crypto ikev2 fragmentation mtu 1400 (or mtu 1300 on constrained tactical links) [190, 249, 354]. This forces the router to perform fragmentation natively at the IKEv2 protocol layer rather than the IP layer, splitting the large key payloads into multiple standard, unfragmented UDP packets that pass cleanly through intermediate firewalls and ISPs.

Phased parallel migration strategy

To transition VeriVault's highly sensitive operations to the new SRv6 fabric with zero downtime and absolute data security, Cisco deployed a phased Parallel Build Migration approach rather than a risky " flash cut" .

Figure 29. Phased parallel migration strategy

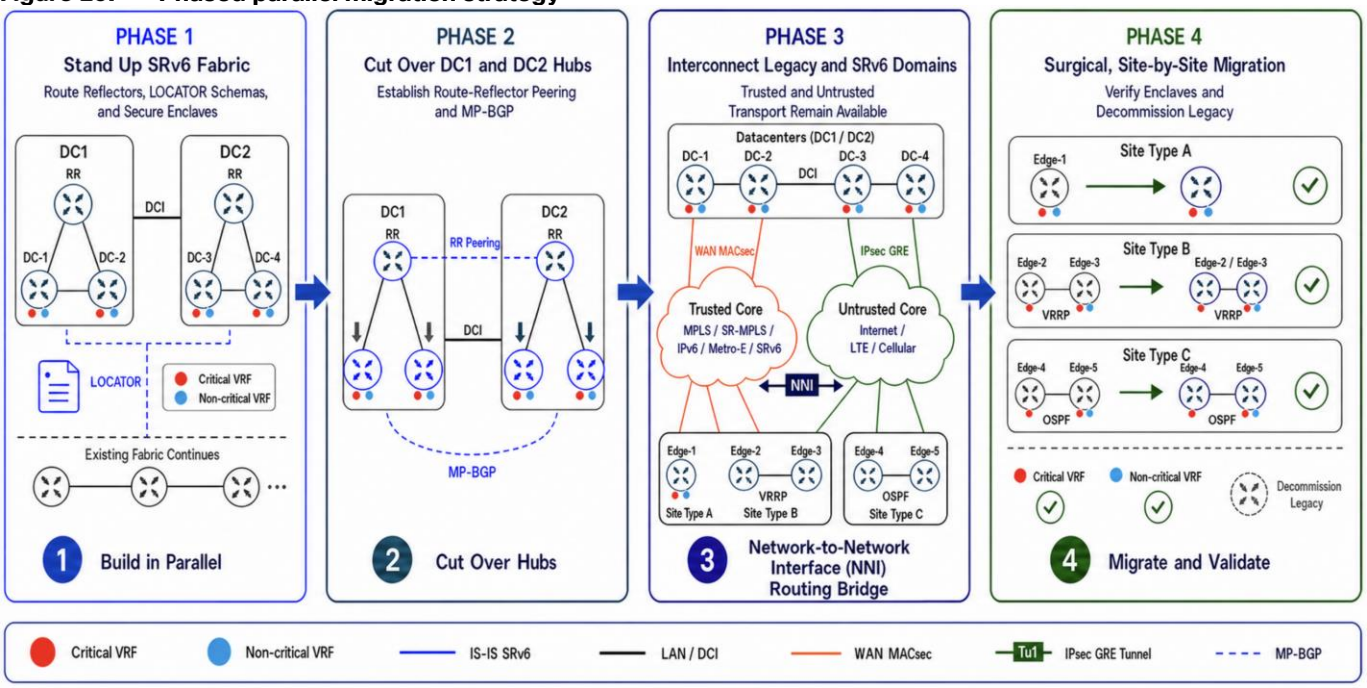

Phase 1: parallel fabric deployment

- **Execution:** The new SRv6 fabric (IS-IS SRTE-OVLY process, loopbacks, and F3216 locator blocks) was stood up in parallel with the legacy WAN.
- **Impact:** Because the fabric operates natively on a clean IPv6 addressing scheme, there was zero routing overlap or packet conflict with the legacy IPv4/MPLS transport core, allowing engineers to validate underlay reachability without impacting production traffic.

Phase 2: datacenter aggregation core cutover

- **Execution:** The Datacenter Hubs (DC1 and DC2) were migrated first.

- **Detail:** Engineers established the centralized ASR 1001-HX Route Reflector, activated the VRF definition mappings, and validated that the DC core routers could successfully receive and process both standard IPv6 unicast and encapsulated SRv6 Service SIDs.

Phase 3: standing up the Network-to-Network Interface (NNI) bridge

- **Execution:** To allow migrated and non-migrated enclaves to communicate seamlessly during the transition, a Network-to-Network Interface (NNI) routing bridge was established at the DC aggregation layer.
- **Detail:** Dynamic routing protocols (BGP/OSPF) were enabled between the legacy core and the new SRv6 hubs. When a migrated SRv6 spoke transmitted traffic to a legacy spoke, the packet was sent to the DC Hub, decapsulated, routed across the NNI bridge into the legacy VRF table, re-encapsulated using classical MPLS/IPsec, and forwarded to its destination—maintaining unbroken end-to-end communication throughout the multi-week migration window.

Phase 4: surgical, site-by-site spoke cutover

- **Execution:** Field Sites and Secure Processing Centers were migrated individually during scheduled, non-disruptive maintenance windows.
- **Detail:** At each site:
    - The legacy edge equipment was physically replaced with the Cisco Secure Router platforms (C8161-G2, C8375-E-G2, or C8475-G2).
    - The unified SRTE-OVLY IS-IS process, uSID locators, and ePBR/NBAR policies were activated.
    - Tunnels were brought up, validating the hybrid ML-KEM-1024 IPsec and WAN MACsec session establishment.
    - Control-plane route propagation via the ASR Route Reflector was verified, confirming that the remote edge routers successfully accepted and programmed the local VRF Service SIDs.
    - Traffic was cleanly swung onto the new fabric, and legacy connectivity was systematically decommissioned.

This structured, phased transition allowed VeriVault to modernize its global WAN incrementally, validating failover, latency steering, and post-quantum encryption performance at every milestone without ever placing active national security payloads at risk.

Unified technical mapping

The business values and technical accomplishments of VeriVault Defense's WAN modernization are summarized in the master mapping matrix below:

Table 4. Unified technical mapping

Business requirement	Technical challenge	Cisco SRv6 validation solution
Data sovereignty and quantum immunity	Retrospective HNDL attacks by nation-states armed with future CRQCs.	Deployed hybrid ML-KEM-1024 (Kyber) key encapsulation for untrusted internet transits, mathematically protecting historical keys.
Strict enclave isolation	Complex VRF-lite underlay configuration and heavy per-prefix label allocation.	Implemented native BGP-based SRv6 L3VPNs with highly scalable, per-VRF Service SIDs (uDT4/uDT6).
Survivability in contested environments	Single points of failure in centralized SD-WAN controllers under jammed/degraded DDIL links.	Pushed path-steering intelligence directly to the edge using local head-end CSPF computation and stateless core routing.

Business requirement	Technical challenge	Cisco SRv6 validation solution
Sub-50ms optical resiliency	Cascading packet drops, slow convergence, and RSVP-TE signaling storms during physical cuts.	Activated underlay Topology Independent Loop-Free Alternate (TI-LFA) to execute local fast-reroute without controller dependency.
Noisy neighbor protection	Bulk background transfers (administrative data) congesting and dropping SCADA and life-safety voice.	Built parent-child Per-Flow Policies (PFPs) driven by ePBR/NBAR deep packet inspection to steer voice to low-latency paths.
Simplified operations	Fragmented and heterogeneous configurations across multiple regional branches.	Standardized on three modular site profiles executing a single, unified Cisco secure routing operating model.

Appendices: Validated IOS-XE configurations

In a zero-trust, mission-critical Wide Area Network (WAN), configuration standardization is the cornerstone of logical isolation, deterministic pathing, and cryptographic integrity. Any configuration drift across edge devices can lead to fragmented Interior Gateway Protocol (IGP) databases, asymmetric routing loops, or silent packet drops within encrypted transit tunnels. This appendix provides the complete, production-grade Cisco IOS XE configuration blueprints for the dual-datacenter, multi-site reference architecture validated within this Cisco Validated Design (CVD).

By providing copy-pasteable command blocks directly derived from our physical lab testbed, this appendix serves as the definitive engineering baseline for deploying secure Segment Routing over IPv6 (SRv6) overlays across both high-speed MACsec-encrypted private E-Lines and post-quantum IPsec-secured public internet transits.

This Appendix Includes:

- Reference Lab Topology and Infrastructure Scope

  Documents the physical-to-logical interface mappings, Autonomous System Boundaries (ASNs), and IPv6 locator schemas that bind the datacenter aggregation hubs and branch edge platforms into a unified transport fabric.

- Datacenter Aggregation Hub Configuration

  Datacenter: Provides the master configuration for the Cisco Secure Router 8570-G2 acting as the primary cryptographic aggregation headend, Route Reflector (RR), and Anycast Rendezvous Point (RP).

- Site Type A Secure Branch Edge Configuration (EDGE1_1)

  Small Site: Blueprints the single-router, dual-transport branch profile (Cisco Secure Router 8161-G2) enforcing strict, logical isolation between critical enclaves and administrative transits using Per-Flow Traffic Engineering.

- Site Type B Medium-Site Primary Edge Configuration (EDGE2_1)

  ◦ Medium Site: Blueprints the primary secure router (Cisco Secure Router 8375-E-G2) executing WAN MACsec line-rate encryption, VRRP active LAN-side tracking, and automated failure-sensing EEM scripts.

  ◦ Medium-Site Backup Edge Configuration (EDGE2_2)
  Blueprints the secondary unsecure router (Cisco Secure Router 8375-E-G2) terminating public internet transits and acting as the active LAN gateway for standard administrative traffic.

- Site Type C Large-Site Primary Edge Configuration (EDGE3_1)

  Blueprints the primary secure router (Cisco Secure Router 8375-E-G2) with dual WAN links with L3 routing towards the LAN router.

- Site Type C Large-Site Primary Edge Configuration (EDGE3_1)

- Blueprints the backup secure router (Cisco Secure Router 8375-E-G2) with dual WAN links with L3 routing towards the LAN router.

Reference lab topology and infrastructure scope

The configurations provided in this appendix are validated against the master multi-site reference lab topology below. The infrastructure is split into a Trusted Core (operating private Layer 2 E-Line transits secured with high-speed WAN MACsec) and an Untrusted Core (operating public internet uplinks secured end-to-end with post-quantum IPsec tunnels).

Validated for code version IOS-XE release 26.1.1.

Figure 30. Reference Lab Topology

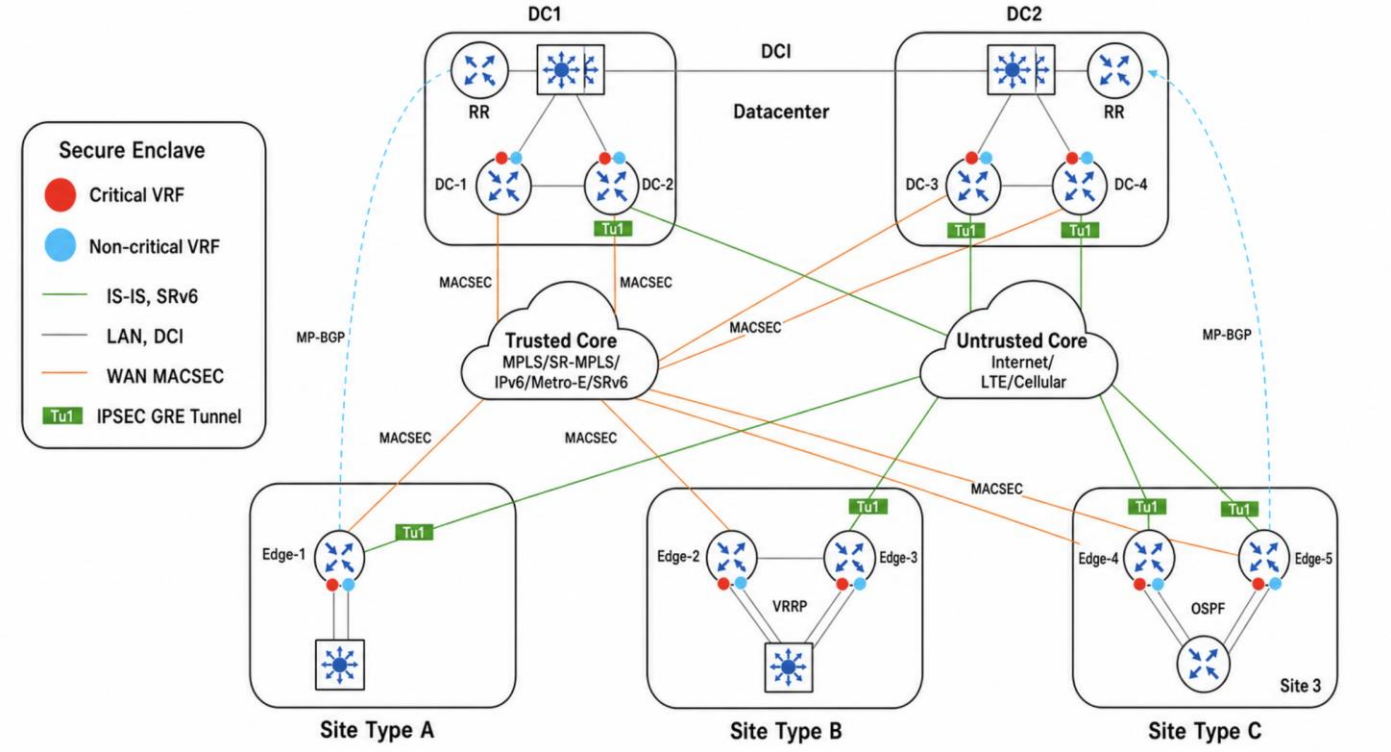

Hub 1_1 specifications

Table 5. Hub 1_1 Specifications

Parameter	HUB1_1
WAN Interfaces (MACSEC)	Te0/0/1 parent .4078 to EDGE3_2 (172.31.78.2/30, 2001:DB8:4078:1::2/64) .4084 to EDGE3_1 .4093 to EDGE2_1 .4094 to EDGE1_1 secured E-Line exception
WAN Interfaces (IPsec)	Te0/0/2 172.16.11.2/30 Tu200 to EDGE2_2 Tu213 to EDGE1_1 Tu253 to EDGE3_1 Tu262 to EDGE3_2
LAN Interfaces (Critical VRF)	Te0/0/4.101 198.18.11.1/24 Lo100 10.11.11.1/24 IPv6 2001:DB8:11:11::1/64
LAN Interfaces (Non-Critical VRF)	Te0/0/4.201 198.19.111.1/24 Lo210 10.111.111.1/24 IPv6 2001:DB8:111:111::1/64
Locator	SLOC 2001:DB8:AC11::/48

Parameter	HUB1_1
BGP ASN	65000

Hub 1_2 specifications

Table 6. Hub 1_2 Specifications

Parameter	HUB1_2
WAN Interfaces (MACSEC)	Te0/0/1 parent .4081 to EDGE1_1 secured E-Line exception .4089 to EDGE2_1 .4091 to EDGE3_1 .4092 to EDGE3_2
WAN Interfaces (IPsec)	Te0/0/2 172.16.12.2/30 Tu210 to EDGE1_1 Tu220 to EDGE2_2 Tu250 to EDGE3_1 Tu260 to EDGE3_2
LAN Interfaces (Critical VRF)	Te0/0/4.101 198.18.12.1/24 Lo100 10.12.12.1/24 IPv6 2001:DB8:12:12::1/64
LAN Interfaces (Non-Critical VRF)	Te0/0/4.201 198.19.112.1/24 Lo210 10.112.112.1/24 IPv6 2001:DB8:112:112::1/64
Locator	SLOC 2001:DB8:AC12::/48
BGP ASN	65000

Hub 2_1 specifications

Table 7. Hub 2_1 Specifications

Parameter	HUB2_1
WAN Interfaces (MACSEC)	Te0/0/1 parent .4079 to EDGE3_2 (172.31.79.2/30, 2001:DB8:4079:1::2/64) .4082 to EDGE1_1 secured E-Line exception .4085 to EDGE3_1 .4088 to EDGE2_1
WAN Interfaces (IPsec)	Te0/0/2 172.16.21.2/30 Tu211 to EDGE1_1 Tu221 to EDGE2_2 Tu251 to EDGE3_1 Tu263 to EDGE3_2

Parameter	HUB2_1
LAN Interfaces (Critical VRF)	Te0/0/4.101 198.18.21.1/24 Lo100 10.21.21.1/24 IPv6 2001:DB8:21:21::1/64
LAN Interfaces (Non-Critical VRF)	Te0/0/4.201 198.19.121.1/24 Lo210 10.121.121.1/24 IPv6 2001:DB8:121:121::1/64
Locator	SLOC 2001:DB8:AC21::/48
BGP ASN	65000

Hub 2_2 specifications

Table 8. Hub 2_2 Specifications

Parameter	HUB2_2
WAN Interfaces (MACSEC)	Te0/0/1 parent .4083 to EDGE1_1 secured E-Line exception .4086 to EDGE3_1 .4087 to EDGE2_1 .4090 to EDGE3_2
WAN Interfaces (IPsec)	Te0/0/2 172.16.23.2/30 Tu212 to EDGE1_1 Tu222 to EDGE2_2 Tu252 to EDGE3_1 Tu261 to EDGE3_2
LAN Interfaces (Critical VRF)	Te0/0/4.101 198.18.22.1/24 Lo100 10.22.22.1/24 IPv6 2001:DB8:22:22::1/64
LAN Interfaces (Non-Critical VRF)	Te0/0/4.201 198.19.122.1/24 Lo210 10.122.122.1/24 IPv6 2001:DB8:122:122::1/64
Locator	SLOC 2001:DB8:AC22::/48
BGP ASN	65000

Edge1_1 specifications

Table 9. Edge1_1 Specifications

Parameter	Edge1_1
WAN Interfaces (MACSEC)	No native WAN MACsec endpoint. Uses secured E-Line exception on Gi0/0/0 with VLANs .4081, .4082, .4083, .4094. GRE/IPsec overlay protection is used for the secured exception tunnels.
WAN Interfaces (IPsec)	Gi0/0/1 172.16.31.2/30 to INET. Tu210 to HUB1_2, Tu211 to HUB2_1, Tu212 to HUB2_2, Tu213 to HUB1_1.
LAN Interfaces (Critical VRF)	Gi0/1/2 trunk to LAN switch. Vlan101 198.18.31.1/24. Lo100 10.31.31.1/24. Multicast receiver scale: Vlan102-150, 198.18.102.1-198.18.150.1/24.
LAN Interfaces (Non-Critical VRF)	Gi0/1/2 trunk to LAN switch. Vlan201 198.19.31.1/24. Lo210 10.131.131.1/24. Multicast receiver scale: Vlan202-250, 198.19.202.1-198.19.250.1/24.
Locator	SLOC 2001:DB8:AC31::/48
BGP ASN	65000

Edge 2_1 specifications

Table 10. Edge 2_1 Specifications

Parameter	Edge2_1
WAN Interfaces (MACsec)	Te0/0/4 parent. .4087 to HUB2_2, .4088 to HUB2_1, .4089 to HUB1_2, .4093 to HUB1_1.
WAN Interfaces (IPsec)	Not Applicable
WAN Cross-Connect Link	Te0/0/5 reserved/direct inter-edge link; no routed IP in current config.
LAN Interfaces (Critical VRF)	Tw0/0/0.101 198.18.41.2/24. VRRP41 VIP 198.18.41.1, priority 150. Lo100 10.41.41.1/24.
LAN Interfaces (Non-Critical VRF)	Tw0/0/0.201 198.19.42.2/24. VRRP42 VIP 198.19.42.1, backup. Lo210 10.142.142.1/24.
Locator	SLOC 2001:DB8:AC41::/48
VRRP	Critical VRF primary. Non-critical VRF backup.
BGP ASN	65000

Edge 2_2 specifications

Table 11. Edge 2_2 Specifications

Parameter	Edge2_2
WAN Interfaces (MACsec)	Not Applicable
WAN Interfaces (IPsec)	Te0/0/4 172.16.22.2/30 to INET. Tu200 to HUB1_1, Tu220 to HUB1_2, Tu221 to HUB2_1, Tu222 to HUB2_2.
WAN Cross-Connect Link	Te0/0/5 reserved/direct inter-edge link; no routed IP in current config.
LAN Interfaces (Critical VRF)	Tw0/0/0.101 198.18.41.3/24. VRRP41 VIP 198.18.41.1, default priority.
LAN Interfaces (Non-Critical VRF)	Tw0/0/0.201 198.19.42.3/24. VRRP42 VIP 198.19.42.1, priority 150. Lo210 10.42.42.1/24.
Locator	SLOC 2001:DB8:AC42::/48
VRRP	Critical VRF secondary. Non-critical VRF primary.
BGP ASN	65000

Edge 3_1 specifications

Table 12. Edge 3_1 Specifications

Parameter	Edge3_1
WAN Interfaces (MACsec)	Te0/0/8 parent. .4084 to HUB1_1, .4085 to HUB2_1, .4086 to HUB2_2, .4091 to HUB1_2.
WAN Interfaces (IPsec)	Te0/0/9 172.16.51.2/30 to INET. Tu250 to HUB1_2, Tu251 to HUB2_1, Tu252 to HUB2_2, Tu253 to HUB1_1.
LAN Interfaces (Critical VRF)	Te0/0/11 198.18.51.1/24 to EDGE3_X Te0/1/0. EDGE3_X client VLAN101 is 198.18.53.1/24 via OSPF.
LAN Interfaces (Non-Critical VRF)	Tunnel320 198.19.151.1/30 to EDGE3_X 198.19.151.2. Lo210 10.151.151.1/24. EDGE3_X client VLAN201 is 198.19.53.1/24.
Locator	SLOC 2001:DB8:AC51::/48
BGP ASN	65000

Edge 3_2 specifications

Table 13. Edge 3_2 Specifications

Parameter	Edge3_2
WAN Interfaces (MACsec)	Te0/1/0 parent. .4078 to HUB1_1, .4079 to HUB2_1, .4090 to HUB2_2, .4092 to HUB1_2.
WAN Interfaces (IPsec)	Te0/1/1 172.16.52.2/30 to INET. Tu260 to HUB1_2, Tu261 to HUB2_2, Tu262 to HUB1_1, Tu263 to HUB2_1.

Parameter	Edge3_2
LAN Interfaces (Critical VRF)	Te0/1/3 198.18.52.1/24 to EDGE3_X Te0/1/1. EDGE3_X client VLAN101 is 198.18.53.1/24 via OSPF.
LAN Interfaces (Non-Critical VRF)	Tunnel321 198.19.152.1/30 to EDGE3_X 198.19.152.2. Lo210 10.152.152.1/24. EDGE3_X client VLAN201 is 198.19.53.1/24.
Locator	SLOC 2001:DB8:AC52::/48
BGP ASN	65000

