BGP LP SEQUENCING · ZERO-DOWNTIME CUTOVER · IPODWDM

SRv6 Core Migration Blueprint

Dual-Plane BGP Local-Pref Migration — Two Ships in the Night, Zero Downtime

0	<30s	LP 40/50/100
SR-MPLS Steps	Rollback	Migration Sequence

This white paper presents the definitive production blueprint for migrating legacy RSVP-TE/LDP MPLS backbones directly to SRv6 uSID — covering the complete technical architecture of the dual-plane 'two ships in the night' migration strategy, the precise BGP Local Preference sequencing that enables seamless incremental traffic cutover without touching eBGP PE–CE local routes, and the VRF-by-VRF, E-LINE-by-E-LINE, E-LAN-by-E-LAN service migration sequence from lowest to highest impact. Drawing on the SRv6 TEK Innovations LinkedIn blog series (DWDM and IPoDWDM integration) and validated against a 200-node ContainerLab backbone with 900K+ BGP routes, this paper documents a migration path that maintains full rollback capability at every step, eliminates the SR-MPLS intermediate architecture that creates an irresolvable IPv6 plane conflict, and achieves zero planned downtime for production services throughout the entire transition.

1. SRv6 uSID Architecture vs. Legacy Technologies

Before presenting the migration blueprint, it is essential to understand precisely why SRv6 uSID replaces — rather than coexists alongside — every major legacy networking architecture operating in a Tier-1 service provider core today. The comparison is not a feature-set debate; it is a fundamental architectural convergence: SRv6 uSID eliminates an entire class of protocol complexity that exists solely because earlier architectures had no native mechanism to encode forwarding intent in a packet header without a separate label or tunnel encapsulation.

Dimension	Legacy RSVP-TE / LDP	SR-MPLS	SRv6 uSID (F3216)
Control Plane	LDP + RSVP-TE — two separate label distribution and TE signaling protocols	IS-IS/OSPF SR extensions — eliminates LDP/RSVP-TE, still on MPLS data plane	IS-IS/OSPF SRv6 locator — single unified IPv6 control plane, zero MPLS
Data Plane	MPLS label stack, 4 bytes/label, no native IPv6 integration	MPLS label stack (v4) or MPLS-in-IPv6 (v6) — still a label stack	Native IPv6 DA carries all 6 uSIDs in 128-bit field — no separate data plane
Header Overhead	4 bytes/label, grows with hop count, MTU planning required	4 bytes/label, same MTU impact as LDP	Fixed 40-byte IPv6 header regardless of hop count (≤6 hops, no SRH)
P-Node Upgrade	Every steered-path node must support MPLS label swap at line rate	All nodes need SR-MPLS capability — software upgrade minimum	Zero — P-nodes forward IPv6 DA blindly (LPM) without any modification
Network Slicing	Not natively supported — requires separate VPN overlays per slice	Improved via SR-TE policies, but still MPLS-label based	Native per-uSID Flex-Algo plane, DSCP mapping, eBPF/Cilium integration
5G / Cloud Native	No native integration — GTP-U tunnel required for mobile user plane	Hybrid MPLS+SRv6 — complex interworking at mobile core boundary	Native SRv6 MUP (TR 23.700-40), eBPF kernel integration, Kubernetes CRD
Migration Risk	Baseline — existing production state	Low initial risk, but creates IPv6 plane conflict on SRv6 introduction	Direct path: RSVP-TE/LDP → SRv6, zero SR-MPLS intermediate state
End-State	Two protocols (LDP + RSVP-TE) + MPLS plane — permanent complexity	Unified SR control, still on MPLS data plane — an interim architecture	One data plane (IPv6), one control plane (IS-IS + SRv6), zero MPLS

1.1 Why SR-MPLS Is a Dead End, Not a Stepping Stone

The textbook migration path has been MPLS → SR-MPLS → SRv6. This path was recommended because SR-MPLS is a software-only upgrade from LDP/RSVP-TE — no hardware replacement, no P-node changes, and it eliminates the two-protocol complexity of LDP+RSVP-TE while preserving the MPLS data plane. However, when SRv6 is later

introduced into an SR-MPLS domain, a structural conflict appears: both SR-MPLS Segment IDs and SRv6 SIDs are advertised in IS-IS as extensions on the same IS-IS LSDB, but they operate on different data planes — SR-MPLS on the IPv4 label-forwarding plane, SRv6 on the IPv6 native-forwarding plane. The conflict is not just operational; it is a rollback impossibility: removing SR-MPLS state requires touching both the IPv4 forwarding tables (MPLS LFIB) and the IPv6 forwarding tables (SRv6 FIB) simultaneously on every node in the network, with no clean boundary between 'before' and 'after.'

> **The Architecture Decision:** The 6-hop uSID sweet spot — RFC 9800 NEXT-CSID — eliminates the reason SR-MPLS was ever attractive. P-nodes require zero upgrade under SRv6 uSID (exactly the same proposition that made SR-MPLS attractive), so the entire justification for SR-MPLS as an intermediate step dissolves. MPLS → SRv6 uSID direct is now the production-proven path.

1.2 SRv6 uSID F3216 Encoding — The Technical Foundation

32-bit Locator Block	Network domain identifier shared by all SIDs in the operator's SRv6 domain. Advertised via IS-IS SRv6 Locator TLV to all routers. Establishes the /32 prefix from which all /64 node locators and /128 SIDs are derived.
16-bit uSID Function	Per-node per-behavior micro-segment: End (plain routing), End.X (cross-connect to adjacency), End.DT4/DT6/DT46 (L3VPN table lookup), End.DX2/uDX2V (L2 EVPN cross-connect), uDT2U/uDT2M (E-LAN unicast/multicast).
32-bit GIB / Argument	Group Identifier Block for container /64 locators, per-session identifiers (MUP/GTP-U replacement), or VPN discriminators. Provides 4.2 billion simultaneous /64 assignments per domain.
6-Hop DA-Only Forwarding	When ≤6 uSIDs are packed into the 128-bit IPv6 DA with overlay/underlay merge disabled (vendor default), no SRH is required. P-nodes see a standard IPv6 unicast packet and forward on LPM at full line rate — zero hardware modification.
NEXT-CSID Compression (RFC 9800)	Each 16-bit uSID function is shifted into the active position by updating only the IPv6 DA at each SRv6 endpoint — the Shift-and-Forward operation that executes natively on Jericho2+, Tomahawk4+, Cisco Silicon One, Nokia FP4/FP5, and Juniper Trio5+.

2. The Dual-Plane Migration Architecture — Two Ships in the Night

The core insight of the SRv6 TEK Innovations migration blueprint is that the IPv4 MPLS plane and the IPv6 SRv6 plane are, by nature, completely orthogonal forwarding domains on dual-stack routers. They do not share LFIB entries, they do not share FIB state, and they do not interact with each other at the forwarding ASIC level. This orthogonality is not a limitation to work around — it is the architectural gift that enables zero-impact parallel operation and clean, reversible cutover. The MPLS plane carries all existing production traffic. The SRv6 plane is built out alongside it, tested, and validated before a single production byte is moved. The two planes are like two ships passing in the night: they share the same physical fiber infrastructure but never interact, never conflict, and either can be independently decommissioned without affecting the other.

> ***"The IPv4 MPLS plane and the IPv6 SRv6 plane are fully orthogonal forwarding domains — they share fiber but never touch each other's state."***

2.1 Dual-Stack PE Node Configuration

Every Provider Edge (PE) node in the migration-ready network runs both forwarding planes simultaneously from day one of Phase 2. On the IPv4 plane: the existing RSVP-TE LSP mesh (or LDP in access domains) carries all production MPLS-encapsulated VPN, EVPN, and transit traffic exactly as before, with zero configuration change. On the IPv6 plane: IS-IS is extended with SRv6 Locator TLV advertisements, each PE node receives a /64 SRv6 locator prefix, and BGP is extended with RFC 9252 SRv6 L3VPN and EVPN NLRIs advertising the same VPN prefixes as the existing MPLS path but with SRv6 next-hops. The critical point: both plane's VPN prefixes land in the same VRF RIB on every PE simultaneously, and BGP Local Preference is the sole mechanism that determines which forwarding plane carries production traffic at any given moment.

