Carrier-Grade Routing Solution Brief: Deploying Dynamic Segment Routing (SRv6) Deployment Architecture in ISP Transport

TARGET NETWORK POSITIONING

Dynamic Segment Routing over IPv6 (SRv6) represents the definitive architectural evolution for Internet Service Provider (ISP) transport networks, engineered to dissolve the rigid constraints of traditional MPLS and RSVP-TE frameworks. This solution brief articulates the deployment of a carrier-grade routing platform specifically optimized for the dynamic SRv6 data plane, positioning the hardware at the nexus of aggregation, edge, and core transport layers. The platform is designed to address the exponential demand for scalable, low-latency, and policy-driven traffic engineering across metro and long-haul backbones. By leveraging the native source-routing paradigm of IPv6, the architecture eliminates the need for per-flow state in the core, enabling unprecedented network slicing, service chaining, and deterministic latency guarantees essential for 5G transport, cloud leased lines, and wholesale broadband aggregation.



HARDWARE-LEVEL FAILOVER REDUNDANCY

The proposed carrier-grade chassis is architected from the silicon up to ensure zero-touch, sub-second failover in SRv6 transport environments. The system employs a fully redundant control plane with dual Route Processor Modules (RPMs) operating in 1:1 active-standby configuration, synchronized via a dedicated high-speed heartbeat channel. The forwarding plane is underpinned by a distributed, non-blocking switch fabric with cell-based load balancing across multiple fabric cards, ensuring that the failure of any single fabric element does not degrade the line-rate capacity of the SRv6 segment endpoint. Furthermore, the platform supports TI-LFA (Topology Independent Loop-Free Alternate) FRR in hardware, providing 50ms protection for any SRv6 policy path. Hot-swappable, load-sharing AC/DC power supplies and redundant variable-speed fan trays complete the hardware redundancy model,

guaranteeing carrier-grade 99.999% availability for critical ISP transport services.

PROTOCOL SUPPORT MATRIX

The platform delivers comprehensive, hardware-accelerated support for the SRv6 ecosystem and its underlying transport protocols. The forwarding ASIC is purpose-built to execute SRv6 Endpoint Behaviors (End, End.X, End.T, End.DX4, End.DX6, End.DT4, End.DT6, and End.B6.Encaps) at line rate, ensuring no performance penalty for advanced traffic engineering. Key protocol support includes:

- IGP Extensions: IS-IS and OSPFv3 with SRv6 locator advertisements and flexible algorithm (Flex-Algo) support.
- BGP-EPE: BGP Egress Peer Engineering for inter-domain SRv6 policy creation.
- BGP-LU: BGP Labeled Unicast for seamless interworking with legacy MPLS domains.
- PCEP & BGP SR-Policy: For dynamic, controller-based SRv6 policy instantiation and optimization.
- EVPN-VPWS & EVPN-MPLS: For Layer 2 and Layer 3 VPN services over an SRv6 transport underlay.
- Telemetry: Model-Driven Telemetry (MDT) with gRPC and gNMI for real-time SRv6 path visibility and SLA monitoring.

CORE PARAMETERS

The following core parameters define the baseline performance and physical characteristics of the SRv6-optimized routing platform, establishing the hardware foundation for large-scale ISP transport deployments.

Parameter	**Specification**
Form Factor	16RU Modular Chassis
Switching Capacity	Up to 768 Tbps (Full Duplex)
Forwarding Throughput	Up to 384 Bpps
Power Supply	Dual Hot-Swappable Redundant AC/DC
SRv6 Endpoint Behaviors	End, End.X, End.T, End.DX4/6, End.DT4/6, End.B6.Encaps (Line Rate)
SRv6 Traffic Engineering	SR-Policy, PCEP, BGP-EPE, Flex-Algo
Redundancy	Dual RPMs, N+1 Fabric Cards, N+N Power, TI-LFA FRR
MTBF	> 500,000 Hours (Telcordia SR-332)

LIFECYCLE ASSURANCE (MTBF)

A cornerstone of this ISP transport solution is its lifecycle assurance, validated through rigorous reliability engineering. The platform is designed for a Mean Time Between Failures (MTBF) exceeding 500,000 hours, calculated using Telcordia SR-332 Issue 4 methodology at a 40 °C operating ambient. This exceptional reliability metric is achieved through the use of high-grade industrial components, advanced thermal management, and a shock-mounted chassis design. The system supports In-Service Software Upgrades (ISSU) for both control and forwarding planes, allowing for software maintenance windows without any traffic-affecting outages. A comprehensive 10-year lifecycle commitment ensures long-term investment protection for ISP transport infrastructure.

TARGET BACKBONE TOPOLOGIES

The platform is engineered for seamless integration into the most demanding ISP backbone topologies. Its primary deployment scenarios include:

- Metro Aggregation Ring: As a high-density SRv6 aggregation node, it terminates thousands of subscriber sessions and applies SRv6 policies for differentiated services.

- Core Transport Mesh: In a full-mesh or partial-mesh core, the platform serves as an SRv6 ingress/egress and transit node, enabling end-to-end traffic

engineering with optimized path computation and low-latency forwarding.

- Data Center Interconnect (DCI): The platform excels in DCI environments, creating SRv6-based overlays for secure, multi-tenant cloud connectivity and workload mobility.

- 5G Transport Slicing: It acts as a pivotal node for network slicing, using SRv6 SIDs to guarantee distinct SLA characteristics for enhanced Mobile Broadband (eMBB), Ultra-Reliable Low-Latency Communication (URLLC), and massive Machine-Type Communication (mMTC) services.