Datacenter aggregation hub configuration

Hub1_1

```
HUB1_1#sh run
Building configuration...

Current configuration : 32488 bytes
!
! Last configuration change at 03:40:23 UTC Fri Jul 24 2026 by admin
!
version 26.1
system mode insecure
service timestamps debug datetime msec
service timestamps log datetime msec
service internal
platform qfp utilization monitor load 80
!
hostname HUB1_1
!
boot-start-marker
boot system bootflash:c8000aep-universalk9.26.01.02prd6.SPA.bin
! Warning: Booting with bundle mode will be deprecated in the near future. Migration to install mode is required.
boot system bootflash:/c8000aep-universalk9.BLD_V261_1_THROTTLE_LATEST_20260305_183234_V26_1_0_54.SSA.bin
! Warning: Booting with bundle mode will be deprecated in the near future. Migration to install mode is required.
boot-end-marker
!
!
vrf definition Mgmt-intf
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
 exit-address-family
```

```
!
vrf definition CRITICAL-VRF
 rd 65000:11
 route-target export 65000:1001
 route-target import 65000:1001
 !
 address-family ipv4
  srv6-mcast ingress-replication partitioned
  route-target export 65000:1001
  route-target import 65000:1001
  route-target export 65000:1001 stitching
  route-target import 65000:1001 stitching
 exit-address-family
 !
 address-family ipv6
  route-target export 65000:1001
  route-target import 65000:1001
 exit-address-family
!
vrf definition NON-CRITICAL-VRF
 rd 65000:111
 route-target export 65000:2001
 route-target import 65000:2001
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
 exit-address-family
!
no logging console
aaa new-model
!
!
aaa authentication login default local
aaa authentication enable default enable
aaa authorization console
aaa authorization exec default local
!
!
aaa session-id common
!
!
subscriber templating
!
ip multicast-routing vrf CRITICAL-VRF distributed
ip multicast-routing vrf NON-CRITICAL-VRF distributed
```

```
!
!
!
!
!
!
!
!
!
!
login on-success log
!
!
!
!
!
!
ipv6 unicast-routing
!
!
!
!
!
!
!
!
key chain KC-WAN-MACSEC macsec
 key 01
  cryptographic-algorithm aes-256-cmac
  key-string 3637fc51ed2f71b5a354e1aca9a91adb0f5a6744f1cda2cc7e18417cbd025e51
  lifetime local 00:00:00 May 13 2026 infinite
product-analytics
!
crypto pki trustpoint TP-self-signed-3951335788
 enrollment selfsigned
 subject-name cn=IOS-Self-Signed-Certificate-3951335788
 revocation-check none
 rsakeypair TP-self-signed-3951335788
 hash sha512
!
crypto pki trustpoint SLA-TrustPoint
 enrollment pkcs12
 revocation-check crl
 hash sha512
!
!
```

```
crypto pki certificate chain TP-self-signed-3951335788
 certificate self-signed 01
  30820330 30820218 A0030201 02020101 300D0609 2A864886 F70D0101 0D050030
  31312F30 2D060355 04030C26 494F532D 53656C66 2D536967 6E65642D 43657274
  69666963 6174652D 33393531 33333537 3838301E 170D3236 30353131 30393334
  35335A17 0D333630 35313030 39333435 335A3031 312F302D 06035504 030C2649
  4F532D53 656C662D 5369676E 65642D43 65727469 66696361 74652D33 39353133
  33353738 38308201 22300D06 092A8648 86F70D01 01010500 0382010F 00308201
  0A028201 010093DC 02B11083 F34593AD 7246A921 317E476F C397DC87 36BA8E37
  2A112D3F 44EF05E4 B0788F61 01E6BFF7 83700EB7 137DB918 CBFB6B77 1E9DE456
  AC9D29BD E31FBC5B 47567499 65ED3A26 36A5A161 BE338EC7 A0631FA6 B9EA1EF8
  9F442371 4F5EE2C6 11D136C7 C0F41BAA EFB9C072 CB2C03B3 A67E876D 8E4DBE65
  8E1F8827 29BC031F 1160AACD 8AFC09B1 C518F01C E8CCF6B9 27107DBF 27B35E4D
  32FF053D FB23E125 D74514F5 46156DF8 52B19A42 B6C15A28 4B9B957A 569FB610
  1D0ECE95 A1CBD2EC 63A57646 E77A79AA BDB9B2FF 2E268186 31DF81E2 B77ABE98
  951D822C 26BDA272 98B2CE55 4E779141 7281F0FC 6B5EF06E 9C0F0D96 AE95FD83
  BB221849 E52B0203 010001A3 53305130 1D060355 1D0E0416 0414C2A7 6C57986F
  2FFD74D1 DC42AE4D 6BBADA23 4DD8301F 0603551D 23041830 168014C2 A76C5798
  6F2FFD74 D1DC42AE 4D6BBADA 234DD830 0F060355 1D130101 FF040530 030101FF
  300D0609 2A864886 F70D0101 0D050003 82010100 48E121E1 62CD621F 6FFBC95C
  266C1450 8BCFF269 BFF572F5 9F900520 B1321A9F 8CA3CC40 930BB7D6 7CFCFD83
  CBDB8723 8067366D 9EDC6C10 482370BF C758602A 955B7846 2CE011F7 1A5CFAB8
  4FB3CDAC 04A0607A D31C6419 08FE70E2 9F5BE7AF 8E139A26 0D19843C 06B6CBCB
  1C60FF96 9BE0467E 0DB924EE 076818EC BC8CCA52 BDCDEB21 44924A0E CFAFCD13
  2E373FD4 EFF92006 B2A022AC F08874F7 481FF2AA 68D73414 88D0CC75 B198C611
  13BD79C5 3E19F501 402B35BB 5231F848 633EF9C8 62AF5B57 96D78EFB 5E6EE071
  8CF879F6 D01F95BB C4B9288A F73F509B 2F9B09A5 EA264260 15560426 59F3A7E4
  61145E95 DF08E7C1 89056475 571E22C5 4A7FD922
          quit
crypto pki certificate chain SLA-TrustPoint
 certificate ca 01
  30820321 30820209 A0030201 02020101 300D0609 2A864886 F70D0101 0B050030
  32310E30 0C060355 040A1305 43697363 6F312030 1E060355 04031317 43697363
  6F204C69 63656E73 696E6720 526F6F74 20434130 1E170D31 33303533 30313934
  3834375A 170D3338 30353330 31393438 34375A30 32310E30 0C060355 040A1305
  43697363 6F312030 1E060355 04031317 43697363 6F204C69 63656E73 696E6720
  526F6F74 20434130 82012230 0D06092A 864886F7 0D010101 05000382 010F0030
  82010A02 82010100 A6BCBD96 131E05F7 145EA72C 2CD686E6 17222EA1 F1EFF64D
  CBB4C798 212AA147 C655D8D7 9471380D 8711441E 1AAF071A 9CAE6388 8A38E520
  1C394D78 462EF239 C659F715 B98C0A59 5BBB5CBD 0CFEBEA3 700A8BF7 D8F256EE
  4AA4E80D DB6FD1C9 60B1FD18 FFC69C96 6FA68957 A2617DE7 104FDC5F EA2956AC
  7390A3EB 2B5436AD C847A2C5 DAB553EB 69A9A535 58E9F3E3 C0BD23CF 58BD7188
  68E69491 20F320E7 948E71D7 AE3BCC84 F10684C7 4BC8E00F 539BA42B 42C68BB7
  C7479096 B4CB2D62 EA2F505D C7B062A4 6811D95B E8250FC4 5D5D5FB8 8F27D191
  C55F0D76 61F9A4CD 3D992327 A8BB03BD 4E6D7069 7CBADF8B DF5F4368 95135E44
  DFC7C6CF 04DD7FD1 02030100 01A34230 40300E06 03551D0F 0101FF04 04030201
  06300F06 03551D13 0101FF04 05300301 01FF301D 0603551D 0E041604 1449DC85
```

```
  4B3D31E5 1B3E6A17 606AF333 3D3B4C73 E8300D06 092A8648 86F70D01 010B0500
  03820101 00507F24 D3932A66 86025D9F E838AE5C 6D4DF6B0 49631C78 240DA905
  604EDCDE FF4FED2B 77FC460E CD636FDB DD44681E 3A5673AB 9093D3B1 6C9E3D8B
  D98987BF E40CBD9E 1AECA0C2 2189BB5C 8FA85686 CD98B646 5575B146 8DFC66A8
  467A3DF4 4D565700 6ADF0F0D CF835015 3C04FF7C 21E878AC 11BA9CD2 55A9232C
  7CA7B7E6 C1AF74F6 152E99B7 B1FCF9BB E973DE7F 5BDDEB86 C71E3B49 1765308B
  5FB0DA06 B92AFE7F 494E8A9E 07B85737 F3A58BE1 1A48A229 C37C1E69 39F08678
  80DDCD16 D6BACECA EEBC7CF9 8428787B 35202CDC 60E4616A B623CDBD 230E3AFB
  418616A9 4093E049 4D10AB75 27E86F73 932E35B5 8862FDAE 0275156F 719BB2F0
  D697DF7F 28
           quit
!
!
!
!
!
!
!
!
!
license udi pid C8570-G2 sn FCB3013D188
license boot level advantage
!
memory free low-watermark processor 3199455
hw-module subslot 0/1 mode 10G
hw-module subslot 0/2 mode 40G
hw-module subslot 0/2 breakout none port all
diagnostic bootup level minimal
!
spanning-tree extend system-id
!
mka policy POL-WAN-MACSEC
!
!
!
!
username lab privilege 15 secret 9 $9$Al2Po6UnO/2meU$X4dU4QYCwO73hO2DbvRWSJmkive/00s2VfHE8UcpQ.g
username admin privilege 15 secret 9
$9$7hi99LSMVCqkYU$o1ZoIDj4ZIAO/kxaszJpy3JMKeVzzMNdUGgu.06Z3cM
!
redundancy
 mode none
!
!
!
crypto ikev2 proposal U-IKEV2-PROP
 pqc mlkem1024 optional
```

```
 encryption aes-gcm-256
 prf sha256
 group 19
crypto ikev2 proposal UC2-IKEV2-PROP
 pqc mlkem1024 optional
 encryption aes-gcm-256
 prf sha256
 group 19
!
crypto ikev2 policy U-IKEV2-POLICY
 proposal U-IKEV2-PROP
crypto ikev2 policy UC2-IKEV2-POLICY
 proposal UC2-IKEV2-PROP
!
crypto ikev2 keyring U-PSK-KR
 peer EDGE2_2
  address 172.16.22.2
  pre-shared-key cisco123!
 !
 peer EDGE3_1
  address 172.16.51.2
  pre-shared-key cisco123!
 !
 peer EDGE3_2
  address 172.16.52.2
  pre-shared-key cisco123!
 !
 peer EDGE1_1
  address 172.16.31.2
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring U-PPK-KR
 peer EDGE2_2
  address 172.16.22.2
  ppk manual id UNSEC-E22-H11 key pqcReady123! required
 !
 peer EDGE3_1
  address 172.16.51.2
  ppk manual id UNSEC-E31-H11 key pqcReady123! required
 !
 peer EDGE3_2
  address 172.16.52.2
  ppk manual id UNSEC-E32-H11 key pqcReady123! required
 !
 peer EDGE1_1
  address 172.16.31.2
```

```
  ppk manual id UNSEC-E11-H11 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-E11-PSK-KR
 peer EDGE1_1
  address 2001:DB8:3100:1::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-E11-PPK-KR
 peer EDGE1_1
  address 2001:DB8:3100:1::2/128
  ppk manual id UC2-E11-H11 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-E31-PSK-KR
 peer EDGE1_1
  address 2001:DB8:3100:1::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-E31-PPK-KR
 peer EDGE1_1
  address 2001:DB8:3100:1::2/128
  ppk manual id UC2-E31-H11 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-E41-PSK-KR
 peer EDGE2_1
  address 2001:DB8:3100:2::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-E41-PPK-KR
 peer EDGE2_1
  address 2001:DB8:3100:2::2/128
  ppk manual id UC2-E41-H11 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-E51-PSK-KR
 peer EDGE3_1
  address 2001:DB8:3100:3::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-E51-PPK-KR
 peer EDGE3_1
```

```
  address 2001:DB8:3100:3::2/128
  ppk manual id UC2-E51-H11 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-E52-PSK-KR
 peer EDGE3_2
  address 2001:DB8:3100:4::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-E52-PPK-KR
 peer EDGE3_2
  address 2001:DB8:3100:4::2/128
  ppk manual id UC2-E52-H11 key pqcReady123! required
 !
!
crypto ikev2 keyring UC4-E31-PSK-KR
 peer EDGE1_1
  address 172.31.94.1
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC4-E31-PPK-KR
 peer EDGE1_1
  address 172.31.94.1
  ppk manual id UC4-E31-H11 key pqcReady123! required
 !
!
!
crypto ikev2 profile U-IKEV2-PROFILE
 match identity remote address 172.16.22.2 255.255.255.255
 match identity remote address 172.16.51.2 255.255.255.255
 match identity remote address 172.16.52.2 255.255.255.255
 match identity remote address 172.16.31.2 255.255.255.255
 identity local address 172.16.11.2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk U-PPK-KR
 keyring local U-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-E11-IKEV2-PROFILE
 match identity remote address 2001:DB8:3100:1::2/128
 identity local address 2001:DB8:2100:1::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-E11-PPK-KR
```

```
 keyring local UC2-E11-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-E31-IKEV2-PROFILE
 match identity remote address 2001:DB8:3100:1::2/128
 identity local address 2001:DB8:2100:1::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-E31-PPK-KR
 keyring local UC2-E31-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-E41-IKEV2-PROFILE
 match identity remote address 2001:DB8:3100:2::2/128
 identity local address 2001:DB8:2100:1::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-E41-PPK-KR
 keyring local UC2-E41-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-E51-IKEV2-PROFILE
 match identity remote address 2001:DB8:3100:3::2/128
 identity local address 2001:DB8:2100:1::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-E51-PPK-KR
 keyring local UC2-E51-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-E52-IKEV2-PROFILE
 match identity remote address 2001:DB8:3100:4::2/128
 identity local address 2001:DB8:2100:1::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-E52-PPK-KR
 keyring local UC2-E52-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC4-E31-IKEV2-PROFILE
 match identity remote address 172.31.94.1 255.255.255.255
 identity local address 172.31.94.2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC4-E31-PPK-KR
 keyring local UC4-E31-PSK-KR
 dpd 10 2 on-demand
```

```
!
crypto ikev2 fragmentation mtu 1400
!
!
cdp run
!
!
class-map match-any CM-BUSINESS
 match dscp af31
class-map match-any CM-PILOT-VIDEO
 match dscp af41
class-map match-any CM-PILOT-VOICE
 match dscp ef
class-map match-any CM-CRITICAL
 match dscp ef
class-map type inspect match-any CM-PILOT-NGFW-ALLOW
 match protocol icmp
 match protocol tcp
 match protocol udp
 match access-group name IPV4-PILOT-NGFW-ALLOW
 match access-group name IPV6-PILOT-NGFW-ALLOW
class-map match-any CM-PILOT-CRITICAL
 match dscp af31
class-map match-any CM-PILOT-BULK
 match dscp af11
!
policy-map PM-PILOT-SECURED-QOS
 class CM-PILOT-VOICE
  priority percent 10
 class CM-PILOT-VIDEO
  bandwidth percent 25
 class CM-PILOT-CRITICAL
  bandwidth percent 30
 class CM-PILOT-BULK
  bandwidth percent 10
 class class-default
  fair-queue
policy-map type epbr PM-PFP
 class CM-CRITICAL
  set forward-class 1
 class CM-BUSINESS
  set forward-class 2
 class class-default
  set forward-class 0
policy-map type inspect PM-PILOT-NGFW
 class type inspect CM-PILOT-NGFW-ALLOW
  inspect
```

```
 class class-default
  drop log
policy-map type inspect PM-PILOT-NGFW-PASS
 class type inspect CM-PILOT-NGFW-ALLOW
  pass
 class class-default
  drop log
!
!
zone security Z-PILOT-LAN
zone security Z-PILOT-SRV6-WAN
zone security Z-PILOT-LAN-U
zone-pair security ZP-LAN-TO-SRV6-WAN source Z-PILOT-LAN destination Z-PILOT-SRV6-WAN
 service-policy type inspect PM-PILOT-NGFW
zone-pair security ZP-LANU-TO-SRV6-WAN source Z-PILOT-LAN-U destination Z-PILOT-SRV6-WAN
 service-policy type inspect PM-PILOT-NGFW
zone-pair security ZP-SRV6-WAN-TO-LAN source Z-PILOT-SRV6-WAN destination Z-PILOT-LAN
 service-policy type inspect PM-PILOT-NGFW
zone-pair security ZP-SRV6-WAN-TO-LANU source Z-PILOT-SRV6-WAN destination Z-PILOT-LAN-U
 service-policy type inspect PM-PILOT-NGFW
!
!
!
!
!
!
!
crypto ipsec transform-set U-TS esp-gcm 256
 mode transport
crypto ipsec transform-set UC2-TS esp-gcm 256
 mode transport
!
crypto ipsec profile U-E32-H11-IPSEC
 set transform-set U-TS
 set ikev2-profile U-IKEV2-PROFILE
!
crypto ipsec profile U-IPSEC
 set transform-set U-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile U-IKEV2-PROFILE
!
crypto ipsec profile UC2-E11-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-E11-IKEV2-PROFILE
!
crypto ipsec profile UC2-E31-IPSEC
```

```
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-E31-IKEV2-PROFILE
!
crypto ipsec profile UC2-E41-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-E41-IKEV2-PROFILE
!
crypto ipsec profile UC2-E51-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-E51-IKEV2-PROFILE
!
crypto ipsec profile UC2-E52-IPSEC
 set transform-set UC2-TS
 set pfs group19
 set ikev2-profile UC2-E52-IKEV2-PROFILE
!
crypto ipsec profile UC4-E31-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC4-E31-IKEV2-PROFILE
!
!
!
!
!
!
!
!
!
!
interface Loopback0
 description CNI_UC2_BGP_UPDATE_SOURCE
 ip address 10.255.0.11 255.255.255.255
 ip proxy-arp
 ipv6 address 2001:DB8:100:11::1/128
 ipv6 router isis SRTE-OVLY
!
interface Loopback100
 vrf forwarding CRITICAL_VRF
 ip address 10.11.11.1 255.255.255.0
 ip proxy-arp
 ipv6 address 2001:DB8:11:11::1/64
!
interface Loopback110
```

```
 description ANYCAST_RP_SHARED_HUB1_PAIR_CRITICAL_VRF
 vrf forwarding CRITICAL_VRF
 ip address 10.255.11.100 255.255.255.255
 ip proxy-arp
 ip pim sparse-mode
!
interface Loopback111
 description ANYCAST_RP_MSDP_UNIQUE_HUB1_1_CRITICAL_VRF
 vrf forwarding CRITICAL_VRF
 ip address 10.255.11.11 255.255.255.255
 ip proxy-arp
 ip pim sparse-mode
!
interface Loopback200
 no ip address
 ip proxy-arp
 ipv6 address 2001:DB8:111:200::1/128
 ipv6 router isis SRTE-OVLY
!
interface Loopback210
 description HUB1_1 unsecured l3vpn service loopback
 vrf forwarding NON-CRITICAL_VRF
 ip address 10.111.111.1 255.255.255.0
 ip proxy-arp
 ipv6 address 2001:DB8:111:111::1/64
!
interface Tunnel200
 description HUB1_1 to EDGE2_2 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:200::1/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/0/2
 tunnel destination 172.16.22.2
 tunnel protection ipsec profile U-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel213
 description HUB1_1 to EDGE1_1 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
```

```
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:213::2/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/0/2
 tunnel destination 172.16.31.2
 tunnel protection ipsec profile U-IPSEC
 clns mtu 1300
 isis circuit-type level-2-only
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel253
 description HUB1_1 to EDGE3_1 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:253::2/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/0/2
 tunnel destination 172.16.51.2
 tunnel protection ipsec profile U-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel262
 description HUB1_1 to EDGE3_2 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:262::2/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/0/2
 tunnel destination 172.16.52.2
 tunnel protection ipsec profile U-E32-H11-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel381
 description ASM_PIM_HUB1_1_TO_RR_CRITICAL-VRF
```

```
 vrf forwarding CRITICAL-VRF
 ip address 172.31.131.2 255.255.255.252
 ip proxy-arp
 ip pim sparse-mode
 tunnel source Loopback0
 tunnel mode gre ipv6
 tunnel destination 2001:DB8:100:10::1
 tunnel path-mtu-discovery
!
interface Tunnel382
 description ASM_PIM_HUB1_1_TO_RR_NON-CRITICAL-VRF
 vrf forwarding NON-CRITICAL-VRF
 ip address 172.31.132.2 255.255.255.252
 ip proxy-arp
 ip pim sparse-mode
 tunnel source Loopback0
 tunnel mode gre ipv6
 tunnel destination 2001:DB8:100:10::1
 tunnel key 382
 tunnel path-mtu-discovery
!
interface Tunnel383
 description ASM_PIM_HUB1_1_TO_EDGE1_1_CRITICAL-VRF
 vrf forwarding CRITICAL-VRF
 ip address 172.31.133.1 255.255.255.252
 ip proxy-arp
 ip pim sparse-mode
 tunnel source Loopback0
 tunnel mode gre ipv6
 tunnel destination 2001:DB8:100:31::1
 tunnel path-mtu-discovery
!
interface Tunnel384
 description ASM_PIM_HUB1_1_TO_EDGE1_1_NON-CRITICAL-VRF
 vrf forwarding NON-CRITICAL-VRF
 ip address 172.31.134.1 255.255.255.252
 ip proxy-arp
 ip pim sparse-mode
 tunnel source Loopback0
 tunnel mode gre ipv6
 tunnel destination 2001:DB8:100:31::1
 tunnel key 384
 tunnel path-mtu-discovery
!
interface Tunnel4094
 description UC4_EDGE1_EXCEPTION_GRE_OVER_IPSEC_SECURED_ELINE
 no ip address
```

```
 ip proxy-arp
 ip mtu 1300
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5094:1::2/64
 ipv6 mtu 1300
 ipv6 router isis SRTE-OVLY
 tunnel source 172.31.94.2
 tunnel destination 172.31.94.1
 tunnel path-mtu-discovery
 tunnel protection ipsec profile UC4-E31-IPSEC
 clns mtu 1300
 isis circuit-type level-2-only
 isis affinity flex-algo
  name SECURED
  !
!
interface TenGigabitEthernet0/0/0
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/0/1
 description CNI_UC4_HUB1_1_WAN_MACSEC_PARENT
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
 ipv6 address 2001:DB8:2100:1::2/64
 macsec dot1q-in-clear 1
 macsec access-control should-secure
 service-policy output PM-PILOT-SECURED-QOS
!
interface TenGigabitEthernet0/0/1.4078
 description UC4_TAGGED_WAN_MACSEC_NATIVE_ISIS_TO_EDGE3_2
 encapsulation dot1Q 4078
 ip address 172.31.78.2 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:4078:1::2/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address 28af.fdb3.5188
 eapol eth-type 876F
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
```

```
 clns mtu 1400
 isis circuit-type level-2-only
 isis network point-to-point
 isis affinity flex-algo
  name SECURED
  !
!
interface TenGigabitEthernet0/0/1.4084
 description UC4_TAGGED_WAN_MACSEC_GRE_ISIS_TO_EDGE3_1
 encapsulation dot1Q 4084
 ip address 172.31.84.2 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:4084:1::2/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address ec19.2e55.5f08
 eapol eth-type 876F
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
 clns mtu 1400
 isis circuit-type level-2-only
 isis network point-to-point
 isis affinity flex-algo
  name SECURED
  !
!
interface TenGigabitEthernet0/0/1.4093
 description UC4_WAN_MACSEC_PILOT_TO_EDGE2_1_VLAN4093
 encapsulation dot1Q 4093
 ip address 172.31.93.2 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:4093:1::2/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address 5000.e06c.b2a4
 eapol eth-type 876F
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
 clns mtu 1400
 isis circuit-type level-2-only
 isis network point-to-point
 isis affinity flex-algo
  name SECURED
```

```
  !
!
interface TenGigabitEthernet0/0/1.4094
 description UC4_ELINE_PILOT_TEST_TO_EDGE1_1
 encapsulation dot1Q 4094
 ip address 172.31.94.2 255.255.255.252
 ip proxy-arp
 ipv6 address 2001:DB8:4094:1::2/64
!
interface TenGigabitEthernet0/0/2
 description CNI_UC2_INET_UNDERLAY_BLR_TUNNEL200
 ip address 172.16.11.2 255.255.255.252
 ip proxy-arp
 no negotiation auto
 cdp enable
 service-policy output PM-PILOT-SECURED-QOS
!
interface TenGigabitEthernet0/0/3
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/0/4
 description HUB1_1 WAN switch/DCI segment
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
 ipv6 address 2001:DB8:4100:1::11/64
!
interface TenGigabitEthernet0/0/4.101
 description SECURED_CRITICAL_VRF_HUB1_1
 encapsulation dot1Q 101
 vrf forwarding CRITICAL-VRF
 ip address 198.18.11.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 ip nbar protocol-discovery
 zone-member security Z-PILOT-LAN
 service-policy type epbr input PM-PFP
!
interface TenGigabitEthernet0/0/4.201
 description UNSECURED_NON_CRITICAL_VRF_HUB1_1
 encapsulation dot1Q 201
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.111.1 255.255.255.0
```

```
 ip proxy-arp
 ip pim sparse-mode
 ip nbar protocol-discovery
 zone-member security Z-PILOT-LAN-U
 service-policy type epbr input PM-PFP
!
interface TenGigabitEthernet0/0/4.511
 description WAN_MGMT511_HUB1_1
 encapsulation dot1Q 511
 ip address 192.168.12.11 255.255.255.0
 ip proxy-arp
 ipv6 address 2001:DB8:511:12::11/64
!
interface TenGigabitEthernet0/0/5
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/0/6
 description ANYCAST_RP_MSDP_DIRECT_TO_HUB1_2
 vrf forwarding CRITICAL-VRF
 ip address 10.255.11.17 255.255.255.252
 ip proxy-arp
 ip pim sparse-mode
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/0/7
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/1/0
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/1/1
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/1/2
```

```
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/1/3
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface FortyGigabitEthernet0/2/0
 no ip address
 ip proxy-arp
 shutdown
 no negotiation auto
!
interface FortyGigabitEthernet0/2/4
 no ip address
 ip proxy-arp
 shutdown
 no negotiation auto
!
interface FortyGigabitEthernet0/2/8
 no ip address
 ip proxy-arp
 shutdown
 no negotiation auto
!
interface GigabitEthernet0
 vrf forwarding Mgmt-intf
 ip address 175.13.8.32 255.255.255.0
 ip proxy-arp
 negotiation auto
!
segment-routing traffic-eng
 !
 on-demand color 100
  authorize
  candidate-paths
   preference 100
    constraints
     segments
      dataplane srv6
     !
     affinity
      include-all
```

```
       name SECURED
      !
     !
    !
    dynamic
     metric
      type te
     !
    !
   !
  !
 !
 on-demand color 200
  authorize
  candidate-paths
   preference 100
    constraints
     segments
      dataplane srv6
     !
     affinity
      include-all
       name UNSECURED
      !
     !
    !
    dynamic
     metric
      type te
     !
    !
   !
  !
 !
 on-demand color 300
  authorize
  candidate-paths
   preference 100
    per-flow
     forward-class 0 color 200
     forward-class 1 color 100
     forward-class 2 color 100
    !
   !
  !
 !
!
```

```
segment-routing srv6
 encapsulation
  source-address 2001:DB8:AC11::
 locators
  locator SLOC
   prefix 2001:DB8:AC11::/48
   format usid-f3216
!
router isis SRTE-OVLY
 net 49.0099.0099.0000.1011.00
 is-type level-2-only
 metric-style wide
 lsp-mtu 1300
 distribute link-state instance-id 99
 affinity-map SECURED bit-position 0
 affinity-map UNSECURED bit-position 1
 !
 address-family ipv6
  router-id Loopback0
  segment-routing srv6
   locator SLOC
  fast-reroute per-prefix level-2 all
  fast-reroute ti-lfa level-2
 exit-address-family
!
router bgp 65000
 bgp router-id 10.255.0.11
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 2001:DB8:100:10::1 remote-as 65000
 neighbor 2001:DB8:100:10::1 update-source Loopback0
 neighbor 2001:DB8:100:31::1 remote-as 65000
 neighbor 2001:DB8:100:31::1 update-source Loopback0
 neighbor 2001:DB8:100:41::1 remote-as 65000
 neighbor 2001:DB8:100:41::1 update-source Loopback0
 neighbor 2001:DB8:100:42::1 remote-as 65000
 neighbor 2001:DB8:100:42::1 update-source Loopback0
 neighbor 2001:DB8:100:51::1 remote-as 65000
 neighbor 2001:DB8:100:51::1 update-source Loopback0
 neighbor 2001:DB8:100:52::1 remote-as 65000
 neighbor 2001:DB8:100:52::1 update-source Loopback0
 neighbor 2001:DB8:511:12::208 remote-as 65000
 neighbor 192.168.255.208 remote-as 65000
 !
 address-family ipv4
 exit-address-family
 !
```

```
address-family ipv4 mvpn
 neighbor 2001:DB8:100:51::1 activate
 neighbor 2001:DB8:100:51::1 send-community both
 neighbor 2001:DB8:100:52::1 activate
 neighbor 2001:DB8:100:52::1 send-community both
 neighbor 2001:DB8:100:52::1 route-reflector-client
exit-address-family
!
address-family vpnv4
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
 neighbor 2001:DB8:100:10::1 activate
 neighbor 2001:DB8:100:10::1 send-community both
 neighbor 2001:DB8:100:10::1 next-hop-self all
 neighbor 2001:DB8:100:10::1 route-map RM_DCI_RR_NH_OUT out
 neighbor 2001:DB8:100:31::1 activate
 neighbor 2001:DB8:100:31::1 send-community both
 neighbor 2001:DB8:100:31::1 route-reflector-client
 neighbor 2001:DB8:100:41::1 activate
 neighbor 2001:DB8:100:41::1 send-community both
 neighbor 2001:DB8:100:41::1 route-reflector-client
 neighbor 2001:DB8:100:42::1 activate
 neighbor 2001:DB8:100:42::1 send-community both
 neighbor 2001:DB8:100:42::1 route-reflector-client
 neighbor 2001:DB8:100:51::1 activate
 neighbor 2001:DB8:100:51::1 send-community both
 neighbor 2001:DB8:100:51::1 route-reflector-client
 neighbor 2001:DB8:100:52::1 activate
 neighbor 2001:DB8:100:52::1 send-community both
 neighbor 2001:DB8:100:52::1 route-reflector-client
exit-address-family
!
address-family vpnv6
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
 neighbor 2001:DB8:100:10::1 activate
 neighbor 2001:DB8:100:10::1 send-community both
 neighbor 2001:DB8:100:10::1 next-hop-self all
 neighbor 2001:DB8:100:10::1 route-map RM_DCI_RR_NH_OUT out
```

```
 neighbor 2001:DB8:100:31::1 activate
 neighbor 2001:DB8:100:31::1 send-community both
 neighbor 2001:DB8:100:31::1 route-reflector-client
 neighbor 2001:DB8:100:41::1 activate
 neighbor 2001:DB8:100:41::1 send-community both
 neighbor 2001:DB8:100:41::1 route-reflector-client
 neighbor 2001:DB8:100:42::1 activate
 neighbor 2001:DB8:100:42::1 send-community both
 neighbor 2001:DB8:100:42::1 route-reflector-client
 neighbor 2001:DB8:100:51::1 activate
 neighbor 2001:DB8:100:51::1 send-community both
 neighbor 2001:DB8:100:51::1 route-reflector-client
 neighbor 2001:DB8:100:52::1 activate
 neighbor 2001:DB8:100:52::1 send-community both
 neighbor 2001:DB8:100:52::1 route-reflector-client
exit-address-family
!
address-family link-state link-state
 neighbor 2001:DB8:511:12::208 activate
 neighbor 2001:DB8:511:12::208 send-community both
 neighbor 192.168.255.208 activate
 neighbor 192.168.255.208 send-community both
exit-address-family
!
address-family ipv4 vrf CRITICAL-VRF
 network 10.11.11.0 mask 255.255.255.0 route-map RM_ODN_COLOR_300
 network 10.255.11.11 mask 255.255.255.255
 network 10.255.11.100 mask 255.255.255.255
 network 198.18.11.0 route-map RM_ODN_COLOR_300
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
exit-address-family
!
address-family ipv6 vrf CRITICAL-VRF
 network 2001:DB8:11:11::/64 route-map RM_ODN_COLOR_300
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
exit-address-family
!
```

```
 address-family ipv4 vrf NON-CRITICAL-VRF
  network 10.111.111.0 mask 255.255.255.0
  network 198.19.111.0
  !
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
 exit-address-family
 !
 address-family ipv6 vrf NON-CRITICAL-VRF
  network 2001:DB8:111:111::/64
  !
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
 exit-address-family
!
ip rcmd domain-lookup
ip forward-protocol nd
ip forward-protocol udp tftp
!
ip telnet comport enable
ip pim vrf CRITICAL-VRF rp-address 10.255.13.100 ACL-ASM-PILOT-GROUPS
ip pim vrf CRITICAL-VRF rp-address 10.255.11.100 ACL-ASM-ANYCAST-RP-PILOT
ip pim vrf NON-CRITICAL-VRF rp-address 10.255.13.200 ACL-ASM-PILOT-GROUPS
ip msdp vrf CRITICAL-VRF peer 10.255.11.12 connect-source Loopback111 remote-as 65000
ip msdp vrf CRITICAL-VRF originator-id Loopback111
ip tftp source-interface GigabitEthernet0
ip tftp blocksize 512
ip ftp passive
ip http server
ip http authentication local
ip http secure-server
ip ssh bulk-mode 131072
ip route 172.16.22.2 255.255.255.255 172.16.11.1
ip route 172.16.31.2 255.255.255.255 172.16.11.1
ip route 172.16.51.2 255.255.255.255 172.16.11.1
ip route 172.16.52.2 255.255.255.255 172.16.11.1
ip route 192.168.255.208 255.255.255.255 192.168.12.208
ip route vrf Mgmt-intf 0.0.0.0 0.0.0.0 175.13.8.1
ip route vrf Mgmt-intf 202.153.144.0 255.255.255.0 175.13.8.1
ip route vrf CRITICAL-VRF 10.255.11.12 255.255.255.255 10.255.11.18 name MSDP_HUB12_DIRECT
ip route vrf CRITICAL-VRF 10.255.13.100 255.255.255.255 Tunnel381
```

```
ip route vrf NON-CRITICAL-VRF 10.255.13.200 255.255.255.255 Tunnel382
ip route vrf NON-CRITICAL-VRF 198.19.0.0 255.255.0.0 Tunnel384
!
ip access-list standard ACL-ASM-ANYCAST-RP-PILOT
 10 permit 239.19.10.0 0.0.0.255
ip access-list standard ACL-ASM-PILOT-GROUPS
 10 permit 239.18.0.0 0.0.255.255
!
ip access-list extended IPV4-PILOT-NGFW-ALLOW
 10 permit ip any any
!
ip prefix-list PL_ODN_PILOT_HUB11 seq 10 permit 198.18.11.0/24
ip prefix-list PL_ODN_PILOT_HUB11 seq 20 permit 10.255.11.11/32
ip prefix-list PL_ODN_PILOT_HUB11 seq 30 permit 10.255.11.100/32
ipv6 route 2001:DB8:100:10::1/128 2001:DB8:4100:1::10
ipv6 route 2001:DB8:511:255::208/128 2001:DB8:511:12::208
ipv6 route 2001:DB8:3100:1::/64 2001:DB8:2100:1::1
ipv6 route 2001:DB8:3100:2::/64 2001:DB8:2100:1::1
ipv6 route 2001:DB8:3100:3::/64 2001:DB8:2100:1::1
ipv6 route 2001:DB8:3100:4::/64 2001:DB8:2100:1::1
ipv6 route 2001:DB8:AC12::/48 2001:DB8:4100:1::12 250
ipv6 route 2001:DB8:AC21::/48 2001:DB8:4100:1::21 250
ipv6 route 2001:DB8:AC22::/48 2001:DB8:4100:1::22 250
ipv6 route 2001:DB8:AC51::/48 2001:DB8:4100:1::12 250
ipv6 route 2001:DB8:AC52::/48 2001:DB8:4100:1::22 250
route-map RM_DCI_RR_NH_OUT permit 10
 set ipv6 next-hop 2001:DB8:100:11::1
!
route-map RM_ODN_COLOR_300 permit 10
 match ip address prefix-list PL_ODN_PILOT_HUB11
 set extcommunity color 300 additive
!
route-map RM_SRV6_NH_T4094_OUT permit 10
 set ipv6 next-hop 2001:DB8:5094:1::2
!
!
!
performance-measurement
!
!
!
!
ipv6 access-list IPV6-PILOT-NGFW-ALLOW
 sequence 10 permit ipv6 any any
!
control-plane
!
```

```
!
!
!
!
!
line con 0
 exec-timeout 0 0
 activation-character 13
 transport preferred ssh
 transport output ssh
 stopbits 1
line aux 0
 activation-character 13
line vty 0 4
 exec-timeout 0 0
 activation-character 13
 transport input ssh
line vty 5 15
 activation-character 13
 transport input ssh
!
ntp allow mode control 3
!
!
!
!
!
!
end
```

Hub1_2

```
HUB1_2#sh run
Building configuration...

Current configuration : 28999 bytes
!
! Last configuration change at 10:07:52 UTC Thu Jul 23 2026 by admin
!
version 26.1
system mode insecure
service timestamps debug datetime msec
service timestamps log datetime msec
platform qfp utilization monitor load 80
!
hostname HUB1_2
!
boot-start-marker
```

```
boot system bootflash:/c8000aep-
universalk9.BLD_V261_1_THROTTLE_LATEST_20260305_183234_V26_1_0_54.SSA.bin
! Warning: Booting with bundle mode will be deprecated in the near future. Migration to install
mode is required.
boot-end-marker
!
!
vrf definition Mgmt-intf
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
 exit-address-family
!
vrf definition CRITICAL-VRF
 rd 65000:12
 route-target export 65000:1001
 route-target import 65000:1001
 !
 address-family ipv4
  route-target export 65000:1001
  route-target import 65000:1001
 exit-address-family
 !
 address-family ipv6
  route-target export 65000:1001
  route-target import 65000:1001
 exit-address-family
!
vrf definition NON-CRITICAL-VRF
 rd 65000:112
 route-target export 65000:2001
 route-target import 65000:2001
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
 exit-address-family
!
no logging console
aaa new-model
!
!
aaa authentication login default local
aaa authentication enable default enable
aaa authorization console
```

```
aaa authorization exec default local
!
!
aaa session-id common
!
!
subscriber templating
!
ip multicast-routing vrf CRITICAL-VRF distributed
!
!
!
!
!
!
!
!
!
login on-success log
!
!
!
!
!
!
ipv6 unicast-routing
!
!
!
!
!
!
!
!
key chain KC-WAN-MACSEC macsec
 key 01
  cryptographic-algorithm aes-256-cmac
  key-string 3637fc51ed2f71b5a354e1aca9a91adb0f5a6744f1cda2cc7e18417cbd025e51
  lifetime local 00:00:00 Jan 1 2026 infinite
product-analytics
!
crypto pki trustpoint TP-self-signed-903902264
 enrollment selfsigned
 subject-name cn=IOS-Self-Signed-Certificate-903902264
 revocation-check none
 rsakeypair TP-self-signed-903902264
 hash sha512
```

```
!
crypto pki trustpoint SLA-TrustPoint
 enrollment pkcs12
 revocation-check crl
 hash sha512
!
!
crypto pki certificate chain TP-self-signed-903902264
 certificate self-signed 01
  3082032E 30820216 A0030201 02020101 300D0609 2A864886 F70D0101 0D050030
  30312E30 2C060355 04030C25 494F532D 53656C66 2D536967 6E65642D 43657274
  69666963 6174652D 39303339 30323236 34301E17 0D323630 35313130 39333531
  365A170D 33363035 31303039 33353136 5A303031 2E302C06 03550403 0C25494F
  532D5365 6C662D53 69676E65 642D4365 72746966 69636174 652D3930 33393032
  32363430 82012230 0D06092A 864886F7 0D010101 05000382 010F0030 82010A02
  82010100 D0968951 00E50372 0EB77C83 F736151A 4810C654 4E48E368 F004E886
  4A2D7237 8146DF02 9B269FF9 5D063F8B 3A20E3DF D3CF9AAE 94C27F48 1A4A45CE
  FA67524C 1653F452 6AD1A697 932F01BB 073438C7 3814D364 A516A6AA E7BDF205
  6594A00E E534DA3A 717042C3 8B4E2733 9692ADF1 6A4F9EEB 7C3BDE1C 16CA435D
  9938A1B9 56D4DC18 9FA31C68 0DA00D3D E4D8BB97 A52C6AAE 0C44D866 B9E39FF6
  97F1462B 51C47809 35F0531C B64C7CE8 69ADBAC2 A726C561 FABE625A 89374BF2
  8D15F101 A3162D70 1B1FA847 671C07E7 4ED3253A 50702529 380E43AE 51E5073A
  940DADE0 A7B3DF24 E97904ED 9E937744 7C759A49 68B1DAE5 F0C21BFF E7B1932D
  106F89C3 02030100 01A35330 51301D06 03551D0E 04160414 5AD71D74 29E8A722
  B02051A1 8A3C38A3 057CC806 301F0603 551D2304 18301680 145AD71D 7429E8A7
  22B02051 A18A3C38 A3057CC8 06300F06 03551D13 0101FF04 05300301 01FF300D
  06092A86 4886F70D 01010D05 00038201 01005566 3A4ADD89 CA568587 38DEB69C
  92A00D31 7FC0F8E1 22E7946B 56485BD5 21D2C432 375C127B A817EDF1 0DAE97DA
  0455E78B 31D112A9 AB512853 CD5D1EC0 B5CF60BA BDF5AAB5 4BAC3FEE 752A8940
  3F76DF7A C4EF70BD 28162059 145A55AC 61F1FEA5 8666EBAB 66B872CA C8970C62
  B1B7F6BA 6D9468E8 3F0CED1D D6CD66D7 1E66D2FE 01DB1247 B8B94C44 A34A6BA0
  87D8DA44 CBD197BB 24775FF2 AACEDBB3 B063427E 91D71D5F F083A992 D322CCB0
  0B8E100F 13879A6E 1766877D F1A538FB 015CB49D F11F00C5 60B8236A 1240D552
  FFF280F5 BFB888C4 60354B76 F176326D F25ACC0A 52F06BA5 DBCB57E0 4887BFB7
  BF261596 8441B1E3 0B0F790E D193E323 7559
          quit
crypto pki certificate chain SLA-TrustPoint
 certificate ca 01
  30820321 30820209 A0030201 02020101 300D0609 2A864886 F70D0101 0B050030
  32310E30 0C060355 040A1305 43697363 6F312030 1E060355 04031317 43697363
  6F204C69 63656E73 696E6720 526F6F74 20434130 1E170D31 33303533 30313934
  3834375A 170D3338 30353330 31393438 34375A30 32310E30 0C060355 040A1305
  43697363 6F312030 1E060355 04031317 43697363 6F204C69 63656E73 696E6720
  526F6F74 20434130 82012230 0D06092A 864886F7 0D010101 05000382 010F0030
  82010A02 82010100 A6BCBD96 131E05F7 145EA72C 2CD686E6 17222EA1 F1EFF64D
  CBB4C798 212AA147 C655D8D7 9471380D 8711441E 1AAF071A 9CAE6388 8A38E520
  1C394D78 462EF239 C659F715 B98C0A59 5BBB5CBD 0CFEBEA3 700A8BF7 D8F256EE
```

```
  4AA4E80D DB6FD1C9 60B1FD18 FFC69C96 6FA68957 A2617DE7 104FDC5F EA2956AC
  7390A3EB 2B5436AD C847A2C5 DAB553EB 69A9A535 58E9F3E3 C0BD23CF 58BD7188
  68E69491 20F320E7 948E71D7 AE3BCC84 F10684C7 4BC8E00F 539BA42B 42C68BB7
  C7479096 B4CB2D62 EA2F505D C7B062A4 6811D95B E8250FC4 5D5D5FB8 8F27D191
  C55F0D76 61F9A4CD 3D992327 A8BB03BD 4E6D7069 7CBADF8B DF5F4368 95135E44
  DFC7C6CF 04DD7FD1 02030100 01A34230 40300E06 03551D0F 0101FF04 04030201
  06300F06 03551D13 0101FF04 05300301 01FF301D 0603551D 0E041604 1449DC85
  4B3D31E5 1B3E6A17 606AF333 3D3B4C73 E8300D06 092A8648 86F70D01 010B0500
  03820101 00507F24 D3932A66 86025D9F E838AE5C 6D4DF6B0 49631C78 240DA905
  604EDCDE FF4FED2B 77FC460E CD636FDB DD44681E 3A5673AB 9093D3B1 6C9E3D8B
  D98987BF E40CBD9E 1AECA0C2 2189BB5C 8FA85686 CD98B646 5575B146 8DFC66A8
  467A3DF4 4D565700 6ADF0F0D CF835015 3C04FF7C 21E878AC 11BA9CD2 55A9232C
  7CA7B7E6 C1AF74F6 152E99B7 B1FCF9BB E973DE7F 5BDDEB86 C71E3B49 1765308B
  5FB0DA06 B92AFE7F 494E8A9E 07B85737 F3A58BE1 1A48A229 C37C1E69 39F08678
  80DDCD16 D6BACECA EEBC7CF9 8428787B 35202CDC 60E4616A B623CDBD 230E3AFB
  418616A9 4093E049 4D10AB75 27E86F73 932E35B5 8862FDAE 0275156F 719BB2F0
  D697DF7F 28
          quit
!
!
!
!
!
!
!
!
!
license udi pid C8550-G2 sn FCB3013D190
license boot level advantage
!
memory free low-watermark processor 3199455
diagnostic bootup level minimal
!
spanning-tree extend system-id
!
mka policy POL-WAN-MACSEC
!
!
!
!
username lab privilege 15 secret 9 $9$Al2Po6UnO/2meU$X4dU4QYCwO73hO2DbvRWSJmkive/00s2VfHE8UcpQ.g
username admin privilege 15 secret 9
$9$J7q4Qwprgf85Ck$UIB8iLjPFOmttOb7uLFU9SKGijvCCUfvp06k0fyz0Bc
!
redundancy
 mode none
!
```

```
!
!
crypto ikev2 proposal U-IKEV2-PROP
 pqc mlkem1024 optional
 encryption aes-gcm-256
 prf sha256
 group 19
crypto ikev2 proposal UC2-IKEV2-PROP
 pqc mlkem1024 optional
 encryption aes-gcm-256
 prf sha256
 group 19
!
crypto ikev2 policy U-IKEV2-POLICY
 proposal U-IKEV2-PROP
crypto ikev2 policy UC2-IKEV2-POLICY
 proposal UC2-IKEV2-PROP
!
crypto ikev2 keyring U-H12-PSK-KR
 peer EDGE1_1
  address 172.16.31.2
  pre-shared-key cisco123!
 !
 peer EDGE3_1
  address 172.16.51.2
  pre-shared-key cisco123!
 !
 peer EDGE3_2
  address 172.16.52.2
  pre-shared-key cisco123!
 !
 peer EDGE2_2
  address 172.16.22.2
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring U-H12-PPK-KR
 peer EDGE1_1
  address 172.16.31.2
  ppk manual id UNSEC-E11-H12 key pqcReady123! required
 !
 peer EDGE3_1
  address 172.16.51.2
  ppk manual id UNSEC-E31-H12 key pqcReady123! required
 !
 peer EDGE3_2
  address 172.16.52.2
```

```
  ppk manual id UNSEC-E32-H12 key pqcReady123! required
 !
 peer EDGE2_2
  address 172.16.22.2
  ppk manual id UNSEC-E22-H12 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-E31-PSK-KR
 peer EDGE1_1
  address 2001:DB8:3100:1::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-E31-PPK-KR
 peer EDGE1_1
  address 2001:DB8:3100:1::2/128
  ppk manual id UC2-E31-H12 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-E41-PSK-KR
 peer EDGE2_1
  address 2001:DB8:3100:2::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-E41-PPK-KR
 peer EDGE2_1
  address 2001:DB8:3100:2::2/128
  ppk manual id UC2-E41-H12 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-E51-PSK-KR
 peer EDGE3_1
  address 2001:DB8:3100:3::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-E51-PPK-KR
 peer EDGE3_1
  address 2001:DB8:3100:3::2/128
  ppk manual id UC2-E51-H12 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-E52-PSK-KR
 peer EDGE3_2
  address 2001:DB8:3100:4::2/128
  pre-shared-key cisco123!
```

```
 !
!
crypto ikev2 keyring UC2-E52-PPK-KR
 peer EDGE3_2
  address 2001:DB8:3100:4::2/128
  ppk manual id UC2-E52-H12 key pqcReady123! required
 !
!
crypto ikev2 keyring UC4-E31-H12-PSK-KR
 peer EDGE1_1
  address 172.31.81.1
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC4-E31-H12-PPK-KR
 peer EDGE1_1
  address 172.31.81.1
  ppk manual id UC4-E31-H12 key pqcReady123! required
 !
!
!
crypto ikev2 profile U-H12-IKEV2-PROFILE
 match identity remote address 172.16.31.2 255.255.255.255
 match identity remote address 172.16.51.2 255.255.255.255
 match identity remote address 172.16.52.2 255.255.255.255
 match identity remote address 172.16.22.2 255.255.255.255
 identity local address 172.16.12.2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk U-H12-PPK-KR
 keyring local U-H12-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-E31-IKEV2-PROFILE
 match identity remote address 2001:DB8:3100:1::2/128
 identity local address 2001:DB8:2100:2::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-E31-PPK-KR
 keyring local UC2-E31-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-E41-IKEV2-PROFILE
 match identity remote address 2001:DB8:3100:2::2/128
 identity local address 2001:DB8:2100:2::2
 authentication remote pre-share
 authentication local pre-share
```

```
 keyring ppk UC2-E41-PPK-KR
 keyring local UC2-E41-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-E51-IKEV2-PROFILE
 match identity remote address 2001:DB8:3100:3::2/128
 identity local address 2001:DB8:2100:2::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-E51-PPK-KR
 keyring local UC2-E51-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-E52-IKEV2-PROFILE
 match identity remote address 2001:DB8:3100:4::2/128
 identity local address 2001:DB8:2100:2::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-E52-PPK-KR
 keyring local UC2-E52-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC4-E31-H12-IKEV2-PROFILE
 match identity remote address 172.31.81.1 255.255.255.255
 identity local address 172.31.81.2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC4-E31-H12-PPK-KR
 keyring local UC4-E31-H12-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 fragmentation mtu 1400
!
!
cdp run
!
!
class-map match-any CM-BUSINESS
 match dscp af31
class-map match-any CM-PILOT-VIDEO
 match dscp af41
class-map match-any CM-PILOT-VOICE
 match dscp ef
class-map match-any CM-CRITICAL
 match dscp ef
class-map type inspect match-any CM-PILOT-NGFW-ALLOW
 match protocol icmp
```

```
 match protocol tcp
 match protocol udp
 match access-group name IPV4-PILOT-NGFW-ALLOW
 match access-group name IPV6-PILOT-NGFW-ALLOW
class-map match-any CM-PILOT-CRITICAL
 match dscp af31
class-map match-any CM-PILOT-BULK
 match dscp af11
!
policy-map PM-PILOT-SECURED-QOS
 class CM-PILOT-VOICE
  priority percent 10
 class CM-PILOT-VIDEO
  bandwidth percent 25
 class CM-PILOT-CRITICAL
  bandwidth percent 30
 class CM-PILOT-BULK
  bandwidth percent 10
 class class-default
  fair-queue
policy-map type epbr PM-PFP
 class CM-CRITICAL
  set forward-class 1
 class CM-BUSINESS
  set forward-class 2
 class class-default
  set forward-class 0
policy-map type inspect PM-PILOT-NGFW
 class type inspect CM-PILOT-NGFW-ALLOW
  inspect
 class class-default
  drop log
policy-map type inspect PM-PILOT-NGFW-PASS
 class type inspect CM-PILOT-NGFW-ALLOW
  pass
 class class-default
  drop log
!
!
zone security Z-PILOT-LAN
zone security Z-PILOT-SRV6-WAN
zone security Z-PILOT-LAN-U
zone-pair security ZP-LAN-TO-SRV6-WAN source Z-PILOT-LAN destination Z-PILOT-SRV6-WAN
 service-policy type inspect PM-PILOT-NGFW
zone-pair security ZP-LANU-TO-SRV6-WAN source Z-PILOT-LAN-U destination Z-PILOT-SRV6-WAN
 service-policy type inspect PM-PILOT-NGFW
zone-pair security ZP-SRV6-WAN-TO-LAN source Z-PILOT-SRV6-WAN destination Z-PILOT-LAN
```

```
 service-policy type inspect PM-PILOT-NGFW
zone-pair security ZP-SRV6-WAN-TO-LANU source Z-PILOT-SRV6-WAN destination Z-PILOT-LAN-U
 service-policy type inspect PM-PILOT-NGFW
!
!
!
!
!
!
!
crypto ipsec transform-set U-TS esp-gcm 256
 mode transport
crypto ipsec transform-set UC2-TS esp-gcm 256
 mode transport
!
crypto ipsec profile U-H12-IPSEC
 set transform-set U-TS
 set ikev2-profile U-H12-IKEV2-PROFILE
!
crypto ipsec profile U-H12-MLKEM-IPSEC
 set transform-set U-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile U-H12-IKEV2-PROFILE
!
crypto ipsec profile UC2-E31-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-E31-IKEV2-PROFILE
!
crypto ipsec profile UC2-E41-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-E41-IKEV2-PROFILE
!
crypto ipsec profile UC2-E51-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-E51-IKEV2-PROFILE
!
crypto ipsec profile UC2-E52-IPSEC
 set transform-set UC2-TS
 set pfs group19
 set ikev2-profile UC2-E52-IKEV2-PROFILE
!
crypto ipsec profile UC4-E31-H12-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
```

```
 set ikev2-profile UC4-E31-H12-IKEV2-PROFILE
!
!
!
!
!
!
!
!
!
!
interface Loopback0
 description CNI_UC2_BGP_UPDATE_SOURCE
 no ip address
 ip proxy-arp
 ipv6 address 2001:DB8:100:12::1/128
 ipv6 router isis SRTE-OVLY
!
interface Loopback100
 vrf forwarding CRITICAL-VRF
 ip address 10.12.12.1 255.255.255.0
 ip proxy-arp
 ipv6 address 2001:DB8:12:12::1/64
!
interface Loopback110
 description ANYCAST_RP_SHARED_HUB1_PAIR_CRITICAL_VRF
 vrf forwarding CRITICAL-VRF
 ip address 10.255.11.100 255.255.255.255
 ip proxy-arp
 ip pim sparse-mode
!
interface Loopback112
 description ANYCAST_RP_MSDP_UNIQUE_HUB1_2_CRITICAL_VRF
 vrf forwarding CRITICAL-VRF
 ip address 10.255.11.12 255.255.255.255
 ip proxy-arp
 ip pim sparse-mode
!
interface Loopback200
 no ip address
 ip proxy-arp
 ipv6 address 2001:DB8:112:200::1/128
 ipv6 router isis SRTE-OVLY
!
interface Loopback210
 vrf forwarding NON-CRITICAL-VRF
 ip address 10.112.112.1 255.255.255.0
```

```
 ip proxy-arp
 ipv6 address 2001:DB8:112:112::1/64
!
interface Tunnel210
 description HUB1_2 to EDGE1_1 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:210::2/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/0/2
 tunnel destination 172.16.31.2
 tunnel protection ipsec profile U-H12-MLKEM-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel220
 description HUB1_2 to EDGE2_2 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:220::1/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/0/2
 tunnel destination 172.16.22.2
 tunnel protection ipsec profile U-H12-MLKEM-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel250
 description HUB1_2 to EDGE3_1 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:250::2/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/0/2
 tunnel destination 172.16.51.2
 tunnel protection ipsec profile U-H12-MLKEM-IPSEC
```

```
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel260
 description HUB1_2 to EDGE3_2 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:260::2/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/0/2
 tunnel destination 172.16.52.2
 tunnel protection ipsec profile U-H12-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel4081
 description UC4_EDGE1_EXCEPTION_GRE_OVER_IPSEC_SECURED_ELINE_TO_EDGE1_1
 no ip address
 ip proxy-arp
 ip mtu 1300
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5081:1::2/64
 ipv6 mtu 1300
 ipv6 router isis SRTE-OVLY
 tunnel source 172.31.81.2
 tunnel destination 172.31.81.1
 tunnel path-mtu-discovery
 tunnel protection ipsec profile UC4-E31-H12-IPSEC
 clns mtu 1300
 isis circuit-type level-2-only
 isis affinity flex-algo
  name SECURED
  !
!
interface TenGigabitEthernet0/0/0
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/0/1
 description CNI_UC4_HUB1_2_WAN_MACSEC_PARENT
```

```
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
 macsec dot1q-in-clear 1
 macsec access-control should-secure
 service-policy output PM-PILOT-SECURED-QOS
!
interface TenGigabitEthernet0/0/1.4081
 description UC4_ELINE_EXCEPTION_TO_EDGE1_1
 encapsulation dot1Q 4081
 ip address 172.31.81.2 255.255.255.252
 ip proxy-arp
 ipv6 address 2001:DB8:4081:1::2/64
!
interface TenGigabitEthernet0/0/1.4089
 description UC4_HUB1_2_TO_EDGE2_1_WAN_MACSEC_GRE_ISIS_VLAN4089
 encapsulation dot1Q 4089
 ip address 172.31.89.2 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:4089:1::2/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address 5000.e06c.b2a4
 eapol eth-type 876F
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
 clns mtu 1400
 isis circuit-type level-2-only
 isis network point-to-point
 isis affinity flex-algo
  name SECURED
  !
!
interface TenGigabitEthernet0/0/1.4091
 description UC4_TAGGED_WAN_MACSEC_GRE_ISIS_TO_PE
 encapsulation dot1Q 4091
 ip address 172.31.91.2 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:4091:1::2/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address ec19.2e55.5f08
 eapol eth-type 876F
```

```
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
 bfd interval 500 min_rx 500 multiplier 3
 clns mtu 1400
 isis circuit-type level-2-only
 isis network point-to-point
 isis bfd
 isis affinity flex-algo
  name SECURED
  !
!
interface TenGigabitEthernet0/0/1.4092
 description UC4_TAGGED_WAN_MACSEC_GRE_ISIS_TO_EDGE3_2
 encapsulation dot1Q 4092
 ip address 172.31.92.2 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:4092:1::2/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address 28af.fdb3.5188
 eapol eth-type 876F
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
 clns mtu 1400
 isis circuit-type level-2-only
 isis network point-to-point
 isis affinity flex-algo
  name SECURED
  !
!
interface TenGigabitEthernet0/0/2
 description CNI_UC2_INET_UNDERLAY_BLR_TUNNEL220
 ip address 172.16.12.2 255.255.255.252
 ip proxy-arp
 no negotiation auto
 cdp enable
 ipv6 address 2001:DB8:2100:2::2/64
 service-policy output PM-PILOT-SECURED-QOS
!
interface TenGigabitEthernet0/0/3
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
```

```
!
interface TenGigabitEthernet0/0/4
 description HUB1_2 WAN switch/DCI segment
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
 ipv6 address 2001:DB8:4100:1::12/64
!
interface TenGigabitEthernet0/0/4.101
 description SECURED_CRITICAL_VRF_HUB1_2
 encapsulation dot1Q 101
 vrf forwarding CRITICAL-VRF
 ip address 198.18.12.1 255.255.255.0
 ip proxy-arp
 ip nbar protocol-discovery
 zone-member security Z-PILOT-LAN
 service-policy type epbr input PM-PFP
!
interface TenGigabitEthernet0/0/4.201
 description UNSECURED_NON_CRITICAL_VRF_HUB1_2
 encapsulation dot1Q 201
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.112.1 255.255.255.0
 ip proxy-arp
 ip nbar protocol-discovery
 zone-member security Z-PILOT-LAN-U
 service-policy type epbr input PM-PFP
!
interface TenGigabitEthernet0/0/4.511
 description WAN_MGMT511_HUB1_2
 encapsulation dot1Q 511
 ip address 192.168.12.12 255.255.255.0
 ip proxy-arp
!
interface TenGigabitEthernet0/0/5
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/0/6
 description ANYCAST_RP_MSDP_DIRECT_TO_HUB1_1
 vrf forwarding CRITICAL-VRF
 ip address 10.255.11.18 255.255.255.252
 ip proxy-arp
 ip pim sparse-mode
```

```
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/0/7
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/0/8
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/0/9
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/0/10
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/0/11
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface GigabitEthernet0
 vrf forwarding Mgmt-intf
 ip address 175.13.8.33 255.255.255.0
 ip proxy-arp
 negotiation auto
!
segment-routing traffic-eng
 !
 on-demand color 100
  authorize
  candidate-paths
   preference 100
    constraints
     segments
```

```
     dataplane srv6
     !
     affinity
      include-all
       name SECURED
      !
     !
    !
    dynamic
     metric
      type te
     !
    !
   !
  !
 !
on-demand color 200
 authorize
 candidate-paths
  preference 100
   constraints
    segments
     dataplane srv6
    !
    affinity
     include-all
      name UNSECURED
     !
    !
   !
   dynamic
    metric
     type te
    !
   !
  !
 !
!
on-demand color 300
 authorize
 candidate-paths
  preference 100
   per-flow
    forward-class 0 color 200
    forward-class 1 color 100
    forward-class 2 color 100
   !
```

```
   !
  !
 !
!
segment-routing srv6
 encapsulation
  source-address 2001:DB8:AC12::
 locators
  locator SLOC
   prefix 2001:DB8:AC12::/48
   format usid-f3216
!
router isis SRTE-OVLY
 net 49.0099.0099.0000.1012.00
 is-type level-2-only
 metric-style wide
 lsp-mtu 1300
 distribute link-state instance-id 99
 affinity-map SECURED bit-position 0
 affinity-map UNSECURED bit-position 1
 bfd all-interfaces
 !
 address-family ipv6
  bfd all-interfaces
  router-id Loopback0
  segment-routing srv6
   locator SLOC
  fast-reroute per-prefix level-2 all
  fast-reroute ti-lfa level-2
 exit-address-family
!
router bgp 65000
 bgp router-id 10.255.0.12
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 2001:DB8:100:10::1 remote-as 65000
 neighbor 2001:DB8:100:10::1 update-source Loopback0
 neighbor 2001:DB8:100:31::1 remote-as 65000
 neighbor 2001:DB8:100:31::1 update-source Loopback0
 neighbor 2001:DB8:100:41::1 remote-as 65000
 neighbor 2001:DB8:100:41::1 update-source Loopback0
 neighbor 2001:DB8:100:42::1 remote-as 65000
 neighbor 2001:DB8:100:42::1 update-source Loopback0
 neighbor 2001:DB8:100:51::1 remote-as 65000
 neighbor 2001:DB8:100:51::1 update-source Loopback0
 neighbor 2001:DB8:100:52::1 remote-as 65000
 neighbor 2001:DB8:100:52::1 update-source Loopback0
```

```
neighbor 2001:DB8:131:200::1 remote-as 65000
neighbor 2001:DB8:131:200::1 update-source Loopback200
!
address-family ipv4
exit-address-family
!
address-family vpnv4
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
 neighbor 2001:DB8:100:10::1 activate
 neighbor 2001:DB8:100:10::1 send-community both
 neighbor 2001:DB8:100:10::1 next-hop-self all
 neighbor 2001:DB8:100:10::1 route-map RM_DCI_RR_NH_OUT out
 neighbor 2001:DB8:100:31::1 activate
 neighbor 2001:DB8:100:31::1 send-community both
 neighbor 2001:DB8:100:31::1 route-reflector-client
 neighbor 2001:DB8:100:41::1 activate
 neighbor 2001:DB8:100:41::1 send-community both
 neighbor 2001:DB8:100:41::1 route-reflector-client
 neighbor 2001:DB8:100:42::1 activate
 neighbor 2001:DB8:100:42::1 send-community both
 neighbor 2001:DB8:100:42::1 route-reflector-client
 neighbor 2001:DB8:100:51::1 activate
 neighbor 2001:DB8:100:51::1 send-community both
 neighbor 2001:DB8:100:51::1 route-reflector-client
 neighbor 2001:DB8:100:52::1 activate
 neighbor 2001:DB8:100:52::1 send-community both
 neighbor 2001:DB8:100:52::1 route-reflector-client
 neighbor 2001:DB8:131:200::1 activate
 neighbor 2001:DB8:131:200::1 send-community both
 neighbor 2001:DB8:131:200::1 route-reflector-client
 neighbor 2001:DB8:131:200::1 next-hop-self all
exit-address-family
!
address-family ipv6
exit-address-family
!
address-family vpnv6
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
```

```
 !
 neighbor 2001:DB8:100:10::1 activate
 neighbor 2001:DB8:100:10::1 send-community both
 neighbor 2001:DB8:100:10::1 next-hop-self all
 neighbor 2001:DB8:100:10::1 route-map RM_DCI_RR_NH_OUT out
 neighbor 2001:DB8:100:31::1 activate
 neighbor 2001:DB8:100:31::1 send-community both
 neighbor 2001:DB8:100:31::1 route-reflector-client
 neighbor 2001:DB8:100:41::1 activate
 neighbor 2001:DB8:100:41::1 send-community both
 neighbor 2001:DB8:100:41::1 route-reflector-client
 neighbor 2001:DB8:100:42::1 activate
 neighbor 2001:DB8:100:42::1 send-community both
 neighbor 2001:DB8:100:42::1 route-reflector-client
 neighbor 2001:DB8:100:51::1 activate
 neighbor 2001:DB8:100:51::1 send-community both
 neighbor 2001:DB8:100:51::1 route-reflector-client
 neighbor 2001:DB8:100:52::1 activate
 neighbor 2001:DB8:100:52::1 send-community both
 neighbor 2001:DB8:100:52::1 route-reflector-client
 neighbor 2001:DB8:131:200::1 activate
 neighbor 2001:DB8:131:200::1 send-community both
 neighbor 2001:DB8:131:200::1 route-reflector-client
 neighbor 2001:DB8:131:200::1 next-hop-self all
exit-address-family
!
address-family ipv4 vrf CRITICAL-VRF
 network 10.12.12.0 mask 255.255.255.0 route-map RM_ODN_COLOR_300
 network 10.255.11.12 mask 255.255.255.255
 network 10.255.11.100 mask 255.255.255.255
 network 198.18.12.0 route-map RM_ODN_COLOR_300
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
exit-address-family
!
address-family ipv6 vrf CRITICAL-VRF
 network 2001:DB8:12:12::/64 route-map RM_ODN_COLOR_300
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
```

```
 exit-address-family
 !
 address-family ipv4 vrf NON-CRITICAL-VRF
  network 10.112.112.0 mask 255.255.255.0
  network 198.19.112.0
  !
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
 exit-address-family
 !
 address-family ipv6 vrf NON-CRITICAL-VRF
  network 2001:DB8:112:112::/64
  !
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
 exit-address-family
!
ip rcmd domain-lookup
ip forward-protocol nd
ip forward-protocol udp tftp
!
ip telnet comport enable
ip pim vrf CRITICAL-VRF rp-address 10.255.11.100 ACL-ASM-ANYCAST-RP-PILOT
ip msdp vrf CRITICAL-VRF peer 10.255.11.11 connect-source Loopback112 remote-as 65000
ip msdp vrf CRITICAL-VRF originator-id Loopback112
ip tftp source-interface GigabitEthernet0
ip tftp blocksize 512
ip ftp passive
ip http server
ip http authentication local
ip http secure-server
ip ssh bulk-mode 131072
ip route 172.16.22.2 255.255.255.255 172.16.12.1
ip route 172.16.31.2 255.255.255.255 172.16.12.1
ip route 172.16.51.2 255.255.255.255 172.16.12.1
ip route 172.16.52.2 255.255.255.255 172.16.12.1
ip route vrf Mgmt-intf 0.0.0.0 0.0.0.0 175.13.8.1
ip route vrf Mgmt-intf 202.153.144.0 255.255.255.0 175.13.8.1
ip route vrf CRITICAL-VRF 10.255.11.11 255.255.255.255 10.255.11.17 name MSDP_HUB11_DIRECT
!
ip access-list standard ACL-ASM-ANYCAST-RP-PILOT
```

```
 10 permit 239.19.10.0 0.0.0.255
!
ip access-list extended IPV4-PILOT-NGFW-ALLOW
 10 permit ip any any
ipv6 route 2001:DB8:100:10::1/128 2001:DB8:4100:1::10
ipv6 route 2001:DB8:3100:1::/64 2001:DB8:2100:2::1
ipv6 route 2001:DB8:3100:2::/64 2001:DB8:2100:2::1
ipv6 route 2001:DB8:3100:3::/64 2001:DB8:2100:2::1
ipv6 route 2001:DB8:3100:4::/64 2001:DB8:2100:2::1
ipv6 route 2001:DB8:AC11::/48 2001:DB8:4100:1::11 250
ipv6 route 2001:DB8:AC21::/48 2001:DB8:4100:1::21 250
ipv6 route 2001:DB8:AC22::/48 2001:DB8:4100:1::22 250
ipv6 route 2001:DB8:AC31::/48 2001:DB8:4100:1::11 250
ipv6 route 2001:DB8:AC41::/48 2001:DB8:4100:1::11 250
ipv6 route 2001:DB8:AC42::/48 2001:DB8:4100:1::11 250
ipv6 route 2001:DB8:AC52::/48 2001:DB8:4100:1::22 250
route-map RM_DCI_RR_NH_OUT permit 10
 set ipv6 next-hop 2001:DB8:100:12::1
!
route-map RM_ODN_COLOR_300 permit 10
 set extcommunity color 300 additive
!
!
!
!
!
!
ipv6 access-list IPV6-PILOT-NGFW-ALLOW
 sequence 10 permit ipv6 any any
!
control-plane
!
!
!
!
!
!
line con 0
 exec-timeout 0 0
 activation-character 13
 transport preferred ssh
 transport output ssh
 stopbits 1
line aux 0
 activation-character 13
line vty 0 4
 exec-timeout 0 0
```

```
 activation-character 13
 transport input ssh
line vty 5 15
 activation-character 13
 transport input ssh
!
ntp allow mode control 3
!
!
!
!
!
!
End
```