IPv4/MPLS Plane (Existing)	RSVP-TE LSP mesh fully operational. LDP in access. BGP/MPLS L3VPN (RFC 4364) with MPLS next-hops. EVPN/VPLS with MPLS transport. Zero changes to this plane throughout the migration.
IPv6/SRv6 Plane (New)	IS-IS SRv6 Locator TLV enabled. Every PE assigned a /64 locator from the 32-bit block. BGP SRv6 L3VPN (RFC 9252) NLRIs advertising same VRF prefixes with SRv6 next-hops. EVPN SRv6 NLRIs for E-LINE/E-LAN services.
Shared Infrastructure	Same physical routers, same IS-IS adjacencies, same BGP sessions. Both planes run concurrently with zero physical topology change. P-nodes forward IPv6 SRv6 traffic blindly via LPM on their existing IPv6 FIB.
VRF RIB Behavior	BGP installs both MPLS-next-hop and SRv6-next-hop paths for the same VPN prefix into the same VRF RIB. Local Preference determines which path is selected as best — the other path remains in RIB as a standby, available for instant reactivation.

3. BGP Local Preference Migration Sequencing

BGP Local Preference (LOCAL_PREF) is the cleanest possible migration control mechanism for this dual-plane architecture because it operates entirely within the iBGP domain, affects only the best-path selection among iBGP-learned routes, and has no interaction whatsoever with eBGP PE–CE routes. This is the critical safety property: the migration never touches or risks displacing any PE–CE eBGP learned prefix that carries a customer's default LOCAL_PREF of 100. The only routes being manipulated are the iBGP VPN routes that carry the transport plane information — specifically, the choice between the MPLS next-hop and the SRv6 next-hop for the same VPN prefix.

3.1 The Three-State Local Preference Model

The migration uses exactly three Local Preference values to manage the transition across the entire network, with one invariant constraint: no value must exceed the default eBGP import LOCAL_PREF (100), because any value above 100 on an iBGP VPN route would make it prefer over eBGP PE–CE customer routes, potentially black-holing customer traffic.

Local Pref Value	Applied To	Effect	Phase
LP = 40	SRv6 VPN NLRI (iBGP)	SRv6 path is de-preferred — MPLS path (LP=100 default) wins. SRv6 plane is built, advertised, and validated but carries zero production traffic.	Phase 2: Build-Out
LP = 100 (default)	Both MPLS and SRv6 VPN NLRIs (initial)	Default BGP behavior — tie-broken by other attributes (next-hop IGP cost, router-ID). Establishes the pre-migration baseline.	Phase 1: Audit
LP = 50	MPLS VPN NLRI (iBGP)	MPLS path is de-preferred below default (100) but above the SRv6 build-out value (40). SRv6 path at default LP=100 now wins — production traffic migrates to SRv6 plane.	Phase 3: Cutover
LP = 100 (SRv6)	SRv6 VPN NLRI (iBGP)	SRv6 path at default LP=100 is preferred over MPLS path at LP=50. Production traffic on SRv6 plane. MPLS path remains in RIB for instant rollback.	Phase 3: Post-Cutover
LP removed (MPLS)	MPLS VPN NLRI withdrawn	RSVP-TE LSPs torn down gracefully after full verification. MPLS plane decommissioned VRF by VRF.	Phase 4: Decommission

3.2 Why LP=50 for MPLS De-Preference Is Precise

The value 50 is not arbitrary. The constraint is that the MPLS de-preference value must satisfy three conditions simultaneously: (1) it must be lower than the SRv6 path's LP value at cutover time (LP=100 default) so that SRv6 wins; (2) it must be higher than the SRv6 build-out LP value (LP=40) so that if both paths are somehow at their respective migration-phase values simultaneously, the MPLS path still wins over an accidentally un-promoted SRv6 path; and (3) it must be lower than any eBGP PE–CE import policy value (typically LP=100 for internet routes,

LP=150 for VPN preferred routes in some operator designs) to avoid accidentally de-preferring customer-facing eBGP routes. LP=50 satisfies all three constraints with a comfortable margin in both directions.

> **Safety Invariant:** eBGP PE–CE routes are never touched. The LP manipulation applies only to iBGP VPN NLRIs carrying the transport next-hop choice between MPLS and SRv6. Customer prefixes learned via eBGP PE–CE sessions retain their import LP (typically 100 or above) and are never candidates for this manipulation. The migration is structurally invisible to every customer CE router in the network.

3.3 Why LP=100 Is the Safe Ceiling — LOCAL_PREF, eBGP, and the Dual-Attached CE Loop

The most critical constraint in the entire migration design is this: the SRv6 iBGP VPN NLRI LOCAL_PREF must never be set higher than 100. Understanding precisely why requires a clear explanation of how LOCAL_PREF behaves on eBGP PE–CE sessions, and why the loop risk is specific to dual-attached CE configurations and does not exist for single-attached CEs.

LOCAL_PREF Is Optional Non-Transitive — The CE Never Sets It

LOCAL_PREF is an optional non-transitive BGP path attribute: it is meaningful only within a single autonomous system and is never carried across eBGP session boundaries. When a CE router sends a BGP UPDATE to its PE on the eBGP PE–CE session, it cannot set LOCAL_PREF outbound — and even if it could, the PE would strip it on receipt. The PE therefore always sees the CE-advertised prefix arrive at the eBGP session with no LOCAL_PREF attribute, and BGP assigns it the default value of 100. This is not configurable at the CE. It is a BGP protocol invariant: every CE-to-PE eBGP route lands on the PE at exactly LOCAL_PREF=100, always. LOCAL_PREF is set only inbound on the PE side via import route-policy, and only outbound from the PE on iBGP sessions to reflect migration state.

Single-Attached CE — No Loop Risk

When a CE is single-homed to exactly one PE, that PE holds the only eBGP route for that CE's prefixes in the entire network. No other PE has a local eBGP PE–CE route for those same prefixes — they can only learn them via iBGP VPN NLRIs redistributed from the single-attached PE. In this topology, there is no competing iBGP VPN NLRI at the attached PE that could be preferred over the local eBGP route, because the iBGP VPN NLRI for that prefix originates from this same PE and is not re-imported locally. Single-attached CE prefixes are therefore immune to the loop condition regardless of the LOCAL_PREF value set on outbound iBGP VPN advertisements.

Dual-Attached CE — The Precise Loop Condition

The loop risk is specific and precisely defined: it requires a CE dual-homed to two different PEs (CE dual-attached to PE-A and PE-B), with the same customer prefix advertised via eBGP from CE to both PE-A and PE-B simultaneously. Both PE-A and PE-B receive the prefix at LP=100 on their respective local eBGP PE–CE sessions. Both then redistribute the prefix into the iBGP VPN mesh as a VPNv4/VPNv6 NLRI with their own MPLS or SRv6 next-hop.

Now consider what happens on PE-A if the iBGP VPN NLRI for that prefix — originating from PE-B — is advertised with LOCAL_PREF higher than 100, say LP=150. BGP best-path selection on PE-A runs: local eBGP route at LP=100 vs. iBGP VPN NLRI from PE-B at LP=150. LOCAL_PREF is the third BGP decision criterion, evaluated

before the eBGP-over-iBGP tiebreak. LP=150 wins. PE-A selects PE-B's iBGP VPN NLRI as best path and installs PE-B as the next-hop for a prefix that the CE is advertising directly and locally to PE-A. Traffic arriving at PE-A for that customer prefix is encapsulated and forwarded to PE-B. PE-B faces the identical condition in reverse: it also sees PE-A's iBGP VPN NLRI at LP=150, which beats its own local eBGP CE route at LP=100, and forwards the traffic back to PE-A. The traffic bounces between PE-A and PE-B until TTL expires. Both directly-attached CEs receive nothing. This is the forwarding loop.

> **The Precise Loop-Prevention Rule:** LOCAL_PREF on iBGP VPN NLRIs must never exceed 100. At LP=100 equal to the locally-learned eBGP PE–CE route at LP=100, BGP best-path selection reaches the eBGP-over-iBGP tiebreak — because LOCAL_PREF comparison is a tie at 100. The eBGP path wins that tiebreak unconditionally. The local CE-attached route is preferred at every PE that has a direct CE attachment for that prefix. The SRv6 iBGP VPN NLRI at LP=100 wins only on PEs that have no local eBGP CE attachment — exactly the PEs where the iBGP route should win because they have no better local path. No loop can form.

