

The SRv6 Network Disruptor

Paradigm Shift

Architectural Viability & Business Transformation
C-Level Executives · Enterprise Boards · Private Equity Partners

LATENCY REDUCTION –25% production Tier-1 provider	CPU OVERHEAD –60% control-plane MPLS eliminated	CAPEX SAVINGS –40–60% white-box disaggregation
--	--	---

Why Legacy MPLS Is a Commercial Liability

Legacy MPLS/RSVP-TE architectures, engineered for a narrowband, voice-centric telecommunications paradigm of the early 2000s, are structurally incompatible with the operational and financial realities of 2026 and beyond. The commercial liability is measurable across four critical business dimensions:

- **Operational Complexity Tax:** RSVP-TE requires per-LSP state maintenance at every node — consuming CPU, memory, and engineering hours on a protocol adding zero incremental business value.
- **Cloud-Native Incompatibility:** Modern containerized workloads and Kubernetes-native service meshes require sub-millisecond path programmability that MPLS/RSVP-TE cannot deliver without proprietary controller overlays.
- **5G/6G Network Slicing Barrier:** 3GPP Release 17+ slicing architectures mandate transport-layer isolation with guaranteed SLA enforcement per slice — making dynamic AI-driven slice orchestration operationally impossible on MPLS at commercial scale.
- **Competitive Obsolescence Risk:** Tier-1 carriers (AT&T, Deutsche Telekom, NTT, China Mobile) and hyperscalers (Google, Microsoft, Amazon) have already committed to SRv6 as their strategic forwarding plane.

The Economics of SRv6 uSID

SRv6 micro-SID (uSID) is not an incremental protocol update. It represents the architectural unification of network transport, service delivery, and cloud orchestration into a single IPv6-native forwarding plane — eliminating the protocol layer proliferation (RSVP-TE, LDP, BGP-LU, MPLS, GRE, VXLAN) that drives 70–80% of carrier network operational expenditure.

In production deployment at a Tier-1 provider national backbone, the SRv6 uSID migration delivered a **25% reduction in end-to-end latency** and near-zero traffic loss during cutover — achieved through a Flex-Algo based traffic engineering architecture that eliminated RSVP-TE tunnel provisioning entirely.

SRv6 uSID replaces MPLS, RSVP-TE, LDP, and segment routing label stacks with a compact 128-bit IPv6 address encoding — reducing encapsulation overhead to 40 bytes of native IPv6 header.

SRv6 uSID TCO/ROI Impact Matrix

COST DIMENSION	LEGACY MPLS/RSVP-TE	SRv6 uSID ARCHITECTURE	IMPACT
Control-plane protocols	RSVP-TE + LDP + BGP-LU + SR	SR-Policy + service route color signaling	–60% CPU overhead
PE hardware CapEx	Chassis-based IOS/JUNOS	Disaggregated white-box	–40–60% CapEx
Provisioning time (per TE path)	Hours (RSVP tunnel config)	Seconds (SR-Policy API)	–98% time-to-activate
Network slicing support	Manual, static MPLS VPN	Dynamic, per-SID slices	5G/6G commercial enablement
Latency (backbone)	Baseline	–25% (production validated)	SLA tier expansion
Engineering skill requirement	Multi-protocol specialists	Unified SRv6 team	–30–40% OpEx per engineer
Fast reroute coverage	LFA/RLFA — partial topology coverage, multi-hop gaps	TI-LFA — 100% topology coverage, hitless multi-hop backup path	Zero traffic loss on any link/node failure
Route summarization	Exact-match only — LPM summarization not supported; loopbacks must be flooded inter-area/AS; results in TCAM exhaustion risk	LPM summary routes — no TCAM expansion issues; SRv6 steers aggregates natively	Eliminates TCAM pressure & inter-domain flooding
TE hop limit (> 6 hops)	MPLS label stack depth — ~12 SIDs max at linerate (NCS-5700/J2); recirculation required beyond limit	SRv6 SRH — 26 uSIDs at linerate in single pass, no recirculation (Bell Canada/Rakuten/Cisco, MPLS WC Paris 2022, segment-routing.net)	2–3× HW scale advantage — 26 vs ~12 SIDs, zero recirculation penalty

Engagement Scope & Executive Deliverables

This engagement is structured as a high-intensity, confidential advisory partnership delivered across three coordinated workstreams. All sessions are conducted under mutual NDA and tailored exclusively to the client organization's infrastructure profile, competitive landscape, and strategic transformation timeline.

WORKSTREAM	DESCRIPTION	DELIVERABLE	INVESTMENT
A — Executive Strategy Briefing	2-day intensive: SRv6 architecture mapped to board outcomes, ROI modeling workshop, competitive disruption analysis.	C-suite slide deck (20 slides), recorded sessions, board-ready executive summary report.	\$8,500
B — Custom ROI/TCO Migration Matrix	5-year TCO model: legacy MPLS vs SRv6 uSID. CapEx/OpEx savings with sensitivity analysis. Revenue opportunity sizing for new service tiers.	Financial model (Excel + PDF), risk-adjusted NPV, payback period analysis tailored to client vertical.	\$9,000
C — Architectural Look-Ahead Report	3-to-5 year architecture roadmap, vendor selection guidance, IETF standards trajectory, migration phasing with rollback procedures.	Written architectural report, migration phasing plan, cloud-native integration blueprint (Kubernetes/Cilium, SR-Policy color signaling, EVPN).	\$7,500
ALL THREE WORKSTREAMS	Complete engagement — all deliverables across A, B, and C. Fixed fee, fully deliverable-guaranteed. 30-day revision commitment included.	Full deliverable package: briefing deck, TCO model, architectural report, IC defense session.	\$25,000 Flat Fee

Pricing and Engagement Terms

- Payment Terms:** 100% due upon SOW execution for engagements under \$30,000; structured payment options available for enterprise retainer extensions.
- Exclusivity:** Engagement includes exclusive advisory access for the duration of the project scope.
- Confidentiality:** Full mutual NDA executed prior to any infrastructure data exchange or briefing session.
- Follow-On Services:** Extended advisory retainer available at \$12,000/month, covering quarterly architecture reviews and IETF standards monitoring.
- Deliverable Guarantee:** All workstream deliverables backed by a 30-day revision commitment at no additional cost.
- Governing Law:** State of Maryland, United States of America.