Hub2_1

```
HUB2_1#sh run
Building configuration...

Current configuration : 27999 bytes
!
! Last configuration change at 03:40:43 UTC Fri Jul 24 2026 by admin
!
version 26.1
system mode insecure
service timestamps debug datetime msec
service timestamps log datetime msec
platform qfp utilization monitor load 80
!
hostname HUB2_1
!
boot-start-marker
boot system bootflash:/c8000aep-
universalk9.BLD_V261_1_THROTTLE_LATEST_20260305_183234_V26_1_0_54.SSA.bin
! Warning: Booting with bundle mode will be deprecated in the near future. Migration to install
mode is required.
boot-end-marker
!
!
vrf definition Mgmt-intf
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
 exit-address-family
!
```

```
vrf definition CRITICAL-VRF
 rd 65000:21
 route-target export 65000:1001
 route-target import 65000:1001
 !
 address-family ipv4
  route-target export 65000:1001
  route-target import 65000:1001
 exit-address-family
 !
 address-family ipv6
  route-target export 65000:1001
  route-target import 65000:1001
 exit-address-family
!
vrf definition NON-CRITICAL-VRF
 rd 65000:121
 route-target export 65000:2001
 route-target import 65000:2001
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
 exit-address-family
!
no logging console
aaa new-model
!
!
aaa authentication login default local
aaa authentication enable default enable
aaa authorization console
aaa authorization exec default local
!
!
aaa session-id common
!
!
subscriber templating
!
!
!
!
!
!
!
```

```
!
!
!
login on-success log
!
!
!
!
!
!
ipv6 unicast-routing
!
!
!
!
!
!
!
!
key chain KC-WAN-MACSEC macsec
 key 01
  cryptographic-algorithm aes-256-cmac
  key-string 3637fc51ed2f71b5a354e1aca9a91adb0f5a6744f1cda2cc7e18417cbd025e51
  lifetime local 00:00:00 May 14 2026 infinite
product-analytics
!
crypto pki trustpoint TP-self-signed-2127709243
 enrollment selfsigned
 subject-name cn=IOS-Self-Signed-Certificate-2127709243
 revocation-check none
 rsakeypair TP-self-signed-2127709243
 hash sha512
!
crypto pki trustpoint SLA-TrustPoint
 enrollment pkcs12
 revocation-check crl
 hash sha512
!
!
crypto pki certificate chain TP-self-signed-2127709243
 certificate self-signed 01
  30820330 30820218 A0030201 02020101 300D0609 2A864886 F70D0101 0D050030
  31312F30 2D060355 04030C26 494F532D 53656C66 2D536967 6E65642D 43657274
  69666963 6174652D 32313237 37303932 3433301E 170D3236 30353131 30393335
  33365A17 0D333630 35313030 39333533 365A3031 312F302D 06035504 030C2649
  4F532D53 656C662D 5369676E 65642D43 65727469 66696361 74652D32 31323737
  30393234 33308201 22300D06 092A8648 86F70D01 01010500 0382010F 00308201
```

```
  0A028201 0100CB56 5418BBD9 C5B92AE7 3F415A96 48345256 924AD7D2 EAFAA8C5
  CE041A2C CA516230 8D7B48BD BC59650A 15CA2398 B6349C33 C17C069C 44D406A8
  0CC49C60 B4C6181E 0B3CAD25 52898C73 7F0ECB28 E8D1F78F 07B1086A 9BA4A47B
  0B9E1F21 DB0F9056 39B29BE2 1E2604C0 197CA645 F311DBCD F88B26E7 887BAFF2
  2AF486D9 0897289C FA8BECC5 6C9188F6 F039E290 281CF683 BA2CC8AF D6C7D627
  26478E54 F79FFBA6 4F2950DF D04D4BD4 2CC145BE 220AABDC 78A780D1 E4C8B7F8
  CBD1D688 EB49B4E9 FBD948D2 C17EF6C9 919B0AA5 8821F99F 2D86063C B4314786
  47AA9FAB 8C5BBF0D 33AEBD69 B377264E 81A66C3B 65AB660D A9098517 D31DAA45
  97754865 20130203 010001A3 53305130 1D060355 1D0E0416 0414F63D 2F181C81
  47EF5DAF D7D5BAE0 81128FF9 7925301F 0603551D 23041830 168014F6 3D2F181C
  8147EF5D AFD7D5BA E081128F F9792530 0F060355 1D130101 FF040530 030101FF
  300D0609 2A864886 F70D0101 0D050003 82010100 C29E176F 1B102538 FF061232
  BD526CE7 955D6481 08C08761 0FB5D57C 9BC8417C 4B9C808D C4F75A54 D2E1FA57
  24F09B9A 55ACC1E3 89B3FC41 957CB0B9 4CBA59EC 68935FE0 96493BC2 8004FEF5
  3082F2F4 3B52C2F3 F5ED271A 5AC06788 46AF6988 B07DD288 4A272797 822CCF5E
  A42B96C7 BC371685 47D6154B 13F9A4B4 6823200B C78E1C79 EB4A5BB6 31F117F1
  75C89FD3 868BFFCC 2AEFCD23 FF9048CA B46BA810 E84AA8A3 87DA211F 595E489E
  D6096617 A69D8929 BE04521E 99B3E611 CFFA1EC3 F48483AC 1482B909 24CF591D
  9AF65492 FD85181D 36F14DE1 1CCD1B9A 6FEEAAD2 5ADF36E8 89CCD133 3244DBDB
  FAF9E506 877650E6 4CC95A5D 1BAD96DF 63FD9A55
          quit
crypto pki certificate chain SLA-TrustPoint
 certificate ca 01
  30820321 30820209 A0030201 02020101 300D0609 2A864886 F70D0101 0B050030
  32310E30 0C060355 040A1305 43697363 6F312030 1E060355 04031317 43697363
  6F204C69 63656E73 696E6720 526F6F74 20434130 1E170D31 33303533 30313934
  3834375A 170D3338 30353330 31393438 34375A30 32310E30 0C060355 040A1305
  43697363 6F312030 1E060355 04031317 43697363 6F204C69 63656E73 696E6720
  526F6F74 20434130 82012230 0D06092A 864886F7 0D010101 05000382 010F0030
  82010A02 82010100 A6BCBD96 131E05F7 145EA72C 2CD686E6 17222EA1 F1EFF64D
  CBB4C798 212AA147 C655D8D7 9471380D 8711441E 1AAF071A 9CAE6388 8A38E520
  1C394D78 462EF239 C659F715 B98C0A59 5BBB5CBD 0CFEBEA3 700A8BF7 D8F256EE
  4AA4E80D DB6FD1C9 60B1FD18 FFC69C96 6FA68957 A2617DE7 104FDC5F EA2956AC
  7390A3EB 2B5436AD C847A2C5 DAB553EB 69A9A535 58E9F3E3 C0BD23CF 58BD7188
  68E69491 20F320E7 948E71D7 AE3BCC84 F10684C7 4BC8E00F 539BA42B 42C68BB7
  C7479096 B4CB2D62 EA2F505D C7B062A4 6811D95B E8250FC4 5D5D5FB8 8F27D191
  C55F0D76 61F9A4CD 3D992327 A8BB03BD 4E6D7069 7CBADF8B DF5F4368 95135E44
  DFC7C6CF 04DD7FD1 02030100 01A34230 40300E06 03551D0F 0101FF04 04030201
  06300F06 03551D13 0101FF04 05300301 01FF301D 0603551D 0E041604 1449DC85
  4B3D31E5 1B3E6A17 606AF333 3D3B4C73 E8300D06 092A8648 86F70D01 010B0500
  03820101 00507F24 D3932A66 86025D9F E838AE5C 6D4DF6B0 49631C78 240DA905
  604EDCDE FF4FED2B 77FC460E CD636FDB DD44681E 3A5673AB 9093D3B1 6C9E3D8B
  D98987BF E40CBD9E 1AECA0C2 2189BB5C 8FA85686 CD98B646 5575B146 8DFC66A8
  467A3DF4 4D565700 6ADF0F0D CF835015 3C04FF7C 21E878AC 11BA9CD2 55A9232C
  7CA7B7E6 C1AF74F6 152E99B7 B1FCF9BB E973DE7F 5BDDEB86 C71E3B49 1765308B
  5FB0DA06 B92AFE7F 494E8A9E 07B85737 F3A58BE1 1A48A229 C37C1E69 39F08678
  80DDCD16 D6BACECA EEBC7CF9 8428787B 35202CDC 60E4616A B623CDBD 230E3AFB
```

```
  418616A9 4093E049 4D10AB75 27E86F73 932E35B5 8862FDAE 0275156F 719BB2F0
  D697DF7F 28
          quit
!
!
!
!
!
!
!
!
!
license udi pid C8550-G2 sn FCB3013D18Z
license boot level advantage
!
memory free low-watermark processor 3199455
diagnostic bootup level minimal
!
spanning-tree extend system-id
!
mka policy POL-WAN-MACSEC
!
!
!
!
username lab privilege 15 secret 9 $9$Al2Po6UnO/2meU$X4dU4QYCwO73hO2DbvRWSJmkive/00s2VfHE8UcpQ.g
username admin privilege 15 secret 9
$9$K9NmZsQJLACYyk$NlXoh.Q.nz4ujgbA5dzMBukbMBWTHmXigspGnDvV.js
!
redundancy
 mode none
!
!
!
crypto ikev2 proposal U-IKEV2-PROP
 pqc mlkem1024 optional
 encryption aes-gcm-256
 prf sha256
 group 19
crypto ikev2 proposal UC2-IKEV2-PROP
 pqc mlkem1024 optional
 encryption aes-gcm-256
 prf sha256
 group 19
!
crypto ikev2 policy U-IKEV2-POLICY
 proposal U-IKEV2-PROP
```

```
crypto ikev2 policy UC2-IKEV2-POLICY
 proposal UC2-IKEV2-PROP
!
crypto ikev2 keyring U-H21-PSK-KR
 peer EDGE1_1
  address 172.16.31.2
  pre-shared-key cisco123!
 !
 peer EDGE3_1
  address 172.16.51.2
  pre-shared-key cisco123!
 !
 peer EDGE2_2
  address 172.16.22.2
  pre-shared-key cisco123!
 !
 peer EDGE3_2
  address 172.16.52.2
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring U-H21-PPK-KR
 peer EDGE1_1
  address 172.16.31.2
  ppk manual id UNSEC-E11-H21 key pqcReady123! required
 !
 peer EDGE3_1
  address 172.16.51.2
  ppk manual id UNSEC-E31-H21 key pqcReady123! required
 !
 peer EDGE2_2
  address 172.16.22.2
  ppk manual id UNSEC-E22-H21 key pqcReady123! required
 !
 peer EDGE3_2
  address 172.16.52.2
  ppk manual id UNSEC-E32-H21 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-E31-PSK-KR
 peer EDGE1_1
  address 2001:DB8:3100:1::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-E31-PPK-KR
 peer EDGE1_1
```

```
  address 2001:DB8:3100:1::2/128
  ppk manual id UC2-E31-H21 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-E41-PSK-KR
 peer EDGE2_1
  address 2001:DB8:3100:2::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-E41-PPK-KR
 peer EDGE2_1
  address 2001:DB8:3100:2::2/128
  ppk manual id UC2-E41-H21 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-E51-PSK-KR
 peer EDGE3_1
  address 2001:DB8:3100:3::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-E51-PPK-KR
 peer EDGE3_1
  address 2001:DB8:3100:3::2/128
  ppk manual id UC2-E51-H21 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-E52-PSK-KR
 peer EDGE3_2
  address 2001:DB8:3100:4::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-E52-PPK-KR
 peer EDGE3_2
  address 2001:DB8:3100:4::2/128
  ppk manual id UC2-E52-H21 key pqcReady123! required
 !
!
crypto ikev2 keyring UC4-E31-H21-PSK-KR
 peer EDGE1_1
  address 172.31.82.1
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC4-E31-H21-PPK-KR
```

```
 peer EDGE1_1
  address 172.31.82.1
  ppk manual id UC4-E31-H21 key pqcReady123! required
 !
!
!
crypto ikev2 profile U-H21-IKEV2-PROFILE
 match identity remote address 172.16.31.2 255.255.255.255
 match identity remote address 172.16.51.2 255.255.255.255
 match identity remote address 172.16.22.2 255.255.255.255
 match identity remote address 172.16.52.2 255.255.255.255
 identity local address 172.16.21.2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk U-H21-PPK-KR
 keyring local U-H21-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-E31-IKEV2-PROFILE
 match identity remote address 2001:DB8:3100:1::2/128
 identity local address 2001:DB8:2100:3::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-E31-PPK-KR
 keyring local UC2-E31-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-E41-IKEV2-PROFILE
 match identity remote address 2001:DB8:3100:2::2/128
 identity local address 2001:DB8:2100:3::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-E41-PPK-KR
 keyring local UC2-E41-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-E51-IKEV2-PROFILE
 match identity remote address 2001:DB8:3100:3::2/128
 identity local address 2001:DB8:2100:3::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-E51-PPK-KR
 keyring local UC2-E51-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-E52-IKEV2-PROFILE
 match identity remote address 2001:DB8:3100:4::2/128
```

```
 identity local address 2001:DB8:2100:3::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-E52-PPK-KR
 keyring local UC2-E52-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC4-E31-H21-IKEV2-PROFILE
 match identity remote address 172.31.82.1 255.255.255.255
 identity local address 172.31.82.2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC4-E31-H21-PPK-KR
 keyring local UC4-E31-H21-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 fragmentation mtu 1400
!
!
cdp run
!
!
class-map match-any CM-BUSINESS
 match dscp af31
class-map match-any CM-PILOT-VIDEO
 match dscp af41
class-map match-any CM-PILOT-VOICE
 match dscp ef
class-map match-any CM-CRITICAL
 match dscp ef
class-map type inspect match-any CM-PILOT-NGFW-ALLOW
 match protocol icmp
 match protocol tcp
 match protocol udp
 match access-group name IPV4-PILOT-NGFW-ALLOW
 match access-group name IPV6-PILOT-NGFW-ALLOW
class-map match-any CM-PILOT-CRITICAL
 match dscp af31
class-map match-any CM-PILOT-BULK
 match dscp af11
!
policy-map PM-PILOT-SECURED-QOS
 class CM-PILOT-VOICE
  priority percent 10
 class CM-PILOT-VIDEO
  bandwidth percent 25
 class CM-PILOT-CRITICAL
```

```
  bandwidth percent 30
 class CM-PILOT-BULK
  bandwidth percent 10
 class class-default
  fair-queue
policy-map type epbr PM-PFP
 class CM-CRITICAL
  set forward-class 1
 class CM-BUSINESS
  set forward-class 2
 class class-default
  set forward-class 0
policy-map type inspect PM-PILOT-NGFW
 class type inspect CM-PILOT-NGFW-ALLOW
  inspect
 class class-default
  drop log
policy-map type inspect PM-PILOT-NGFW-PASS
 class type inspect CM-PILOT-NGFW-ALLOW
  pass
 class class-default
  drop log
!
!
zone security Z-PILOT-LAN
zone security Z-PILOT-SRV6-WAN
zone security Z-PILOT-LAN-U
zone-pair security ZP-LAN-TO-SRV6-WAN source Z-PILOT-LAN destination Z-PILOT-SRV6-WAN
 service-policy type inspect PM-PILOT-NGFW
zone-pair security ZP-LANU-TO-SRV6-WAN source Z-PILOT-LAN-U destination Z-PILOT-SRV6-WAN
 service-policy type inspect PM-PILOT-NGFW
zone-pair security ZP-SRV6-WAN-TO-LAN source Z-PILOT-SRV6-WAN destination Z-PILOT-LAN
 service-policy type inspect PM-PILOT-NGFW
zone-pair security ZP-SRV6-WAN-TO-LANU source Z-PILOT-SRV6-WAN destination Z-PILOT-LAN-U
 service-policy type inspect PM-PILOT-NGFW
!
!
!
!
!
!
!
crypto ipsec transform-set U-TS esp-gcm 256
 mode transport
crypto ipsec transform-set UC2-TS esp-gcm 256
 mode transport
!
```

```
crypto ipsec profile U-E32-H21-IPSEC
 set transform-set U-TS
 set ikev2-profile U-H21-IKEV2-PROFILE
!
crypto ipsec profile U-H21-IPSEC
 set transform-set U-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile U-H21-IKEV2-PROFILE
!
crypto ipsec profile UC2-E31-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-E31-IKEV2-PROFILE
!
crypto ipsec profile UC2-E41-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-E41-IKEV2-PROFILE
!
crypto ipsec profile UC2-E51-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-E51-IKEV2-PROFILE
!
crypto ipsec profile UC2-E52-IPSEC
 set transform-set UC2-TS
 set pfs group19
 set ikev2-profile UC2-E52-IKEV2-PROFILE
!
crypto ipsec profile UC4-E31-H21-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC4-E31-H21-IKEV2-PROFILE
!
!
!
!
!
!
!
!
!
!
interface Loopback0
 description CNI_UC2_BGP_UPDATE_SOURCE
 no ip address
 ip proxy-arp
```

```
 ipv6 address 2001:DB8:100:21::1/128
 ipv6 router isis SRTE-OVLY
!
interface Loopback100
 vrf forwarding CRITICAL-VRF
 ip address 10.21.21.1 255.255.255.0
 ip proxy-arp
 ipv6 address 2001:DB8:21:21::1/64
!
interface Loopback200
 no ip address
 ip proxy-arp
 ipv6 address 2001:DB8:121:200::1/128
 ipv6 router isis SRTE-OVLY
!
interface Loopback210
 vrf forwarding NON-CRITICAL-VRF
 ip address 10.121.121.1 255.255.255.0
 ip proxy-arp
 ipv6 address 2001:DB8:121:121::1/64
!
interface Tunnel211
 description HUB2_1 to EDGE1_1 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:211::2/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/0/2
 tunnel destination 172.16.31.2
 tunnel protection ipsec profile U-H21-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel221
 description HUB2_1 to EDGE2_2 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:221::1/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/0/2
```

```
 tunnel destination 172.16.22.2
 tunnel protection ipsec profile U-H21-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel251
 description HUB2_1 to EDGE3_1 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:251::2/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/0/2
 tunnel destination 172.16.51.2
 tunnel protection ipsec profile U-H21-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel263
 description HUB2_1 to EDGE3_2 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:263::2/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/0/2
 tunnel destination 172.16.52.2
 tunnel protection ipsec profile U-E32-H21-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel4082
 description UC4_EDGE1_EXCEPTION_GRE_OVER_IPSEC_SECURED_ELINE_TO_EDGE1_1
 no ip address
 ip proxy-arp
 ip mtu 1300
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5082:1::2/64
 ipv6 mtu 1300
 ipv6 router isis SRTE-OVLY
```

```
 tunnel source 172.31.82.2
 tunnel destination 172.31.82.1
 tunnel path-mtu-discovery
 tunnel protection ipsec profile UC4-E31-H21-IPSEC
 clns mtu 1300
 isis circuit-type level-2-only
 isis affinity flex-algo
  name SECURED
  !
!
interface TenGigabitEthernet0/0/0
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/0/1
 description CNI_UC4_HUB2_1_WAN_MACSEC_PARENT
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
 ipv6 address 2001:DB8:2100:3::2/64
 macsec dot1q-in-clear 1
 macsec access-control should-secure
 service-policy output PM-PILOT-SECURED-QOS
!
interface TenGigabitEthernet0/0/1.4079
 description UC4_TAGGED_WAN_MACSEC_NATIVE_ISIS_TO_EDGE3_2
 encapsulation dot1Q 4079
 ip address 172.31.79.2 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:4079:1::2/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address 28af.fdb3.5188
 eapol eth-type 876F
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
 clns mtu 1400
 isis circuit-type level-2-only
 isis network point-to-point
 isis affinity flex-algo
  name SECURED
  !
```

```
!
interface TenGigabitEthernet0/0/1.4082
 description UC4_ELINE_EXCEPTION_TO_EDGE1_1
 encapsulation dot1Q 4082
 ip address 172.31.82.2 255.255.255.252
 ip proxy-arp
 ipv6 address 2001:DB8:4082:1::2/64
!
interface TenGigabitEthernet0/0/1.4085
 description UC4_TAGGED_WAN_MACSEC_GRE_ISIS_TO_EDGE3_1
 encapsulation dot1Q 4085
 ip address 172.31.85.2 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:4085:1::2/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address ec19.2e55.5f08
 eapol eth-type 876F
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
 clns mtu 1400
 isis circuit-type level-2-only
 isis network point-to-point
 isis affinity flex-algo
  name SECURED
  !
!
interface TenGigabitEthernet0/0/1.4088
 description UC4_HUB2_1_TO_EDGE2_1_WAN_MACSEC_GRE_ISIS_VLAN4088
 encapsulation dot1Q 4088
 ip address 172.31.88.2 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:4088:1::2/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address 5000.e06c.b2a4
 eapol eth-type 876F
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
 clns mtu 1400
 isis circuit-type level-2-only
 isis network point-to-point
 isis affinity flex-algo
```

```
  name SECURED
  !
!
interface TenGigabitEthernet0/0/2
 description CNI_UC2_INET_UNDERLAY_BLR_TUNNEL221
 ip address 172.16.21.2 255.255.255.252
 ip proxy-arp
 no negotiation auto
 cdp enable
 service-policy output PM-PILOT-SECURED-QOS
!
interface TenGigabitEthernet0/0/3
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/0/4
 description HUB2_1 WAN switch/DCI segment
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
 ipv6 address 2001:DB8:4100:1::21/64
!
interface TenGigabitEthernet0/0/4.101
 description SECURED_CRITICAL_VRF_HUB2_1
 encapsulation dot1Q 101
 vrf forwarding CRITICAL-VRF
 ip address 198.18.21.1 255.255.255.0
 ip proxy-arp
 ip nbar protocol-discovery
 zone-member security Z-PILOT-LAN
 service-policy type epbr input PM-PFP
!
interface TenGigabitEthernet0/0/4.201
 description UNSECURED_NON_CRITICAL_VRF_HUB2_1
 encapsulation dot1Q 201
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.121.1 255.255.255.0
 ip proxy-arp
 ip nbar protocol-discovery
 zone-member security Z-PILOT-LAN-U
 service-policy type epbr input PM-PFP
!
interface TenGigabitEthernet0/0/4.511
 description WAN_MGMT511_HUB2_1
```

```
 encapsulation dot1Q 511
 ip address 192.168.12.21 255.255.255.0
 ip proxy-arp
!
interface TenGigabitEthernet0/0/5
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/0/6
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/0/7
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/0/8
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/0/9
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/0/10
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/0/11
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface GigabitEthernet0
```

```
 vrf forwarding Mgmt-intf
 ip address 175.13.8.35 255.255.255.0
 ip proxy-arp
 negotiation auto
!
segment-routing traffic-eng
 !
 on-demand color 100
  authorize
  candidate-paths
   preference 100
    constraints
     segments
      dataplane srv6
     !
     affinity
      include-all
       name SECURED
      !
     !
    !
    dynamic
     metric
      type te
     !
    !
   !
  !
 !
 on-demand color 200
  authorize
  candidate-paths
   preference 100
    constraints
     segments
      dataplane srv6
     !
     affinity
      include-all
       name UNSECURED
      !
     !
    !
    dynamic
     metric
      type te
     !
```

```
    !
   !
  !
 !
 on-demand color 300
  authorize
  candidate-paths
   preference 100
    per-flow
     forward-class 0 color 200
     forward-class 1 color 100
     forward-class 2 color 100
    !
   !
  !
 !
!
segment-routing srv6
 encapsulation
  source-address 2001:DB8:AC21::
 locators
  locator SLOC
   prefix 2001:DB8:AC21::/48
   format usid-f3216
!
router isis SRTE-OVLY
 net 49.0099.0099.0000.1021.00
 is-type level-2-only
 metric-style wide
 lsp-mtu 1300
 distribute link-state instance-id 99
 affinity-map SECURED bit-position 0
 affinity-map UNSECURED bit-position 1
 !
 address-family ipv6
  router-id Loopback0
  segment-routing srv6
   locator SLOC
  fast-reroute per-prefix level-2 all
  fast-reroute ti-lfa level-2
 exit-address-family
!
router bgp 65000
 bgp router-id 10.255.0.21
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 2001:DB8:100:10::1 remote-as 65000
```

```
neighbor 2001:DB8:100:10::1 update-source Loopback0
neighbor 2001:DB8:100:31::1 remote-as 65000
neighbor 2001:DB8:100:31::1 update-source Loopback0
neighbor 2001:DB8:100:41::1 remote-as 65000
neighbor 2001:DB8:100:41::1 update-source Loopback0
neighbor 2001:DB8:100:42::1 remote-as 65000
neighbor 2001:DB8:100:42::1 update-source Loopback0
neighbor 2001:DB8:100:51::1 remote-as 65000
neighbor 2001:DB8:100:51::1 update-source Loopback0
neighbor 2001:DB8:100:52::1 remote-as 65000
neighbor 2001:DB8:100:52::1 update-source Loopback0
neighbor 2001:DB8:131:200::1 remote-as 65000
neighbor 2001:DB8:131:200::1 update-source Loopback200
!
address-family ipv4
exit-address-family
!
address-family vpnv4
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
 neighbor 2001:DB8:100:10::1 activate
 neighbor 2001:DB8:100:10::1 send-community both
 neighbor 2001:DB8:100:10::1 next-hop-self all
 neighbor 2001:DB8:100:10::1 route-map RM_DCI_RR_NH_OUT out
 neighbor 2001:DB8:100:31::1 activate
 neighbor 2001:DB8:100:31::1 send-community both
 neighbor 2001:DB8:100:31::1 route-reflector-client
 neighbor 2001:DB8:100:31::1 next-hop-self all
 neighbor 2001:DB8:100:41::1 activate
 neighbor 2001:DB8:100:41::1 send-community both
 neighbor 2001:DB8:100:41::1 route-reflector-client
 neighbor 2001:DB8:100:42::1 activate
 neighbor 2001:DB8:100:42::1 send-community both
 neighbor 2001:DB8:100:42::1 route-reflector-client
 neighbor 2001:DB8:100:51::1 activate
 neighbor 2001:DB8:100:51::1 send-community both
 neighbor 2001:DB8:100:51::1 route-reflector-client
 neighbor 2001:DB8:100:52::1 activate
 neighbor 2001:DB8:100:52::1 send-community both
 neighbor 2001:DB8:100:52::1 route-reflector-client
 neighbor 2001:DB8:131:200::1 activate
 neighbor 2001:DB8:131:200::1 send-community both
 neighbor 2001:DB8:131:200::1 route-reflector-client
```

```
 neighbor 2001:DB8:131:200::1 next-hop-self all
exit-address-family
!
address-family vpnv6
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
 neighbor 2001:DB8:100:10::1 activate
 neighbor 2001:DB8:100:10::1 send-community both
 neighbor 2001:DB8:100:10::1 next-hop-self all
 neighbor 2001:DB8:100:10::1 route-map RM_DCI_RR_NH_OUT out
 neighbor 2001:DB8:100:31::1 activate
 neighbor 2001:DB8:100:31::1 send-community both
 neighbor 2001:DB8:100:31::1 route-reflector-client
 neighbor 2001:DB8:100:31::1 next-hop-self all
 neighbor 2001:DB8:100:41::1 activate
 neighbor 2001:DB8:100:41::1 send-community both
 neighbor 2001:DB8:100:41::1 route-reflector-client
 neighbor 2001:DB8:100:42::1 activate
 neighbor 2001:DB8:100:42::1 send-community both
 neighbor 2001:DB8:100:42::1 route-reflector-client
 neighbor 2001:DB8:100:51::1 activate
 neighbor 2001:DB8:100:51::1 send-community both
 neighbor 2001:DB8:100:51::1 route-reflector-client
 neighbor 2001:DB8:100:52::1 activate
 neighbor 2001:DB8:100:52::1 send-community both
 neighbor 2001:DB8:100:52::1 route-reflector-client
 neighbor 2001:DB8:131:200::1 activate
 neighbor 2001:DB8:131:200::1 send-community both
 neighbor 2001:DB8:131:200::1 route-reflector-client
 neighbor 2001:DB8:131:200::1 next-hop-self all
exit-address-family
!
address-family ipv4 vrf CRITICAL-VRF
 network 10.21.21.0 mask 255.255.255.0 route-map RM_ODN_COLOR_300
 network 198.18.21.0 route-map RM_ODN_COLOR_300
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
exit-address-family
!
```

```
 address-family ipv6 vrf CRITICAL-VRF
  network 2001:DB8:21:21::/64 route-map RM_ODN_COLOR_300
  !
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
 exit-address-family
 !
 address-family ipv4 vrf NON-CRITICAL-VRF
  network 10.121.121.0 mask 255.255.255.0
  network 198.19.121.0
  !
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
 exit-address-family
 !
 address-family ipv6 vrf NON-CRITICAL-VRF
  network 2001:DB8:121:121::/64
  !
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
 exit-address-family
!
ip rcmd domain-lookup
ip forward-protocol nd
ip forward-protocol udp tftp
!
ip telnet comport enable
ip tftp source-interface GigabitEthernet0
ip tftp blocksize 512
ip ftp passive
ip http server
ip http authentication local
ip http secure-server
ip ssh bulk-mode 131072
ip route 172.16.22.2 255.255.255.255 172.16.21.1
ip route 172.16.31.2 255.255.255.255 172.16.21.1
ip route 172.16.51.2 255.255.255.255 172.16.21.1
ip route 172.16.52.2 255.255.255.255 172.16.21.1
```

```
ip route vrf Mgmt-intf 0.0.0.0 0.0.0.0 175.13.8.1
ip route vrf Mgmt-intf 202.153.144.0 255.255.255.0 175.13.8.1
!
ip access-list extended IPV4-PILOT-NGFW-ALLOW
 10 permit ip any any
ipv6 route 2001:DB8:100:10::1/128 2001:DB8:4100:1::10
ipv6 route 2001:DB8:3100:1::/64 2001:DB8:2100:3::1
ipv6 route 2001:DB8:3100:2::/64 2001:DB8:2100:3::1
ipv6 route 2001:DB8:3100:3::/64 2001:DB8:2100:3::1
ipv6 route 2001:DB8:3100:4::/64 2001:DB8:2100:3::1
ipv6 route 2001:DB8:AC11::/48 2001:DB8:4100:1::11 250
ipv6 route 2001:DB8:AC12::/48 2001:DB8:4100:1::12 250
ipv6 route 2001:DB8:AC22::/48 2001:DB8:4100:1::22 250
ipv6 route 2001:DB8:AC31::/48 2001:DB8:4100:1::11 250
ipv6 route 2001:DB8:AC41::/48 2001:DB8:4100:1::11 250
ipv6 route 2001:DB8:AC42::/48 2001:DB8:4100:1::11 250
ipv6 route 2001:DB8:AC51::/48 2001:DB8:4100:1::12 250
ipv6 route 2001:DB8:AC52::/48 2001:DB8:4100:1::22 250
route-map RM_DCI_RR_NH_OUT permit 10
 set ipv6 next-hop 2001:DB8:100:21::1
!
route-map RM_ODN_COLOR_300 permit 10
 set extcommunity color 300 additive
!
!
!
!
!
!
ipv6 access-list IPV6-PILOT-NGFW-ALLOW
 sequence 10 permit ipv6 any any
!
control-plane
!
!
!
!
!
!
line con 0
 exec-timeout 0 0
 activation-character 13
 transport preferred ssh
 transport output ssh
 stopbits 1
line aux 0
 activation-character 13
```

```
line vty 0 4
 exec-timeout 0 0
 activation-character 13
 transport input ssh
line vty 5 15
 activation-character 13
 transport input ssh
!
ntp allow mode control 3
!
!
!
!
!
!
End
```