3.4 The LP Mechanics Per-PE at Migration — Outbound Advertisement Sequence

The key operational mechanism is that LOCAL_PREF is set outbound on the iBGP VPN NLRI advertisement from each PE as it is migrated. This is the precise sequence:

Phase 2: Build-Out (All PEs)	Each PE advertises its SRv6 VPN NLRIs outbound into the iBGP mesh with LP=40. All other PEs receive and install the SRv6 path at LP=40 in their RIB but do not select it (MPLS path at LP=100 default wins). SRv6 plane is live, built, and visible but carries zero production traffic. Local eBGP PE–CE routes at LP=100 are unaffected — LP=40 on iBGP never threatens them.
PE-A Cutover: Unset LP=40, Advertise SRv6 Outbound at LP=100	PE-A removes the LP=40 policy on its SRv6 VPN NLRIs and advertises them outbound at LP=100 (the BGP default). Simultaneously, PE-A sets its MPLS VPN NLRIs outbound to LP=50. All other PEs in the core now see: SRv6 path to PE-A's prefixes at LP=100, MPLS path to PE-A's prefixes at LP=50. Every non-migrated PE selects SRv6 as the best path toward PE-A immediately via iBGP convergence. PE-A's own local eBGP CE routes remain at LP=100 and are still preferred over any iBGP route via the eBGP-over-iBGP tiebreak.
Non-Migrated PEs (PE-B, PE-C…): Forward Path to PE-A = SRv6	All non-migrated PEs select SRv6 as the forward path toward PE-A's customer prefixes (SRv6 NLRI at LP=100 beats MPLS NLRI at LP=50). Traffic from PE-B's customers to PE-A's customers flows over the SRv6 IPv6 plane.
Migrated PE-A: Return Path to Non-Migrated PEs = MPLS	PE-A still receives LP=40 on SRv6 NLRIs from non-migrated PE-B, PE-C etc. PE-A selects MPLS (LP=100) over SRv6 (LP=40) for traffic destined toward non-migrated PEs. Return path from PE-A's customers to PE-B's/PE-C's customers flows over the MPLS IPv4 plane. This is intentional asymmetry.
Asymmetry Is Valid and Expected	Forward path to migrated PE-A: SRv6. Return path from PE-A to non-migrated peers: MPLS. Both planes deliver correctly. TCP sessions are not disrupted. ECMP across planes is avoided because each PE selects exactly one path per prefix based on LOCAL_PREF. QoS classification must cover both planes.

PE-B Migrated: Symmetry Grows	As PE-B unsets LP=40 and advertises SRv6 NLRIs outbound at LP=100, all remaining PEs prefer SRv6 to both PE-A and PE-B. PE-A and PE-B now see each other's SRv6 NLRIs at LP=100 and select SRv6 bidirectionally — full symmetry between migrated PEs. Asymmetry persists only toward remaining non-migrated PEs.
Final PE Migrated: Full SRv6 Symmetry	When the last PE unsets LP=40 on its SRv6 NLRIs and advertises LP=100 outbound, every PE in the network selects SRv6 as best path toward every other PE. All traffic, in all directions, flows over the SRv6 IPv6 plane. The MPLS plane carries zero traffic. Bidirectional SRv6 symmetry is complete. RSVP-TE sunset can now proceed.
MPLS Decommission Gate	RSVP-TE LSP teardown executes only after 72-hour gNMI telemetry confirms zero MPLS traffic counters on all PE-to-PE interfaces. Decommission is VRF-by-VRF, PE-by-PE, following the same lowest-impact-first sequence as the migration itself.

IMPACT: **The complete LP sequence: LP=40 outbound (SRv6 de-preferred during build-out) → LP=40 unset + LP=100 advertised outbound on SRv6, LP=50 set outbound on MPLS (cutover — all other PEs now prefer SRv6 to the migrated PE via iBGP convergence) → each PE migrated incrementally until full SRv6 symmetry → MPLS dark → RSVP-TE sunset.**

3.5 BGP Route Policy Implementation

The LP values are applied via route-policy on the BGP route-reflectors (RRs) or directly on the PEs, depending on the operator's BGP topology. The policy matches on the VPN NLRI's transport attribute (MPLS next-hop vs SRv6 next-hop, distinguishable via the BGP next-hop address family or the SRv6 Service TLV presence in RFC 9252 NLRIs) and sets LOCAL_PREF accordingly. The same policy mechanism is available on all major platforms:

Cisco IOS-XR	route-policy with if extcommunity color matches SRv6 service TLV → set local-preference 40/50/100. Applied inbound on iBGP RR client sessions.
Juniper JunOS	policy-statement with term matching next-hop inet6 address family (SRv6 paths) or MPLS next-hop → then local-preference 40/50/100.
Nokia SR OS	vprn policy-statement with bgp-next-hop criteria distinguishing SRv6 locator prefix from MPLS tunnel endpoint → local-preference assignment.
Arista EOS	route-map with match ip next-hop prefix-list (MPLS loopbacks) or match ipv6 next-hop prefix-list (SRv6 locators) → set local-preference.
FRRouting	route-map with match community (SRv6 color community) or match ipv6 next-hop prefix-list → set local-preference. Applied via neighbor route-map in-policy on RR.

4. Incremental L3VPN Migration — VRF by VRF, Prefix by Prefix

The production migration follows a strict lowest-impact-first discipline across all service types. For L3VPN, this means beginning with a dedicated test VRF carrying no production traffic, progressing through increasingly larger prefix aggregates within production VRFs, then cutting PE pairs, then cutting entire VRFs, and finally decommissioning the MPLS plane entirely once every VRF has been verified on SRv6. At every step, the MPLS path remains in the BGP RIB at LP=50 and can be restored to LP=100 within seconds by a single route-policy change, providing instant rollback without any RSVP-TE LSP renegotiation or LDP re-convergence.

4.1 L3VPN Migration Sequence — Step by Step

Step	Action	Traffic Impact	Rollback
1. Test VRF	Create a dedicated test VRF with a loopback prefix (e.g., 192.0.2.0/24). Set SRv6 NLRI LP=100, MPLS NLRI LP=50. Verify reachability via SRv6 path using traceroute and gNMI telemetry.	Zero — test prefix only, no customer traffic	Remove test VRF; no production impact
2. Small Aggregate (Pilot)	Select lowest-traffic production VRF. Set SRv6 NLRI LP=100, MPLS LP=50 for a single /24 or /25 aggregate within the VRF. Monitor via 5-minute gNMI streaming window.	Minimal — one small prefix, rollback in <30s	Set MPLS LP back to 100 for that prefix
3. Larger Aggregates	Expand LP=100 SRv6 preference to all prefixes within the pilot VRF, one aggregate at a time. Typical sequence: /24 → /23 → /22 → full VRF prefix set.	Low — one VRF, verified incrementally	Per-prefix LP reversion via route-policy
4. Full Pilot VRF Cutover	Apply LP=100 SRv6 / LP=50 MPLS to all prefixes in the pilot VRF simultaneously. 24-hour observation window. Verify MPLS remains in RIB as standby.	Medium — one production VRF fully on SRv6	Single policy change restores MPLS LP=100
5. PE Pair Cutover	Select one pair of PEs serving the pilot VRF. Cut that PE pair's VRF traffic to SRv6 while leaving other PEs on MPLS. Validates per-PE SRv6 locator and SR-TE policies.	Scoped to one PE pair's customer set	Per-PE route-policy LP reversion
6. VRF-by-VRF Progression	Repeat steps 2–5 for each remaining VRF in order of ascending traffic volume and business criticality. One VRF per change window.	Controlled — one VRF per window	Per-VRF policy reversion at any step
7. Final PE Pairs	After all VRFs cut, execute final PE-by-PE MPLS LP deprecation. Verify zero MPLS traffic remains on each PE before decommission.	Low — all VRFs already on SRv6	Full LP reversion per PE in <60s
8. RSVP-TE Sunset	After 72-hour clean observation on all PEs, gracefully tear down RSVP-TE LSPs. Remove LDP from access. Decommission MPLS data plane.	Zero — all traffic already on SRv6	N/A — rollback complete only if needed pre-sunset

4.2 Verification at Each Step

Each step uses a structured three-part verification before any subsequent step proceeds: BGP RIB verification confirms both the SRv6 path (LP=100, selected) and MPLS path (LP=50, in RIB but not active) are present on every PE serving the migrated VRF; data-plane verification via traceroute with SRv6 segment-display confirms traffic is traversing the expected SRv6 uSID path rather than the MPLS LSP; and telemetry verification via gNMI streaming shows zero packet loss and normal latency distribution across a minimum 5-minute observation window before the next step executes.

IMPACT: **The BGP RIB dual-presence guarantee — SRv6 path selected at LP=100, MPLS path present at LP=50 — means rollback at any step is a single route-policy line change propagating via iBGP in under 30 seconds, with no RSVP-TE renegotiation, no LDP re-convergence, and no customer-visible impact.**

5. L2 VPN Migration — EVPN E-LINE and E-LAN Incremental Cutover

L2 VPN migration follows the same lowest-impact-first discipline as L3VPN, but the service types require distinct migration sequences because E-LINE (point-to-point Ethernet pseudowire replacement) and E-LAN (multipoint VPLS replacement) have fundamentally different failure modes. An E-LINE failure affects one customer's private line; an E-LAN failure can affect all participants in a multipoint broadcast domain simultaneously. The migration therefore keeps E-LINE and E-LAN migrations structurally independent and begins each with dedicated test circuits before any production service is touched.

5.1 SRv6 EVPN L2 Service Type Mapping (RFC 9252)

uDX2 — EVPN E-LINE	Ethernet cross-connect, all-to-one forwarding. Direct replacement for VPWS MPLS pseudowire. Single SID per AC encodes the remote PE and the attachment circuit. No MAC learning required at the transport layer.
uDX2V — E-LINE VLAN-aware	VLAN-aware bundled service. Each VLAN tag maps to a distinct uSID, preserving per-VLAN forwarding semantics across the SRv6 transport. Replaces VLAN-bundled VPWS pseudowires.
uDT2U — EVPN E-LAN Unicast	MAC/IP advertisement unicast forwarding. Each remote MAC address learned via BGP EVPN Type-2 route maps to a per-MAC uSID at the egress PE. Replaces VPLS unicast data plane.
uDT2M — EVPN E-LAN BUM Flooding	Unknown unicast, broadcast, and multicast (BUM) flooding. Replaces the VPLS flooding data plane. BUM traffic is replicated by the ingress PE to a BUM SID list derived from the EVPN Inclusive Multicast Route (Type-3).
End.DX2 — L2 Cross-Connect	Single-home attachment circuit cross-connect. Used for EVPN single-active or all-active multihoming at the PE level. Encodes both the physical port and the EVI identifier in the uSID function.

5.2 E-LINE Migration Sequence

Step	Action	Impact	Rollback
1. Test E-LINE	Provision a dedicated test EVPN E-LINE (uDX2) between two PEs using SRv6 transport at LP=100, MPLS transport at LP=50. Send test L2 frames. Verify MAC learning, VLAN tagging, and BFD session on SRv6 path.	Zero — test circuit only	Delete test E-LINE
2. Lowest-Traffic E-LINE	Select the single lowest-traffic production E-LINE (e.g., management circuit, OOB link). Set SRv6 NLRI LP=100. Monitor for 30 minutes.	Minimal — one low-priority private line	LP reversion <30s
3. E-LINE Batch (Non-Critical)	Group all non-critical E-LINEs (e.g., inter-DC management, test/dev circuits). Apply LP=100 SRv6 preference in a single change window.	Low — no SLA-protected circuits	Batch LP reversion via single route-policy
4. SLA-Protected E-LINEs	Migrate production SLA-protected E-LINEs one at a time during off-peak windows. Verify BFD, OAM, and latency SLA compliance on SRv6 path before proceeding.	Controlled — SLA circuits, verified first	Per-circuit LP reversion
5. E-LINE Sunset	After all E-LINEs verified on SRv6 for 72h, remove VPWS MPLS pseudowire configurations. Decommission E-LINE MPLS transport.	Zero — all E-LINEs already on SRv6	N/A

5.3 E-LAN Migration Sequence

E-LAN migration is sequenced separately from E-LINE because multipoint broadcast domains have amplified failure blast radius: a misconfiguration in a large E-LAN can simultaneously affect all CE sites in that broadcast domain, potentially disrupting dozens of enterprise customers at once. The migration therefore treats each E-LAN instance as an independent atomic unit, begins with a test E-LAN before touching production, and proceeds from smallest to largest instance by participant count.

Step	Action	Impact	Rollback
1. Test E-LAN	Provision a dedicated test EVPN E-LAN (uDT2U + uDT2M) between 3 test PEs. Send unicast and BUM frames. Verify MAC learning (uDT2U), BUM replication (uDT2M), split-horizon, and loop prevention on SRv6 path.	Zero — test instance only	Delete test E-LAN instance
2. Smallest Production E-LAN	Select the production E-LAN with fewest participants (e.g., 2-PE E-LAN). Set SRv6 NLRI LP=100. Verify all participants reachable via SRv6 MAC learning. Monitor for 60 minutes.	Minimal — smallest broadcast domain	LP reversion, MPLS path resumes

3. Medium E-LAN Instances	Migrate E-LAN instances with 3–8 PE participants, one per change window. Verify BUM replication tree rebuilds correctly on SRv6 EVPN Type-3 Inclusive Multicast Routes.	Controlled — medium broadcast domains	Per-instance LP reversion
4. Large E-LAN / Enterprise	Migrate largest E-LAN instances (10+ PEs, enterprise VPLS replacements). These are highest-impact and require dedicated maintenance windows with customer notification.	High — largest blast radius, full verification required	LP reversion in <30s, MPLS path resumes
5. E-LAN Sunset	After all E-LAN instances verified on SRv6 for 72h, decommission VPLS/EVPN-MPLS transport. Remove MPLS E-LAN forwarding state.	Zero — all E-LANs already on SRv6	N/A

6. DWDM and IPoDWDM Integration — The Complete Optical Migration

The SRv6 TEK Innovations LinkedIn blog series on DWDM and IPoDWDM integration presents the optical transport dimension of the same architectural transformation that this white paper addresses at the IP/MPLS layer. The 'IPoDWDM Revolution') identifies the structural cost of maintaining two completely separate infrastructure layers: an IP routing plane managed by network engineers and an optical transport DWDM plane managed by optical engineers. The SRv6 uSID migration creates the precise precondition for collapsing this two-tier model: once the IP core is running on a unified SRv6 IPv6 plane with Flex-Algo for traffic engineering, the optical layer can be collapsed into it using Coherent IPoDWDM — 400G/800G ZR+ pluggables installed directly in router line cards, driven by the same SRv6 Flex-Algo TE policies that govern IP forwarding.

6.1 IPoDWDM Architecture — Eliminating the Transponder Tier

Before: Two-Tier Architecture	IP routing plane (routers, IS-IS, BGP) managed by network engineers. Optical DWDM plane (transponders, ROADMs, optical controllers) managed by optical engineers. Gray optics connect routers to external transponder chassis ($200K–$800K per chassis). Two management domains, two toolchains, two sets of engineers.
After: Coherent IPoDWDM	400G/800G ZR+ coherent pluggables installed directly in router line cards. Router OS manages optical performance parameters (chromatic dispersion, OSNR, pre-FEC BER) natively via gNMI/YANG. SRv6 Flex-Algo manages wavelength path selection. One management domain, one toolchain, one set of engineers.
CapEx Reduction	40–65% on metro and regional P2P connectivity. Elimination of external optical transponder chassis, ROADM aggregation for Option B (disaggregated λ channels), and optical domain controller software licensing.

The SRv6 Connection	SRv6 Flex-Algo plane assignments map directly to wavelength paths in the IPoDWDM architecture: FA-128 (low-latency) rides the shortest optical path, FA-129 (bandwidth) rides the highest-capacity wavelength. The same Flex-Algo constraints that govern IP forwarding govern optical path selection.

6.2 Flex-Algo Channel Grouping vs. Static Fat-Pipe Aggregation

The Flex-Algo Channel Grouping analysis identifies the optical analog of the ECMP spray-and-pray problem: static fat-pipe aggregation that groups all traffic types into a single rigid optical wavelength through static ROADM channels. Low-latency financial transactions, high-volume cloud backups, and standard internet traffic all compete for the same physical wavelength path, with no mechanism for the optical layer to differentiate between them. SRv6 Flex-Algo channel grouping solves this by partitioning optical bandwidth dynamically: each Flex-Algo plane (FA-128 Latency, FA-129 BW, FA-130 TE, FA-131 Backup) maps to its own spectral allocation in Option B's disaggregated λ channel architecture, providing native optical-layer traffic engineering without any ROADM controller.