Hub2_2

```
HUB2_2#sh run
Building configuration...

Current configuration : 26621 bytes
!
! Last configuration change at 10:07:22 UTC Thu Jul 23 2026 by admin
!
version 26.1
system mode insecure
service timestamps debug datetime msec
service timestamps log datetime msec
platform qfp utilization monitor load 80
!
hostname HUB2_2
!
boot-start-marker
boot system bootflash:/c8000aep-
universalk9.BLD_V261_1_THROTTLE_LATEST_20260305_183234_V26_1_0_54.SSA.bin
! Warning: Booting with bundle mode will be deprecated in the near future. Migration to install
mode is required.
boot-end-marker
!
!
vrf definition Mgmt-intf
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
 exit-address-family
```

```
!
vrf definition CRITICAL-VRF
 rd 65000:22
 route-target export 65000:1001
 route-target import 65000:1001
 !
 address-family ipv4
  route-target export 65000:1001
  route-target import 65000:1001
 exit-address-family
 !
 address-family ipv6
  route-target export 65000:1001
  route-target import 65000:1001
 exit-address-family
!
vrf definition NON-CRITICAL-VRF
 rd 65000:122
 route-target export 65000:2001
 route-target import 65000:2001
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
 exit-address-family
!
no logging console
aaa new-model
!
!
aaa authentication login default local
aaa authentication enable default enable
aaa authorization console
aaa authorization exec default local
!
!
aaa session-id common
!
!
subscriber templating
!
!
!
!
!
!
```

```
!
!
!
!
login on-success log
!
!
!
!
!
!
ipv6 unicast-routing
!
!
!
!
!
!
!
!
key chain KC-WAN-MACSEC macsec
 key 01
  cryptographic-algorithm aes-256-cmac
  key-string 3637fc51ed2f71b5a354e1aca9a91adb0f5a6744f1cda2cc7e18417cbd025e51
  lifetime local 00:00:00 May 14 2026 infinite
product-analytics
!
crypto pki trustpoint TP-self-signed-1103863800
 enrollment selfsigned
 subject-name cn=IOS-Self-Signed-Certificate-1103863800
 revocation-check none
 rsakeypair TP-self-signed-1103863800
 hash sha512
!
crypto pki trustpoint SLA-TrustPoint
 enrollment pkcs12
 revocation-check crl
 hash sha512
!
!
crypto pki certificate chain TP-self-signed-1103863800
 certificate self-signed 01
  30820330 30820218 A0030201 02020101 300D0609 2A864886 F70D0101 0D050030
  31312F30 2D060355 04030C26 494F532D 53656C66 2D536967 6E65642D 43657274
  69666963 6174652D 31313033 38363338 3030301E 170D3236 30353131 30393335
  34355A17 0D333630 35313030 39333534 355A3031 312F302D 06035504 030C2649
  4F532D53 656C662D 5369676E 65642D43 65727469 66696361 74652D31 31303338
```

```
  36333830 30308201 22300D06 092A8648 86F70D01 01010500 0382010F 00308201
  0A028201 0100C960 47E062A4 99E61AC0 3286558B D40EFC04 719E37A1 1EA051F4
  B3C6202D 9C6C2709 8417FEA9 E5219E06 EE36D0B8 71CFC365 D137A194 3FA82749
  DC292704 7451D35D 8D4D3B30 8219C241 81E858AD 1D29EE4B BED5ECA3 FBF078F8
  685D89B8 112FDF1D 624F4AEA 8AD0831E FFEE57A3 C39203D7 FA255B7D BD6A8C60
  179217DB B8CADD21 9DEB8501 FE537829 CBAA8149 5CBB755D 3AF8C023 00608D52
  186F7E3D 0C99EF74 4470F269 34933827 FDAD85BC A6745689 A94F014D 967259A1
  2DA04871 B1B0D3E6 14B7FB12 72E5C765 263639D3 71552CAC 72FFC710 D5D78537
  F6455408 F5AC64C4 FD6BE51A 72293E76 A6983356 DF7D5C8F D903A286 D3CDDB81
  64A7324B CD550203 010001A3 53305130 1D060355 1D0E0416 04149CDB 4FF2C7C1
  05BF6FB1 9A16AFDD A2542C42 6F67301F 0603551D 23041830 1680149C DB4FF2C7
  C105BF6F B19A16AF DDA2542C 426F6730 0F060355 1D130101 FF040530 030101FF
  300D0609 2A864886 F70D0101 0D050003 82010100 0C092664 01E9D33E 006D60FA
  E4C46A93 4A6D9EB7 9BFD8FC9 8B678169 CC5DA48D FEA9CE0A 9277F28D 7CCC74B2
  CFE8F413 4A52BA2D 831766E2 9BBCBB38 2D866F30 AA34BAF8 E427B3BB 3C28DCCA
  6DC452E4 2D5AA558 1B001DBD BAB8E8A4 4D4B4141 3C86B372 05FE3A96 A59B8608
  6F7A6848 E41890BE CFF1F8A5 8A575672 75326346 93E20105 31D34C79 2220E888
  81342014 5CC00B53 A75E42C8 60325B56 DC444E94 D807E986 0691761A 7D864DF2
  D46E4E89 A5C65C10 EFFDA980 C4A5F16D 67228B08 B1F5BCB3 CCBDAE32 7A576A26
  55197416 0FD0CF33 95E36D03 DE8F9A2C 9189070C AD7817D5 65690168 B86134E0
  7EC1229C 03F90410 1E8F0FC6 C732CB69 FAB2C867
          quit
crypto pki certificate chain SLA-TrustPoint
 certificate ca 01
  30820321 30820209 A0030201 02020101 300D0609 2A864886 F70D0101 0B050030
  32310E30 0C060355 040A1305 43697363 6F312030 1E060355 04031317 43697363
  6F204C69 63656E73 696E6720 526F6F74 20434130 1E170D31 33303533 30313934
  3834375A 170D3338 30353330 31393438 34375A30 32310E30 0C060355 040A1305
  43697363 6F312030 1E060355 04031317 43697363 6F204C69 63656E73 696E6720
  526F6F74 20434130 82012230 0D06092A 864886F7 0D010101 05000382 010F0030
  82010A02 82010100 A6BCBD96 131E05F7 145EA72C 2CD686E6 17222EA1 F1EFF64D
  CBB4C798 212AA147 C655D8D7 9471380D 8711441E 1AAF071A 9CAE6388 8A38E520
  1C394D78 462EF239 C659F715 B98C0A59 5BBB5CBD 0CFEBEA3 700A8BF7 D8F256EE
  4AA4E80D DB6FD1C9 60B1FD18 FFC69C96 6FA68957 A2617DE7 104FDC5F EA2956AC
  7390A3EB 2B5436AD C847A2C5 DAB553EB 69A9A535 58E9F3E3 C0BD23CF 58BD7188
  68E69491 20F320E7 948E71D7 AE3BCC84 F10684C7 4BC8E00F 539BA42B 42C68BB7
  C7479096 B4CB2D62 EA2F505D C7B062A4 6811D95B E8250FC4 5D5D5FB8 8F27D191
  C55F0D76 61F9A4CD 3D992327 A8BB03BD 4E6D7069 7CBADF8B DF5F4368 95135E44
  DFC7C6CF 04DD7FD1 02030100 01A34230 40300E06 03551D0F 0101FF04 04030201
  06300F06 03551D13 0101FF04 05300301 01FF301D 0603551D 0E041604 1449DC85
  4B3D31E5 1B3E6A17 606AF333 3D3B4C73 E8300D06 092A8648 86F70D01 010B0500
  03820101 00507F24 D3932A66 86025D9F E838AE5C 6D4DF6B0 49631C78 240DA905
  604EDCDE FF4FED2B 77FC460E CD636FDB DD44681E 3A5673AB 9093D3B1 6C9E3D8B
  D98987BF E40CBD9E 1AECA0C2 2189BB5C 8FA85686 CD98B646 5575B146 8DFC66A8
  467A3DF4 4D565700 6ADF0F0D CF835015 3C04FF7C 21E878AC 11BA9CD2 55A9232C
  7CA7B7E6 C1AF74F6 152E99B7 B1FCF9BB E973DE7F 5BDDEB86 C71E3B49 1765308B
  5FB0DA06 B92AFE7F 494E8A9E 07B85737 F3A58BE1 1A48A229 C37C1E69 39F08678
```

```
  80DDCD16 D6BACECA EEBC7CF9 8428787B 35202CDC 60E4616A B623CDBD 230E3AFB
  418616A9 4093E049 4D10AB75 27E86F73 932E35B5 8862FDAE 0275156F 719BB2F0
  D697DF7F 28
          quit
!
!
!
!
!
!
!
!
!
license udi pid C8570-G2 sn FCB3013D187
license boot level advantage
!
memory free low-watermark processor 3199455
hw-module subslot 0/1 mode 10G
hw-module subslot 0/2 mode 40G
hw-module subslot 0/2 breakout none port all
diagnostic bootup level minimal
!
spanning-tree extend system-id
!
mka policy POL-WAN-MACSEC
!
!
!
!
username lab privilege 15 secret 9 $9$Al2Po6UnO/2meU$X4dU4QYCwO73hO2DbvRWSJmkive/00s2VfHE8UcpQ.g
username admin privilege 15 secret 9
$9$S23ynXWeM8rpiE$ar4S.RnVZuzBkohiJHsPL1nZOvj3DO/cju0i/9Ntoy.
!
redundancy
 mode none
!
!
!
crypto ikev2 proposal U-IKEV2-PROP
 pqc mlkem1024 optional
 encryption aes-gcm-256
 prf sha256
 group 19
crypto ikev2 proposal UC2-IKEV2-PROP
 pqc mlkem1024 optional
 encryption aes-gcm-256
 prf sha256
```

```
 group 19
!
crypto ikev2 policy U-IKEV2-POLICY
 proposal U-IKEV2-PROP
crypto ikev2 policy UC2-IKEV2-POLICY
 proposal UC2-IKEV2-PROP
!
crypto ikev2 keyring U-H22-PSK-KR
 peer EDGE3_2
  address 172.16.52.2
  pre-shared-key cisco123!
 !
 peer EDGE2_2
  address 172.16.22.2
  pre-shared-key cisco123!
 !
 peer EDGE3_1
  address 172.16.51.2
  pre-shared-key cisco123!
 !
 peer EDGE1_1
  address 172.16.31.2
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring U-H22-PPK-KR
 peer EDGE3_2
  address 172.16.52.2
  ppk manual id UNSEC-E32-H22 key pqcReady123! required
 !
 peer EDGE2_2
  address 172.16.22.2
  ppk manual id UNSEC-E22-H22 key pqcReady123! required
 !
 peer EDGE3_1
  address 172.16.51.2
  ppk manual id UNSEC-E31-H22 key pqcReady123! required
 !
 peer EDGE1_1
  address 172.16.31.2
  ppk manual id UNSEC-E11-H22 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-E31-PSK-KR
 peer EDGE1_1
  address 2001:DB8:3100:1::2/128
  pre-shared-key cisco123!
```

```
 !
!
crypto ikev2 keyring UC2-E31-PPK-KR
 peer EDGE1_1
  address 2001:DB8:3100:1::2/128
  ppk manual id UC2-E31-H22 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-E41-PSK-KR
 peer EDGE2_1
  address 2001:DB8:3100:2::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-E41-PPK-KR
 peer EDGE2_1
  address 2001:DB8:3100:2::2/128
  ppk manual id UC2-E41-H22 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-E51-PSK-KR
 peer EDGE3_1
  address 2001:DB8:3100:3::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-E51-PPK-KR
 peer EDGE3_1
  address 2001:DB8:3100:3::2/128
  ppk manual id UC2-E51-H22 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-E52-PSK-KR
 peer EDGE3_2
  address 2001:DB8:3100:4::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-E52-PPK-KR
 peer EDGE3_2
  address 2001:DB8:3100:4::2/128
  ppk manual id UC2-E52-H22 key pqcReady123! required
 !
!
!
crypto ikev2 profile U-H22-IKEV2-PROFILE
 match identity remote address 172.16.52.2 255.255.255.255
```

```
 match identity remote address 172.16.22.2 255.255.255.255
 match identity remote address 172.16.51.2 255.255.255.255
 match identity remote address 172.16.31.2 255.255.255.255
 identity local address 172.16.23.2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk U-H22-PPK-KR
 keyring local U-H22-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-E31-IKEV2-PROFILE
 match identity remote address 2001:DB8:3100:1::2/128
 identity local address 2001:DB8:2100:4::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-E31-PPK-KR
 keyring local UC2-E31-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-E41-IKEV2-PROFILE
 match identity remote address 2001:DB8:3100:2::2/128
 identity local address 2001:DB8:2100:4::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-E41-PPK-KR
 keyring local UC2-E41-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-E51-IKEV2-PROFILE
 match identity remote address 2001:DB8:3100:3::2/128
 identity local address 2001:DB8:2100:4::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-E51-PPK-KR
 keyring local UC2-E51-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-E52-IKEV2-PROFILE
 match identity remote address 2001:DB8:3100:4::2/128
 identity local address 2001:DB8:2100:4::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-E52-PPK-KR
 keyring local UC2-E52-PSK-KR
 dpd 10 2 on-demand
!
!
```

```
!
cdp run
!
!
class-map match-any CM-BUSINESS
 match dscp af31
class-map match-any CM-PILOT-VIDEO
 match dscp af41
class-map match-any CM-PILOT-VOICE
 match dscp ef
class-map match-any CM-CRITICAL
 match dscp ef
class-map type inspect match-any CM-PILOT-NGFW-ALLOW
 match protocol icmp
 match protocol tcp
 match protocol udp
 match access-group name IPV4-PILOT-NGFW-ALLOW
 match access-group name IPV6-PILOT-NGFW-ALLOW
class-map match-any CM-PILOT-CRITICAL
 match dscp af31
class-map match-any CM-PILOT-BULK
 match dscp af11
!
policy-map PM-PILOT-SECURED-QOS
 class CM-PILOT-VOICE
  priority percent 10
 class CM-PILOT-VIDEO
  bandwidth percent 25
 class CM-PILOT-CRITICAL
  bandwidth percent 30
 class CM-PILOT-BULK
  bandwidth percent 10
 class class-default
  fair-queue
policy-map type epbr PM-PFP
 class CM-CRITICAL
  set forward-class 1
 class CM-BUSINESS
  set forward-class 2
 class class-default
  set forward-class 0
policy-map type inspect PM-PILOT-NGFW
 class type inspect CM-PILOT-NGFW-ALLOW
  inspect
 class class-default
  drop log
policy-map type inspect PM-PILOT-NGFW-PASS
```

```
 class type inspect CM-PILOT-NGFW-ALLOW
  pass
 class class-default
  drop log
!
!
zone security Z-PILOT-LAN
zone security Z-PILOT-SRV6-WAN
zone security Z-PILOT-LAN-U
zone-pair security ZP-LAN-TO-SRV6-WAN source Z-PILOT-LAN destination Z-PILOT-SRV6-WAN
 service-policy type inspect PM-PILOT-NGFW
zone-pair security ZP-LANU-TO-SRV6-WAN source Z-PILOT-LAN-U destination Z-PILOT-SRV6-WAN
 service-policy type inspect PM-PILOT-NGFW
zone-pair security ZP-SRV6-WAN-TO-LAN source Z-PILOT-SRV6-WAN destination Z-PILOT-LAN
 service-policy type inspect PM-PILOT-NGFW
zone-pair security ZP-SRV6-WAN-TO-LANU source Z-PILOT-SRV6-WAN destination Z-PILOT-LAN-U
 service-policy type inspect PM-PILOT-NGFW
!
!
!
!
!
!
!
crypto ipsec transform-set U-TS esp-gcm 256
 mode transport
crypto ipsec transform-set UC2-TS esp-gcm 256
 mode transport
!
crypto ipsec profile U-H22-IPSEC
 set transform-set U-TS
 set ikev2-profile U-H22-IKEV2-PROFILE
!
crypto ipsec profile U-H22-MLKEM-IPSEC
 set transform-set U-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile U-H22-IKEV2-PROFILE
!
crypto ipsec profile UC2-E31-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-E31-IKEV2-PROFILE
!
crypto ipsec profile UC2-E41-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-E41-IKEV2-PROFILE
```

```
!
crypto ipsec profile UC2-E51-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-E51-IKEV2-PROFILE
!
crypto ipsec profile UC2-E52-IPSEC
 set transform-set UC2-TS
 set pfs group19
 set ikev2-profile UC2-E52-IKEV2-PROFILE
!
!
!
!
!
!
!
!
!
!
interface Loopback0
 description CNI_UC2_BGP_UPDATE_SOURCE
 no ip address
 ip proxy-arp
 ipv6 address 2001:DB8:100:22::1/128
 ipv6 router isis SRTE-OVLY
!
interface Loopback100
 vrf forwarding CRITICAL-VRF
 ip address 10.22.22.1 255.255.255.0
 ip proxy-arp
 ipv6 address 2001:DB8:22:22::1/64
!
interface Loopback200
 no ip address
 ip proxy-arp
 ipv6 address 2001:DB8:122:200::1/128
 ipv6 router isis SRTE-OVLY
!
interface Loopback210
 vrf forwarding NON-CRITICAL-VRF
 ip address 10.122.122.1 255.255.255.0
 ip proxy-arp
 ipv6 address 2001:DB8:122:122::1/64
!
interface Tunnel212
 description HUB2_2 to EDGE1_1 unsecured GRE over IPsec
```

```
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:212::2/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/0/2
 tunnel destination 172.16.31.2
 tunnel protection ipsec profile U-H22-MLKEM-IPSEC
 clns mtu 1300
 isis circuit-type level-2-only
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel222
 description HUB2_2 to EDGE2_2 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:222::1/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/0/2
 tunnel destination 172.16.22.2
 tunnel protection ipsec profile U-H22-MLKEM-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel252
 description HUB2_2 to EDGE3_1 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:252::2/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/0/2
 tunnel destination 172.16.51.2
 tunnel protection ipsec profile U-H22-MLKEM-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
```

```
!
interface Tunnel261
 description HUB2_2 to EDGE3_2 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:261::2/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/0/2
 tunnel destination 172.16.52.2
 tunnel protection ipsec profile U-H22-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface TenGigabitEthernet0/0/0
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/0/1
 description HUB2_2_SECURED_TAGGED_WAN_MACSEC_TO_PE2
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
 macsec dot1q-in-clear 1
 macsec access-control should-secure
 service-policy output PM-PILOT-SECURED-QOS
!
interface TenGigabitEthernet0/0/1.4083
 description UC4_ELINE_EFP_TEST_TO_EDGE1_1
 encapsulation dot1Q 4083
 ip address 172.31.83.2 255.255.255.252
 ip proxy-arp
 ipv6 address 2001:DB8:4083:1::2/64
!
interface TenGigabitEthernet0/0/1.4086
 description UC4_TAGGED_WAN_MACSEC_GRE_ISIS_TO_EDGE3_1
 encapsulation dot1Q 4086
 ip address 172.31.86.2 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
```

```
 ipv6 address 2001:DB8:4086:1::2/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address ec19.2e55.5f08
 eapol eth-type 876F
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
 clns mtu 1400
 isis circuit-type level-2-only
 isis network point-to-point
 isis affinity flex-algo
  name SECURED
  !
!
interface TenGigabitEthernet0/0/1.4087
 description UC4_HUB2_2_TO_EDGE2_1_WAN_MACSEC_GRE_ISIS_VLAN4087
 encapsulation dot1Q 4087
 ip address 172.31.87.2 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:4087:1::2/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address 5000.e06c.b2a4
 eapol eth-type 876F
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
 clns mtu 1400
 isis circuit-type level-2-only
 isis network point-to-point
 isis affinity flex-algo
  name SECURED
  !
!
interface TenGigabitEthernet0/0/1.4090
 description UC4_TAGGED_WAN_MACSEC_GRE_ISIS_TO_EDGE3_2
 encapsulation dot1Q 4090
 ip address 172.31.90.2 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:4090:1::2/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address 28af.fdb3.5188
 eapol eth-type 876F
 mka policy POL-WAN-MACSEC
```

```
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
 clns mtu 1400
 isis circuit-type level-2-only
 isis network point-to-point
 isis affinity flex-algo
  name SECURED
  !
!
interface TenGigabitEthernet0/0/2
 description CNI_UC2_INET_UNDERLAY_BLR_TUNNEL222
 ip address 172.16.23.2 255.255.255.252
 ip proxy-arp
 no negotiation auto
 cdp enable
 ipv6 address 2001:DB8:2100:4::2/64
 service-policy output PM-PILOT-SECURED-QOS
!
interface TenGigabitEthernet0/0/3
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/0/4
 description HUB2_2 WAN switch/DCI segment
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
 ipv6 address 2001:DB8:4100:1::22/64
!
interface TenGigabitEthernet0/0/4.101
 description SECURED_CRITICAL_VRF_HUB2_2
 encapsulation dot1Q 101
 vrf forwarding CRITICAL-VRF
 ip address 198.18.22.1 255.255.255.0
 ip proxy-arp
 ip nbar protocol-discovery
 zone-member security Z-PILOT-LAN
 service-policy type epbr input PM-PFP
!
interface TenGigabitEthernet0/0/4.201
 description UNSECURED_NON_CRITICAL_VRF_HUB2_2
 encapsulation dot1Q 201
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.122.1 255.255.255.0
```

```
 ip proxy-arp
 ip nbar protocol-discovery
 zone-member security Z-PILOT-LAN-U
 service-policy type epbr input PM-PFP
!
interface TenGigabitEthernet0/0/4.511
 description WAN_MGMT511_HUB2_2
 encapsulation dot1Q 511
 ip address 192.168.12.22 255.255.255.0
 ip proxy-arp
!
interface TenGigabitEthernet0/0/5
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/0/6
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/0/7
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/1/0
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/1/1
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface TenGigabitEthernet0/1/2
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
```

```
interface TenGigabitEthernet0/1/3
 no ip address
 ip proxy-arp
 no negotiation auto
 cdp enable
!
interface FortyGigabitEthernet0/2/0
 no ip address
 ip proxy-arp
 shutdown
 no negotiation auto
!
interface FortyGigabitEthernet0/2/4
 no ip address
 ip proxy-arp
 shutdown
 no negotiation auto
!
interface FortyGigabitEthernet0/2/8
 no ip address
 ip proxy-arp
 shutdown
 no negotiation auto
!
interface GigabitEthernet0
 vrf forwarding Mgmt-intf
 ip address 175.13.8.36 255.255.255.0
 ip proxy-arp
 negotiation auto
!
segment-routing traffic-eng
 !
 on-demand color 100
  authorize
  candidate-paths
   preference 100
    constraints
     segments
      dataplane srv6
     !
     affinity
      include-all
       name SECURED
      !
     !
    !
    dynamic
```

```
     metric
      type te
     !
    !
   !
  !
 !
 on-demand color 200
  authorize
  candidate-paths
   preference 100
    constraints
     segments
      dataplane srv6
     !
     affinity
      include-all
       name UNSECURED
      !
     !
    !
    dynamic
     metric
      type te
     !
    !
   !
  !
 !
 on-demand color 300
  authorize
  candidate-paths
   preference 100
    per-flow
     forward-class 0 color 200
     forward-class 1 color 100
     forward-class 2 color 100
    !
   !
  !
 !
!
segment-routing srv6
 encapsulation
  source-address 2001:DB8:AC22::
 locators
  locator SLOC
```

```
   prefix 2001:DB8:AC22::/48
   format usid-f3216
!
router isis SRTE-OVLY
 net 49.0099.0099.0000.1022.00
 is-type level-2-only
 metric-style wide
 lsp-mtu 1300
 distribute link-state instance-id 99
 affinity-map SECURED bit-position 0
 affinity-map UNSECURED bit-position 1
 !
 address-family ipv6
  router-id Loopback0
  segment-routing srv6
   locator SLOC
  fast-reroute per-prefix level-2 all
  fast-reroute ti-lfa level-2
 exit-address-family
!
router bgp 65000
 bgp router-id 10.255.0.22
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 2001:DB8:100:10::1 remote-as 65000
 neighbor 2001:DB8:100:10::1 update-source Loopback0
 neighbor 2001:DB8:100:31::1 remote-as 65000
 neighbor 2001:DB8:100:31::1 update-source Loopback0
 neighbor 2001:DB8:100:41::1 remote-as 65000
 neighbor 2001:DB8:100:41::1 update-source Loopback0
 neighbor 2001:DB8:100:42::1 remote-as 65000
 neighbor 2001:DB8:100:42::1 update-source Loopback0
 neighbor 2001:DB8:100:51::1 remote-as 65000
 neighbor 2001:DB8:100:51::1 update-source Loopback0
 neighbor 2001:DB8:100:52::1 remote-as 65000
 neighbor 2001:DB8:100:52::1 update-source Loopback0
 !
 address-family ipv4
 exit-address-family
 !
 address-family vpnv4
  !
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
```

```
 neighbor 2001:DB8:100:10::1 activate
 neighbor 2001:DB8:100:10::1 send-community both
 neighbor 2001:DB8:100:10::1 next-hop-self all
 neighbor 2001:DB8:100:10::1 route-map RM_DCI_RR_NH_OUT out
 neighbor 2001:DB8:100:31::1 activate
 neighbor 2001:DB8:100:31::1 send-community both
 neighbor 2001:DB8:100:31::1 route-reflector-client
 neighbor 2001:DB8:100:41::1 activate
 neighbor 2001:DB8:100:41::1 send-community both
 neighbor 2001:DB8:100:41::1 route-reflector-client
 neighbor 2001:DB8:100:42::1 activate
 neighbor 2001:DB8:100:42::1 send-community both
 neighbor 2001:DB8:100:42::1 route-reflector-client
 neighbor 2001:DB8:100:51::1 activate
 neighbor 2001:DB8:100:51::1 send-community both
 neighbor 2001:DB8:100:51::1 route-reflector-client
 neighbor 2001:DB8:100:52::1 activate
 neighbor 2001:DB8:100:52::1 send-community both
 neighbor 2001:DB8:100:52::1 route-reflector-client
exit-address-family
!
address-family vpnv6
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
 neighbor 2001:DB8:100:10::1 activate
 neighbor 2001:DB8:100:10::1 send-community both
 neighbor 2001:DB8:100:10::1 next-hop-self all
 neighbor 2001:DB8:100:10::1 route-map RM_DCI_RR_NH_OUT out
 neighbor 2001:DB8:100:31::1 activate
 neighbor 2001:DB8:100:31::1 send-community both
 neighbor 2001:DB8:100:31::1 route-reflector-client
 neighbor 2001:DB8:100:41::1 activate
 neighbor 2001:DB8:100:41::1 send-community both
 neighbor 2001:DB8:100:41::1 route-reflector-client
 neighbor 2001:DB8:100:42::1 activate
 neighbor 2001:DB8:100:42::1 send-community both
 neighbor 2001:DB8:100:42::1 route-reflector-client
 neighbor 2001:DB8:100:51::1 activate
 neighbor 2001:DB8:100:51::1 send-community both
 neighbor 2001:DB8:100:51::1 route-reflector-client
 neighbor 2001:DB8:100:52::1 activate
 neighbor 2001:DB8:100:52::1 send-community both
 neighbor 2001:DB8:100:52::1 route-reflector-client
```

```
 exit-address-family
 !
 address-family ipv4 vrf CRITICAL-VRF
  network 10.22.22.0 mask 255.255.255.0 route-map RM_ODN_COLOR_300
  network 198.18.22.0 route-map RM_ODN_COLOR_300
  !
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
 exit-address-family
 !
 address-family ipv6 vrf CRITICAL-VRF
  network 2001:DB8:22:22::/64 route-map RM_ODN_COLOR_300
  !
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
 exit-address-family
 !
 address-family ipv4 vrf NON-CRITICAL-VRF
  network 10.122.122.0 mask 255.255.255.0
  network 198.19.122.0
  !
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
 exit-address-family
 !
 address-family ipv6 vrf NON-CRITICAL-VRF
  network 2001:DB8:122:122::/64
  !
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
 exit-address-family
!
ip rcmd domain-lookup
ip forward-protocol nd
ip forward-protocol udp tftp
```

```
!
ip telnet comport enable
ip tftp source-interface GigabitEthernet0
ip tftp blocksize 512
ip ftp passive
ip http server
ip http authentication local
ip http secure-server
ip ssh bulk-mode 131072
ip route 172.16.22.2 255.255.255.255 172.16.23.1
ip route 172.16.31.2 255.255.255.255 172.16.23.1
ip route 172.16.51.2 255.255.255.255 172.16.23.1
ip route 172.16.52.2 255.255.255.255 172.16.23.1
ip route vrf Mgmt-intf 0.0.0.0 0.0.0.0 175.13.8.1
ip route vrf Mgmt-intf 202.153.144.0 255.255.255.0 175.13.8.1
!
ip access-list extended IPV4-PILOT-NGFW-ALLOW
 10 permit ip any any
ipv6 route 2001:DB8:100:10::1/128 2001:DB8:4100:1::10
ipv6 route 2001:DB8:3100:1::/64 2001:DB8:2100:4::1
ipv6 route 2001:DB8:3100:2::/64 2001:DB8:2100:4::1
ipv6 route 2001:DB8:3100:3::/64 2001:DB8:2100:4::1
ipv6 route 2001:DB8:3100:4::/64 2001:DB8:2100:4::1
ipv6 route 2001:DB8:AC11::/48 2001:DB8:4100:1::11 250
ipv6 route 2001:DB8:AC12::/48 2001:DB8:4100:1::12 250
ipv6 route 2001:DB8:AC21::/48 2001:DB8:4100:1::21 250
ipv6 route 2001:DB8:AC31::/48 2001:DB8:4100:1::11 250
ipv6 route 2001:DB8:AC41::/48 2001:DB8:4100:1::11 250
ipv6 route 2001:DB8:AC42::/48 2001:DB8:4100:1::11 250
ipv6 route 2001:DB8:AC51::/48 2001:DB8:4100:1::12 250
route-map RM_DCI_RR_NH_OUT permit 10
 set ipv6 next-hop 2001:DB8:100:22::1
!
route-map RM_ODN_COLOR_300 permit 10
 set extcommunity color 300 additive
!
!
!
!
!
!
ipv6 access-list IPV6-PILOT-NGFW-ALLOW
 sequence 10 permit ipv6 any any
!
control-plane
!
!
```

```
!
!
!
!
line con 0
 exec-timeout 0 0
 activation-character 13
 transport preferred ssh
 transport output ssh
 stopbits 1
line aux 0
 activation-character 13
line vty 0 4
 exec-timeout 0 0
 activation-character 13
 transport input ssh
line vty 5 15
 activation-character 13
 transport input ssh
!
ntp allow mode control 3
!
!
!
!
!
!
end
```

Site Type A Secure branch edge configuration

Single-router, dual-transport branch profile (Cisco Secure Router 8161-G2) enforcing strict, logical isolation between critical enclaves and administrative transits.