6.3 Gold/Silver/Bronze Optical Tiering, Dark Fiber, and Alien Wavelengths

Gold/Silver/Bronze Optical Network Slicing	Native optical network slicing at the protocol layer. Gold: lowest-latency shortest-reach fiber spans via Flex-Algo min-latency metric. Silver: protection-mapped channels with TI-LFA on fiber. Bronze: highest-capacity longer-reach paths for non-latency-sensitive bulk transfer.
P2P POI over Dark Fiber	For financial enterprises requiring data sovereignty: coherent pluggables drive raw dark fiber directly from the router port, managed by SRv6 TI-LFA path protection. No external optical controller, no vendor lock-in, no intermediate encapsulation layer.
Alien Wavelengths over Lit Fiber	On leased lit fiber: custom optical frequencies generated from standard coherent transceivers at the edge router bypass the fiber provider's transponder requirement. SRv6 Flex-Algo manages alien wavelength path selection across the third-party ROADM infrastructure without interference.

6.4 Long-Span Tuning, Spectral Coexistence, and Controller Elimination

Long-Span Amplitude Tuning	Router OS adjusts modulation parameters (constellation shaping, launch power) natively for each span's physical fiber characteristics. Extends 400G/800G reach without mid-point regenerators — directly driven by the SRv6 routing OS that already manages the IP layer.
Spectral Coexistence with Legacy OTN	During optical migration, 400G/800G coherent wavelengths deploy alongside active legacy OTN frequencies in the same fiber. Automated laser power shaping prevents high-speed channels from overwhelming legacy transponders — the same zero-downtime philosophy as the IP/MPLS→SRv6 migration.

Coherent DSP Packaging in QSFP-DD	The entire DSP engine packaged inside QSFP-DD. Router OS controls chromatic dispersion and SNR directly. Eliminates the proprietary transponder shelf management silo that previously required optical engineers for every wavelength change.
Eliminating Standalone Optical Controllers	Single-engine routing logic (SRv6 IS-IS + Flex-Algo) replaces separate Optical Domain Controller + IP Network Controller cross-vendor API stack. Direct Layer-3 visibility into fiber path conditions enables instantaneous path re-computation — the optical layer becomes an extension of the IP control plane, not a separate domain.

7. Complete Protocol Stack Coverage

Protocol	RFC / Standard	SRv6 Migration Role
IS-IS SRv6	RFC 9252 / draft-ietf-lsr-isis-srv6	SRv6 Locator TLV, MT-ISIS for IPv6 topology, Flex-Algo FAD TLV for per-class SID allocation. Advertises both MPLS SR-prefix-SID and SRv6 locators during dual-plane phase.
OSPFv2/v3 SRv6	RFC 9513 / draft-ietf-lsr-ospfv3-srv6	SRv6 Capabilities LSA, End SID sub-TLVs. Used in access domains where OSPF is preferred over IS-IS.
BGP SRv6 L3VPN	RFC 9252	BGP-LU/L3VPN NLRIs with SRv6 Service TLV carrying uDT4/uDT6/uDT46 SIDs. Dual-advertised with MPLS NLRIs for parallel-plane operation.
BGP EVPN SRv6	RFC 9252, RFC 8365	EVPN Type-2/3 routes with SRv6 Service TLV carrying uDX2/uDX2V/uDT2U/uDT2M SIDs. Parallel to EVPN MPLS transport.
BGP-LS	RFC 9552	Topology advertisement to SR-PCE. Both MPLS and SRv6 topology exported. Enables PCE path computation across dual-plane migration state.
BGP SR-TE Policy	RFC 9256	SR Policy NLRI with color communities for SR-TE explicit path steering. Used for post-migration TE optimization on SRv6 plane.
PCE / PCEP	RFC 8664	SRv6-ERO/RRO with uSID sub-objects. PCE computes optimal SRv6 paths and delegates via PCEP to headend PEs.
RSVP-TE (Existing)	RFC 3209	Remains fully operational on IPv4 plane throughout migration. Gracefully decommissioned in Phase 4 after all services verified on SRv6.
LDP (Existing)	RFC 5036	Remains in access domains. Decommissioned after L3VPN and EVPN access services fully migrated to SRv6.

BFD for SRv6	RFC 9127	BFD discriminator-to-SID mapping. Sub-second failure detection for SRv6 SR-TE paths. Replaces RSVP-TE hello/dead-interval failure detection.
TI-LFA	RFC 8355	Pre-computed SRv6 backup segment lists for sub-50ms local repair. Requires P-node uSID capability (Phase 2). Phase 1: standard IP-FRR fallback.
RPKI / ROV	RFC 8210	Route Origin Validation applied to all BGP updates including SRv6 VPN NLRIs. Routinator/GoRTR integration validated in ContainerLab simulation.

8. Hardware Breakpoint Matrix — What Needs Upgrading and What Doesn't

One of the most persistent migration objections is the belief that SRv6 requires a full hardware forklift. The 6-hop uSID sweet spot (RFC 9800 NEXT-CSID) eliminates this belief entirely for transit nodes. The only hardware that requires specific SRv6 capability is the PE (Provider Edge) nodes performing SID endpoint behaviors — and even here, most platforms already in production require only a software/microcode update.

Category	Platforms	Migration Role
Blind Transit — Zero Upgrade	Trident 2/3, Tomahawk 1/2/3, Jericho 1, Qumran-MX; ASR9000 Typhoon, NCS5500 Gen-1; Juniper MX Trio 1–4; Nokia FP3	Forward SRv6 uSID packets as standard IPv6 LPM — no SRH parsing, no DA manipulation. Serve as P-nodes indefinitely.
SW Upgrade Only — PE Capable	Trident 4/5, Jericho 2/2c/3; Cisco NCS5500/5700 NC57 (IOS-XR 7.3.2+), ASR9000 TH Gen-3; Juniper MX Trio5/6, ACX7000; Nokia 7250 IXR, FP4/FP5 (SR OS 22.5.R1+); Arista Jericho2 (EOS 4.28+)	Full uSID Endpoint: End, End.X, End.DT4/6/46, uDX2/uDT2U/uDT2M. Activate via firmware/software update — no hardware replacement.
Greenfield Native	Cisco Silicon One Q100/Q200/P100/P200 (Cisco 8000, Nexus SONiC); Broadcom Tomahawk 4/5/6; NVIDIA Spectrum-2/3/4; Jericho 3 AI	Built for NEXT-CSID from day one. All SRv6 behaviors, TI-LFA, Flex-Algo at line rate.

9. Phase-by-Phase Migration Timeline

Phase	Activities	Duration	Rollback
Phase 1 Audit & Plan	RSVP-TE LSP inventory. Traffic matrix capture per VRF and E-LINE/E-LAN. SPOF blast-radius analysis. Hardware breakpoint audit per node. SRv6 locator addressing plan. ContainerLab simulation build. No SR-MPLS deployment.	4–8 weeks	N/A — no production change
Phase 2 Dual-Plane Build	IS-IS SRv6 locator enable on all PEs. BGP RFC 9252 SRv6 NLRIs advertised with LP=40. SRv6 plane carries zero traffic (MPLS LP=100 wins). IPoDWDM coherent pluggables deployed on Phase 1 priority links. ContainerLab CI/CD validates every config change.	4–6 weeks	Remove IS-IS SRv6 locator and BGP SRv6 NLRIs — zero traffic impact
Phase 3 Incremental Cutover	SRv6 LP raised to 100 per VRF/service as per Sections 4–5. MPLS LP set to 50 at cutover. Test VRF → small agg → large agg → PE pair → full VRF. E-LINE: test → non-critical → SLA. E-LAN: test → small → medium → large. 5-minute gNMI verification at each step.	8–16 weeks	Single LP reversion per VRF or service, <30s, no RSVP renegotiation
Phase 4 Decommission	72h clean SRv6 operation on all services. Graceful RSVP-TE LSP tear-down. LDP removal from access. IPoDWDM controller decommission. BGP MPLS NLRIs withdrawn. MPLS LFIB cleared. IS-IS SR-MPLS extensions removed.	2–4 weeks	N/A — decommission only after full Phase 3 verification

10. 200-Node ContainerLab Digital-Twin Validation

Every configuration change proposed in Phases 1–4 is first executed in the SRv6 TEK Innovations 200-node ContainerLab simulation before any production router is touched. The simulation runs the exact same container images that will execute in production — Cisco xRd (IOS-XR), Juniper cRPD, Nokia SR Linux, Arista cEOS, FRRouting — wired together by a YAML topology definition that mirrors the production backbone.

BGP Table Scale	Route Views BGP full table injection: 900K+ IPv4 unicast prefixes, 150K+ IPv6 prefixes, 50K+ VPN prefixes across test VRFs. Validates BGP convergence time and RIB memory under production-scale prefix loads.
RPKI / ROV	Live Routinator/GoRTR RPKI validator connected to the simulation. Every BGP update including SRv6 VPN NLRIs validated against Route Origin Authorizations. Invalid prefixes rejected at ingress.
Dual-Plane LP Testing	All three LP states (40/50/100) tested per VRF. Automated verification confirms correct path selection (SRv6 vs MPLS) at each LP value. Rollback simulation: LP reversion from 100 to 50 verified under live traffic injection.

EVPN L2 Service Tests	test E-LINE (uDX2), test E-LAN (uDT2U + uDT2M) provisioned in simulation. MAC learning, BUM replication, split-horizon, and loop prevention verified for both SRv6 and MPLS transport paths simultaneously.
TI-LFA Sub-50ms	Fiber-pull simulation (link shutdown) on every P-node link. TI-LFA backup activation time measured via gNMI streaming. All paths verified <50ms recovery. Pre-FEC BER optical degradation injection via ContainerLab custom links.
CI/CD GitHub Actions	Every topology change — IS-IS locator enable, BGP LP policy change, VRF cutover — triggers automated regression: BGP convergence, EVPN MAC learning, TI-LFA recovery, RPKI validation. Rollback tested on every run. Self-service sandbox handed off to customer engineering team.

11. Production Migration Outcomes

0	<30s	25%	≤6
SR-MPLS Intermediate Steps	Rollback at Any Phase	Latency Reduction	uSID Hops, No SRH

Metric	Before Migration	After Migration
Control Plane Protocols	LDP + RSVP-TE + IS-IS + BGP/MPLS + PCE	IS-IS SRv6 + BGP RFC 9252 + SR-PCE — 3 protocols vs 5+
Forwarding Plane	MPLS LFIB + IPv4/IPv6 FIB — two planes	IPv6 FIB only with uSID — one plane
Header Overhead	4 bytes/label × hop count; grows with depth	Fixed 40-byte IPv6 header, 6 hops in DA, no SRH
P-Node Hardware	MPLS-capable required throughout	Any IPv6-capable silicon — never needs upgrading
Network Slicing	Per-VPN RSVP-TE tunnel proliferation	Flex-Algo FA-128/129/130/131 native planes
OAM / Telemetry	RSVP-TE hello/dead, SNMP polling	BFD RFC 9127, IOAM, gNMI 100ms streaming
MTTR on Failure	RSVP-TE re-signaling: 1–60s	TI-LFA pre-computed backup: <50ms
Optical Integration	Separate OTN/DWDM domain, external controllers	IPoDWDM: coherent pluggables driven by SRv6 IS-IS, zero external optical controller

Automation Toolchain	Vendor CLI + RSVP-TE provisioning scripts	GitOps: NETCONF/gNMI + ArgoCD + ContainerLab CI/CD

12. SRv6 TEK Innovations Blog Series — DWDM and IPoDWDM Reference

The following technical articles published on the SRv6 TEK Innovations website (srv6tek-innovations.com) provide the optical transport design and architecture foundation that extends the IP/MPLS→SRv6 migration blueprint into the physical optical domain. Each article presents a self-contained deep-dive into a specific aspect of collapsing the IP-and-optical two-tier architecture into a single SRv6 IPv6 forwarding plane. This white paper serves as the comprehensive technical synthesis and detailed design analysis for all articles in the DWDM and IPoDWDM integration series.

12.1 The IPoDWDM Revolution — Collapsing IP and Optical into One Plane

Published on srv6tek-innovations.com. For decades, network operators maintained two completely separate infrastructure layers: an IP routing plane managed by network engineers and an optical transport DWDM plane managed by optical engineers. This two-tier architecture requires massive arrays of short-reach gray optics connecting routers to external transponder chassis (costing $200K–$800K per chassis), heavy mid-tier cabling, separate management systems, and two engineering disciplines operating independently with no shared toolchain. The SRv6 TEK Innovations IPoDWDM revolution blog establishes the architectural collapse: 400G/800G ZR+ coherent pluggables installed directly in router line cards drive DWDM wavelengths natively from the routing OS. SRv6 Flex-Algo manages physical wavelength path selection using the same IS-IS metric extensions that govern IP TE. The entire optical management domain — chromatic dispersion, OSNR monitoring, pre-FEC BER, wavelength assignment — is controlled via standard gNMI/YANG from the same automation stack that manages BGP and IS-IS. The IPoDWDM migration is the natural sequel to the IP/MPLS→SRv6 migration: once the IP core runs on a unified SRv6 IPv6 Flex-Algo plane, the optical layer collapses into it, completing the stack consolidation from five layers (Application/IP, MPLS/RSVP-TE, OTN, SONET/SDH, per-λ transponder) to three (SRv6 uSID/IP, coherent pluggable, fiber).

12.2 Flex-Algo Channel Grouping vs. Static Fat-Pipe Aggregation

Published on srv6tek-innovations.com. Traditional optical transport relies on fat-pipe aggregation: all traffic types grouped into a single rigid optical wavelength through static ROADM channels, with low-latency financial transactions, high-volume cloud backups, and standard internet traffic competing for the same physical path. This is the optical analog of ECMP spray-and-pray: a static, traffic-class-blind assignment that wastes expensive spectral capacity. The SRv6 TEK Innovations Flex-Algo channel grouping design partitions optical bandwidth dynamically: FA-128 (min-latency) rides the shortest optical path on the lowest-latency fiber spans; FA-129 (max-bandwidth) rides the highest-capacity wavelength; FA-130 (TE policy) uses explicit optical path constraints; FA-131 (backup) pre-computes protection paths. Each Flex-Algo plane maps to its own spectral allocation in Option B disaggregated λ channel architecture, providing native optical-layer traffic engineering with zero ROADM controller. The same IS-IS Flex-Algo FAD TLV constraints that govern IP TE path selection also govern which wavelength carries which traffic class.

12.3 Gold, Silver, Bronze Optical Network Slicing over Fiber

Published on srv6tek-innovations.com. Carrier revenue maximization requires selling performance-tiered service classes rather than undifferentiated bandwidth, but implementing strict tiering over raw fiber has historically required a complex Layer 2 provisioning layer that is labor-intensive to configure and maintain. The SRv6 TEK Innovations optical tiering design implements Gold/Silver/Bronze slicing natively at the SRv6 Flex-Algo protocol layer: Gold class traffic steered over the absolute lowest-latency, shortest-reach fiber spans via Flex-Algo min-latency metric; Silver class over reliable protection-mapped channels with full TI-LFA on fiber, balancing distance and latency; Bronze class over highest-capacity, longer-reach paths optimized for cost efficiency on non-latency-sensitive bulk transfers. Because tiering is expressed natively at the SRv6 Flex-Algo layer — not via a separate provisioning system — the physical transport network becomes a directly monetizable, tiered business asset. The same Flex-Algo plane assignment that governs IP forwarding also governs the optical tier each customer's traffic rides.

12.4 P2P POI over Dark Fiber — Data Sovereignty without Vendor Lock-In

Published on srv6tek-innovations.com. Global financial enterprises and critical infrastructure operators requiring absolute data sovereignty deploy dedicated Point-to-Point Packet-Optical Integration (POI) directly over private dark fiber. Traditional P2P POI requires purchasing expensive proprietary software suites from optical hardware manufacturers to control the underlying laser paths — a vendor lock-in risk compounding the capital cost of private fiber. The SRv6 TEK Innovations dark fiber P2P design eliminates this: coherent pluggables installed directly in routing switches drive raw dark fiber natively, with SRv6 TI-LFA providing path protection without any external optical controller. The routing data is contained entirely within private, unencapsulated IPv6 infrastructure. The same open-standards principle that governs the IP/MPLS→SRv6 migration — no proprietary controllers, no vendor SDK dependencies, full BGP/IS-IS open protocol stack — extends identically to the optical layer. Dark fiber P2P POI is the logical extension of the IPoDWDM architecture for operators who own their fiber plant and require full stack control.