Edge1_1

```
EDGE1_1#sh run
Building configuration...

Current configuration : 50514 bytes
!
! Last configuration change at 13:11:27 UTC Fri Jul 17 2026 by admin
!
version 26.1
system mode insecure
service timestamps debug datetime msec
service timestamps log datetime msec
platform qfp utilization monitor load 80
!
```

```
hostname EDGE1_1
!
boot-start-marker
boot system bootflash:c81g2be-
universalk9.BLD_V261_1_THROTTLE_LATEST_20260305_183234_V26_1_0_54.SSA.bin
! Warning: Booting with bundle mode will be deprecated in the near future. Migration to install
mode is required.
boot-end-marker
!
!
vrf definition Mgmt-intf
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
 exit-address-family
!
vrf definition CRITICAL-VRF
 rd 65000:31
 route-target export 65000:1001
 route-target import 65000:1001
 !
 address-family ipv4
  route-target export 65000:1001
  route-target import 65000:1001
 exit-address-family
 !
 address-family ipv6
  route-target export 65000:1001
  route-target import 65000:1001
 exit-address-family
!
vrf definition NON-CRITICAL-VRF
 rd 65000:131
 route-target export 65000:2001
 route-target import 65000:2001
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
 exit-address-family
!
no logging console
aaa new-model
!
!
```

```
aaa authentication login default local
aaa authentication enable default enable
aaa authorization console
aaa authorization exec default local
!
!
aaa session-id common
!
ip multicast-routing vrf CRITICAL-VRF distributed
ip multicast-routing vrf NON-CRITICAL-VRF distributed
!
!
!
!
!
!
!
!
!
login on-success log
!
!
!
!
!
!
ipv6 unicast-routing
!
!
subscriber templating
!
!
!
!
!
!
!
!
!
!
!
!
!
!
!
key chain KC-WAN-MACSEC macsec
 key 01
  cryptographic-algorithm aes-256-cmac
```

```
  key-string 3637fc51ed2f71b5a354e1aca9a91adb0f5a6744f1cda2cc7e18417cbd025e51
  lifetime local 00:00:00 May 14 2026 infinite
product-analytics
!
!
crypto pki trustpoint SLA-TrustPoint
 enrollment pkcs12
 revocation-check crl
 hash sha512
!
crypto pki trustpoint TP-self-signed-1409266032
 enrollment selfsigned
 subject-name cn=IOS-Self-Signed-Certificate-1409266032
 revocation-check none
 rsakeypair TP-self-signed-1409266032
 hash sha512
!
!
crypto pki certificate chain SLA-TrustPoint
 certificate ca 01
  30820321 30820209 A0030201 02020101 300D0609 2A864886 F70D0101 0B050030
  32310E30 0C060355 040A1305 43697363 6F312030 1E060355 04031317 43697363
  6F204C69 63656E73 696E6720 526F6F74 20434130 1E170D31 33303533 30313934
  3834375A 170D3338 30353330 31393438 34375A30 32310E30 0C060355 040A1305
  43697363 6F312030 1E060355 04031317 43697363 6F204C69 63656E73 696E6720
  526F6F74 20434130 82012230 0D06092A 864886F7 0D010101 05000382 010F0030
  82010A02 82010100 A6BCBD96 131E05F7 145EA72C 2CD686E6 17222EA1 F1EFF64D
  CBB4C798 212AA147 C655D8D7 9471380D 8711441E 1AAF071A 9CAE6388 8A38E520
  1C394D78 462EF239 C659F715 B98C0A59 5BBB5CBD 0CFEBEA3 700A8BF7 D8F256EE
  4AA4E80D DB6FD1C9 60B1FD18 FFC69C96 6FA68957 A2617DE7 104FDC5F EA2956AC
  7390A3EB 2B5436AD C847A2C5 DAB553EB 69A9A535 58E9F3E3 C0BD23CF 58BD7188
  68E69491 20F320E7 948E71D7 AE3BCC84 F10684C7 4BC8E00F 539BA42B 42C68BB7
  C7479096 B4CB2D62 EA2F505D C7B062A4 6811D95B E8250FC4 5D5D5FB8 8F27D191
  C55F0D76 61F9A4CD 3D992327 A8BB03BD 4E6D7069 7CBADF8B DF5F4368 95135E44
  DFC7C6CF 04DD7FD1 02030100 01A34230 40300E06 03551D0F 0101FF04 04030201
  06300F06 03551D13 0101FF04 05300301 01FF301D 0603551D 0E041604 1449DC85
  4B3D31E5 1B3E6A17 606AF333 3D3B4C73 E8300D06 092A8648 86F70D01 010B0500
  03820101 00507F24 D3932A66 86025D9F E838AE5C 6D4DF6B0 49631C78 240DA905
  604EDCDE FF4FED2B 77FC460E CD636FDB DD44681E 3A5673AB 9093D3B1 6C9E3D8B
  D98987BF E40CBD9E 1AECA0C2 2189BB5C 8FA85686 CD98B646 5575B146 8DFC66A8
  467A3DF4 4D565700 6ADF0F0D CF835015 3C04FF7C 21E878AC 11BA9CD2 55A9232C
  7CA7B7E6 C1AF74F6 152E99B7 B1FCF9BB E973DE7F 5BDDEB86 C71E3B49 1765308B
  5FB0DA06 B92AFE7F 494E8A9E 07B85737 F3A58BE1 1A48A229 C37C1E69 39F08678
  80DDCD16 D6BACECA EEBC7CF9 8428787B 35202CDC 60E4616A B623CDBD 230E3AFB
  418616A9 4093E049 4D10AB75 27E86F73 932E35B5 8862FDAE 0275156F 719BB2F0
  D697DF7F 28
          quit
```

```
crypto pki certificate chain TP-self-signed-1409266032
 certificate self-signed 01
  30820330 30820218 A0030201 02020101 300D0609 2A864886 F70D0101 0D050030
  31312F30 2D060355 04031326 494F532D 53656C66 2D536967 6E65642D 43657274
  69666963 6174652D 31343039 32363630 3332301E 170D3236 30343031 30383333
  32335A17 0D333630 33333130 38333332 335A3031 312F302D 06035504 03132649
  4F532D53 656C662D 5369676E 65642D43 65727469 66696361 74652D31 34303932
  36363033 32308201 22300D06 092A8648 86F70D01 01010500 0382010F 00308201
  0A028201 0100C3E7 03FDB300 B600B0A8 5F9DADCD BF66F270 B398FEDE 818BA58D
  726F6F47 D13944BA C6C3E3D1 74FDEEB1 E354A7B0 582C130E 7C400568 502E6B73
  D412D232 4DF2F7E4 F0F88918 E17144C9 D6F43681 28D70DAF 716DCF96 5F9EA2C4
  E34BB47F 70DA409C 5B23F6DC 4ADCC73A 8FB99692 154E01A4 5A815E3F 7D191F3A
  6B68BA8A 47484865 FC9B25EC 38BBC3A7 1D5B89EF 611D5C0D B88288D4 FC327171
  AD9C5FB5 11ECA69D 9C1F4BE2 744B4F34 0D3CE7EF 502EC47F 293C7872 E848D4F4
  6C69A802 C064BCAF 1BDCE8B8 D0F1997A F5AE802A 58494C01 83429CC1 14C53C2D
  2B960183 F27C6AB7 D8533347 56BEDFC9 2B9F291C BB26ACC7 E03B9F02 DFA1A370
  A7F4EA8A A88F0203 010001A3 53305130 1D060355 1D0E0416 0414FE51 D0439546
  08871F59 78122284 CF8634B7 6FE8301F 0603551D 23041830 168014FE 51D04395
  4608871F 59781222 84CF8634 B76FE830 0F060355 1D130101 FF040530 030101FF
  300D0609 2A864886 F70D0101 0D050003 82010100 728D0652 F979B85A AB14C38C
  00F9F70B 028E8956 65024223 907EEC70 821F0B5E 4F7D55E0 9B00AF41 F566EEB5
  C8B7F9DB 6822C698 B49F2064 4BADB6E8 05658BF2 58643A11 3AFD170E 95D1B257
  D1F393AC 24D13F8F 4A5C9DF1 632F9B42 E928DFDD ECB873E1 9079BAB0 C25F91B3
  D2B91B89 53D5614C 55C1BBA7 91D863CA 1B3350A2 6BDFED24 5A797533 F0C8C48E
  755EDAA8 CA13077D 068644B8 C022165D E2B04EE5 7BD4D2EC E8601250 E3B26853
  CFE1445D B7F43B3D B5935E6D 46C34C0E B241F0D3 4CB03DCC 5DFFCA11 2B00355D
  C44BBE2A 32F62027 5D23629B D717074B 72E11A21 FC9B5B18 8C9F1E03 E95D682F
  37E0774B 8D633390 4CF3B541 7E43CF0B DEE74FB5
          quit
!
!
!
!
!
!
!
!
!
diagnostic bootup level minimal
!
license udi pid C8161-G2 sn FGL2946L0LF
memory free low-watermark processor 62531
!
spanning-tree extend system-id
!
mka policy POL-WAN-MACSEC
!
```

```
!
!
enable password admin
!
username lab privilege 15 secret 9 $9$Iax15YbT2GKmzE$QtZMcL0HHGhjbjshJUlOYcKHw/naPWaxnTn0EYuy8LE
username admin privilege 15 secret 9
$9$kjgIYI7raHX36.$8nP64zGnEzQQ6B0BUMHejbtNbKvsXTQgSDyyu7hc7G2
!
redundancy
 mode none
!
!
!
crypto ikev2 proposal U-IKEV2-PROP
 pqc mlkem1024 optional
 encryption aes-gcm-256
 prf sha256
 group 19
crypto ikev2 proposal UC2-IKEV2-PROP
 pqc mlkem1024 optional
 encryption aes-gcm-256
 prf sha256
 group 19
!
crypto ikev2 policy U-IKEV2-POLICY
 proposal U-IKEV2-PROP
crypto ikev2 policy UC2-IKEV2-POLICY
 proposal UC2-IKEV2-PROP
!
crypto ikev2 keyring U-H12-PSK-KR
 peer HUB1_2
  address 172.16.12.2
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring U-H12-PPK-KR
 peer HUB1_2
  address 172.16.12.2
  ppk manual id UNSEC-E11-H12 key pqcReady123! required
 !
!
crypto ikev2 keyring U-H21-PSK-KR
 peer HUB2_1
  address 172.16.21.2
  pre-shared-key cisco123!
 !
!
```

```
crypto ikev2 keyring U-H21-PPK-KR
 peer HUB2_1
  address 172.16.21.2
  ppk manual id UNSEC-E11-H21 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-H11-PSK-KR
 peer HUB1_1
  address 2001:DB8:2100:1::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-H11-PPK-KR
 peer HUB1_1
  address 2001:DB8:2100:1::2/128
  ppk manual id UC2-E31-H11 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-H12-PSK-KR
 peer HUB1_2
  address 2001:DB8:2100:2::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-H12-PPK-KR
 peer HUB1_2
  address 2001:DB8:2100:2::2/128
  ppk manual id UC2-E31-H12 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-H21-PSK-KR
 peer HUB2_1
  address 2001:DB8:2100:3::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-H21-PPK-KR
 peer HUB2_1
  address 2001:DB8:2100:3::2/128
  ppk manual id UC2-E31-H21 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-H22-PSK-KR
 peer HUB2_2
  address 2001:DB8:2100:4::2/128
  pre-shared-key cisco123!
 !
```

```
!
crypto ikev2 keyring UC2-H22-PPK-KR
 peer HUB2_2
  address 2001:DB8:2100:4::2/128
  ppk manual id UC2-E31-H22 key pqcReady123! required
 !
!
crypto ikev2 keyring UC4-H11-PSK-KR
 peer HUB1_1
  address 172.31.94.2
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC4-H11-PPK-KR
 peer HUB1_1
  address 172.31.94.2
  ppk manual id UC4-E31-H11 key pqcReady123! required
 !
!
crypto ikev2 keyring U-H22-PSK-KR
 peer HUB2_2
  address 172.16.23.2
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring U-H22-PPK-KR
 peer HUB2_2
  address 172.16.23.2
  ppk manual id UNSEC-E11-H22 key pqcReady123! required
 !
!
crypto ikev2 keyring UC4-H12-PSK-KR
 peer HUB1_2
  address 172.31.81.2
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC4-H12-PPK-KR
 peer HUB1_2
  address 172.31.81.2
  ppk manual id UC4-E31-H12 key pqcReady123! required
 !
!
crypto ikev2 keyring UC4-H21-PSK-KR
 peer HUB2_1
  address 172.31.82.2
  pre-shared-key cisco123!
```

```
 !
!
crypto ikev2 keyring UC4-H21-PPK-KR
 peer HUB2_1
  address 172.31.82.2
  ppk manual id UC4-E31-H21 key pqcReady123! required
 !
!
crypto ikev2 keyring U-H11-PSK-KR
 peer HUB1_1
  address 172.16.11.2
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring U-H11-PPK-KR
 peer HUB1_1
  address 172.16.11.2
  ppk manual id UNSEC-E11-H11 key pqcReady123! required
 !
!
!
crypto ikev2 profile U-H12-IKEV2-PROFILE
 match identity remote address 172.16.12.2 255.255.255.255
 identity local address 172.16.31.2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk U-H12-PPK-KR
 keyring local U-H12-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile U-H21-IKEV2-PROFILE
 match identity remote address 172.16.21.2 255.255.255.255
 identity local address 172.16.31.2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk U-H21-PPK-KR
 keyring local U-H21-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-H11-IKEV2-PROFILE
 match identity remote address 2001:DB8:2100:1::2/128
 identity local address 2001:DB8:3100:1::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-H11-PPK-KR
 keyring local UC2-H11-PSK-KR
 dpd 10 2 on-demand
```

```
!
crypto ikev2 profile UC2-H12-IKEV2-PROFILE
 match identity remote address 2001:DB8:2100:2::2/128
 identity local address 2001:DB8:3100:1::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-H12-PPK-KR
 keyring local UC2-H12-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-H21-IKEV2-PROFILE
 match identity remote address 2001:DB8:2100:3::2/128
 identity local address 2001:DB8:3100:1::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-H21-PPK-KR
 keyring local UC2-H21-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-H22-IKEV2-PROFILE
 match identity remote address 2001:DB8:2100:4::2/128
 identity local address 2001:DB8:3100:1::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-H22-PPK-KR
 keyring local UC2-H22-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC4-H11-IKEV2-PROFILE
 match identity remote address 172.31.94.2 255.255.255.255
 identity local address 172.31.94.1
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC4-H11-PPK-KR
 keyring local UC4-H11-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile U-H22-IKEV2-PROFILE
 match identity remote address 172.16.23.2 255.255.255.255
 identity local address 172.16.31.2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk U-H22-PPK-KR
 keyring local U-H22-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC4-H12-IKEV2-PROFILE
```

```
 match identity remote address 172.31.81.2 255.255.255.255
 identity local address 172.31.81.1
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC4-H12-PPK-KR
 keyring local UC4-H12-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC4-H21-IKEV2-PROFILE
 match identity remote address 172.31.82.2 255.255.255.255
 identity local address 172.31.82.1
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC4-H21-PPK-KR
 keyring local UC4-H21-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile U-H11-IKEV2-PROFILE
 match identity remote address 172.16.11.2 255.255.255.255
 identity local address 172.16.31.2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk U-H11-PPK-KR
 keyring local U-H11-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 fragmentation mtu 1400
!
!
vlan internal allocation policy ascending
!
vlan 3101
 name CNI_UC2_INET_UNDERLAY
!
vlan 4094
 name CNI_UC4_EDGE1_PE_HANDOFF
!
!
class-map match-any CM-PILOT-VIDEO
 match dscp af41
class-map match-any CM-BUSINESS
 match dscp af31
class-map match-any CM-PILOT-VOICE
 match dscp ef
class-map match-any CM-CRITICAL
 match dscp ef
class-map type inspect match-any CM-PILOT-NGFW-ALLOW
```

```
 match protocol icmp
 match protocol tcp
 match protocol udp
 match access-group name IPV4-PILOT-NGFW-ALLOW
 match access-group name IPV6-PILOT-NGFW-ALLOW
class-map match-any CM-PILOT-CRITICAL
 match dscp af31
class-map match-any CM-PILOT-BULK
 match dscp af11
!
policy-map PM-PILOT-SECURED-QOS
 class CM-PILOT-VOICE
  priority percent 10
 class CM-PILOT-VIDEO
  bandwidth percent 25
 class CM-PILOT-CRITICAL
  bandwidth percent 30
 class CM-PILOT-BULK
  bandwidth percent 10
 class class-default
  fair-queue
policy-map type epbr PM-PFP
 class CM-CRITICAL
  set forward-class 1
 class CM-BUSINESS
  set forward-class 2
 class class-default
  set forward-class 0
policy-map type inspect PM-PILOT-NGFW
 class type inspect CM-PILOT-NGFW-ALLOW
  inspect
 class class-default
  drop log
!
zone security Z-PILOT-LAN
zone security Z-PILOT-SRV6-WAN
zone security Z-PILOT-LAN-U
zone-pair security ZP-LAN-TO-SRV6-WAN source Z-PILOT-LAN destination Z-PILOT-SRV6-WAN
 service-policy type inspect PM-PILOT-NGFW
zone-pair security ZP-LANU-TO-SRV6-WAN source Z-PILOT-LAN-U destination Z-PILOT-SRV6-WAN
 service-policy type inspect PM-PILOT-NGFW
zone-pair security ZP-SRV6-WAN-TO-LAN source Z-PILOT-SRV6-WAN destination Z-PILOT-LAN
 service-policy type inspect PM-PILOT-NGFW
zone-pair security ZP-SRV6-WAN-TO-LANU source Z-PILOT-SRV6-WAN destination Z-PILOT-LAN-U
 service-policy type inspect PM-PILOT-NGFW
!
!
```

```
!
!
!
!
!
!
!
crypto ipsec transform-set U-TS esp-gcm 256
 mode transport
crypto ipsec transform-set UC2-TS esp-gcm 256
 mode transport
!
crypto ipsec profile U-H11-IPSEC
 set transform-set U-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile U-H11-IKEV2-PROFILE
!
crypto ipsec profile U-H12-IPSEC
 set transform-set U-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile U-H12-IKEV2-PROFILE
!
crypto ipsec profile U-H21-IPSEC
 set transform-set U-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile U-H21-IKEV2-PROFILE
!
crypto ipsec profile U-H22-MLKEM-IPSEC
 set transform-set U-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile U-H22-IKEV2-PROFILE
!
crypto ipsec profile UC2-H11-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-H11-IKEV2-PROFILE
!
crypto ipsec profile UC2-H12-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-H12-IKEV2-PROFILE
!
crypto ipsec profile UC2-H21-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-H21-IKEV2-PROFILE
```

```
!
crypto ipsec profile UC2-H22-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-H22-IKEV2-PROFILE
!
crypto ipsec profile UC4-H11-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC4-H11-IKEV2-PROFILE
!
crypto ipsec profile UC4-H12-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC4-H12-IKEV2-PROFILE
!
crypto ipsec profile UC4-H21-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC4-H21-IKEV2-PROFILE
!
!
!
!
!
!
!
!
!
!
interface Loopback0
 description CNI_UC2_BGP_UPDATE_SOURCE
 no ip address
 ip proxy-arp
 ipv6 address 2001:DB8:100:31::1/128
 ipv6 router isis SRTE-OVLY
!
interface Loopback100
 vrf forwarding CRITICAL-VRF
 ip address 10.31.31.1 255.255.255.0
 ip proxy-arp
 ipv6 address 2001:DB8:31:31::1/64
!
interface Loopback200
 description CNI_UC2_SERVICE_LOOPBACK
 vrf forwarding CRITICAL-VRF
 no ip address
```

```
 ip proxy-arp
!
interface Loopback210
 description EDGE1_1 unsecured service loopback
 vrf forwarding NON-CRITICAL-VRF
 ip address 10.131.131.1 255.255.255.0
 ip proxy-arp
 ipv6 address 2001:DB8:131:131::1/64
!
interface Tunnel210
 description EDGE1_1 to HUB1_2 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:210::1/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source GigabitEthernet0/0/1
 tunnel destination 172.16.12.2
 tunnel protection ipsec profile U-H12-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel211
 description EDGE1_1 to HUB2_1 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:211::1/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source GigabitEthernet0/0/1
 tunnel destination 172.16.21.2
 tunnel protection ipsec profile U-H21-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel212
 description EDGE1_1 to HUB2_2 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
```

```
 ipv6 address 2001:DB8:5100:212::1/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source GigabitEthernet0/0/1
 tunnel destination 172.16.23.2
 tunnel protection ipsec profile U-H22-MLKEM-IPSEC
 clns mtu 1300
 isis circuit-type level-2-only
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel213
 description EDGE1_1 to HUB1_1 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:213::1/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source GigabitEthernet0/0/1
 tunnel destination 172.16.11.2
 tunnel protection ipsec profile U-H11-IPSEC
 clns mtu 1300
 isis circuit-type level-2-only
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel383
 description ASM_PIM_EDGE1_1_TO_HUB1_1_CRITICAL-VRF
 vrf forwarding CRITICAL-VRF
 ip address 172.31.133.2 255.255.255.252
 ip proxy-arp
 ip pim sparse-mode
 tunnel source Loopback0
 tunnel mode gre ipv6
 tunnel destination 2001:DB8:100:11::1
 tunnel path-mtu-discovery
!
interface Tunnel384
 description ASM_PIM_EDGE1_1_TO_HUB1_1_NON-CRITICAL-VRF
 vrf forwarding NON-CRITICAL-VRF
 ip address 172.31.134.2 255.255.255.252
 ip proxy-arp
 ip pim sparse-mode
```

```
 tunnel source Loopback0
 tunnel mode gre ipv6
 tunnel destination 2001:DB8:100:11::1
 tunnel key 384
 tunnel path-mtu-discovery
!
interface Tunnel4081
 description UC4_EDGE1_EXCEPTION_GRE_OVER_IPSEC_SECURED_ELINE_TO_HUB1_2
 no ip address
 ip proxy-arp
 ip mtu 1300
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5081:1::1/64
 ipv6 mtu 1300
 ipv6 router isis SRTE-OVLY
 tunnel source 172.31.81.1
 tunnel destination 172.31.81.2
 tunnel path-mtu-discovery
 tunnel protection ipsec profile UC4-H12-IPSEC
 clns mtu 1300
 isis circuit-type level-2-only
 isis affinity flex-algo
  name SECURED
  !
!
interface Tunnel4082
 description UC4_EDGE1_EXCEPTION_GRE_OVER_IPSEC_SECURED_ELINE_TO_HUB2_1
 no ip address
 ip proxy-arp
 ip mtu 1300
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5082:1::1/64
 ipv6 mtu 1300
 ipv6 router isis SRTE-OVLY
 tunnel source 172.31.82.1
 tunnel destination 172.31.82.2
 tunnel path-mtu-discovery
 tunnel protection ipsec profile UC4-H21-IPSEC
 clns mtu 1300
 isis circuit-type level-2-only
 isis affinity flex-algo
  name SECURED
  !
!
interface Tunnel4094
 description UC4_EDGE1_EXCEPTION_GRE_OVER_IPSEC_SECURED_ELINE
 no ip address
```

```
 ip proxy-arp
 ip mtu 1300
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5094:1::1/64
 ipv6 mtu 1300
 ipv6 router isis SRTE-OVLY
 tunnel source 172.31.94.1
 tunnel destination 172.31.94.2
 tunnel path-mtu-discovery
 tunnel protection ipsec profile UC4-H11-IPSEC
 clns mtu 1300
 isis circuit-type level-2-only
 isis affinity flex-algo
  name SECURED
  !
!
interface GigabitEthernet0/0/0
 description CNI_UC4_EDGE1_L3_PARENT
 no ip address
 ip proxy-arp
 ip nbar protocol-discovery
 negotiation auto
 ipv6 address 2001:DB8:3100:1::2/64
 service-policy output PM-PILOT-SECURED-QOS
!
interface GigabitEthernet0/0/0.4081
 description UC4_ELINE_EXCEPTION_TO_HUB1_2
 encapsulation dot1Q 4081
 ip address 172.31.81.1 255.255.255.252
 ip proxy-arp
 ipv6 address 2001:DB8:4081:1::1/64
!
interface GigabitEthernet0/0/0.4082
 description UC4_ELINE_EXCEPTION_TO_HUB2_1
 encapsulation dot1Q 4082
 ip address 172.31.82.1 255.255.255.252
 ip proxy-arp
 ipv6 address 2001:DB8:4082:1::1/64
!
interface GigabitEthernet0/0/0.4083
 description UC4_ELINE_EFP_TEST_TO_HUB2_2
 encapsulation dot1Q 4083
 ip address 172.31.83.1 255.255.255.252
 ip proxy-arp
 ipv6 address 2001:DB8:4083:1::1/64
!
interface GigabitEthernet0/0/0.4094
```

```
 description UC4_ELINE_PILOT_TEST_TO_HUB1_1
 encapsulation dot1Q 4094
 ip address 172.31.94.1 255.255.255.252
 ip proxy-arp
 ipv6 address 2001:DB8:4094:1::1/64
!
interface GigabitEthernet0/0/1
 description CNI_UC2_INET_UNDERLAY_BLR_TUNNELS_210_211_212
 ip address 172.16.31.2 255.255.255.252
 no ip proxy-arp
 negotiation auto
 service-policy output PM-PILOT-SECURED-QOS
!
interface GigabitEthernet0/1/0
 shutdown
!
interface GigabitEthernet0/1/1
 shutdown
!
interface GigabitEthernet0/1/2
 description LAN_SWITCH_Twe1/0/1_NATIVE_SERVICE_TAGGED_MGMT511
 switchport mode trunk
!
interface GigabitEthernet0/1/3
 shutdown
!
interface GigabitEthernet0/1/4
 shutdown
!
interface GigabitEthernet0/1/5
 shutdown
!
interface GigabitEthernet0/1/6
 switchport
 shutdown
!
interface GigabitEthernet0/1/7
 switchport
 shutdown
!
interface Vlan1
 no ip address
 ip proxy-arp
!
interface Vlan101
 description SECURED_CRITICAL_VRF_EDGE1_1
 vrf forwarding CRITICAL-VRF
```

```
 ip address 198.18.31.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 ip nbar protocol-discovery
 zone-member security Z-PILOT-LAN
 service-policy type epbr input PM-PFP
!
interface Vlan102
 description SECURED_MCAST_RECEIVER_VLAN102_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.102.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan103
 description SECURED_MCAST_RECEIVER_VLAN103_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.103.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan104
 description SECURED_MCAST_RECEIVER_VLAN104_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.104.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan105
 description SECURED_MCAST_RECEIVER_VLAN105_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.105.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan106
 description SECURED_MCAST_RECEIVER_VLAN106_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.106.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
```

```
interface Vlan107
 description SECURED_MCAST_RECEIVER_VLAN107_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.107.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan108
 description SECURED_MCAST_RECEIVER_VLAN108_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.108.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan109
 description SECURED_MCAST_RECEIVER_VLAN109_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.109.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan110
 description SECURED_MCAST_RECEIVER_VLAN110_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.110.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan111
 description SECURED_MCAST_RECEIVER_VLAN111_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.111.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan112
 description SECURED_MCAST_RECEIVER_VLAN112_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.112.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
```

```
!
interface Vlan113
 description SECURED_MCAST_RECEIVER_VLAN113_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.113.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan114
 description SECURED_MCAST_RECEIVER_VLAN114_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.114.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan115
 description SECURED_MCAST_RECEIVER_VLAN115_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.115.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan116
 description SECURED_MCAST_RECEIVER_VLAN116_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.116.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan117
 description SECURED_MCAST_RECEIVER_VLAN117_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.117.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan118
 description SECURED_MCAST_RECEIVER_VLAN118_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.118.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
```

```
 zone-member security Z-PILOT-LAN
!
interface Vlan119
 description SECURED_MCAST_RECEIVER_VLAN119_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.119.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan120
 description SECURED_MCAST_RECEIVER_VLAN120_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.120.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan121
 description SECURED_MCAST_RECEIVER_VLAN121_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.121.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan122
 description SECURED_MCAST_RECEIVER_VLAN122_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.122.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan123
 description SECURED_MCAST_RECEIVER_VLAN123_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.123.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan124
 description SECURED_MCAST_RECEIVER_VLAN124_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.124.1 255.255.255.0
 ip proxy-arp
```

```
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan125
 description SECURED_MCAST_RECEIVER_VLAN125_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.125.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan126
 description SECURED_MCAST_RECEIVER_VLAN126_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.126.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan127
 description SECURED_MCAST_RECEIVER_VLAN127_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.127.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan128
 description SECURED_MCAST_RECEIVER_VLAN128_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.128.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan129
 description SECURED_MCAST_RECEIVER_VLAN129_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.129.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan130
 description SECURED_MCAST_RECEIVER_VLAN130_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.130.1 255.255.255.0
```

```
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan131
 description SECURED_MCAST_RECEIVER_VLAN131_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.131.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan132
 description SECURED_MCAST_RECEIVER_VLAN132_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.132.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan133
 description SECURED_MCAST_RECEIVER_VLAN133_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.133.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan134
 description SECURED_MCAST_RECEIVER_VLAN134_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.134.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan135
 description SECURED_MCAST_RECEIVER_VLAN135_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.135.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan136
 description SECURED_MCAST_RECEIVER_VLAN136_EDGE1_1
 vrf forwarding CRITICAL-VRF
```

```
 ip address 198.18.136.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan137
 description SECURED_MCAST_RECEIVER_VLAN137_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.137.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan138
 description SECURED_MCAST_RECEIVER_VLAN138_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.138.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan139
 description SECURED_MCAST_RECEIVER_VLAN139_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.139.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan140
 description SECURED_MCAST_RECEIVER_VLAN140_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.140.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan141
 description SECURED_MCAST_RECEIVER_VLAN141_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.141.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan142
 description SECURED_MCAST_RECEIVER_VLAN142_EDGE1_1
```

```
 vrf forwarding CRITICAL-VRF
 ip address 198.18.142.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan143
 description SECURED_MCAST_RECEIVER_VLAN143_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.143.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan144
 description SECURED_MCAST_RECEIVER_VLAN144_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.144.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan145
 description SECURED_MCAST_RECEIVER_VLAN145_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.145.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan146
 description SECURED_MCAST_RECEIVER_VLAN146_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.146.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan147
 description SECURED_MCAST_RECEIVER_VLAN147_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.147.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan148
```

```
 description SECURED_MCAST_RECEIVER_VLAN148_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.148.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan149
 description SECURED_MCAST_RECEIVER_VLAN149_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.149.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan150
 description SECURED_MCAST_RECEIVER_VLAN150_EDGE1_1
 vrf forwarding CRITICAL-VRF
 ip address 198.18.150.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN
!
interface Vlan201
 description UNSECURED_NON_CRITICAL_VRF_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.31.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 ip nbar protocol-discovery
 zone-member security Z-PILOT-LAN-U
 service-policy type epbr input PM-PFP
!
interface Vlan202
 description UNSECURED_MCAST_RECEIVER_VLAN202_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.202.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan203
 description UNSECURED_MCAST_RECEIVER_VLAN203_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.203.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
```

```
 zone-member security Z-PILOT-LAN-U
!
interface Vlan204
 description UNSECURED_MCAST_RECEIVER_VLAN204_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.204.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan205
 description UNSECURED_MCAST_RECEIVER_VLAN205_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.205.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan206
 description UNSECURED_MCAST_RECEIVER_VLAN206_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.206.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan207
 description UNSECURED_MCAST_RECEIVER_VLAN207_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.207.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan208
 description UNSECURED_MCAST_RECEIVER_VLAN208_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.208.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan209
 description UNSECURED_MCAST_RECEIVER_VLAN209_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.209.1 255.255.255.0
 ip proxy-arp
```

```
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan210
 description UNSECURED_MCAST_RECEIVER_VLAN210_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.210.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan211
 description UNSECURED_MCAST_RECEIVER_VLAN211_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.211.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan212
 description UNSECURED_MCAST_RECEIVER_VLAN212_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.212.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan213
 description UNSECURED_MCAST_RECEIVER_VLAN213_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.213.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan214
 description UNSECURED_MCAST_RECEIVER_VLAN214_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.214.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan215
 description UNSECURED_MCAST_RECEIVER_VLAN215_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.215.1 255.255.255.0
```

```
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan216
 description UNSECURED_MCAST_RECEIVER_VLAN216_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.216.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan217
 description UNSECURED_MCAST_RECEIVER_VLAN217_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.217.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan218
 description UNSECURED_MCAST_RECEIVER_VLAN218_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.218.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan219
 description UNSECURED_MCAST_RECEIVER_VLAN219_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.219.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan220
 description UNSECURED_MCAST_RECEIVER_VLAN220_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.220.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan221
 description UNSECURED_MCAST_RECEIVER_VLAN221_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
```

```
 ip address 198.19.221.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan222
 description UNSECURED_MCAST_RECEIVER_VLAN222_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.222.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan223
 description UNSECURED_MCAST_RECEIVER_VLAN223_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.223.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan224
 description UNSECURED_MCAST_RECEIVER_VLAN224_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.224.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan225
 description UNSECURED_MCAST_RECEIVER_VLAN225_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.225.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan226
 description UNSECURED_MCAST_RECEIVER_VLAN226_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.226.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan227
 description UNSECURED_MCAST_RECEIVER_VLAN227_EDGE1_1
```

```
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.227.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan228
 description UNSECURED_MCAST_RECEIVER_VLAN228_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.228.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan229
 description UNSECURED_MCAST_RECEIVER_VLAN229_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.229.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan230
 description UNSECURED_MCAST_RECEIVER_VLAN230_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.230.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan231
 description UNSECURED_MCAST_RECEIVER_VLAN231_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.231.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan232
 description UNSECURED_MCAST_RECEIVER_VLAN232_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.232.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan233
```

```
 description UNSECURED_MCAST_RECEIVER_VLAN233_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.233.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan234
 description UNSECURED_MCAST_RECEIVER_VLAN234_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.234.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan235
 description UNSECURED_MCAST_RECEIVER_VLAN235_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.235.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan236
 description UNSECURED_MCAST_RECEIVER_VLAN236_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.236.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan237
 description UNSECURED_MCAST_RECEIVER_VLAN237_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.237.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan238
 description UNSECURED_MCAST_RECEIVER_VLAN238_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.238.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
```

```
interface Vlan239
 description UNSECURED_MCAST_RECEIVER_VLAN239_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.239.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan240
 description UNSECURED_MCAST_RECEIVER_VLAN240_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.240.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan241
 description UNSECURED_MCAST_RECEIVER_VLAN241_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.241.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan242
 description UNSECURED_MCAST_RECEIVER_VLAN242_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.242.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan243
 description UNSECURED_MCAST_RECEIVER_VLAN243_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.243.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan244
 description UNSECURED_MCAST_RECEIVER_VLAN244_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.244.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
```

```
!
interface Vlan245
 description UNSECURED_MCAST_RECEIVER_VLAN245_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.245.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan246
 description UNSECURED_MCAST_RECEIVER_VLAN246_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.246.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan247
 description UNSECURED_MCAST_RECEIVER_VLAN247_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.247.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan248
 description UNSECURED_MCAST_RECEIVER_VLAN248_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.248.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan249
 description UNSECURED_MCAST_RECEIVER_VLAN249_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.249.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 zone-member security Z-PILOT-LAN-U
!
interface Vlan250
 description UNSECURED_MCAST_RECEIVER_VLAN250_EDGE1_1
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.250.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
```

```
 zone-member security Z-PILOT-LAN-U
!
interface Vlan511
 description LAN_MGMT511_EDGE1_1
 ip address 192.168.11.11 255.255.255.0
 ip proxy-arp
!
segment-routing traffic-eng
 !
 on-demand color 100
  authorize
  candidate-paths
   preference 100
    constraints
     segments
      dataplane srv6
     !
     affinity
      include-all
       name SECURED
      !
     !
    !
    dynamic
     metric
      type te
     !
    !
   !
  !
 !
 on-demand color 200
  authorize
  candidate-paths
   preference 100
    constraints
     segments
      dataplane srv6
     !
     affinity
      include-all
       name UNSECURED
      !
     !
    !
    dynamic
     metric
```

```
      type te
     !
    !
   !
  !
 !
 on-demand color 300
  authorize
  candidate-paths
   preference 100
    per-flow
     forward-class 0 color 200
     forward-class 1 color 100
     forward-class 2 color 100
    !
   !
  !
 !
!
segment-routing srv6
 encapsulation
  source-address 2001:DB8:AC31::
 locators
  locator SLOC
   prefix 2001:DB8:AC31::/48
   format usid-f3216
!
router isis SRTE-OVLY
 net 49.0099.0099.0000.1031.00
 is-type level-2-only
 metric-style wide
 lsp-mtu 1300
 distribute link-state instance-id 99
 affinity-map SECURED bit-position 0
 affinity-map UNSECURED bit-position 1
 !
 address-family ipv6
  router-id Loopback0
  segment-routing srv6
   locator SLOC
  fast-reroute per-prefix level-2 all
  fast-reroute ti-lfa level-2
 exit-address-family
!
router bgp 65000
 bgp router-id 10.255.0.31
 bgp log-neighbor-changes
```

```
no bgp default ipv4-unicast
neighbor 2001:DB8:100:11::1 remote-as 65000
neighbor 2001:DB8:100:11::1 update-source Loopback0
neighbor 2001:DB8:100:12::1 remote-as 65000
neighbor 2001:DB8:100:12::1 update-source Loopback0
neighbor 2001:DB8:100:21::1 remote-as 65000
neighbor 2001:DB8:100:21::1 update-source Loopback0
neighbor 2001:DB8:100:22::1 remote-as 65000
neighbor 2001:DB8:100:22::1 update-source Loopback0
neighbor 2001:DB8:112:200::1 remote-as 65000
neighbor 2001:DB8:112:200::1 update-source Loopback200
neighbor 2001:DB8:121:200::1 remote-as 65000
neighbor 2001:DB8:121:200::1 update-source Loopback200
!
address-family ipv4
exit-address-family
!
address-family vpnv4
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
 neighbor 2001:DB8:100:11::1 activate
 neighbor 2001:DB8:100:11::1 send-community both
 neighbor 2001:DB8:100:12::1 activate
 neighbor 2001:DB8:100:12::1 send-community both
 neighbor 2001:DB8:100:21::1 activate
 neighbor 2001:DB8:100:21::1 send-community both
 neighbor 2001:DB8:100:22::1 activate
 neighbor 2001:DB8:100:22::1 send-community both
 neighbor 2001:DB8:112:200::1 activate
 neighbor 2001:DB8:112:200::1 send-community both
 neighbor 2001:DB8:112:200::1 next-hop-self
 neighbor 2001:DB8:121:200::1 activate
 neighbor 2001:DB8:121:200::1 send-community both
 neighbor 2001:DB8:121:200::1 next-hop-self
exit-address-family
!
address-family vpnv6
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
```

```
 neighbor 2001:DB8:100:11::1 activate
 neighbor 2001:DB8:100:11::1 send-community both
 neighbor 2001:DB8:100:12::1 activate
 neighbor 2001:DB8:100:12::1 send-community both
 neighbor 2001:DB8:100:21::1 activate
 neighbor 2001:DB8:100:21::1 send-community both
 neighbor 2001:DB8:100:22::1 activate
 neighbor 2001:DB8:100:22::1 send-community both
 neighbor 2001:DB8:112:200::1 activate
 neighbor 2001:DB8:112:200::1 send-community both
 neighbor 2001:DB8:112:200::1 next-hop-self
 neighbor 2001:DB8:121:200::1 activate
 neighbor 2001:DB8:121:200::1 send-community both
 neighbor 2001:DB8:121:200::1 next-hop-self
exit-address-family
!
address-family ipv4 vrf CRITICAL-VRF
 network 10.31.31.0 mask 255.255.255.0 route-map RM_ODN_COLOR_300
 network 198.18.31.0 route-map RM_ODN_COLOR_300
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
exit-address-family
!
address-family ipv6 vrf CRITICAL-VRF
 network 2001:DB8:31:31::/64 route-map RM_ODN_COLOR_300
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
exit-address-family
!
address-family ipv4 vrf NON-CRITICAL-VRF
 network 10.131.131.0 mask 255.255.255.0
 network 198.19.31.0
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
exit-address-family
```

```
 !
 address-family ipv6 vrf NON-CRITICAL-VRF
  network 2001:DB8:131:131::/64
  !
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
 exit-address-family
!
ip rcmd domain-lookup
ip forward-protocol nd
ip forward-protocol udp tftp
ip pim vrf CRITICAL-VRF rp-address 10.255.13.100 ACL-ASM-PILOT-GROUPS
ip pim vrf NON-CRITICAL-VRF rp-address 10.255.13.200 ACL-ASM-PILOT-GROUPS
ip telnet comport enable
ip tftp blocksize 512
ip ftp passive
no ip http server
ip http secure-server
!
ip route 172.16.11.2 255.255.255.255 172.16.31.1
ip route 172.16.12.2 255.255.255.255 172.16.31.1
ip route 172.16.21.2 255.255.255.255 172.16.31.1
ip route 172.16.23.2 255.255.255.255 172.16.31.1
ip route vrf Mgmt-intf 0.0.0.0 0.0.0.0 175.13.8.1
ip route vrf Mgmt-intf 202.153.144.0 255.255.255.0 175.13.8.1
ip route vrf CRITICAL-VRF 10.255.13.100 255.255.255.255 Tunnel383
ip route vrf NON-CRITICAL-VRF 10.255.13.200 255.255.255.255 Tunnel384
ip route vrf NON-CRITICAL-VRF 198.19.111.0 255.255.255.0 Tunnel384
ip ssh bulk-mode 131072
!
ip access-list standard ACL-ASM-PILOT-GROUPS
 10 permit 239.18.0.0 0.0.255.255
!
ip access-list extended IPV4-PILOT-NGFW-ALLOW
 10 permit ip any any
ipv6 route 2001:DB8:2100:1::/64 2001:DB8:3100:1::1
ipv6 route 2001:DB8:2100:2::/64 2001:DB8:3100:1::1
ipv6 route 2001:DB8:2100:3::/64 2001:DB8:3100:1::1
ipv6 route 2001:DB8:2100:4::/64 2001:DB8:3100:1::1
ipv6 route 2001:DB8:AC11::/48 Tunnel4094 250
ipv6 route 2001:DB8:AC12::/48 Tunnel4094 250
ipv6 route 2001:DB8:AC21::/48 Tunnel4094 250
ipv6 route 2001:DB8:AC22::/48 Tunnel4094 250
ipv6 route 2001:DB8:AC41::/48 Tunnel4094 250
```

```
ipv6 route 2001:DB8:AC42::/48 Tunnel4094 250
ipv6 route 2001:DB8:AC51::/48 Tunnel4094 250
ipv6 route 2001:DB8:AC52::/48 Tunnel4094 250
route-map RM_ODN_COLOR_300 permit 10
 set extcommunity color 300 additive
!
route-map RM_SRV6_NH_T4094_OUT permit 10
 set ipv6 next-hop 2001:DB8:5094:1::1
!
!
!
!
!
!
ipv6 access-list IPV6-PILOT-NGFW-ALLOW
 sequence 10 permit ipv6 any any
!
control-plane
!
!
mgcp behavior rsip-range tgcp-only
mgcp behavior comedia-role none
mgcp behavior comedia-check-media-src disable
mgcp behavior comedia-sdp-force disable
!
mgcp profile default
!
!
!
!
!
!
line con 0
 exec-timeout 0 0
 activation-character 13
 transport preferred ssh
 transport output ssh
 stopbits 1
line vty 0 4
 exec-timeout 0 0
 activation-character 13
 transport input ssh
line vty 5 15
 activation-character 13
 transport input ssh
!
ntp allow mode control 3
```

```
!
!
!
!
!
!
end
```

Site Type B Medium-site primary edge configuration

The primary secure router (Cisco Secure Router 8375-E-G2) executing WAN MACsec line-rate encryption, VRRP active LAN-side tracking..

Edge2_1

```
EDGE2_1#sh run
Building configuration...

Current configuration : 23860 bytes
!
! Last configuration change at 05:32:11 UTC Sat Aug 8 2026 by admin
!
version 26.1
system mode insecure
service timestamps debug datetime msec
service timestamps log datetime msec
platform qfp utilization monitor load 80
!
hostname EDGE2_1
!
boot-start-marker
boot system bootflash:c8kg2be-
universalk9.BLD_V261_1_THROTTLE_LATEST_20260305_183234_V26_1_0_54.SSA.bin
! Warning: Booting with bundle mode will be deprecated in the near future. Migration to install
mode is required.
boot-end-marker
!
!
vrf definition Mgmt-intf
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
 exit-address-family
!
vrf definition CRITICAL-VRF
 rd 65000:41
 route-target export 65000:1001
```

```
 route-target import 65000:1001
 !
 address-family ipv4
  route-target export 65000:1001
  route-target import 65000:1001
 exit-address-family
 !
 address-family ipv6
  route-target export 65000:1001
  route-target import 65000:1001
 exit-address-family
!
vrf definition NON-CRITICAL-VRF
 rd 65000:141
 route-target export 65000:2001
 route-target import 65000:2001
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
 exit-address-family
!
no logging console
aaa new-model
!
!
aaa authentication login default local
aaa authentication enable default enable
aaa authorization console
aaa authorization exec default local
!
!
aaa session-id common
!
!
!
!
!
!
!
!
!
!
login on-success log
!
!
```

```
!
!
!
!
ipv6 unicast-routing
!
!
subscriber templating
!
!
!
!
!
!
!
!
!
!
!
!
!
key chain KC-WAN-MACSEC macsec
 key 01
  cryptographic-algorithm aes-256-cmac
  key-string 3637fc51ed2f71b5a354e1aca9a91adb0f5a6744f1cda2cc7e18417cbd025e51
  lifetime local 00:00:00 May 13 2026 infinite
product-analytics
!
!
crypto pki trustpoint SLA-TrustPoint
 enrollment pkcs12
 revocation-check crl
 hash sha512
!
crypto pki trustpoint TP-self-signed-598932909
 enrollment selfsigned
 subject-name cn=IOS-Self-Signed-Certificate-598932909
 revocation-check none
 rsakeypair TP-self-signed-598932909
 hash sha512
!
!
crypto pki certificate chain SLA-TrustPoint
 certificate ca 01
  30820321 30820209 A0030201 02020101 300D0609 2A864886 F70D0101 0B050030
  32310E30 0C060355 040A1305 43697363 6F312030 1E060355 04031317 43697363
```

```
  6F204C69 63656E73 696E6720 526F6F74 20434130 1E170D31 33303533 30313934
  3834375A 170D3338 30353330 31393438 34375A30 32310E30 0C060355 040A1305
  43697363 6F312030 1E060355 04031317 43697363 6F204C69 63656E73 696E6720
  526F6F74 20434130 82012230 0D06092A 864886F7 0D010101 05000382 010F0030
  82010A02 82010100 A6BCBD96 131E05F7 145EA72C 2CD686E6 17222EA1 F1EFF64D
  CBB4C798 212AA147 C655D8D7 9471380D 8711441E 1AAF071A 9CAE6388 8A38E520
  1C394D78 462EF239 C659F715 B98C0A59 5BBB5CBD 0CFEBEA3 700A8BF7 D8F256EE
  4AA4E80D DB6FD1C9 60B1FD18 FFC69C96 6FA68957 A2617DE7 104FDC5F EA2956AC
  7390A3EB 2B5436AD C847A2C5 DAB553EB 69A9A535 58E9F3E3 C0BD23CF 58BD7188
  68E69491 20F320E7 948E71D7 AE3BCC84 F10684C7 4BC8E00F 539BA42B 42C68BB7
  C7479096 B4CB2D62 EA2F505D C7B062A4 6811D95B E8250FC4 5D5D5FB8 8F27D191
  C55F0D76 61F9A4CD 3D992327 A8BB03BD 4E6D7069 7CBADF8B DF5F4368 95135E44
  DFC7C6CF 04DD7FD1 02030100 01A34230 40300E06 03551D0F 0101FF04 04030201
  06300F06 03551D13 0101FF04 05300301 01FF301D 0603551D 0E041604 1449DC85
  4B3D31E5 1B3E6A17 606AF333 3D3B4C73 E8300D06 092A8648 86F70D01 010B0500
  03820101 00507F24 D3932A66 86025D9F E838AE5C 6D4DF6B0 49631C78 240DA905
  604EDCDE FF4FED2B 77FC460E CD636FDB DD44681E 3A5673AB 9093D3B1 6C9E3D8B
  D98987BF E40CBD9E 1AECA0C2 2189BB5C 8FA85686 CD98B646 5575B146 8DFC66A8
  467A3DF4 4D565700 6ADF0F0D CF835015 3C04FF7C 21E878AC 11BA9CD2 55A9232C
  7CA7B7E6 C1AF74F6 152E99B7 B1FCF9BB E973DE7F 5BDDEB86 C71E3B49 1765308B
  5FB0DA06 B92AFE7F 494E8A9E 07B85737 F3A58BE1 1A48A229 C37C1E69 39F08678
  80DDCD16 D6BACECA EEBC7CF9 8428787B 35202CDC 60E4616A B623CDBD 230E3AFB
  418616A9 4093E049 4D10AB75 27E86F73 932E35B5 8862FDAE 0275156F 719BB2F0
  D697DF7F 28
           quit
crypto pki certificate chain TP-self-signed-598932909
 certificate self-signed 01
  3082032E 30820216 A0030201 02020101 300D0609 2A864886 F70D0101 0D050030
  30312E30 2C060355 04031325 494F532D 53656C66 2D536967 6E65642D 43657274
  69666963 6174652D 35393839 33323930 39301E17 0D323630 34303130 38313930
  335A170D 33363033 33313038 31393033 5A303031 2E302C06 03550403 1325494F
  532D5365 6C662D53 69676E65 642D4365 72746966 69636174 652D3539 38393332
  39303930 82012230 0D06092A 864886F7 0D010101 05000382 010F0030 82010A02
  82010100 D5511398 61DED442 6DD0038A 5060F999 70C9B266 E75F7E9B E8839EF4
  06850D16 F90FCDD3 972E3545 DB5CDBBD 097DE8A0 D58F71D1 8F00B5ED 5FD8EE01
  F2C943D9 DB4A1958 36A24014 9A4672CB 2C9D2F06 E516719C D67044D9 FDC09706
  8416B3B7 0E07AACA 4C540A24 803EC22C 068112C8 825416FB 3B55BE29 EA3EA289
  883852DC 2E5B86EE 599D09BD 71CE63B6 9F949C0A 562719C3 C72C164E CF3AD9F9
  8EAC93AF 67011109 2DF01ACB 4E36FAFA AA893249 06673EA6 B8446DFD 535CC1AE
  9A7EC7FB E7483717 0C589A55 2285FFEB B3AFDF4E 8A3AE88F 7DFC11CE 89253ABD
  52795271 BA70A239 B5B73867 20CD93B2 3CC3EF1C 0BCD3A0F DF56F0AE 4E90CDD4
  A81DD283 02030100 01A35330 51301D06 03551D0E 04160414 8F6A8912 667CA96C
  B50EA7CC 5362FD17 5B0BC601 301F0603 551D2304 18301680 148F6A89 12667CA9
  6CB50EA7 CC5362FD 175B0BC6 01300F06 03551D13 0101FF04 05300301 01FF300D
  06092A86 4886F70D 01010D05 00038201 0100A8B6 4F85A71D EBDFE461 9F9B2395
  AAF3FB8C A6ABD5EC BEE0339C 47FEFD09 98777937 D4A54159 49E5AD86 B68F7D45
  EDE40652 0D36563D FC7FA865 3CFC6BDC 4767D467 181B62EE C46A0CE6 7553AEBF
```

```
  D4A3DBAA 9C42EF38 84102832 BF6E589A 82120552 B1A7A34E 8CD8E48B 53B502FD
  2D4311C3 32C824FD F68BF234 985B650F 979909F9 7B4765F6 7154BD9A CF04C565
  0ABF3BB8 860BA9D1 8C8C9245 0CC7A0E3 13C2ADDF 43054798 248940D1 954D4C49
  49D28E9D 604C7B2E A4455A47 CEEDBFEF 23E8BB5A 42EA76A8 4CD76B89 BFAFE857
  C060A579 7344DE5F 1EDE41DC E3DD2F3B FE132AC5 F98DB941 70763914 39AA56C9
  BCD52154 35B4A879 E06A0706 F3FAFCA3 8427
          quit
!
!
!
!
!
!
!
!
!
!
voice-card 0/4
!
diagnostic bootup level minimal
!
license udi pid C8375-E-G2 sn FDO2922M0B7
memory free low-watermark processor 62496
!
spanning-tree extend system-id
!
mka policy POL-WAN-MACSEC
!
!
!
!
username lab privilege 15 secret 9 $9$pljcNe68rquL0E$5IDc7Mo4LJe6AJ.iIP85WW4rPoWilUSaCXrAzNeXOwg
username admin privilege 15 secret 9
$9$4xresnAmOJOHmk$HbXnyVu9QbvTsXNN8Ma2HMM.zRCTv76j/55SBPdt/wc
!
redundancy
 mode none
!
!
!
crypto ikev2 proposal UC2-IKEV2-PROP
 pqc mlkem1024 optional
 encryption aes-gcm-256
 prf sha256
 group 19
!
crypto ikev2 policy UC2-IKEV2-POLICY
```

```
 proposal UC2-IKEV2-PROP
!
crypto ikev2 keyring UC2-H11-PSK-KR
 peer HUB1_1
  address 2001:DB8:2100:1::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-H11-PPK-KR
 peer HUB1_1
  address 2001:DB8:2100:1::2/128
  ppk manual id UC2-E41-H11 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-H12-PSK-KR
 peer HUB1_2
  address 2001:DB8:2100:2::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-H12-PPK-KR
 peer HUB1_2
  address 2001:DB8:2100:2::2/128
  ppk manual id UC2-E41-H12 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-H21-PSK-KR
 peer HUB2_1
  address 2001:DB8:2100:3::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-H21-PPK-KR
 peer HUB2_1
  address 2001:DB8:2100:3::2/128
  ppk manual id UC2-E41-H21 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-H22-PSK-KR
 peer HUB2_2
  address 2001:DB8:2100:4::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-H22-PPK-KR
 peer HUB2_2
  address 2001:DB8:2100:4::2/128
```

```
  ppk manual id UC2-E41-H22 key pqcReady123! required
 !
!
!
crypto ikev2 profile UC2-H11-IKEV2-PROFILE
 match identity remote address 2001:DB8:2100:1::2/128
 identity local address 2001:DB8:3100:2::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-H11-PPK-KR
 keyring local UC2-H11-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-H12-IKEV2-PROFILE
 match identity remote address 2001:DB8:2100:2::2/128
 identity local address 2001:DB8:3100:2::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-H12-PPK-KR
 keyring local UC2-H12-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-H21-IKEV2-PROFILE
 match identity remote address 2001:DB8:2100:3::2/128
 identity local address 2001:DB8:3100:2::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-H21-PPK-KR
 keyring local UC2-H21-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-H22-IKEV2-PROFILE
 match identity remote address 2001:DB8:2100:4::2/128
 identity local address 2001:DB8:3100:2::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-H22-PPK-KR
 keyring local UC2-H22-PSK-KR
 dpd 10 2 on-demand
!
!
!
!
!
track 41 ip sla 41 reachability
!
class-map match-any CM-BUSINESS
```

```
 match dscp af31
class-map match-any CM-PILOT-VIDEO
 match dscp af41
class-map match-any CM-PILOT-VOICE
 match dscp ef
class-map match-any CM-CRITICAL
 match dscp ef
class-map type inspect match-any CM-PILOT-NGFW-ALLOW
 match protocol icmp
 match protocol tcp
 match protocol udp
 match access-group name IPV6-PILOT-NGFW-ALLOW
 match access-group name IPV4-PILOT-NGFW-ALLOW
class-map match-any CM-PILOT-CRITICAL
 match dscp af31
class-map match-any CM-PILOT-BULK
 match dscp af11
!
policy-map PM-PILOT-SECURED-QOS
 class CM-PILOT-VOICE
  priority percent 10
 class CM-PILOT-VIDEO
  bandwidth percent 25
 class CM-PILOT-CRITICAL
  bandwidth percent 30
 class CM-PILOT-BULK
  bandwidth percent 10
 class class-default
  fair-queue
policy-map type epbr PM-PFP
 class CM-CRITICAL
  set forward-class 1
 class CM-BUSINESS
  set forward-class 2
 class class-default
  set forward-class 0
policy-map type inspect PM-PILOT-NGFW
 class type inspect CM-PILOT-NGFW-ALLOW
  inspect
 class class-default
  drop log
policy-map type inspect PM-PILOT-NGFW-PASS
 class type inspect CM-PILOT-NGFW-ALLOW
  pass
 class class-default
  drop log
policy-map PM-PILOT-TUNNEL-QOS
```

```
 class CM-PILOT-VOICE
  bandwidth percent 10
 class CM-PILOT-VIDEO
  bandwidth percent 25
 class CM-PILOT-CRITICAL
  bandwidth percent 30
 class CM-PILOT-BULK
  bandwidth percent 10
 class class-default
  bandwidth percent 25
!
zone security Z-PILOT-LAN
zone security Z-PILOT-SRV6-WAN
zone security Z-PILOT-LAN-U
zone-pair security ZP-LAN-TO-SRV6-WAN source Z-PILOT-LAN destination Z-PILOT-SRV6-WAN
 service-policy type inspect PM-PILOT-NGFW
zone-pair security ZP-LANU-TO-SRV6-WAN source Z-PILOT-LAN-U destination Z-PILOT-SRV6-WAN
 service-policy type inspect PM-PILOT-NGFW
zone-pair security ZP-SRV6-WAN-TO-LAN source Z-PILOT-SRV6-WAN destination Z-PILOT-LAN
 service-policy type inspect PM-PILOT-NGFW
zone-pair security ZP-SRV6-WAN-TO-LANU source Z-PILOT-SRV6-WAN destination Z-PILOT-LAN-U
 service-policy type inspect PM-PILOT-NGFW
!
!
!
!
!
!
!
!
!
!
crypto ipsec transform-set UC2-TS esp-gcm 256
 mode transport
!
crypto ipsec profile UC2-H11-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-H11-IKEV2-PROFILE
!
crypto ipsec profile UC2-H12-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-H12-IKEV2-PROFILE
!
crypto ipsec profile UC2-H21-IPSEC
 set transform-set UC2-TS
```

```
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-H21-IKEV2-PROFILE
!
crypto ipsec profile UC2-H22-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-H22-IKEV2-PROFILE
!
!
!
!
!
!
!
!
!
interface Loopback0
 ip address 10.255.0.41 255.255.255.255
 ip proxy-arp
 ipv6 address 2001:DB8:100:41::1/128
 ipv6 router isis SRTE-OVLY
!
interface Loopback100
 vrf forwarding CRITICAL-VRF
 ip address 10.41.41.1 255.255.255.0
 ip proxy-arp
 ipv6 address 2001:DB8:41:41::1/64
!
interface Loopback200
 no ip address
 ip proxy-arp
!
interface Loopback210
 description EDGE2_1 backup unsecured l3vpn service loopback
 vrf forwarding NON-CRITICAL-VRF
 ip address 10.142.142.1 255.255.255.0
 ip proxy-arp
 ipv6 address 2001:DB8:142:142::1/64
!
interface TwoGigabitEthernet0/0/0
 no ip address
 ip proxy-arp
 negotiation auto
!
interface TwoGigabitEthernet0/0/0.101
 description SECURED_CRITICAL_VRF_EDGE2_1
 encapsulation dot1Q 101
```

```
 vrf forwarding CRITICAL-VRF
 ip address 198.18.41.2 255.255.255.0
 ip proxy-arp
 ip nbar protocol-discovery
 zone-member security Z-PILOT-LAN
 vrrp 41 ip 198.18.41.1
 vrrp 41 priority 150
 vrrp 41 track 41 decrement 80
 service-policy type epbr input PM-PFP
!
interface TwoGigabitEthernet0/0/0.201
 description UNSECURED_NON_CRITICAL_VRF_EDGE2_1_BACKUP
 encapsulation dot1Q 201
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.42.2 255.255.255.0
 ip proxy-arp
 ip nbar protocol-discovery
 zone-member security Z-PILOT-LAN-U
 vrrp 42 ip 198.19.42.1
 vrrp 42 track 41 decrement 80
 service-policy type epbr input PM-PFP
!
interface TwoGigabitEthernet0/0/0.511
 description LAN_MGMT511_EDGE2_1
 encapsulation dot1Q 511
 ip address 192.168.11.12 255.255.255.0
 ip proxy-arp
 ipv6 address 2001:DB8:511:11::12/64
!
interface TwoGigabitEthernet0/0/1
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface TwoGigabitEthernet0/0/2
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface TwoGigabitEthernet0/0/3
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
```

```
interface TenGigabitEthernet0/0/4
 description CNI_UC4_EDGE2_WAN_MACSEC_PARENT
 no ip address
 ip proxy-arp
 ipv6 address 2001:DB8:3100:2::2/64
 macsec dot1q-in-clear 1
 macsec access-control should-secure
 service-policy output PM-PILOT-SECURED-QOS
!
interface TenGigabitEthernet0/0/4.4087
 description UC4_EDGE2_1_TO_HUB2_2_WAN_MACSEC_GRE_ISIS_VLAN4087
 encapsulation dot1Q 4087
 ip address 172.31.87.1 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:4087:1::1/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address 10e6.769b.9881
 eapol eth-type 876F
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
 clns mtu 1400
 isis circuit-type level-2-only
 isis network point-to-point
 isis affinity flex-algo
  name SECURED
  !
!
interface TenGigabitEthernet0/0/4.4088
 description UC4_EDGE2_1_TO_HUB2_1_WAN_MACSEC_GRE_ISIS_VLAN4088
 encapsulation dot1Q 4088
 ip address 172.31.88.1 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:4088:1::1/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address d862.ca95.da81
 eapol eth-type 876F
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
 clns mtu 1400
 isis circuit-type level-2-only
 isis network point-to-point
```

```
 isis affinity flex-algo
  name SECURED
  !
!
interface TenGigabitEthernet0/0/4.4089
 description UC4_EDGE2_1_TO_HUB1_2_WAN_MACSEC_GRE_ISIS_VLAN4089
 encapsulation dot1Q 4089
 ip address 172.31.89.1 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:4089:1::1/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address d862.ca95.d101
 eapol eth-type 876F
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
 clns mtu 1400
 isis circuit-type level-2-only
 isis network point-to-point
 isis affinity flex-algo
  name SECURED
  !
!
interface TenGigabitEthernet0/0/4.4093
 description UC4_WAN_MACSEC_PILOT_TO_HUB1_1_VLAN4093
 encapsulation dot1Q 4093
 ip address 172.31.93.1 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:4093:1::1/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address d862.caba.a401
 eapol eth-type 876F
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
 clns mtu 1400
 isis circuit-type level-2-only
 isis network point-to-point
 isis affinity flex-algo
  name SECURED
  !
!
interface TenGigabitEthernet0/0/4.4095
```

```
 ip proxy-arp
!
interface TenGigabitEthernet0/0/5
 description MEDIUM_SITE_EDGE2_INTER_EDGE_XCONNECT_TO_EDGE2_2
 no ip address
 ip proxy-arp
 ipv6 address 2001:DB8:245:12::1/64
 ipv6 router isis SRTE-OVLY
 isis network point-to-point
 isis metric 1000 level-2
!
interface Service-Engine0/4/0
!
interface GigabitEthernet0
 vrf forwarding Mgmt-intf
 ip address 175.13.8.22 255.255.255.0
 ip proxy-arp
 negotiation auto
!
interface Vlan1
 no ip address
 ip proxy-arp
!
segment-routing traffic-eng
 !
 on-demand color 100
  authorize
  candidate-paths
   preference 100
    constraints
     segments
      dataplane srv6
     !
     affinity
      include-all
       name SECURED
      !
     !
    !
    dynamic
     metric
      type te
     !
    !
   !
  !
 !
```

```
 on-demand color 200
  authorize
  candidate-paths
   preference 100
    constraints
     segments
      dataplane srv6
     !
     affinity
      include-all
       name UNSECURED
      !
     !
    !
    dynamic
     metric
      type te
     !
    !
   !
  !
 !
 on-demand color 300
  authorize
  candidate-paths
   preference 100
    per-flow
     forward-class 0 color 200
     forward-class 1 color 100
     forward-class 2 color 100
    !
   !
  !
 !
!
segment-routing srv6
 encapsulation
  source-address 2001:DB8:AC41::
 locators
  locator SLOC
   prefix 2001:DB8:AC41::/48
   format usid-f3216
!
router isis SRTE-OVLY
 net 49.0099.0099.0000.1041.00
 is-type level-2-only
 metric-style wide
```

```
 lsp-mtu 1300
 distribute link-state instance-id 99
 affinity-map SECURED bit-position 0
 affinity-map UNSECURED bit-position 1
 !
 address-family ipv6
  router-id Loopback0
  segment-routing srv6
   locator SLOC
  fast-reroute per-prefix level-2 all
  fast-reroute ti-lfa level-2
 exit-address-family
!
router bgp 65000
 !
 segment-routing srv6
  locator SLOC
 exit-srv6
 !
 bgp router-id 10.255.0.41
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 2001:DB8:100:11::1 remote-as 65000
 neighbor 2001:DB8:100:11::1 update-source Loopback0
 neighbor 2001:DB8:100:12::1 remote-as 65000
 neighbor 2001:DB8:100:12::1 update-source Loopback0
 neighbor 2001:DB8:100:21::1 remote-as 65000
 neighbor 2001:DB8:100:21::1 update-source Loopback0
 neighbor 2001:DB8:100:22::1 remote-as 65000
 neighbor 2001:DB8:100:22::1 update-source Loopback0
 neighbor 2001:DB8:511:11::208 remote-as 65000
 neighbor 192.168.255.208 remote-as 65000
 !
 address-family ipv4
 exit-address-family
 !
 address-family vpnv4
  !
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
  neighbor 2001:DB8:100:11::1 activate
  neighbor 2001:DB8:100:11::1 send-community both
  neighbor 2001:DB8:100:11::1 route-map RM-E21-UNSEC-A-BACKUP out
  neighbor 2001:DB8:100:12::1 activate
```

```
 neighbor 2001:DB8:100:12::1 send-community both
 neighbor 2001:DB8:100:12::1 route-map RM-E21-UNSEC-A-BACKUP out
 neighbor 2001:DB8:100:21::1 activate
 neighbor 2001:DB8:100:21::1 send-community both
 neighbor 2001:DB8:100:21::1 route-map RM-E21-UNSEC-A-BACKUP out
 neighbor 2001:DB8:100:22::1 activate
 neighbor 2001:DB8:100:22::1 send-community both
 neighbor 2001:DB8:100:22::1 route-map RM-E21-UNSEC-A-BACKUP out
exit-address-family
!
address-family vpnv6
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
 neighbor 2001:DB8:100:11::1 activate
 neighbor 2001:DB8:100:11::1 send-community both
 neighbor 2001:DB8:100:11::1 route-map RM-E21-UNSEC-A-BACKUP-V6 out
 neighbor 2001:DB8:100:12::1 activate
 neighbor 2001:DB8:100:12::1 send-community both
 neighbor 2001:DB8:100:12::1 route-map RM-E21-UNSEC-A-BACKUP-V6 out
 neighbor 2001:DB8:100:21::1 activate
 neighbor 2001:DB8:100:21::1 send-community both
 neighbor 2001:DB8:100:21::1 route-map RM-E21-UNSEC-A-BACKUP-V6 out
 neighbor 2001:DB8:100:22::1 activate
 neighbor 2001:DB8:100:22::1 send-community both
 neighbor 2001:DB8:100:22::1 route-map RM-E21-UNSEC-A-BACKUP-V6 out
exit-address-family
!
address-family link-state link-state
 neighbor 2001:DB8:511:11::208 activate
 neighbor 2001:DB8:511:11::208 send-community both
 neighbor 192.168.255.208 activate
 neighbor 192.168.255.208 send-community both
exit-address-family
!
address-family ipv4 vrf CRITICAL-VRF
 network 10.41.41.0 mask 255.255.255.0 route-map RM_ODN_COLOR_300
 network 198.18.41.0 route-map RM_ODN_COLOR_300
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
```

```
 exit-address-family
 !
 address-family ipv6 vrf CRITICAL-VRF
  network 2001:DB8:41:41::/64 route-map RM_ODN_COLOR_300
  !
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
 exit-address-family
 !
 address-family ipv4 vrf NON-CRITICAL-VRF
  network 10.142.142.0 mask 255.255.255.0
  network 198.19.42.0
  !
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
 exit-address-family
 !
 address-family ipv6 vrf NON-CRITICAL-VRF
  network 2001:DB8:142:142::/64
  !
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
 exit-address-family
!
ip rcmd domain-lookup
ip forward-protocol nd
ip forward-protocol udp tftp
ip telnet comport enable
ip tftp blocksize 512
ip ftp passive
no ip http server
ip http secure-server
!
ip route 192.168.255.208 255.255.255.255 192.168.11.208
ip route vrf Mgmt-intf 0.0.0.0 0.0.0.0 175.13.8.1
ip route vrf Mgmt-intf 202.153.144.0 255.255.255.0 175.13.8.1
ip ssh bulk-mode 131072
!
```

```
ip access-list extended IPV4-PILOT-NGFW-ALLOW
 10 permit ip any any
!
ip prefix-list PL-E21-UNSEC-A-SVC seq 5 permit 198.19.42.0/24
ip prefix-list PL-E21-UNSEC-A-SVC seq 10 permit 10.142.142.0/24
ip sla 41
 icmp-echo 172.31.93.2 source-ip 172.31.93.1
  frequency 5
ip sla schedule 41 life forever start-time now
ipv6 route 2001:DB8:511:255::208/128 2001:DB8:511:11::208
ipv6 route 2001:DB8:2100:1::/64 2001:DB8:3100:2::1
ipv6 route 2001:DB8:2100:2::/64 2001:DB8:3100:2::1
ipv6 route 2001:DB8:2100:3::/64 2001:DB8:3100:2::1
ipv6 route 2001:DB8:2100:4::/64 2001:DB8:3100:2::1
!
ipv6 prefix-list PL-E21-UNSEC-A-SVC6 seq 5 permit 2001:DB8:142:142::/64
route-map RM_ODN_COLOR_300 permit 10
 set extcommunity color 300 additive
!
route-map RM-E21-UNSEC-A-BACKUP permit 10
 match ip address prefix-list PL-E21-UNSEC-A-SVC
 set local-preference 50
!
route-map RM-E21-UNSEC-A-BACKUP permit 100
!
route-map RM-E21-UNSEC-A-BACKUP-V6 permit 10
 match ipv6 address prefix-list PL-E21-UNSEC-A-SVC6
 set local-preference 50
!
route-map RM-E21-UNSEC-A-BACKUP-V6 permit 100
!
!
!
performance-measurement
!
!
!
!
ipv6 access-list IPV6-PILOT-NGFW-ALLOW
 sequence 10 permit ipv6 any any
!
control-plane
!
!
mgcp behavior rsip-range tgcp-only
mgcp behavior comedia-role none
mgcp behavior comedia-check-media-src disable
```

```
mgcp behavior comedia-sdp-force disable
!
mgcp profile default
!
!
!
!
!
!
line con 0
 exec-timeout 0 0
 activation-character 13
 transport preferred ssh
 transport output ssh
 stopbits 1
line aux 0
 activation-character 13
line vty 0 4
 exec-timeout 0 0
 activation-character 13
 transport input ssh
line vty 5 15
 activation-character 13
 transport input ssh
!
ntp allow mode control 3
!
!
!
!
!
event manager session cli username "admin"
event manager applet EDGE2_SEC_PATH_DOWN
 event track 41 state down
 action 1.0 syslog msg "EDGE2_1 secured path down: withdrawing 198.18.41.0/24"
 action 2.0 cli command "enable"
 action 3.0 cli command "configure terminal"
 action 4.0 cli command "router bgp 65000"
 action 5.0 cli command "address-family ipv4 vrf CRITICAL-VRF"
 action 6.0 cli command "no network 198.18.41.0"
 action 7.0 cli command "end"
event manager applet EDGE2_SEC_PATH_UP
 event track 41 state up
 action 1.0 syslog msg "EDGE2_1 secured path up: restoring 198.18.41.0/24"
 action 2.0 cli command "enable"
 action 3.0 cli command "configure terminal"
 action 4.0 cli command "router bgp 65000"
```

```
 action 5.0 cli command "address-family ipv4 vrf CRITICAL-VRF"
 action 6.0 cli command "network 198.18.41.0 route-map RM_ODN_COLOR_300"
 action 7.0 cli command "end"
!
end
```

Site Type B medium-site backup edge configuration

Blueprints the secondary unsecure router (Cisco Secure Router 8375-E-G2) terminating public internet transits and acting as the active LAN gateway for standard administrative traffic.