12.5 Alien Wavelengths over Leased Lit Fiber

Published on srv6tek-innovations.com. When expanding capacity over leased or pre-existing lit fiber, operators traditionally face high costs from the fiber provider's transponder equipment requirement for lighting new channels. The SRv6 TEK Innovations alien wavelength design bypasses this: custom optical frequencies are generated from standard coherent transceivers installed in edge routers, transmitting alien wavelengths safely across the third-party provider's active ROADM infrastructure without interference. SRv6 Flex-Algo manages alien wavelength path selection across the third-party infrastructure. New 400G/800G channels light up on leased fiber in minutes, bypassing vendor provisioning delays and the end-to-end optical service purchase model. The technique applies directly to the migration context: operators building out the SRv6 parallel plane during the dual-plane migration phase can use alien wavelength capacity on leased fiber for the new SRv6 IPv6 paths without requiring new fiber contracts.

12.6 Long-Span Amplitude Tuning and Coherent DSP Management

Published on srv6tek-innovations.com. Connecting regional data centers over long-span fiber requires maximizing optical signal strength without introducing noise. As speeds reach 400G and 800G, signal degradation forces expensive mid-point optical amplification sites every few miles under traditional architectures. The SRv6 TEK Innovations long-span design deploys advanced amplitude tuning and constellation shaping directly from the router OS: optical modulation parameters (launch power, modulation format from DP-16QAM to DP-QPSK, baud rate) are adjusted natively via gNMI/YANG to match each span's physical fiber characteristics. OpenZR+ adaptive modulation (400Gbps, up to 1,200km with EDFA) enables metro-to-regional reach. The Coherent DSP Packaging article in the same series covers the QSFP-DD form factor that packages the entire DSP engine in a hot-pluggable transceiver: chromatic dispersion and SNR monitored directly from the router OS, eliminating the proprietary transponder shelf management silo that previously required optical engineering expertise for every wavelength change. The migration implication: operators deploying IPoDWDM on the new SRv6 IPv6 parallel plane gain coherent optical management natively through the same automation toolchain managing the IP migration.

12.7 Spectral Coexistence with Legacy OTN

Published on srv6tek-innovations.com. The optical migration, like the IP/MPLS migration, rarely has the luxury of empty fiber: new 400G/800G coherent wavelengths must deploy alongside active legacy OTN frequencies in the same fiber bundle. High-power coherent signals generate wide spectral footprints that can bleed into adjacent channels, disrupting existing legacy services. The SRv6 TEK Innovations spectral coexistence design uses the same zero-downtime philosophy as the IP migration: granular grid allocation maps modern coherent wavelengths onto strict, non-overlapping channel grids adjacent to active legacy OTN frequencies; automated laser power shaping prevents high-speed channels from overwhelming sensitive legacy transponders; and SRv6 wraps and routes incoming streams identically regardless of whether the underlying wavelength is legacy or modern. The optical and IP migrations can proceed in parallel: the IP dual-plane migration runs on existing fiber while new IPoDWDM coherent channels are deployed alongside legacy OTN on the same fiber plant.

12.8 Eliminating Standalone Optical Domain Controllers

Published on srv6tek-innovations.com. Traditional POI architectures require a heavy software layer: a separate Optical Domain Controller (ODC) that continuously synchronizes with an IP Network Controller via complex, fragile cross-vendor APIs. An API update from a single hardware vendor can disrupt cross-layer communication and stall automated traffic engineering. The SRv6 TEK Innovations controller elimination design directly parallels the IP migration's elimination of RSVP-TE and PCE separate controller stacks: SRv6 IS-IS Flex-Algo replaces the Optical Domain Controller for wavelength path selection; TI-LFA replaces the optical protection switching controller; gNMI streaming replaces SNMP optical performance monitoring; and IOAM embedded in production packets replaces synthetic optical test traffic. Single-engine routing logic running natively inside the routers maintains complete direct visibility into fiber path conditions — OSNR, pre-FEC BER, Q-factor — enabling instantaneous path re-computation without waiting on an ODC's telemetry cycle. The optical layer becomes an extension of the SRv6 control plane, managed by the same GitOps automation stack as the IP migration, with no separate controller domain, no cross-vendor API surface, and no additional operational toolchain.

13. SRv6 TEK Innovations Blog Series — Core Migration Reference

The following technical articles published on the SRv6 TEK Innovations website (srv6tek-innovations.com) form the architectural and design foundation for the core migration blueprint presented in this white paper. Each article provides a self-contained deep-dive into a specific dimension of the SRv6 uSID architecture, the migration strategy, hardware breakpoints, and the protocol evolution from legacy MPLS to the unified SRv6 IPv6 plane. This white paper serves as the comprehensive technical synthesis and detailed design analysis of the complete core migration blog series.

13.1 The 6-Hop uSID Sweet Spot — Eliminating the SR-MPLS Detour

Published on srv6tek-innovations.com. Establishes the foundational migration argument: the 6-hop SRH-free forwarding model (RFC 9800 NEXT-CSID) eliminates the only reason SR-MPLS was ever considered a necessary intermediate step. With up to six 16-bit uSIDs packed into the standard 128-bit IPv6 Destination Address, P-nodes forward SRv6 packets as ordinary IPv6 unicast via Longest Prefix Match — zero SRH parsing, zero hardware modification, zero upgrade. The article documents the 62% average header compression vs. classic uncompressed SRv6 (RFC 9800 NEXT-CSID), the 100% P-node upgrade avoidance for networks with ≤6 steering hops, and the H.Insert.Red/H.Encaps.Red decoupling that keeps the full 6-hop transport budget available regardless of service overlay complexity. Core conclusion: MPLS → SR-MPLS → SRv6 was the textbook path; MPLS → SRv6 uSID direct is the production path.

13.2 The Hardware Breakpoint Cheat Sheet — What to Upgrade and When

Published on srv6tek-innovations.com. Provides the definitive per-platform guide to SRv6 uSID capability across the full vendor ecosystem. Three categories: Category 1 (Completely Safe Legacy, 0% upgrade) covers Trident 2/2+/3, Tomahawk 1/2/3, Jericho 1, Qumran-MX, ASR9000 Typhoon, NCS5500 Gen-1 Jericho, Juniper MX Trio 1–4, Nokia FP3 — all serve as blind-transit P-nodes indefinitely, reading SRv6 uSID packets as plain IPv6 at full line rate. Category 2 (Software/Microcode Upgrade Only) covers Trident 4/5, Jericho 2/2c/3, Cisco NCS5500/5700 NC57 (IOS-XR 7.3.2+), ASR9000 Tomahawk Gen-3, Juniper MX Trio 5/6, ACX7000, Nokia 7250 IXR, FP4/FP5 (SR OS 22.5.R1+) — already possess the programmable pipeline, firmware update activates uSID endpoint behavior. Category 3 (Greenfield Native) covers Cisco Silicon One Q100/Q200/P100/P200, Broadcom Tomahawk 4/5/6, NVIDIA Spectrum-2/3/4 — built for NEXT-CSID from day one. The article's core economic argument: hardware spend is scoped exclusively to PE endpoints requiring uSID behavior, not the transit core, eliminating the 'forklift' narrative entirely.

13.3 The Paradigm Shift — Skipping SR-MPLS Entirely

Published on srv6tek-innovations.com. Presents the architectural comparison of three forwarding paradigms side by side: Legacy MPLS (LDP + RSVP-TE, two separate protocols, MPLS label stack, no native IPv6 integration), SR-MPLS/SR-MPLSv6 (unified IS-IS/OSPF control plane but still on MPLS data plane — an interim architecture, not an end-state), and SRv6 uSID (≤6 hops, no SRH, native IPv6 DA, zero MPLS). Documents why SR-MPLS became attractive historically (P-node upgrade avoidance, software-only migration from LDP/RSVP-TE) and why uSID makes that justification obsolete — P-nodes require zero upgrade under SRv6 uSID, exactly the same proposition. Covers

the network slicing dimension: SR-MPLS requires MPLS-label-based tunnel proliferation for 5G slicing, while SRv6 uSID provides native per-uSID Flex-Algo planes that integrate directly with eBPF/Cilium Kubernetes environments at zero additional encapsulation cost. The direct MPLS → SRv6 uSID path is not just architecturally cleaner; it eliminates an entire migration phase worth of operational risk.