Edge2_2

```
EDGE2_2#sh run
Building configuration...

Current configuration : 22475 bytes
!
! Last configuration change at 00:37:11 UTC Mon Jul 27 2026 by admin
!
version 26.1
system mode insecure
service timestamps debug datetime msec
service timestamps log datetime msec
platform qfp utilization monitor load 80
!
hostname EDGE2_2
!
boot-start-marker
boot system bootflash:c8kg2be-
universalk9.BLD_V261_1_THROTTLE_LATEST_20260305_183234_V26_1_0_54.SSA.bin
! Warning: Booting with bundle mode will be deprecated in the near future. Migration to install
mode is required.
boot-end-marker
!
!
vrf definition Mgmt-intf
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
 exit-address-family
!
vrf definition CRITICAL-VRF
 rd 65000:142
 route-target export 65000:1001
 route-target import 65000:1001
 !
 address-family ipv4
```

```
  route-target export 65000:1001
  route-target import 65000:1001
 exit-address-family
 !
 address-family ipv6
  route-target export 65000:1001
  route-target import 65000:1001
 exit-address-family
!
vrf definition NON-CRITICAL-VRF
 rd 65000:42
 route-target export 65000:2001
 route-target import 65000:2001
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
 exit-address-family
!
no logging console
aaa new-model
!
!
aaa authentication login default local
aaa authentication enable default enable
aaa authorization console
aaa authorization exec default local
!
!
aaa session-id common
!
!
!
!
!
!
!
!
!
!
login on-success log
!
!
!
!
!
```

```
!
ipv6 unicast-routing
!
!
subscriber templating
!
!
!
!
!
!
!
!
!
!
!
!
!
product-analytics
!
!
crypto pki trustpoint SLA-TrustPoint
 enrollment pkcs12
 revocation-check crl
 hash sha512
!
crypto pki trustpoint TP-self-signed-2581388279
 enrollment selfsigned
 subject-name cn=IOS-Self-Signed-Certificate-2581388279
 revocation-check none
 rsakeypair TP-self-signed-2581388279
 hash sha512
!
!
crypto pki certificate chain SLA-TrustPoint
 certificate ca 01
  3082031B 30820203 A0030201 02020101 300D0609 2A864886 F70D0101 0B050030
  2F310E30 0C060355 040A1305 43697363 6F311D30 1B060355 04031314 4C696365
  6E73696E 6720526F 6F74202D 20444556 301E170D 31333034 32343231 35353433
  5A170D33 33303432 34323135 3534335A 302F310E 300C0603 55040A13 05436973
  636F311D 301B0603 55040313 144C6963 656E7369 6E672052 6F6F7420 2D204445
  56308201 22300D06 092A8648 86F70D01 01010500 0382010F 00308201 0A028201
  01009C56 7101D61E DF2EBCC3 BA7AE0DB B241B3B4 328A9B00 EB8A80D0 2AA86F5E
  F1AEBFDE B67BD6AD 7DAD7B43 F582753B FFCC1CA5 A7841A07 6934D3AF 99078EF6
  179196FA 4FB3F2ED 3942C756 BF1CA0A9 CC98A7A7 F9E43724 D9E61D47 89E9E792
  DD9F27B4 517C2BDE D0EB5B9A 787BA085 D9BBF003 F0563BE0 A4450C8F 127B5583
  3EBC1385 2D9BAD98 68D3AE07 5C27987C 6B814B99 0686B14A 5F61753C 813089E6
```

```
  AEC48C68 F6D45267 0E365F44 B4456E11 96DCB950 233C8ADB 9FEEBAF1 2B5F3BB6
  7CE521B5 F277EBF6 03B7B0A4 958C9C7D 5460C20B CF9CCFC7 14B80F58 B5268947
  6D081172 26916B41 FB07DF42 EB9B9408 EC346138 23FBD8C4 19909697 A30845F3
  01C50203 010001A3 42304030 0E060355 1D0F0101 FF040403 02010630 0F060355
  1D130101 FF040530 030101FF 301D0603 551D0E04 16041443 214521B5 FB217A1A
  4D1BB702 36E664CB EC8B6530 0D06092A 864886F7 0D01010B 05000382 01010085
  F1B1F2AE AE7D2F9C AB0351C3 29E3F1AE 982DF11F 5E3C90F6 00B3CDED 5A1491FB
  DF07E06C AA0F4325 9FB4C4AE 2080F675 8C3B7AC5 4EAAA03E C5B50A2F 670AFF87
  EDA6462F CFC43967 C024AB32 EE3CCDCF A04B9DAE 1BBABBDA C8DF5587 CF51CB1C
  005A282F 8B518A5A 8C6F9B3C AABA3446 32EF3A75 C2F45450 7A9BCFD3 0C8BE54A
  11872DE0 CF1200D0 D1018FD9 AC685968 167E421C 9BC394ED 9BC85463 83B28146
  07B2BDED DFC1605B 4D16007B 68723E25 55908512 4EEB0A70 B2A74C2A CB1EC882
  C3215B87 6FC74304 241E59D7 C7C02C6D BD3042F5 196E8133 7A4446A4 81216E70
  CF52CF22 50A7D23E FA9F6B07 FB0F6386 9DCC3BBC 65250693 38CF6BA6 CB8EFD
          quit
crypto pki certificate chain TP-self-signed-2581388279
 certificate self-signed 01
  30820330 30820218 A0030201 02020101 300D0609 2A864886 F70D0101 0D050030
  31312F30 2D060355 04031326 494F532D 53656C66 2D536967 6E65642D 43657274
  69666963 6174652D 32353831 33383832 3739301E 170D3236 30343031 30383231
  35395A17 0D333630 33333130 38323135 395A3031 312F302D 06035504 03132649
  4F532D53 656C662D 5369676E 65642D43 65727469 66696361 74652D32 35383133
  38383237 39308201 22300D06 092A8648 86F70D01 01010500 0382010F 00308201
  0A028201 0100D862 2FD6CCC1 2F165BD9 0BA4198B C4A74D97 CD0ECF2C 95C5F079
  E23FD832 7DCCAD3B E3B5FB17 60BC3292 DF39D68B DB74F188 68E4A4D7 29515B6C
  32DA09F8 635BED33 E66B8AC8 135544E4 9B20382A BA7169B4 422382E2 F4787607
  A3D8759D B9AF637C 2F0737E9 EAA35084 2CED5ECD FEF2AE1C 9DDBE14B F0EA5475
  1163D120 CB627190 304A1D7C 4C4825A1 C0E60010 6E617727 AE4F5B69 F95B6198
  25932FD6 A26E0ED2 DEFE3A07 B8B576D8 89B34993 64D4E6A8 D252E477 16831C58
  D06F9904 50E3DF9B 1FD9D901 90A91667 542B47F9 75F233B5 18CAC73C 4CF49417
  8D7B6231 ECEBF6B8 436B43A3 78CF2CB7 B28DCE3B EDE00376 3D09E3ED DF9B3F65
  170C5C5F 04630203 010001A3 53305130 1D060355 1D0E0416 04142B18 BC2817B7
  348A9EE8 F739A797 8A1E9E7D B51F301F 0603551D 23041830 1680142B 18BC2817
  B7348A9E E8F739A7 978A1E9E 7DB51F30 0F060355 1D130101 FF040530 030101FF
  300D0609 2A864886 F70D0101 0D050003 82010100 876C6C96 23BC4032 078A8808
  2AC0C056 63647B8E 1BA74469 1B7C6AC7 226EF9B1 468C4D80 D9C19A7D D4A8AF91
  7C14647C 5175CE5C B8998822 3366EA2F 309E293A 515E7117 A7AB6AAD 3B3F17D8
  1A9C1731 4421F1D1 63BE12DA 4696BA44 DB66AD76 4505FBAB 399DF39B E5F75B33
  8558DBE0 E88A7B01 B33AD4CA C5D90C39 190E0B2A F4DB8BED 430F3A96 D91C6049
  F86EA707 9FCFBA4A C3D17FC5 B9CB09E3 3B1376F1 DF37B950 99C14861 B6F6CA6C
  2B8BED8F 90F98B71 ADAA0EE0 489EED7B 05254F76 0E36114B 9D24E09E 0C8EC839
  C7585181 925AE665 86972351 F109D23F C13869F7 7FE5802B 14A161FB 5A4BD555
  5D3DA4F0 6F41EFDE AB8D7019 7E713C72 1ADF02A7
          quit
!
!
!
```

```
!
!
!
!
!
!
!
voice-card 0/4
!
diagnostic bootup level minimal
!
license udi pid C8375-E-G2 sn FDO2922M0AX
memory free low-watermark processor 62496
!
spanning-tree extend system-id
!
!
!
username lab privilege 15 secret 9 $9$GA0B3PDG7Fcw7E$sKFqo6Qg6lqUFnP6D97gkAGKjpdI4UqAmBmmxRna6qI
username admin privilege 15 secret 9
$9$mKAjljsnbv5bT.$yJiA.Nod5GL6HM.aYvKJ7Okvhl6hVn3i0tRBCVwDhYo
!
redundancy
 mode none
!
!
!
crypto ikev2 proposal U-IKEV2-PROP
 pqc mlkem1024 optional
 encryption aes-gcm-256
 prf sha256
 group 19
crypto ikev2 proposal UC2-IKEV2-PROP
 pqc mlkem1024 optional
 encryption aes-gcm-256
 prf sha256
 group 19
!
crypto ikev2 policy U-IKEV2-POLICY
 proposal U-IKEV2-PROP
crypto ikev2 policy UC2-IKEV2-POLICY
 proposal UC2-IKEV2-PROP
!
crypto ikev2 keyring U-PSK-KR
 peer HUB1_1
  address 172.16.11.2
  pre-shared-key cisco123!
```

```
 !
 peer HUB1_2
  address 172.16.12.2
  pre-shared-key cisco123!
 !
 peer HUB2_1
  address 172.16.21.2
  pre-shared-key cisco123!
 !
 peer HUB2_2
  address 172.16.23.2
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring U-PPK-KR
 peer HUB1_1
  address 172.16.11.2
  ppk manual id UNSEC-E22-H11 key pqcReady123! required
 !
 peer HUB1_2
  address 172.16.12.2
  ppk manual id UNSEC-E22-H12 key pqcReady123! required
 !
 peer HUB2_1
  address 172.16.21.2
  ppk manual id UNSEC-E22-H21 key pqcReady123! required
 !
 peer HUB2_2
  address 172.16.23.2
  ppk manual id UNSEC-E22-H22 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-H11-PSK-KR
 peer HUB1_1
  address 2001:DB8:2100:1::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-H11-PPK-KR
 peer HUB1_1
  address 2001:DB8:2100:1::2/128
  ppk manual id UC2-E41-H11 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-H12-PSK-KR
 peer HUB1_2
  address 2001:DB8:2100:2::2/128
```

```
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-H12-PPK-KR
 peer HUB1_2
  address 2001:DB8:2100:2::2/128
  ppk manual id UC2-E41-H12 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-H21-PSK-KR
 peer HUB2_1
  address 2001:DB8:2100:3::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-H21-PPK-KR
 peer HUB2_1
  address 2001:DB8:2100:3::2/128
  ppk manual id UC2-E41-H21 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-H22-PSK-KR
 peer HUB2_2
  address 2001:DB8:2100:4::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-H22-PPK-KR
 peer HUB2_2
  address 2001:DB8:2100:4::2/128
  ppk manual id UC2-E41-H22 key pqcReady123! required
 !
!
!
crypto ikev2 profile U-IKEV2-PROFILE
 match identity remote address 172.16.11.2 255.255.255.255
 match identity remote address 172.16.12.2 255.255.255.255
 match identity remote address 172.16.21.2 255.255.255.255
 match identity remote address 172.16.23.2 255.255.255.255
 identity local address 172.16.22.2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk U-PPK-KR
 keyring local U-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-H11-IKEV2-PROFILE
```

```
 match identity remote address 2001:DB8:2100:1::2/128
 identity local address 2001:DB8:3100:2::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-H11-PPK-KR
 keyring local UC2-H11-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-H12-IKEV2-PROFILE
 match identity remote address 2001:DB8:2100:2::2/128
 identity local address 2001:DB8:3100:2::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-H12-PPK-KR
 keyring local UC2-H12-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-H21-IKEV2-PROFILE
 match identity remote address 2001:DB8:2100:3::2/128
 identity local address 2001:DB8:3100:2::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-H21-PPK-KR
 keyring local UC2-H21-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-H22-IKEV2-PROFILE
 match identity remote address 2001:DB8:2100:4::2/128
 identity local address 2001:DB8:3100:2::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-H22-PPK-KR
 keyring local UC2-H22-PSK-KR
 dpd 10 2 on-demand
!
!
crypto ikev2 ppk PPK1
! PPK Incomplete (MUST have id and secret assigned)
 !
 crypto ikev2 fragmentation mtu 1400
 !
!
!
!
track 42 interface Tunnel200 line-protocol
!
class-map match-any CM-BUSINESS
```

```
 match dscp af31
class-map match-any CM-PILOT-VIDEO
 match dscp af41
class-map match-any CM-PILOT-VOICE
 match dscp ef
class-map match-any CM-CRITICAL
 match dscp ef
class-map type inspect match-any CM-PILOT-NGFW-ALLOW
 match protocol icmp
 match protocol tcp
 match protocol udp
 match access-group name IPV4-PILOT-NGFW-ALLOW
 match access-group name IPV6-PILOT-NGFW-ALLOW
class-map match-any CM-PILOT-CRITICAL
 match dscp af31
class-map match-any CM-PILOT-BULK
 match dscp af11
!
policy-map PM-PILOT-SECURED-QOS
 class CM-PILOT-VOICE
  priority percent 10
 class CM-PILOT-VIDEO
  bandwidth percent 25
 class CM-PILOT-CRITICAL
  bandwidth percent 30
 class CM-PILOT-BULK
  bandwidth percent 10
 class class-default
  fair-queue
policy-map type epbr PM-PFP
 class CM-CRITICAL
  set forward-class 1
 class CM-BUSINESS
  set forward-class 2
 class class-default
  set forward-class 0
policy-map type inspect PM-PILOT-NGFW
 class type inspect CM-PILOT-NGFW-ALLOW
  inspect
 class class-default
  drop log
policy-map type inspect PM-PILOT-NGFW-PASS
 class type inspect CM-PILOT-NGFW-ALLOW
  pass
 class class-default
  drop log
!
```

```
zone security Z-PILOT-LAN
zone security Z-PILOT-SRV6-WAN
zone security Z-PILOT-LAN-U
zone-pair security ZP-LAN-TO-SRV6-WAN source Z-PILOT-LAN destination Z-PILOT-SRV6-WAN
 service-policy type inspect PM-PILOT-NGFW
zone-pair security ZP-LANU-TO-SRV6-WAN source Z-PILOT-LAN-U destination Z-PILOT-SRV6-WAN
 service-policy type inspect PM-PILOT-NGFW
zone-pair security ZP-SRV6-WAN-TO-LAN source Z-PILOT-SRV6-WAN destination Z-PILOT-LAN
 service-policy type inspect PM-PILOT-NGFW
zone-pair security ZP-SRV6-WAN-TO-LANU source Z-PILOT-SRV6-WAN destination Z-PILOT-LAN-U
 service-policy type inspect PM-PILOT-NGFW
!
!
!
!
!
!
!
!
!
!
crypto ipsec transform-set U-TS esp-gcm 256
 mode transport
crypto ipsec transform-set UC2-TS esp-gcm 256
 mode transport
!
crypto ipsec profile U-IPSEC
 set transform-set U-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile U-IKEV2-PROFILE
!
crypto ipsec profile UC2-H11-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-H11-IKEV2-PROFILE
!
crypto ipsec profile UC2-H12-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-H12-IKEV2-PROFILE
!
crypto ipsec profile UC2-H21-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-H21-IKEV2-PROFILE
!
crypto ipsec profile UC2-H22-IPSEC
```

```
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-H22-IKEV2-PROFILE
!
!
!
!
!
!
!
!
!
!
interface Loopback0
 description CNI_UC2_BGP_UPDATE_SOURCE
 no ip address
 ip proxy-arp
 ipv6 address 2001:DB8:100:42::1/128
 ipv6 router isis SRTE-OVLY
!
interface Loopback200
 description EDGE2_2 unsecured inet-sr loopback
 no ip address
 ip proxy-arp
 ipv6 address 2001:DB8:142:200::1/128
 ipv6 router isis SRTE-OVLY
!
interface Loopback210
 description EDGE2_2 unsecured l3vpn service loopback
 vrf forwarding NON-CRITICAL-VRF
 ip address 10.42.42.1 255.255.255.0
 ip proxy-arp
 ipv6 address 2001:DB8:42:42::1/64
!
interface Tunnel200
 description EDGE2_2 to HUB1_1 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:200::2/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/0/4
 tunnel destination 172.16.11.2
 tunnel protection ipsec profile U-IPSEC
 isis affinity flex-algo
```

```
  name UNSECURED
  !
!
interface Tunnel220
 description EDGE2_2 to HUB1_2 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:220::2/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/0/4
 tunnel destination 172.16.12.2
 tunnel protection ipsec profile U-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel221
 description EDGE2_2 to HUB2_1 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:221::2/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/0/4
 tunnel destination 172.16.21.2
 tunnel protection ipsec profile U-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel222
 description EDGE2_2 to HUB2_2 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:222::2/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/0/4
 tunnel destination 172.16.23.2
 tunnel protection ipsec profile U-IPSEC
```

```
 isis affinity flex-algo
  name UNSECURED
  !
!
interface TwoGigabitEthernet0/0/0
 no ip address
 ip proxy-arp
 negotiation auto
!
interface TwoGigabitEthernet0/0/0.101
 description EDGE2_HA_CRITICAL_VRF_BACKUP
 encapsulation dot1Q 101
 vrf forwarding CRITICAL-VRF
 ip address 198.18.41.3 255.255.255.0
 ip proxy-arp
 ip nbar protocol-discovery
 zone-member security Z-PILOT-LAN
 vrrp 41 ip 198.18.41.1
 vrrp 41 track 42 decrement 80
 service-policy type epbr input PM-PFP
!
interface TwoGigabitEthernet0/0/0.201
 description UNSECURED_NON_CRITICAL_VRF_EDGE2_2_PRIMARY
 encapsulation dot1Q 201
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.42.3 255.255.255.0
 ip proxy-arp
 ip nbar protocol-discovery
 zone-member security Z-PILOT-LAN-U
 vrrp 42 ip 198.19.42.1
 vrrp 42 priority 150
 vrrp 42 track 42 decrement 80
 service-policy type epbr input PM-PFP
!
interface TwoGigabitEthernet0/0/0.511
 description LAN_MGMT511_EDGE2_2
 encapsulation dot1Q 511
 ip address 192.168.11.13 255.255.255.0
 ip proxy-arp
!
interface TwoGigabitEthernet0/0/1
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface TwoGigabitEthernet0/0/2
```

```
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface TwoGigabitEthernet0/0/3
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface TenGigabitEthernet0/0/4
 description CNI_UC2_INET_UNDERLAY_BLR_TUNNEL200
 ip address 172.16.22.2 255.255.255.252
 no ip proxy-arp
 ipv6 address 2001:DB8:3100:2::2/64
 service-policy output PM-PILOT-SECURED-QOS
!
interface TenGigabitEthernet0/0/5
 description MEDIUM_SITE_EDGE2_INTER_EDGE_XCONNECT_TO_EDGE2_1
 no ip address
 ip proxy-arp
 ipv6 address 2001:DB8:245:12::2/64
 ipv6 router isis SRTE-OVLY
 isis network point-to-point
 isis metric 1000 level-2
!
interface Service-Engine0/4/0
!
interface GigabitEthernet0
 vrf forwarding Mgmt-intf
 ip address 175.13.8.23 255.255.255.0
 ip proxy-arp
 negotiation auto
!
segment-routing traffic-eng
 !
 on-demand color 100
  authorize
  candidate-paths
   preference 100
    constraints
     segments
      dataplane srv6
     !
    !
    dynamic
```

```
     metric
      type te
     !
    !
   !
  !
 !
 on-demand color 200
  authorize
  candidate-paths
   preference 100
    constraints
     segments
      dataplane srv6
     !
    !
    dynamic
     metric
      type te
     !
    !
   !
  !
 !
 on-demand color 300
  authorize
  candidate-paths
   preference 100
    per-flow
     forward-class 0 color 200
     forward-class 1 color 100
     forward-class 2 color 100
    !
   !
  !
 !
!
segment-routing srv6
 encapsulation
  source-address 2001:DB8:AC42::
 locators
  locator SLOC
   prefix 2001:DB8:AC42::/48
   format usid-f3216
!
router isis SRTE-OVLY
 net 49.0099.0099.0000.1042.00
```

```
 is-type level-2-only
 metric-style wide
 lsp-mtu 1300
 distribute link-state instance-id 99
 affinity-map SECURED bit-position 0
 affinity-map UNSECURED bit-position 1
 !
 address-family ipv6
  router-id Loopback0
  segment-routing srv6
   locator SLOC
  fast-reroute per-prefix level-2 all
  fast-reroute ti-lfa level-2
 exit-address-family
!
router bgp 65000
 bgp router-id 10.255.0.42
 bgp log-neighbor-changes
 neighbor 2001:DB8:100:11::1 remote-as 65000
 neighbor 2001:DB8:100:11::1 update-source Loopback0
 neighbor 2001:DB8:100:12::1 remote-as 65000
 neighbor 2001:DB8:100:12::1 update-source Loopback0
 neighbor 2001:DB8:100:21::1 remote-as 65000
 neighbor 2001:DB8:100:21::1 update-source Loopback0
 neighbor 2001:DB8:100:22::1 remote-as 65000
 neighbor 2001:DB8:100:22::1 update-source Loopback0
 !
 address-family vpnv4
  neighbor 2001:DB8:100:11::1 activate
  neighbor 2001:DB8:100:11::1 send-community both
  neighbor 2001:DB8:100:11::1 next-hop-self
  neighbor 2001:DB8:100:11::1 route-map RM-E2-BRANCH-A-BACKUP out
  neighbor 2001:DB8:100:12::1 activate
  neighbor 2001:DB8:100:12::1 send-community both
  neighbor 2001:DB8:100:12::1 next-hop-self
  neighbor 2001:DB8:100:12::1 route-map RM-E2-BRANCH-A-BACKUP out
  neighbor 2001:DB8:100:21::1 activate
  neighbor 2001:DB8:100:21::1 send-community both
  neighbor 2001:DB8:100:21::1 next-hop-self
  neighbor 2001:DB8:100:21::1 route-map RM-E2-BRANCH-A-BACKUP out
  neighbor 2001:DB8:100:22::1 activate
  neighbor 2001:DB8:100:22::1 send-community both
  neighbor 2001:DB8:100:22::1 next-hop-self
  neighbor 2001:DB8:100:22::1 route-map RM-E2-BRANCH-A-BACKUP out
 exit-address-family
 !
 address-family vpnv6
```

```
 neighbor 2001:DB8:100:11::1 activate
 neighbor 2001:DB8:100:11::1 send-community both
 neighbor 2001:DB8:100:11::1 next-hop-self
 neighbor 2001:DB8:100:12::1 activate
 neighbor 2001:DB8:100:12::1 send-community both
 neighbor 2001:DB8:100:12::1 next-hop-self
 neighbor 2001:DB8:100:21::1 activate
 neighbor 2001:DB8:100:21::1 send-community both
 neighbor 2001:DB8:100:21::1 next-hop-self
 neighbor 2001:DB8:100:22::1 activate
 neighbor 2001:DB8:100:22::1 send-community both
 neighbor 2001:DB8:100:22::1 next-hop-self
exit-address-family
!
address-family ipv4 vrf CRITICAL-VRF
 network 198.18.41.0 route-map RM_ODN_COLOR_300
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
exit-address-family
!
address-family ipv6 vrf CRITICAL-VRF
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
exit-address-family
!
address-family ipv4 vrf NON-CRITICAL-VRF
 network 10.42.42.0 mask 255.255.255.0
 network 198.19.42.0
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
exit-address-family
!
address-family ipv6 vrf NON-CRITICAL-VRF
 network 2001:DB8:42:42::/64
 !
```

```
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
 exit-address-family
!
ip rcmd domain-lookup
ip forward-protocol nd
ip forward-protocol udp tftp
ip telnet comport enable
ip tftp blocksize 512
ip ftp passive
no ip http server
ip http secure-server
!
ip route 172.16.11.2 255.255.255.255 172.16.22.1
ip route 172.16.12.2 255.255.255.255 172.16.22.1
ip route 172.16.21.2 255.255.255.255 172.16.22.1
ip route 172.16.23.2 255.255.255.255 172.16.22.1
ip route vrf Mgmt-intf 0.0.0.0 0.0.0.0 175.13.8.1
ip route vrf Mgmt-intf 202.153.144.0 255.255.255.0 175.13.8.1
ip ssh bulk-mode 131072
!
ip access-list extended IPV4-PILOT-NGFW-ALLOW
 10 permit ip any any
!
ip prefix-list PL-E2-BRANCH-A-SVC seq 5 permit 198.18.41.0/24
ipv6 route 2001:DB8:2100:1::/64 2001:DB8:3100:2::1
ipv6 route 2001:DB8:2100:2::/64 2001:DB8:3100:2::1
ipv6 route 2001:DB8:2100:3::/64 2001:DB8:3100:2::1
ipv6 route 2001:DB8:2100:4::/64 2001:DB8:3100:2::1
ipv6 route 2001:DB8:AC11::/48 Tunnel200 250
ipv6 route 2001:DB8:AC12::/48 Tunnel200 250
ipv6 route 2001:DB8:AC21::/48 Tunnel200 250
ipv6 route 2001:DB8:AC22::/48 Tunnel200 250
ipv6 route 2001:DB8:AC31::/48 Tunnel200 250
ipv6 route 2001:DB8:AC41::/48 Tunnel200 250
ipv6 route 2001:DB8:AC51::/48 Tunnel200 250
ipv6 route 2001:DB8:AC52::/48 Tunnel200 250
route-map RM_ODN_COLOR_300 permit 10
 set extcommunity color 300 additive
!
route-map RM-E2-BRANCH-A-BACKUP permit 10
 match ip address prefix-list PL-E2-BRANCH-A-SVC
 set local-preference 50
!
```

```
route-map RM-E2-BRANCH-A-BACKUP permit 100
!
!
!
!
!
!
ipv6 access-list IPV6-PILOT-NGFW-ALLOW
 sequence 10 permit ipv6 any any
!
control-plane
!
!
mgcp behavior rsip-range tgcp-only
mgcp behavior comedia-role none
mgcp behavior comedia-check-media-src disable
mgcp behavior comedia-sdp-force disable
!
mgcp profile default
!
!
!
!
!
!
line con 0
 exec-timeout 0 0
 activation-character 13
 transport preferred ssh
 transport output ssh
 stopbits 1
line aux 0
 activation-character 13
line vty 0 4
 exec-timeout 0 0
 activation-character 13
 transport input ssh
line vty 5 15
 activation-character 13
 transport input ssh
!
ntp allow mode control 3
!
!
!
!
!
```

```
event manager session cli username "admin"
!
end
```

Site Type C Large-site primary edge configuration

The primary secure router (Cisco Secure Router 8375-E-G2) with dual WAN links with L3 routing towards the LAN router.

Edge3_1

```
EDGE3_1#sh run
Building configuration...

Current configuration : 29484 bytes
!
! Last configuration change at 11:02:44 UTC Thu Jul 23 2026 by admin
!
version 26.1
system mode insecure
service timestamps debug datetime msec
service timestamps log datetime msec
platform qfp utilization monitor load 80
!
hostname EDGE3_1
!
boot-start-marker
boot system bootflash:c84g2aes-
universalk9.BLD_V261_1_THROTTLE_LATEST_20260305_183234_V26_1_0_54.SSA.bin
! Warning: Booting with bundle mode will be deprecated in the near future. Migration to install
mode is required.
boot-end-marker
!
!
vrf definition Mgmt-intf
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
 exit-address-family
!
vrf definition CRITICAL-VRF
 rd 65000:11
 route-target export 65000:1001
 route-target import 65000:1001
 !
 address-family ipv4
  srv6-mcast ingress-replication partitioned
```

```
  route-target export 65000:1001
  route-target import 65000:1001
  route-target export 65000:1001 stitching
  route-target import 65000:1001 stitching
 exit-address-family
 !
 address-family ipv6
  route-target export 65000:1001
  route-target import 65000:1001
 exit-address-family
!
vrf definition NON-CRITICAL-VRF
 rd 65000:151
 route-target export 65000:2001
 route-target import 65000:2001
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
 exit-address-family
!
no logging console
aaa new-model
!
!
aaa authentication login default local
aaa authentication enable default enable
aaa authorization console
aaa authorization exec default local
!
!
aaa session-id common
!
!
subscriber templating
!
ip multicast-routing vrf CRITICAL-VRF distributed
!
!
!
!
!
!
!
!
!
```

```
login on-success log
!
!
!
!
!
!
ipv6 unicast-routing
!
!
!
!
!
!
!
!
key chain KC-WAN-MACSEC macsec
 key 01
  cryptographic-algorithm aes-256-cmac
  key-string 3637fc51ed2f71b5a354e1aca9a91adb0f5a6744f1cda2cc7e18417cbd025e51
  lifetime local 00:00:00 Jan 1 2026 infinite
product-analytics
!
crypto pki trustpoint SLA-TrustPoint
 enrollment pkcs12
 revocation-check crl
 hash sha512
!
crypto pki trustpoint TP-self-signed-4145838372
 enrollment selfsigned
 subject-name cn=IOS-Self-Signed-Certificate-4145838372
 revocation-check none
 rsakeypair TP-self-signed-4145838372
 hash sha512
!
!
crypto pki certificate chain SLA-TrustPoint
 certificate ca 01
  30820321 30820209 A0030201 02020101 300D0609 2A864886 F70D0101 0B050030
  32310E30 0C060355 040A1305 43697363 6F312030 1E060355 04031317 43697363
  6F204C69 63656E73 696E6720 526F6F74 20434130 1E170D31 33303533 30313934
  3834375A 170D3338 30353330 31393438 34375A30 32310E30 0C060355 040A1305
  43697363 6F312030 1E060355 04031317 43697363 6F204C69 63656E73 696E6720
  526F6F74 20434130 82012230 0D06092A 864886F7 0D010101 05000382 010F0030
  82010A02 82010100 A6BCBD96 131E05F7 145EA72C 2CD686E6 17222EA1 F1EFF64D
  CBB4C798 212AA147 C655D8D7 9471380D 8711441E 1AAF071A 9CAE6388 8A38E520
  1C394D78 462EF239 C659F715 B98C0A59 5BBB5CBD 0CFEBEA3 700A8BF7 D8F256EE
```

```
  4AA4E80D DB6FD1C9 60B1FD18 FFC69C96 6FA68957 A2617DE7 104FDC5F EA2956AC
  7390A3EB 2B5436AD C847A2C5 DAB553EB 69A9A535 58E9F3E3 C0BD23CF 58BD7188
  68E69491 20F320E7 948E71D7 AE3BCC84 F10684C7 4BC8E00F 539BA42B 42C68BB7
  C7479096 B4CB2D62 EA2F505D C7B062A4 6811D95B E8250FC4 5D5D5FB8 8F27D191
  C55F0D76 61F9A4CD 3D992327 A8BB03BD 4E6D7069 7CBADF8B DF5F4368 95135E44
  DFC7C6CF 04DD7FD1 02030100 01A34230 40300E06 03551D0F 0101FF04 04030201
  06300F06 03551D13 0101FF04 05300301 01FF301D 0603551D 0E041604 1449DC85
  4B3D31E5 1B3E6A17 606AF333 3D3B4C73 E8300D06 092A8648 86F70D01 010B0500
  03820101 00507F24 D3932A66 86025D9F E838AE5C 6D4DF6B0 49631C78 240DA905
  604EDCDE FF4FED2B 77FC460E CD636FDB DD44681E 3A5673AB 9093D3B1 6C9E3D8B
  D98987BF E40CBD9E 1AECA0C2 2189BB5C 8FA85686 CD98B646 5575B146 8DFC66A8
  467A3DF4 4D565700 6ADF0F0D CF835015 3C04FF7C 21E878AC 11BA9CD2 55A9232C
  7CA7B7E6 C1AF74F6 152E99B7 B1FCF9BB E973DE7F 5BDDEB86 C71E3B49 1765308B
  5FB0DA06 B92AFE7F 494E8A9E 07B85737 F3A58BE1 1A48A229 C37C1E69 39F08678
  80DDCD16 D6BACECA EEBC7CF9 8428787B 35202CDC 60E4616A B623CDBD 230E3AFB
  418616A9 4093E049 4D10AB75 27E86F73 932E35B5 8862FDAE 0275156F 719BB2F0
  D697DF7F 28
          quit
crypto pki certificate chain TP-self-signed-4145838372
 certificate self-signed 01
  30820330 30820218 A0030201 02020101 300D0609 2A864886 F70D0101 0D050030
  31312F30 2D060355 04031326 494F532D 53656C66 2D536967 6E65642D 43657274
  69666963 6174652D 34313435 38333833 3732301E 170D3236 30343031 30383237
  33395A17 0D333630 33333130 38323733 395A3031 312F302D 06035504 03132649
  4F532D53 656C662D 5369676E 65642D43 65727469 66696361 74652D34 31343538
  33383337 32308201 22300D06 092A8648 86F70D01 01010500 0382010F 00308201
  0A028201 010095E7 D9BB2731 4FD23AA7 B866D8E8 8538FA48 73D9FA21 A2EAFFB8
  9DFC8A5E EC8A7A20 78090F28 B415BF33 E88A7795 E16BFA34 3D48C3C1 18F3A2D3
  9FFB0156 8B537162 AF268230 DFF52B9F D3E0B313 AF0A65C0 76754588 42B0B3BF
  7D5A8AB1 8BE6C44F 404F8E58 C62E9228 43A0B088 9A4D978E D7E70AB0 6B9A76AA
  BC1B282E 89DE849C FDDDB63F C6569D91 7A5E354D 56F8D57E 950E8544 7AB7EA04
  F91F9B63 C7314852 984DAE1A BF6F0330 CDE91D5A 59B71757 48ACAE30 5F65B9D4
  61DF3DEF 8A2AB75C 9CC62272 FB8A8CF9 45BB4A85 010D87C7 DC8FD54F 5F6239A3
  4257673E 454C9B4D A2541B7C 8EC97F20 B4467D1F D0159210 D62B240E ED8A4EA4
  B25E0870 36F10203 010001A3 53305130 1D060355 1D0E0416 04143BFB AE8FCB66
  DD06B797 DDA36CAD CDEA2EFC A845301F 0603551D 23041830 1680143B FBAE8FCB
  66DD06B7 97DDA36C ADCDEA2E FCA84530 0F060355 1D130101 FF040530 030101FF
  300D0609 2A864886 F70D0101 0D050003 82010100 840A1CA8 977F3ADB B1B0A0CF
  E2DABD93 C35EF69D D757E1A4 7558E0C0 1E9FC8F6 EB185A03 8160B64B 94245285
  C659A046 5993F080 EE59DC20 B4909160 43790CF9 F679052F 3623B929 B6BFA08B
  E4E3F688 75BE4E2F D1320D60 EBC11961 4747D32D EB87710B F70F6B89 BC3186AF
  F5D6BEF7 47F91412 D6AA44C8 07B6A2D7 00B5AC5F 4E39719C F1DA6550 48D96792
  D30E62E8 21C92FF8 36B6D887 BCBF892A B2A93DB1 8E8F397D 744CAC58 F22278B2
  CF5ABFFA 82A17BD4 AB067275 86A0F424 4B58C4C1 7F6D9BBE 4C5B8202 C77AF502
  548D0F17 CD21F016 75F9F0AF 1B656B52 0E45AD49 C7C17120 B13A07D1 60ECE493
  791F6FDA 582FB8FD 890B593D 4D147AF6 F048B164
          quit
```

```
!
!
!
!
!
!
!
!
!
license udi pid C8475-G2 sn FDO2931M085
!
memory free low-watermark processor 62433
diagnostic bootup level minimal
!
spanning-tree extend system-id
!
mka policy POL-WAN-MACSEC
!
!
!
!
username admin privilege 15 password 0 admin
username lab privilege 15 secret 9 $9$LGuotAGLhJeK0k$BpWmXl8CGoFM5/veRiIBImMjSxsyYkhd9jmXCsCgPYw
!
redundancy
 mode none
!
!
!
crypto ikev2 proposal U-IKEV2-PROP
 pqc mlkem1024 optional
 encryption aes-gcm-256
 prf sha256
 group 19
crypto ikev2 proposal UC2-IKEV2-PROP
 pqc mlkem1024 optional
 encryption aes-gcm-256
 prf sha256
 group 19
!
crypto ikev2 policy U-IKEV2-POLICY
 proposal U-IKEV2-PROP
crypto ikev2 policy UC2-IKEV2-POLICY
 proposal UC2-IKEV2-PROP
!
crypto ikev2 keyring U-H12-PSK-KR
 peer HUB1_2
```

```
  address 172.16.12.2
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring U-H12-PPK-KR
 peer HUB1_2
  address 172.16.12.2
  ppk manual id UNSEC-E31-H12 key pqcReady123! required
 !
!
crypto ikev2 keyring U-H21-PSK-KR
 peer HUB2_1
  address 172.16.21.2
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring U-H21-PPK-KR
 peer HUB2_1
  address 172.16.21.2
  ppk manual id UNSEC-E31-H21 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-H11-PSK-KR
 peer HUB1_1
  address 2001:DB8:2100:1::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-H11-PPK-KR
 peer HUB1_1
  address 2001:DB8:2100:1::2/128
  ppk manual id UC2-E51-H11 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-H12-PSK-KR
 peer HUB1_2
  address 2001:DB8:2100:2::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-H12-PPK-KR
 peer HUB1_2
  address 2001:DB8:2100:2::2/128
  ppk manual id UC2-E51-H12 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-H21-PSK-KR
```

```
 peer HUB2_1
  address 2001:DB8:2100:3::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-H21-PPK-KR
 peer HUB2_1
  address 2001:DB8:2100:3::2/128
  ppk manual id UC2-E51-H21 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-H22-PSK-KR
 peer HUB2_2
  address 2001:DB8:2100:4::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-H22-PPK-KR
 peer HUB2_2
  address 2001:DB8:2100:4::2/128
  ppk manual id UC2-E51-H22 key pqcReady123! required
 !
!
crypto ikev2 keyring U-H11-PSK-KR
 peer HUB1_1
  address 172.16.11.2
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring U-H11-PPK-KR
 peer HUB1_1
  address 172.16.11.2
  ppk manual id UNSEC-E31-H11 key pqcReady123! required
 !
!
crypto ikev2 keyring U-H22-PSK-KR
 peer HUB2_2
  address 172.16.23.2
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring U-H22-PPK-KR
 peer HUB2_2
  address 172.16.23.2
  ppk manual id UNSEC-E31-H22 key pqcReady123! required
 !
!
```

```
!
crypto ikev2 profile U-H12-IKEV2-PROFILE
 match identity remote address 172.16.12.2 255.255.255.255
 identity local address 172.16.51.2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk U-H12-PPK-KR
 keyring local U-H12-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile U-H21-IKEV2-PROFILE
 match identity remote address 172.16.21.2 255.255.255.255
 identity local address 172.16.51.2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk U-H21-PPK-KR
 keyring local U-H21-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-H11-IKEV2-PROFILE
 match identity remote address 2001:DB8:2100:1::2/128
 identity local address 2001:DB8:3100:3::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-H11-PPK-KR
 keyring local UC2-H11-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-H12-IKEV2-PROFILE
 match identity remote address 2001:DB8:2100:2::2/128
 identity local address 2001:DB8:3100:3::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-H12-PPK-KR
 keyring local UC2-H12-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-H21-IKEV2-PROFILE
 match identity remote address 2001:DB8:2100:3::2/128
 identity local address 2001:DB8:3100:3::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-H21-PPK-KR
 keyring local UC2-H21-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-H22-IKEV2-PROFILE
```

```
 match identity remote address 2001:DB8:2100:4::2/128
 identity local address 2001:DB8:3100:3::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-H22-PPK-KR
 keyring local UC2-H22-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile U-H11-IKEV2-PROFILE
 match identity remote address 172.16.11.2 255.255.255.255
 identity local address 172.16.51.2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk U-H11-PPK-KR
 keyring local U-H11-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile U-H22-IKEV2-PROFILE
 match identity remote address 172.16.23.2 255.255.255.255
 identity local address 172.16.51.2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk U-H22-PPK-KR
 keyring local U-H22-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 fragmentation mtu 1400
!
!
!
!
class-map match-any CM-BUSINESS
 match dscp af31
class-map match-any CM-PILOT-VIDEO
 match dscp af41
class-map match-any CM-PILOT-VOICE
 match dscp ef
class-map match-any CM-CRITICAL
 match dscp ef
class-map type inspect match-any CM-PILOT-NGFW-ALLOW
 match protocol icmp
 match protocol tcp
 match protocol udp
 match access-group name IPV4-PILOT-NGFW-ALLOW
 match access-group name IPV6-PILOT-NGFW-ALLOW
class-map match-any CM-PILOT-CRITICAL
 match dscp af31
```

```
class-map match-any CM-PILOT-BULK
 match dscp af11
!
policy-map PM-PILOT-SECURED-QOS
 class CM-PILOT-VOICE
  priority percent 10
 class CM-PILOT-VIDEO
  bandwidth percent 25
 class CM-PILOT-CRITICAL
  bandwidth percent 30
 class CM-PILOT-BULK
  bandwidth percent 10
 class class-default
  fair-queue
policy-map type epbr PM-PFP
 class CM-CRITICAL
  set forward-class 1
 class CM-BUSINESS
  set forward-class 2
 class class-default
  set forward-class 0
policy-map type inspect PM-PILOT-NGFW
 class type inspect CM-PILOT-NGFW-ALLOW
  inspect
 class class-default
  drop log
policy-map type inspect PM-PILOT-NGFW-PASS
 class type inspect CM-PILOT-NGFW-ALLOW
  pass
 class class-default
  drop log
!
!
zone security Z-PILOT-LAN
zone security Z-PILOT-SRV6-WAN
zone security Z-PILOT-LAN-U
zone-pair security ZP-LAN-TO-SRV6-WAN source Z-PILOT-LAN destination Z-PILOT-SRV6-WAN
 service-policy type inspect PM-PILOT-NGFW-PASS
zone-pair security ZP-LANU-TO-SRV6-WAN source Z-PILOT-LAN-U destination Z-PILOT-SRV6-WAN
 service-policy type inspect PM-PILOT-NGFW-PASS
zone-pair security ZP-SRV6-WAN-TO-LAN source Z-PILOT-SRV6-WAN destination Z-PILOT-LAN
 service-policy type inspect PM-PILOT-NGFW-PASS
zone-pair security ZP-SRV6-WAN-TO-LANU source Z-PILOT-SRV6-WAN destination Z-PILOT-LAN-U
 service-policy type inspect PM-PILOT-NGFW-PASS
!
!
!
```

```
!
!
!
!
crypto ipsec transform-set U-TS esp-gcm 256
 mode transport
crypto ipsec transform-set UC2-TS esp-gcm 256
 mode transport
!
crypto ipsec profile U-H11-IPSEC
 set transform-set U-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile U-H11-IKEV2-PROFILE
!
crypto ipsec profile U-H12-IPSEC
 set transform-set U-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile U-H12-IKEV2-PROFILE
!
crypto ipsec profile U-H21-IPSEC
 set transform-set U-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile U-H21-IKEV2-PROFILE
!
crypto ipsec profile U-H22-IPSEC
 set transform-set U-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile U-H22-IKEV2-PROFILE
!
crypto ipsec profile UC2-H11-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-H11-IKEV2-PROFILE
!
crypto ipsec profile UC2-H12-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-H12-IKEV2-PROFILE
!
crypto ipsec profile UC2-H21-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
 set ikev2-profile UC2-H21-IKEV2-PROFILE
!
crypto ipsec profile UC2-H22-IPSEC
 set transform-set UC2-TS
 set pfs group19 pqc mlkem1024
```

```
 set ikev2-profile UC2-H22-IKEV2-PROFILE
!
!
!
!
!
!
!
!
!
!
interface Loopback0
 description CNI_UC2_BGP_UPDATE_SOURCE
 ip address 10.255.0.51 255.255.255.255
 ip proxy-arp
 ipv6 address 2001:DB8:100:51::1/128
 ipv6 router isis SRTE-OVLY
!
interface Loopback100
 vrf forwarding CRITICAL-VRF
 ip address 10.51.51.1 255.255.255.0
 ip proxy-arp
 ipv6 address 2001:DB8:51:51::1/64
!
interface Loopback200
 description EDGE3_1 unsecured loopback
 no ip address
 ip proxy-arp
 ipv6 address 2001:DB8:151:200::1/128
 ipv6 router isis SRTE-OVLY
!
interface Loopback210
 description EDGE3_1 unsecured service loopback
 vrf forwarding NON-CRITICAL-VRF
 ip address 10.151.151.1 255.255.255.0
 ip proxy-arp
 ipv6 address 2001:DB8:151:151::1/64
!
interface Tunnel250
 description EDGE3_1 to HUB1_2 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:250::1/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
```

```
 tunnel source TenGigabitEthernet0/0/9
 tunnel destination 172.16.12.2
 tunnel protection ipsec profile U-H12-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel251
 description EDGE3_1 to HUB2_1 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:251::1/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/0/9
 tunnel destination 172.16.21.2
 tunnel protection ipsec profile U-H21-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel252
 description EDGE3_1 to HUB2_2 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:252::1/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/0/9
 tunnel destination 172.16.23.2
 tunnel protection ipsec profile U-H22-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel253
 description EDGE3_1 to HUB1_1 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:253::1/64
 ipv6 mtu 1400
```

```
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/0/9
 tunnel destination 172.16.11.2
 tunnel protection ipsec profile U-H11-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel320
 description EDGE3_1_TO_EDGE3_X_NON_CRITICAL_VRF_OSPF
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.151.1 255.255.255.252
 ip proxy-arp
 zone-member security Z-PILOT-LAN-U
 ip ospf network point-to-point
 ip ospf 320 area 0
 tunnel source 198.18.51.1
 tunnel destination 198.18.51.2
 tunnel vrf CRITICAL-VRF
!
interface GigabitEthernet0/0/0
 description TO_EDGE3_X_CRITICAL_VRF_OSPF
 vrf forwarding CRITICAL-VRF
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface GigabitEthernet0/0/1
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface GigabitEthernet0/0/2
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface GigabitEthernet0/0/3
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface GigabitEthernet0/0/4
```

```
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface GigabitEthernet0/0/5
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface GigabitEthernet0/0/6
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface GigabitEthernet0/0/7
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface TenGigabitEthernet0/0/8
 description CNI_UC4_EDGE3_1_WAN_MACSEC_PARENT
 no ip address
 ip proxy-arp
 negotiation auto
 macsec dot1q-in-clear 1
 macsec access-control should-secure
 service-policy output PM-PILOT-SECURED-QOS
!
interface TenGigabitEthernet0/0/8.4084
 description UC4_TAGGED_WAN_MACSEC_GRE_ISIS_TO_HUB1_1
 encapsulation dot1Q 4084
 ip address 172.31.84.1 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:4084:1::1/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address d862.caba.a401
 eapol eth-type 876F
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
 clns mtu 1400
```

```
 isis circuit-type level-2-only
 isis network point-to-point
 isis affinity flex-algo
  name SECURED
  !
!
interface TenGigabitEthernet0/0/8.4085
 description UC4_TAGGED_WAN_MACSEC_GRE_ISIS_TO_HUB2_1
 encapsulation dot1Q 4085
 ip address 172.31.85.1 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:4085:1::1/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address d862.ca95.da81
 eapol eth-type 876F
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
 clns mtu 1400
 isis circuit-type level-2-only
 isis network point-to-point
 isis affinity flex-algo
  name SECURED
  !
!
interface TenGigabitEthernet0/0/8.4086
 description UC4_TAGGED_WAN_MACSEC_GRE_ISIS_TO_HUB2_2
 encapsulation dot1Q 4086
 ip address 172.31.86.1 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:4086:1::1/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address 10e6.769b.9881
 eapol eth-type 876F
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
 clns mtu 1400
 isis circuit-type level-2-only
 isis network point-to-point
 isis affinity flex-algo
  name SECURED
  !
```

```
!
interface TenGigabitEthernet0/0/8.4091
 description UC4_TAGGED_WAN_MACSEC_GRE_ISIS_TO_PE
 encapsulation dot1Q 4091
 ip address 172.31.91.1 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:4091:1::1/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address d862.ca95.d101
 eapol eth-type 876F
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
 bfd interval 500 min_rx 500 multiplier 3
 clns mtu 1400
 isis circuit-type level-2-only
 isis network point-to-point
 isis bfd
 isis affinity flex-algo
  name SECURED
  !
!
interface TenGigabitEthernet0/0/9
 description CNI_UC2_INET_UNDERLAY_BLR_TUNNELS_250_251_252_253
 ip address 172.16.51.2 255.255.255.252
 no ip proxy-arp
 negotiation auto
 ipv6 address 2001:DB8:3100:3::2/64
 service-policy output PM-PILOT-SECURED-QOS
!
interface TenGigabitEthernet0/0/10
 no ip address
 ip proxy-arp
 negotiation auto
!
interface TenGigabitEthernet0/0/11
 description CLIENT_EDGE3_X_SECURED_CRITICAL_VRF
 vrf forwarding CRITICAL-VRF
 ip address 198.18.51.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 ip nbar protocol-discovery
 zone-member security Z-PILOT-LAN
 ip ospf network point-to-point
 ip ospf 310 area 0
```

```
 negotiation auto
 service-policy type epbr input PM-PFP
!
interface TenGigabitEthernet0/0/12
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface TenGigabitEthernet0/0/13
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface TenGigabitEthernet0/0/14
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface TenGigabitEthernet0/0/15
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface TwentyFiveGigE0/0/16
 no ip address
 ip proxy-arp
 shutdown
!
interface TwentyFiveGigE0/0/17
 no ip address
 ip proxy-arp
 shutdown
!
interface TwentyFiveGigE0/0/18
 no ip address
 ip proxy-arp
 shutdown
!
interface TwentyFiveGigE0/0/19
 no ip address
 ip proxy-arp
 shutdown
!
```

```
interface GigabitEthernet0
 vrf forwarding Mgmt-intf
 ip address 175.13.8.24 255.255.255.0
 ip proxy-arp
 negotiation auto
!
segment-routing traffic-eng
 !
 on-demand color 100
  authorize
  candidate-paths
   preference 100
    constraints
     segments
      dataplane srv6
     !
     affinity
      include-all
       name SECURED
      !
     !
    !
    dynamic
     metric
      type te
     !
    !
   !
  !
 !
 on-demand color 200
  authorize
  candidate-paths
   preference 100
    constraints
     segments
      dataplane srv6
     !
     affinity
      include-all
       name UNSECURED
      !
     !
    !
    dynamic
     metric
      type te
```

```
     !
    !
   !
  !
 !
 on-demand color 300
  authorize
  candidate-paths
   preference 100
    per-flow
     forward-class 0 color 200
     forward-class 1 color 100
     forward-class 2 color 100
    !
   !
  !
 !
!
segment-routing srv6
 encapsulation
  source-address 2001:DB8:AC51::
 locators
  locator SLOC
   prefix 2001:DB8:AC51::/48
   format usid-f3216
!
router ospf 310 vrf CRITICAL-VRF
 router-id 10.255.0.51
 capability vrf-lite
 passive-interface default
 no passive-interface TenGigabitEthernet0/0/11
 distribute-list prefix EDGE3-OSPF-IN in
!
router ospf 320 vrf NON-CRITICAL-VRF
 router-id 10.255.1.51
 capability vrf-lite
 redistribute bgp 65000 route-map EDGE3-UNSEC-BGP-TO-OSPF
 passive-interface default
 no passive-interface Tunnel320
!
router isis SRTE-OVLY
 net 49.0099.0099.0000.1051.00
 is-type level-2-only
 metric-style wide
 lsp-mtu 1300
 distribute link-state instance-id 99
 affinity-map SECURED bit-position 0
```

```
 affinity-map UNSECURED bit-position 1
 bfd all-interfaces
 !
 address-family ipv6
  bfd all-interfaces
  router-id Loopback0
  segment-routing srv6
   locator SLOC
  fast-reroute per-prefix level-2 all
  fast-reroute ti-lfa level-2
 exit-address-family
!
router bgp 65000
 bgp router-id 10.255.0.51
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 2001:DB8:100:11::1 remote-as 65000
 neighbor 2001:DB8:100:11::1 update-source Loopback0
 neighbor 2001:DB8:100:12::1 remote-as 65000
 neighbor 2001:DB8:100:12::1 update-source Loopback0
 neighbor 2001:DB8:100:21::1 remote-as 65000
 neighbor 2001:DB8:100:21::1 update-source Loopback0
 neighbor 2001:DB8:100:22::1 remote-as 65000
 neighbor 2001:DB8:100:22::1 update-source Loopback0
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv4 mvpn
  neighbor 2001:DB8:100:11::1 activate
  neighbor 2001:DB8:100:11::1 send-community both
 exit-address-family
 !
 address-family vpnv4
  !
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
  neighbor 2001:DB8:100:11::1 activate
  neighbor 2001:DB8:100:11::1 send-community both
  neighbor 2001:DB8:100:11::1 route-map RM-E31-UNSEC-A-BACKUP out
  neighbor 2001:DB8:100:12::1 activate
  neighbor 2001:DB8:100:12::1 send-community both
  neighbor 2001:DB8:100:12::1 route-map RM-E31-UNSEC-A-BACKUP out
  neighbor 2001:DB8:100:21::1 activate
```

```
 neighbor 2001:DB8:100:21::1 send-community both
 neighbor 2001:DB8:100:21::1 route-map RM-E31-UNSEC-A-BACKUP out
 neighbor 2001:DB8:100:22::1 activate
 neighbor 2001:DB8:100:22::1 send-community both
 neighbor 2001:DB8:100:22::1 route-map RM-E31-UNSEC-A-BACKUP out
exit-address-family
!
address-family ipv6
exit-address-family
!
address-family vpnv6
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
 neighbor 2001:DB8:100:11::1 activate
 neighbor 2001:DB8:100:11::1 send-community both
 neighbor 2001:DB8:100:12::1 activate
 neighbor 2001:DB8:100:12::1 send-community both
 neighbor 2001:DB8:100:21::1 activate
 neighbor 2001:DB8:100:21::1 send-community both
 neighbor 2001:DB8:100:22::1 activate
 neighbor 2001:DB8:100:22::1 send-community both
exit-address-family
!
address-family ipv4 vrf CRITICAL-VRF
 network 10.51.51.0 mask 255.255.255.0 route-map RM_ODN_COLOR_300
 network 198.18.51.0 route-map RM_ODN_COLOR_300
 redistribute ospf 310 route-map EDGE3-OSPF-TO-BGP
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
exit-address-family
!
address-family ipv6 vrf CRITICAL-VRF
 network 2001:DB8:51:51::/64 route-map RM_ODN_COLOR_300
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
```

```
 exit-address-family
 !
 address-family ipv4 vrf NON-CRITICAL-VRF
  network 10.151.151.0 mask 255.255.255.0
  redistribute ospf 320 route-map EDGE3-UNSEC-OSPF-TO-BGP
  !
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
 exit-address-family
 !
 address-family ipv6 vrf NON-CRITICAL-VRF
  network 2001:DB8:151:151::/64
  !
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
 exit-address-family
!
ip rcmd domain-lookup
ip forward-protocol nd
ip forward-protocol udp tftp
!
ip telnet comport enable
ip pim vrf CRITICAL-VRF rp-address 10.255.11.100 ACL-ASM-ANYCAST-RP-PILOT
ip tftp blocksize 512
ip ftp passive
no ip http server
ip http secure-server
ip ssh bulk-mode 131072
ip route 172.16.11.2 255.255.255.255 172.16.51.1
ip route 172.16.12.2 255.255.255.255 172.16.51.1
ip route 172.16.21.2 255.255.255.255 172.16.51.1
ip route 172.16.23.2 255.255.255.255 172.16.51.1
ip route vrf Mgmt-intf 0.0.0.0 0.0.0.0 175.13.8.1
ip route vrf Mgmt-intf 202.153.144.0 255.255.255.0 175.13.8.1
!
ip access-list standard ACL-ASM-ANYCAST-RP-PILOT
 10 permit 239.19.10.0 0.0.0.255
!
ip access-list extended IPV4-PILOT-NGFW-ALLOW
 10 permit ip any any
!
```

```
ip prefix-list EDGE3-HUB-LAN seq 5 permit 198.18.11.0/24
ip prefix-list EDGE3-HUB-LAN seq 10 permit 198.18.12.0/24
ip prefix-list EDGE3-HUB-LAN seq 20 permit 198.18.21.0/24
ip prefix-list EDGE3-HUB-LAN seq 22 permit 198.18.22.0/24
ip prefix-list EDGE3-HUB-LAN seq 31 permit 198.18.31.0/24
ip prefix-list EDGE3-HUB-LAN seq 41 permit 198.18.41.0/24
!
ip prefix-list EDGE3-LAN seq 10 permit 198.18.53.0/24
!
ip prefix-list EDGE3-OSPF-IN seq 10 permit 198.18.51.0/24
ip prefix-list EDGE3-OSPF-IN seq 20 permit 198.18.52.0/24
ip prefix-list EDGE3-OSPF-IN seq 30 permit 198.18.53.0/24
ip prefix-list EDGE3-OSPF-IN seq 100 deny 198.18.0.0/16 le 32
ip prefix-list EDGE3-OSPF-IN seq 1000 permit 0.0.0.0/0 le 32
!
ip prefix-list EDGE3-UNSEC-HUB-LAN seq 10 permit 198.19.111.0/24
ip prefix-list EDGE3-UNSEC-HUB-LAN seq 20 permit 198.19.112.0/24
ip prefix-list EDGE3-UNSEC-HUB-LAN seq 30 permit 198.19.121.0/24
ip prefix-list EDGE3-UNSEC-HUB-LAN seq 40 permit 198.19.122.0/24
ip prefix-list EDGE3-UNSEC-HUB-LAN seq 110 permit 10.111.111.0/24
ip prefix-list EDGE3-UNSEC-HUB-LAN seq 120 permit 10.112.112.0/24
ip prefix-list EDGE3-UNSEC-HUB-LAN seq 130 permit 10.121.121.0/24
ip prefix-list EDGE3-UNSEC-HUB-LAN seq 140 permit 10.122.122.0/24
!
ip prefix-list EDGE3-UNSEC-LAN seq 10 permit 198.19.53.0/24
!
ip prefix-list PL-E31-UNSEC-A-BACKUP seq 10 permit 198.19.53.0/24
ipv6 route 2001:DB8:2100:1::/64 2001:DB8:3100:3::1
ipv6 route 2001:DB8:2100:2::/64 2001:DB8:3100:3::1
ipv6 route 2001:DB8:2100:3::/64 2001:DB8:3100:3::1
ipv6 route 2001:DB8:2100:4::/64 2001:DB8:3100:3::1
route-map EDGE3-OSPF-TO-BGP permit 10
 match ip address prefix-list EDGE3-LAN
 set extcommunity color 300 additive
!
route-map EDGE3-UNSEC-BGP-TO-OSPF permit 10
 match ip address prefix-list EDGE3-UNSEC-HUB-LAN
!
route-map RM-E31-UNSEC-A-BACKUP permit 10
 match ip address prefix-list PL-E31-UNSEC-A-BACKUP
 set local-preference 50
!
route-map RM-E31-UNSEC-A-BACKUP permit 100
!
route-map RM_ODN_COLOR_300 permit 10
 set extcommunity color 300 additive
!
```

```
route-map EDGE3-UNSEC-OSPF-TO-BGP permit 10
 match ip address prefix-list EDGE3-UNSEC-LAN
!
route-map EDGE3-BGP-TO-OSPF permit 10
 match ip address prefix-list EDGE3-HUB-LAN
!
!
!
!
!
!
ipv6 access-list IPV6-PILOT-NGFW-ALLOW
 sequence 10 permit ipv6 any any
!
control-plane
!
!
!
!
!
!
line con 0
 exec-timeout 0 0
 activation-character 13
 transport preferred ssh
 transport output ssh
 stopbits 1
line aux 0
 activation-character 13
line vty 0 4
 exec-timeout 0 0
 activation-character 13
 transport input ssh
!
ntp allow mode control 3
!
!
!
!
!
!
end
```

Site Type C Large-site primary edge configuration

The backup secure router (Cisco Secure Router 8375-E-G2) with dual WAN links with L3 routing towards the LAN router.

Edge3_2

```
EDGE3_2#sh run
Building configuration...

Current configuration : 27798 bytes
!
! Last configuration change at 02:57:56 UTC Fri Jul 24 2026 by admin
!
version 26.1
system mode insecure
service timestamps debug datetime msec
service timestamps log datetime msec
platform qfp utilization monitor load 80
!
hostname EDGE3_2
!
boot-start-marker
boot system bootflash:c8000aes-
universalk9.BLD_V261_1_THROTTLE_LATEST_20260305_183234_V26_1_0_54.SSA.bin
! Warning: Booting with bundle mode will be deprecated in the near future. Migration to install
mode is required.
boot-end-marker
!
!
vrf definition Mgmt-intf
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
 exit-address-family
!
vrf definition CRITICAL-VRF
 rd 65000:11
 route-target export 65000:1001
 route-target import 65000:1001
 !
 address-family ipv4
  srv6-mcast ingress-replication partitioned
  route-target export 65000:1001
  route-target import 65000:1001
  route-target export 65000:1001 stitching
  route-target import 65000:1001 stitching
 exit-address-family
 !
 address-family ipv6
  route-target export 65000:1001
  route-target import 65000:1001
```

```
 exit-address-family
!
vrf definition NON-CRITICAL-VRF
 rd 65000:152
 route-target export 65000:2001
 route-target import 65000:2001
 !
 address-family ipv4
 exit-address-family
 !
 address-family ipv6
 exit-address-family
!
no logging console
aaa new-model
!
!
aaa authentication login default local
aaa authentication enable default enable
aaa authorization console
aaa authorization exec default local
!
!
aaa session-id common
!
!
subscriber templating
!
ip multicast-routing vrf CRITICAL-VRF distributed
!
!
!
!
!
!
!
!
!
login on-success log
!
!
!
!
!
!
ipv6 unicast-routing
!
```

```
!
!
!
!
!
!
!
key chain KC-WAN-MACSEC macsec
 key 01
  cryptographic-algorithm aes-256-cmac
  key-string 3637fc51ed2f71b5a354e1aca9a91adb0f5a6744f1cda2cc7e18417cbd025e51
  lifetime local 00:00:00 May 14 2026 infinite
product-analytics
!
crypto pki trustpoint SLA-TrustPoint
 enrollment pkcs12
 revocation-check crl
 hash sha512
!
crypto pki trustpoint TP-self-signed-1676905810
 enrollment selfsigned
 subject-name cn=IOS-Self-Signed-Certificate-1676905810
 revocation-check none
 rsakeypair TP-self-signed-1676905810
 hash sha512
!
!
crypto pki certificate chain SLA-TrustPoint
 certificate ca 01
  30820321 30820209 A0030201 02020101 300D0609 2A864886 F70D0101 0B050030
  32310E30 0C060355 040A1305 43697363 6F312030 1E060355 04031317 43697363
  6F204C69 63656E73 696E6720 526F6F74 20434130 1E170D31 33303533 30313934
  3834375A 170D3338 30353330 31393438 34375A30 32310E30 0C060355 040A1305
  43697363 6F312030 1E060355 04031317 43697363 6F204C69 63656E73 696E6720
  526F6F74 20434130 82012230 0D06092A 864886F7 0D010101 05000382 010F0030
  82010A02 82010100 A6BCBD96 131E05F7 145EA72C 2CD686E6 17222EA1 F1EFF64D
  CBB4C798 212AA147 C655D8D7 9471380D 8711441E 1AAF071A 9CAE6388 8A38E520
  1C394D78 462EF239 C659F715 B98C0A59 5BBB5CBD 0CFEBEA3 700A8BF7 D8F256EE
  4AA4E80D DB6FD1C9 60B1FD18 FFC69C96 6FA68957 A2617DE7 104FDC5F EA2956AC
  7390A3EB 2B5436AD C847A2C5 DAB553EB 69A9A535 58E9F3E3 C0BD23CF 58BD7188
  68E69491 20F320E7 948E71D7 AE3BCC84 F10684C7 4BC8E00F 539BA42B 42C68BB7
  C7479096 B4CB2D62 EA2F505D C7B062A4 6811D95B E8250FC4 5D5D5FB8 8F27D191
  C55F0D76 61F9A4CD 3D992327 A8BB03BD 4E6D7069 7CBADF8B DF5F4368 95135E44
  DFC7C6CF 04DD7FD1 02030100 01A34230 40300E06 03551D0F 0101FF04 04030201
  06300F06 03551D13 0101FF04 05300301 01FF301D 0603551D 0E041604 1449DC85
  4B3D31E5 1B3E6A17 606AF333 3D3B4C73 E8300D06 092A8648 86F70D01 010B0500
  03820101 00507F24 D3932A66 86025D9F E838AE5C 6D4DF6B0 49631C78 240DA905
```

```
  604EDCDE FF4FED2B 77FC460E CD636FDB DD44681E 3A5673AB 9093D3B1 6C9E3D8B
  D98987BF E40CBD9E 1AECA0C2 2189BB5C 8FA85686 CD98B646 5575B146 8DFC66A8
  467A3DF4 4D565700 6ADF0F0D CF835015 3C04FF7C 21E878AC 11BA9CD2 55A9232C
  7CA7B7E6 C1AF74F6 152E99B7 B1FCF9BB E973DE7F 5BDDEB86 C71E3B49 1765308B
  5FB0DA06 B92AFE7F 494E8A9E 07B85737 F3A58BE1 1A48A229 C37C1E69 39F08678
  80DDCD16 D6BACECA EEBC7CF9 8428787B 35202CDC 60E4616A B623CDBD 230E3AFB
  418616A9 4093E049 4D10AB75 27E86F73 932E35B5 8862FDAE 0275156F 719BB2F0
  D697DF7F 28
          quit
crypto pki certificate chain TP-self-signed-1676905810
 certificate self-signed 01
  30820330 30820218 A0030201 02020101 300D0609 2A864886 F70D0101 0D050030
  31312F30 2D060355 04031326 494F532D 53656C66 2D536967 6E65642D 43657274
  69666963 6174652D 31363736 39303538 3130301E 170D3236 30343031 30383036
  30375A17 0D333630 33333130 38303630 375A3031 312F302D 06035504 03132649
  4F532D53 656C662D 5369676E 65642D43 65727469 66696361 74652D31 36373639
  30353831 30308201 22300D06 092A8648 86F70D01 01010500 0382010F 00308201
  0A028201 0100CEEB 9CCD6FCD 17A20A8D E80EBAD6 2CB2FB44 79B406E6 C602D81C
  13F7E0A9 3C899AB8 A4A6F5EA F540F567 F2D20A4E 1808DC87 619AA788 2B866DF5
  53385641 DAE431A8 19172D5A B1CF5C41 4FAB8491 70139218 86F8D54D B0A2AADF
  EE92B53F F585F79C E6C444BF E79B59C1 DF66A519 8C882A5D 7C21E33C E5D0DCDC
  AD1A98A2 5C09E778 2C7114DB 53AF764E 97950DC9 6500AFE0 7424CF3D F161C38D
  2179C58F 65D357DB F5671A0E 6785639E 03049E48 099D05C5 634C8F93 E91D72AB
  C085A8F3 EB84A95E 2566D8B0 F5060C90 43D49699 0AE46819 BA881890 89DE3E70
  B90CA4A8 356CD362 F8C2F18A F4C922FD 1E8BCEAA A1E6095E 92B6B92E 599212C5
  B566DD16 6AFF0203 010001A3 53305130 1D060355 1D0E0416 04146580 9C6BF901
  3A3BD596 51FA7478 37A33D6E 0514301F 0603551D 23041830 16801465 809C6BF9
  013A3BD5 9651FA74 7837A33D 6E051430 0F060355 1D130101 FF040530 030101FF
  300D0609 2A864886 F70D0101 0D050003 82010100 B2D36450 2D32AE8F 742D871A
  5034E30E DEF0BA6C 85F1AB81 346F068B 7FA63D52 D672A53E 5D227708 082B9FB4
  0455D85B FA6AEB2F 63231EA1 F617C474 2E0A6C60 A3C27283 B50765B7 3F1E6D85
  D7824480 E891DA3C 28DD2F68 4D042089 3D0D09E3 C585821F E3AFBDB4 F81FD43B
  8268E32E BD7EDF76 45A1623E 38AFB6D2 CF49EECA 76642CFB 8E96A708 07184323
  8EDAA1ED 7B91B0EC A259CFC8 716DB62B 07066605 BEFBE61C 7B4B5E31 0FA28B33
  A4E5D761 2DD89E01 4FC62AC5 256F2483 2AEA9B96 459D0453 D3AD6D5C A3410329
  53069528 A937E44E E5332696 B56C2027 4B819FB2 E3C1D132 9702D850 E7390B8E
  57BB8676 FF5748EF 4F2E9E3A A610C53B E5912CC7
          quit
!
!
!
!
!
!
!
!
!
```

```
license udi pid C8500L-8S4X sn FLX250402C9
license boot level network-advantage addon dna-advantage
!
memory free low-watermark processor 23524
diagnostic bootup level minimal
!
spanning-tree extend system-id
!
mka policy POL-WAN-MACSEC
!
!
!
!
username admin privilege 15 password 0 admin
username lab privilege 15 secret 9 $9$vfoAzE58gTGlNk$oxxbACHUJZpnRpjr3Cg.jJCrrq9m.9.nv/WrNx8k2K2
!
redundancy
 mode none
!
!
!
crypto ikev2 proposal U-IKEV2-PROP
 encryption aes-gcm-256
 prf sha256
 group 19
crypto ikev2 proposal UC2-IKEV2-PROP
 encryption aes-gcm-256
 prf sha256
 group 19
!
crypto ikev2 policy U-IKEV2-POLICY
 proposal U-IKEV2-PROP
crypto ikev2 policy UC2-IKEV2-POLICY
 proposal UC2-IKEV2-PROP
!
crypto ikev2 keyring U-H12-PSK-KR
 peer HUB1_2
  address 172.16.12.2
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring U-H12-PPK-KR
 peer HUB1_2
  address 172.16.12.2
  ppk manual id UNSEC-E32-H12 key pqcReady123! required
 !
!
```

```
crypto ikev2 keyring U-H22-PSK-KR
 peer HUB2_2
  address 172.16.23.2
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring U-H22-PPK-KR
 peer HUB2_2
  address 172.16.23.2
  ppk manual id UNSEC-E32-H22 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-H11-PSK-KR
 peer HUB1_1
  address 2001:DB8:2100:1::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-H11-PPK-KR
 peer HUB1_1
  address 2001:DB8:2100:1::2/128
  ppk manual id UC2-E52-H11 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-H12-PSK-KR
 peer HUB1_2
  address 2001:DB8:2100:2::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-H12-PPK-KR
 peer HUB1_2
  address 2001:DB8:2100:2::2/128
  ppk manual id UC2-E52-H12 key pqcReady123! required
 !
!
crypto ikev2 keyring UC2-H21-PSK-KR
 peer HUB2_1
  address 2001:DB8:2100:3::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-H21-PPK-KR
 peer HUB2_1
  address 2001:DB8:2100:3::2/128
  ppk manual id UC2-E52-H21 key pqcReady123! required
 !
```

```
!
crypto ikev2 keyring UC2-H22-PSK-KR
 peer HUB2_2
  address 2001:DB8:2100:4::2/128
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring UC2-H22-PPK-KR
 peer HUB2_2
  address 2001:DB8:2100:4::2/128
  ppk manual id UC2-E52-H22 key pqcReady123! required
 !
!
crypto ikev2 keyring U-H11-PSK-KR
 peer HUB1_1
  address 172.16.11.2
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring U-H11-PPK-KR
 peer HUB1_1
  address 172.16.11.2
  ppk manual id UNSEC-E32-H11 key pqcReady123! required
 !
!
crypto ikev2 keyring U-H21-PSK-KR
 peer HUB2_1
  address 172.16.21.2
  pre-shared-key cisco123!
 !
!
crypto ikev2 keyring U-H21-PPK-KR
 peer HUB2_1
  address 172.16.21.2
  ppk manual id UNSEC-E32-H21 key pqcReady123! required
 !
!
!
crypto ikev2 profile U-H12-IKEV2-PROFILE
 match identity remote address 172.16.12.2 255.255.255.255
 identity local address 172.16.52.2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk U-H12-PPK-KR
 keyring local U-H12-PSK-KR
 dpd 10 2 on-demand
!
```

```
crypto ikev2 profile U-H22-IKEV2-PROFILE
 match identity remote address 172.16.23.2 255.255.255.255
 identity local address 172.16.52.2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk U-H22-PPK-KR
 keyring local U-H22-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-H11-IKEV2-PROFILE
 match identity remote address 2001:DB8:2100:1::2/128
 identity local address 2001:DB8:3100:4::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-H11-PPK-KR
 keyring local UC2-H11-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-H12-IKEV2-PROFILE
 match identity remote address 2001:DB8:2100:2::2/128
 identity local address 2001:DB8:3100:4::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-H12-PPK-KR
 keyring local UC2-H12-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-H21-IKEV2-PROFILE
 match identity remote address 2001:DB8:2100:3::2/128
 identity local address 2001:DB8:3100:4::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-H21-PPK-KR
 keyring local UC2-H21-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile UC2-H22-IKEV2-PROFILE
 match identity remote address 2001:DB8:2100:4::2/128
 identity local address 2001:DB8:3100:4::2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk UC2-H22-PPK-KR
 keyring local UC2-H22-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile U-H11-IKEV2-PROFILE
 match identity remote address 172.16.11.2 255.255.255.255
```

```
 identity local address 172.16.52.2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk U-H11-PPK-KR
 keyring local U-H11-PSK-KR
 dpd 10 2 on-demand
!
crypto ikev2 profile U-H21-IKEV2-PROFILE
 match identity remote address 172.16.21.2 255.255.255.255
 identity local address 172.16.52.2
 authentication remote pre-share
 authentication local pre-share
 keyring ppk U-H21-PPK-KR
 keyring local U-H21-PSK-KR
 dpd 10 2 on-demand
!
!
!
!
!
class-map match-any CM-BUSINESS
 match dscp af31
class-map match-any CM-PILOT-VIDEO
 match dscp af41
class-map match-any CM-PILOT-VOICE
 match dscp ef
class-map match-any CM-CRITICAL
 match dscp ef
class-map type inspect match-any CM-PILOT-NGFW-ALLOW
 match protocol icmp
 match protocol tcp
 match protocol udp
 match access-group name IPV4-PILOT-NGFW-ALLOW
 match access-group name IPV6-PILOT-NGFW-ALLOW
class-map match-any CM-PILOT-CRITICAL
 match dscp af31
class-map match-any CM-PILOT-BULK
 match dscp af11
!
policy-map PM-PILOT-SECURED-QOS
 class CM-PILOT-VOICE
  priority percent 10
 class CM-PILOT-VIDEO
  bandwidth percent 25
 class CM-PILOT-CRITICAL
  bandwidth percent 30
 class CM-PILOT-BULK
```

```
  bandwidth percent 10
 class class-default
  fair-queue
policy-map type epbr PM-PFP
 class CM-CRITICAL
  set forward-class 1
 class CM-BUSINESS
  set forward-class 2
 class class-default
  set forward-class 0
policy-map type inspect PM-PILOT-NGFW
 class type inspect CM-PILOT-NGFW-ALLOW
  inspect
 class class-default
  drop log
policy-map type inspect PM-PILOT-NGFW-PASS
 class type inspect CM-PILOT-NGFW-ALLOW
  pass
 class class-default
  drop log
!
!
zone security Z-PILOT-LAN
zone security Z-PILOT-SRV6-WAN
zone security Z-PILOT-LAN-U
zone-pair security ZP-LAN-TO-SRV6-WAN source Z-PILOT-LAN destination Z-PILOT-SRV6-WAN
 service-policy type inspect PM-PILOT-NGFW-PASS
zone-pair security ZP-LANU-TO-SRV6-WAN source Z-PILOT-LAN-U destination Z-PILOT-SRV6-WAN
 service-policy type inspect PM-PILOT-NGFW-PASS
zone-pair security ZP-SRV6-WAN-TO-LAN source Z-PILOT-SRV6-WAN destination Z-PILOT-LAN
 service-policy type inspect PM-PILOT-NGFW-PASS
zone-pair security ZP-SRV6-WAN-TO-LANU source Z-PILOT-SRV6-WAN destination Z-PILOT-LAN-U
 service-policy type inspect PM-PILOT-NGFW-PASS
!
!
!
!
!
!
!
crypto ipsec transform-set U-TS esp-gcm 256
 mode transport
crypto ipsec transform-set UC2-TS esp-gcm 256
 mode transport
!
crypto ipsec profile U-H11-IPSEC
 set transform-set U-TS
```

```
 set ikev2-profile U-H11-IKEV2-PROFILE
!
crypto ipsec profile U-H12-IPSEC
 set transform-set U-TS
 set ikev2-profile U-H12-IKEV2-PROFILE
!
crypto ipsec profile U-H21-IPSEC
 set transform-set U-TS
 set ikev2-profile U-H21-IKEV2-PROFILE
!
crypto ipsec profile U-H22-IPSEC
 set transform-set U-TS
 set ikev2-profile U-H22-IKEV2-PROFILE
!
crypto ipsec profile UC2-H11-IPSEC
 set transform-set UC2-TS
 set pfs group19
 set ikev2-profile UC2-H11-IKEV2-PROFILE
!
crypto ipsec profile UC2-H12-IPSEC
 set transform-set UC2-TS
 set pfs group19
 set ikev2-profile UC2-H12-IKEV2-PROFILE
!
crypto ipsec profile UC2-H21-IPSEC
 set transform-set UC2-TS
 set pfs group19
 set ikev2-profile UC2-H21-IKEV2-PROFILE
!
crypto ipsec profile UC2-H22-IPSEC
 set transform-set UC2-TS
 set pfs group19
 set ikev2-profile UC2-H22-IKEV2-PROFILE
!
!
!
!
!
!
!
!
!
!
interface Loopback0
 description CNI_UC2_BGP_UPDATE_SOURCE
 ip address 10.255.0.52 255.255.255.255
 ip proxy-arp
```

```
 ipv6 address 2001:DB8:100:52::1/128
 ipv6 router isis SRTE-OVLY
!
interface Loopback100
 vrf forwarding CRITICAL-VRF
 ip address 10.52.52.1 255.255.255.0
 ip proxy-arp
 ipv6 address 2001:DB8:52:52::1/64
!
interface Loopback200
 description EDGE3_2 unsecured loopback
 no ip address
 ip proxy-arp
 ipv6 address 2001:DB8:152:200::1/128
 ipv6 router isis SRTE-OVLY
!
interface Loopback210
 description EDGE3_2 unsecured service loopback
 vrf forwarding NON-CRITICAL-VRF
 ip address 10.152.152.1 255.255.255.0
 ip proxy-arp
 ipv6 address 2001:DB8:152:152::1/64
!
interface Tunnel260
 description EDGE3_2 to HUB1_2 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:260::1/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/1/1
 tunnel destination 172.16.12.2
 tunnel protection ipsec profile U-H12-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel261
 description EDGE3_2 to HUB2_2 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:261::1/64
 ipv6 mtu 1400
```

```
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/1/1
 tunnel destination 172.16.23.2
 tunnel protection ipsec profile U-H22-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel262
 description EDGE3_2 to HUB1_1 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:262::1/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/1/1
 tunnel destination 172.16.11.2
 tunnel protection ipsec profile U-H11-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel263
 description EDGE3_2 to HUB2_1 unsecured GRE over IPsec
 no ip address
 ip proxy-arp
 ip mtu 1400
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:5100:263::1/64
 ipv6 mtu 1400
 ipv6 router isis SRTE-OVLY
 tunnel source TenGigabitEthernet0/1/1
 tunnel destination 172.16.21.2
 tunnel protection ipsec profile U-H21-IPSEC
 isis affinity flex-algo
  name UNSECURED
  !
!
interface Tunnel321
 description EDGE3_2_TO_EDGE3_X_NON_CRITICAL_VRF_OSPF
 vrf forwarding NON-CRITICAL-VRF
 ip address 198.19.152.1 255.255.255.252
 ip proxy-arp
 zone-member security Z-PILOT-LAN-U
 ip ospf network point-to-point
```

```
 ip ospf 320 area 0
 tunnel source 198.18.52.1
 tunnel destination 198.18.52.2
 tunnel vrf CRITICAL-VRF
!
interface GigabitEthernet0/0/0
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface GigabitEthernet0/0/1
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface GigabitEthernet0/0/2
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface GigabitEthernet0/0/3
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface GigabitEthernet0/0/4
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface GigabitEthernet0/0/5
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface GigabitEthernet0/0/6
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
```

```
interface GigabitEthernet0/0/7
 no ip address
 ip proxy-arp
 shutdown
 negotiation auto
!
interface TenGigabitEthernet0/1/0
 description EDGE3_2_SECURED_TAGGED_WAN_MACSEC_TO_PE1
 no ip address
 ip proxy-arp
 negotiation auto
 macsec dot1q-in-clear 1
 macsec access-control should-secure
 service-policy output PM-PILOT-SECURED-QOS
!
interface TenGigabitEthernet0/1/0.4078
 description UC4_TAGGED_WAN_MACSEC_NATIVE_ISIS_TO_HUB1_1
 encapsulation dot1Q 4078
 ip address 172.31.78.1 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:4078:1::1/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address d862.caba.a401
 eapol eth-type 876F
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
 clns mtu 1400
 isis circuit-type level-2-only
 isis network point-to-point
 isis affinity flex-algo
  name SECURED
  !
!
interface TenGigabitEthernet0/1/0.4079
 description UC4_TAGGED_WAN_MACSEC_NATIVE_ISIS_TO_HUB2_1
 encapsulation dot1Q 4079
 ip address 172.31.79.1 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:4079:1::1/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address d862.ca95.da81
 eapol eth-type 876F
```

```
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
 clns mtu 1400
 isis circuit-type level-2-only
 isis network point-to-point
 isis affinity flex-algo
  name SECURED
  !
!
interface TenGigabitEthernet0/1/0.4090
 description UC4_TAGGED_WAN_MACSEC_GRE_ISIS_TO_HUB2_2
 encapsulation dot1Q 4090
 ip address 172.31.90.1 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:4090:1::1/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address 10e6.769b.9881
 eapol eth-type 876F
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
 clns mtu 1400
 isis circuit-type level-2-only
 isis network point-to-point
 isis affinity flex-algo
  name SECURED
  !
!
interface TenGigabitEthernet0/1/0.4092
 description UC4_TAGGED_WAN_MACSEC_GRE_ISIS_TO_HUB1_2
 encapsulation dot1Q 4092
 ip address 172.31.92.1 255.255.255.252
 ip proxy-arp
 ip mtu 1468
 zone-member security Z-PILOT-SRV6-WAN
 ipv6 address 2001:DB8:4092:1::1/64
 ipv6 router isis SRTE-OVLY
 eapol destination-address d862.ca95.d101
 eapol eth-type 876F
 mka policy POL-WAN-MACSEC
 mka pre-shared-key key-chain KC-WAN-MACSEC
 macsec
 clns mtu 1400
 isis circuit-type level-2-only
```

```
 isis network point-to-point
 isis affinity flex-algo
  name SECURED
  !
!
interface TenGigabitEthernet0/1/1
 description CNI_UC2_INET_UNDERLAY_BLR_TUNNELS_260_261_262_263
 ip address 172.16.52.2 255.255.255.252
 no ip proxy-arp
 negotiation auto
 ipv6 address 2001:DB8:3100:4::2/64
 service-policy output PM-PILOT-SECURED-QOS
!
interface TenGigabitEthernet0/1/2
 no ip address
 ip proxy-arp
 negotiation auto
!
interface TenGigabitEthernet0/1/3
 description CLIENT_EDGE3_X_SECURED_CRITICAL_VRF
 vrf forwarding CRITICAL-VRF
 ip address 198.18.52.1 255.255.255.0
 ip proxy-arp
 ip pim sparse-mode
 ip nbar protocol-discovery
 zone-member security Z-PILOT-LAN
 ip ospf network point-to-point
 ip ospf 310 area 0
 negotiation auto
 service-policy type epbr input PM-PFP
!
interface GigabitEthernet0
 vrf forwarding Mgmt-intf
 ip address 175.13.8.25 255.255.255.0
 ip proxy-arp
 load-interval 30
 negotiation auto
!
segment-routing traffic-eng
 !
 on-demand color 100
  authorize
  candidate-paths
   preference 100
    constraints
     segments
      dataplane srv6
```

```
    !
    affinity
     include-all
      name SECURED
     !
    !
   !
   dynamic
    metric
     type te
    !
   !
  !
 !
!
on-demand color 200
 authorize
 candidate-paths
  preference 100
   constraints
    segments
     dataplane srv6
    !
    affinity
     include-all
      name UNSECURED
     !
    !
   !
   dynamic
    metric
     type te
    !
   !
  !
 !
!
on-demand color 300
 authorize
 candidate-paths
  preference 100
   per-flow
    forward-class 0 color 200
    forward-class 1 color 100
    forward-class 2 color 100
   !
  !
```

```
  !
 !
!
segment-routing srv6
 encapsulation
  source-address 2001:DB8:AC52::
 locators
  locator SLOC
   prefix 2001:DB8:AC52::/48
   format usid-f3216
!
router ospf 310 vrf CRITICAL-VRF
 router-id 10.255.0.52
 capability vrf-lite
 passive-interface default
 no passive-interface TenGigabitEthernet0/1/3
 distribute-list prefix EDGE3-OSPF-IN in
!
router ospf 320 vrf NON-CRITICAL-VRF
 router-id 10.255.1.52
 capability vrf-lite
 redistribute bgp 65000 route-map EDGE3-UNSEC-BGP-TO-OSPF
 passive-interface default
 no passive-interface Tunnel321
!
router isis SRTE-OVLY
 net 49.0099.0099.0000.1052.00
 is-type level-2-only
 metric-style wide
 lsp-mtu 1300
 distribute link-state instance-id 99
 affinity-map SECURED bit-position 0
 affinity-map UNSECURED bit-position 1
 !
 address-family ipv6
  router-id Loopback0
  segment-routing srv6
   locator SLOC
  fast-reroute per-prefix level-2 all
  fast-reroute ti-lfa level-2
 exit-address-family
!
router bgp 65000
 bgp router-id 10.255.0.52
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 2001:DB8:100:11::1 remote-as 65000
```

```
neighbor 2001:DB8:100:11::1 update-source Loopback0
neighbor 2001:DB8:100:12::1 remote-as 65000
neighbor 2001:DB8:100:12::1 update-source Loopback0
neighbor 2001:DB8:100:21::1 remote-as 65000
neighbor 2001:DB8:100:21::1 update-source Loopback0
neighbor 2001:DB8:100:22::1 remote-as 65000
neighbor 2001:DB8:100:22::1 update-source Loopback0
!
address-family ipv4
exit-address-family
!
address-family ipv4 mvpn
 neighbor 2001:DB8:100:11::1 activate
 neighbor 2001:DB8:100:11::1 send-community both
exit-address-family
!
address-family vpnv4
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
 neighbor 2001:DB8:100:11::1 activate
 neighbor 2001:DB8:100:11::1 send-community both
 neighbor 2001:DB8:100:12::1 activate
 neighbor 2001:DB8:100:12::1 send-community both
 neighbor 2001:DB8:100:21::1 activate
 neighbor 2001:DB8:100:21::1 send-community both
 neighbor 2001:DB8:100:22::1 activate
 neighbor 2001:DB8:100:22::1 send-community both
exit-address-family
!
address-family vpnv6
 !
 segment-routing srv6
  locator SLOC
  alloc-mode per-vrf
 exit-srv6
 !
 neighbor 2001:DB8:100:11::1 activate
 neighbor 2001:DB8:100:11::1 send-community both
 neighbor 2001:DB8:100:12::1 activate
 neighbor 2001:DB8:100:12::1 send-community both
 neighbor 2001:DB8:100:21::1 activate
 neighbor 2001:DB8:100:21::1 send-community both
 neighbor 2001:DB8:100:22::1 activate
```

```
  neighbor 2001:DB8:100:22::1 send-community both
 exit-address-family
 !
 address-family ipv4 vrf CRITICAL-VRF
  network 10.52.52.0 mask 255.255.255.0 route-map RM_ODN_COLOR_300
  network 198.18.52.0 route-map RM_ODN_COLOR_300
  redistribute ospf 310 route-map EDGE3-OSPF-TO-BGP
  !
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
 exit-address-family
 !
 address-family ipv4 vrf NON-CRITICAL-VRF
  network 10.152.152.0 mask 255.255.255.0
  redistribute ospf 320 route-map EDGE3-UNSEC-OSPF-TO-BGP
  !
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
 exit-address-family
 !
 address-family ipv6 vrf NON-CRITICAL-VRF
  network 2001:DB8:152:152::/64
  !
  segment-routing srv6
   locator SLOC
   alloc-mode per-vrf
  exit-srv6
  !
 exit-address-family
!
ip rcmd domain-lookup
ip forward-protocol nd
ip forward-protocol udp tftp
!
ip telnet comport enable
ip pim vrf CRITICAL-VRF rp-address 10.255.11.100 ACL-ASM-ANYCAST-RP-PILOT
ip tftp blocksize 512
ip ftp passive
no ip http server
ip http secure-server
ip ssh bulk-mode 131072
```

```
ip route 172.16.11.2 255.255.255.255 172.16.52.1
ip route 172.16.12.2 255.255.255.255 172.16.52.1
ip route 172.16.21.2 255.255.255.255 172.16.52.1
ip route 172.16.23.2 255.255.255.255 172.16.52.1
ip route vrf Mgmt-intf 0.0.0.0 0.0.0.0 175.13.8.1
ip route vrf Mgmt-intf 202.153.144.0 255.255.255.0 175.13.8.1
!
ip access-list standard ACL-ASM-ANYCAST-RP-PILOT
 10 permit 239.19.10.0 0.0.0.255
!
ip access-list extended IPV4-PILOT-NGFW-ALLOW
 10 permit ip any any
!
ip prefix-list EDGE3-HUB-LAN seq 5 permit 198.18.11.0/24
ip prefix-list EDGE3-HUB-LAN seq 10 permit 198.18.12.0/24
ip prefix-list EDGE3-HUB-LAN seq 20 permit 198.18.22.0/24
ip prefix-list EDGE3-HUB-LAN seq 21 permit 198.18.21.0/24
ip prefix-list EDGE3-HUB-LAN seq 22 permit 198.18.22.0/24
ip prefix-list EDGE3-HUB-LAN seq 31 permit 198.18.31.0/24
ip prefix-list EDGE3-HUB-LAN seq 41 permit 198.18.41.0/24
!
ip prefix-list EDGE3-LAN seq 10 permit 198.18.53.0/24
!
ip prefix-list EDGE3-OSPF-IN seq 10 permit 198.18.51.0/24
ip prefix-list EDGE3-OSPF-IN seq 20 permit 198.18.52.0/24
ip prefix-list EDGE3-OSPF-IN seq 30 permit 198.18.53.0/24
ip prefix-list EDGE3-OSPF-IN seq 100 deny 198.18.0.0/16 le 32
ip prefix-list EDGE3-OSPF-IN seq 1000 permit 0.0.0.0/0 le 32
!
ip prefix-list EDGE3-UNSEC-HUB-LAN seq 10 permit 198.19.111.0/24
ip prefix-list EDGE3-UNSEC-HUB-LAN seq 20 permit 198.19.112.0/24
ip prefix-list EDGE3-UNSEC-HUB-LAN seq 30 permit 198.19.121.0/24
ip prefix-list EDGE3-UNSEC-HUB-LAN seq 40 permit 198.19.122.0/24
ip prefix-list EDGE3-UNSEC-HUB-LAN seq 110 permit 10.111.111.0/24
ip prefix-list EDGE3-UNSEC-HUB-LAN seq 120 permit 10.112.112.0/24
ip prefix-list EDGE3-UNSEC-HUB-LAN seq 130 permit 10.121.121.0/24
ip prefix-list EDGE3-UNSEC-HUB-LAN seq 140 permit 10.122.122.0/24
!
ip prefix-list EDGE3-UNSEC-LAN seq 10 permit 198.19.53.0/24
ipv6 route 2001:DB8:2100:1::/64 2001:DB8:3100:4::1
ipv6 route 2001:DB8:2100:2::/64 2001:DB8:3100:4::1
ipv6 route 2001:DB8:2100:3::/64 2001:DB8:3100:4::1
ipv6 route 2001:DB8:2100:4::/64 2001:DB8:3100:4::1
route-map EDGE3-OSPF-TO-BGP permit 10
 match ip address prefix-list EDGE3-LAN
 set extcommunity color 300 additive
!
```

```
route-map EDGE3-UNSEC-BGP-TO-OSPF permit 10
 match ip address prefix-list EDGE3-UNSEC-HUB-LAN
!
route-map RM_ODN_COLOR_300 permit 10
 set extcommunity color 300 additive
!
route-map EDGE3-UNSEC-OSPF-TO-BGP permit 10
 match ip address prefix-list EDGE3-UNSEC-LAN
!
route-map EDGE3-BGP-TO-OSPF permit 10
 match ip address prefix-list EDGE3-HUB-LAN
!
!
!
!
!
!
ipv6 access-list IPV6-PILOT-NGFW-ALLOW
 sequence 10 permit ipv6 any any
!
control-plane
!
!
!
!
!
!
line con 0
 exec-timeout 0 0
 activation-character 13
 transport preferred ssh
 transport output ssh
 stopbits 1
line aux 0
 activation-character 13
line vty 0 4
 exec-timeout 0 0
 activation-character 13
 transport input ssh
!
ntp allow mode control 3
!
!
!
!
!
!
```

```
end
```