13.4 Blind Transit Migration Strategy — Phase 1 to Phase 2

Published on srv6tek-innovations.com. Details the two-phase migration execution strategy that the BGP Local Preference sequencing in this white paper implements. Phase 1 upgrades only PE (Provider Edge) nodes to SRv6 uSID while every P (core/transit) node remains untouched legacy hardware running vanilla IPv6 forwarding. P-nodes read the uSID-encoded Destination Address as plain IPv6, find the longest-prefix match for the PE's /64 locator prefix in their FIB, and forward — no shift, no SRH parse, no awareness that anything has changed. IS-IS/OSPF simply advertises IPv6 reachability to PE locator prefixes; P-nodes forward SRv6-encapsulated VPN traffic blindly alongside all other IPv6 traffic. Phase 1 immediate results: L2VPN (EVPN-VPWS/VPLS) and L3VPN services work 100% immediately; PEs handle all service encap/decap; the legacy core provides underlay reachability. Phase 1 limitations: no explicit waypoint steering through P-nodes (legacy P cannot Shift-and-Forward), TI-LFA unavailable until Phase 2 P-node refresh (legacy P falls back to standard IP-FRR/LFA). Phase 2 refreshes P-nodes at normal hardware cycles, enabling TI-LFA and explicit SR-TE through the full core.

13.5 L2 VPN EVPN Service Migration — RFC 9252 SRv6 Service Types

Published on srv6tek-innovations.com. Provides the complete technical mapping of every EVPN L2 service type from MPLS transport to SRv6 transport using RFC 9252 SRv6 Service TLV NLRIs. uDX2 (EVPN E-LINE) replaces VPWS MPLS pseudowire: Ethernet cross-connect, all-to-one forwarding, single SID per attachment circuit, no MAC learning at transport layer. uDX2V (E-LINE VLAN-aware bundled) replaces VLAN-bundled VPWS with per-VLAN uSID mapping, preserving VLAN tag forwarding semantics. uDT2U (E-LAN unicast) replaces VPLS unicast: per-remote-MAC BGP EVPN Type-2 route maps to a per-MAC uSID at the egress PE. uDT2M (E-LAN BUM) replaces VPLS flooding: BUM traffic replicated by the ingress PE to a BUM SID list derived from EVPN Inclusive Multicast Route (Type-3). uDT4/uDT6/uDT46 (L3VPN) replaces BGP/MPLS L3VPN label stack with IPv4, IPv6, and dual-stack VRF table lookup SIDs. All service types are dual-advertised during the migration — MPLS NLRI and SRv6 NLRI simultaneously — with LOCAL_PREF governing which transport plane carries production traffic at each step.

13.6 Transport vs. Service SIDs — H.Insert.Red and H.Encaps.Red Decoupling

Published on srv6tek-innovations.com. Explains the critical architectural distinction between transport SIDs (the underlay forwarding path through the core) and service SIDs (the VPN endpoint behavior at the egress PE), and why their decoupling via H.Insert.Red and H.Encaps.Red encapsulation modes is essential to preserving the full 6-hop transport budget. In a merged-encapsulation design (overlay/underlay merge enabled), every service SID consumed from the 6-slot uSID carrier reduces the available transport hop budget: a 2-SID service overlay leaves only 4 entries for transport steering hops. With H.Insert.Red/H.Encaps.Red decoupling — the default on virtually every SRv6 platform — the service identification for L2VPN/L3VPN endpoint behaviors (End.DT4, End.DT6, End.DX2, VPN

function/argument) is placed in a separate, reduced-encapsulation SRH, while the transport uSID carrier in the IPv6 DA carries exclusively the core steering hops. The full 6-hop transport budget remains available regardless of service layer complexity. The article covers per-platform default behavior: overlay/underlay merge is disabled by default across Cisco, Juniper, Nokia, and Arista, making H.Insert.Red/H.Encaps.Red decoupling the production baseline without additional configuration.

13.7 Whitebox Vendors and Open NOS SRv6 uSID Pairing

Published on srv6tek-innovations.com. Maps the disaggregated whitebox platform ecosystem to their SRv6 uSID capability and recommended Open NOS pairing: UfiSpace S9100/S9110 (Qumran 2c, IP Infusion OcNOS or Arrcus ArcOS, Carrier Edge/CSG), UfiSpace S9300-32D (Trident 4, OcNOS, Aggregation/Peering), UfiSpace S9321-64E (Tomahawk 5, ArcOS or SONiC, AI Backend Fabric), Edgecore AS9716-32D (Tomahawk 4, SONiC, DC Fabric), Quanta QuantaMesh T9032 (Tomahawk 5, SONiC 202411+, AI Backend), Celestica Seastone DX010 (TH4/5, SONiC or OcNOS, Hyperscale DC), Ragile RA-B6920-64C (Tomahawk 5, SONiC custom SAI, AI/HPC Fabric). Minimum Open NOS versions: SONiC 202305 (basic SAI SRv6 VPN support), SONiC 202411 PhoenixWing (AI stable, 6-hop DA-only production baseline), Arrcus ArcOS 5.0.1 (industry's first full uSID shift-and-forward and deep SRH parsing), IP Infusion OcNOS 6.3.0 SP Profile (production uSID, full Cisco/Nokia SRv6 domain interop). Note: Cumulus Linux 5.x post-NVIDIA acquisition targets Spectrum ASICs only — Broadcom support frozen at 4.3.x; SONiC, ArcOS, and OcNOS are the production-grade open NOS options for Broadcom whitebox uSID deployments.

13.8 200-Node ContainerLab Digital Twin — The Migration Simulation Platform

Published on srv6tek-innovations.com. Documents the SRv6 TEK Innovations ContainerLab backbone simulation used to validate every migration step before any production router is touched. The 200-node topology runs exact production container images (Cisco xRd IOS-XR, Juniper cRPD, Nokia SR Linux, Arista cEOS, FRRouting) wired by YAML topology definition on a single bare-metal Ubuntu server. Route Views BGP full table injection (900K+ IPv4 prefixes, 150K+ IPv6, 50K+ VPN prefixes across test VRFs) validates convergence time and RIB memory under production-scale load. Live RPKI/Routinator validates all BGP updates including SRv6 VPN NLRIs. All three LOCAL_PREF states (LP=40/50/100) are tested per VRF with automated path-selection verification. EVPN L2 services (test E-LINE uDX2, test E-LAN uDT2U + uDT2M) are provisioned and MAC learning, BUM replication, split-horizon, and loop prevention are verified for both SRv6 and MPLS transport simultaneously. TI-LFA sub-50ms recovery is validated via link-shutdown simulation on every P-node interface. GitHub Actions CI/CD runs the full regression suite on every topology change, handing off a self-service sandbox topology-as-code repository to the customer engineering team for ongoing post-migration testing.

White Paper Summary — SRv6 uSID Technology Conclusion

The dual-plane 'two ships in the night' migration architecture presented in this white paper is the only approach that simultaneously satisfies all four migration requirements: zero production downtime, complete rollback at every step, no intermediate SR-MPLS architecture that creates IPv6 plane conflicts, and incremental traffic cutover from lowest to highest impact. The BGP Local Preference sequencing — SRv6 at LP=40 during build-out, MPLS de-preferred to LP=50 at cutover, SRv6 at default LP=100 as preferred — provides a mathematically precise, operationally safe mechanism that never risks displacing eBGP PE–CE customer routes regardless of the operator's existing BGP policy architecture.

The VRF-by-VRF, prefix-by-prefix, E-LINE-by-E-LINE, E-LAN-by-E-LAN migration sequence transforms what appears to be a massive infrastructure transition into a series of individually reversible, telemetry-verified steps that any operations team can execute safely within normal change windows. The ContainerLab 200-node digital twin provides production-fidelity validation before any production router is touched, and the CI/CD pipeline ensures every configuration change is regression-tested against the full protocol suite — BGP, IS-IS, EVPN, TI-LFA, RPKI — before deployment.

The DWDM/IPoDWDM integration (DWDM and IPoDWDM integration series) represents the natural extension of this migration: once the IP core runs on a unified SRv6 IPv6 Flex-Algo plane, the optical transport layer collapses into it via coherent pluggables driven directly from the routing OS. The same migration philosophy — parallel planes, incremental cutover, no separate controller domain — applies equally to eliminating the optical transponder tier that has been the second major cost and complexity sink in carrier transport networks for two decades.

Ready to Execute Your Zero-Downtime Core Migration?

SRv6 TEK Innovations provides LP sequencing design, ContainerLab digital-twin validation, VRF/EVPN incremental cutover planning, and production cutover support for your RSVP-TE to SRv6 uSID migration.

NAICS: 541512 · 541519 · 541611 · MBE-Eligible Minority-Owned Small Business