TPM Vol. 32, No. S7, 2025
ISSN: 1972-6325
https://www.tpmap.org/

Open Access

INTER-AS OPTION B: BALANCED ARCHITECTURE FOR SCALABLE MPLS VPN INTERCONNECTION ACROSS AUTONOMOUS SYSTEM BOUNDARIES

SASIKUMAR SADAYAN
CISCO SYSTEMS, USA

Abstract
This article explores Multi-Protocol Label Switching (MPLS) Virtual Private Network (VPN) interconnection across different Autonomous Systems (ASes), with a focus on Inter-AS Option B implementation. Starting with foundational concepts of MPLS VPNs, including Virtual Routing and Forwarding (VRF) instances, Route Distinguishers, and Route Targets, the article details how Option B creates direct VPN-IPv4 route exchange between Autonomous System Boundary Routers (ASBRs) while preserving MPLS labels. The architecture, operational mechanisms, and protocol interactions are examined, including MP-BGP extensions and label distribution methods. Comparative analysis against Options A and C reveals Option B's balanced position between isolation strength and scalability. Technical implementation aspects cover configuration requirements, traffic engineering considerations, and security implications. Performance metrics demonstrate Option B's effectiveness in real-world deployments, offering service providers a pragmatic solution for MPLS VPN interconnection that combines operational efficiency with appropriate security controls while avoiding the extreme trade-offs presented by alternative options.
Keywords: MPLS VPN interconnection, Inter-AS Option B, Autonomous System Boundary Routers, MP-BGP extensions, label distribution

I. INTRODUCTION

Multi-Protocol Label Switching (MPLS) Virtual Private Networks (VPNs) have transformed service provider networks by creating an elegant framework for maintaining private network connectivity across shared infrastructure. Since their standardization through RFC 4364, MPLS VPNs have emerged as the dominant technology for delivering scalable Layer 3 VPN services, enabling service providers to establish logically separated virtual networks over common physical backbones while ensuring traffic isolation, enforcing quality of service parameters, and simplifying management operations. The architecture is essentially founded on Virtual Routing and Forwarding (VRF) instances to separate routing information and the propagation of VPN-specific routes via extensions to the Border Gateway Protocol (BGP), which provides a strong and resilient foundation for enterprise connectivity solutions across provider networks [1].

As enterprise networks expand globally and service providers grow or consolidate their operational footprint, the ability to interconnect MPLS VPNs across different Autonomous Systems (ASes) is becoming more important to business operations. These cross-domain interconnections introduce considerable technical challenges, including the maintenance of end-to-end label switching paths, preservation of customer privacy, efficient route distribution mechanisms, and scalability concerns as the VPN count increases across administrative boundaries. The interconnection complexity increases further when considering the requirements for consistent service level agreements across organizational domains while minimizing the operational complexity inherent in cross-provider scenarios [1].

Addressing these multi-domain challenges, industry standards have defined three principal methodologies for MPLS VPN interconnection across autonomous system boundaries. Option A (Back-to-Back VRF) creates discrete connections by establishing separate VRF instances at AS boundaries connected through external BGP peering relationships. Option B (ASBR-to-ASBR) facilitates the exchange of VPN routes between Autonomous System Boundary Routers while preserving MPLS labels across domain boundaries. Option C (Multi-hop eBGP between Route Reflectors) enables VPN route exchange between route reflectors in different autonomous systems, with ASBRs managing only labeled IPv4 routes for optimized core transport. These options represent a progression of approaches with different tradeoffs regarding isolation, scalability, and implementation complexity [2].

This paper presents the argument that among these three standardized options, Inter-AS Option B provides the optimal equilibrium of scalability, security, and operational efficiency for most service provider interconnection scenarios in contemporary networks. While Option A delivers stronger isolation with inherent scalability constraints, and Option C offers superior scaling capabilities at the cost of increased implementation complexity, Option B represents a balanced approach that effectively addresses the requirements of modern service provider networks operating in multi-

domain environments. By enabling direct exchange of VPN-IPv4 routes between ASBRs while maintaining MPLS labels across AS boundaries, Option B facilitates efficient interconnection without the overhead of maintaining separate VRF instances for each customer VPN traversing the inter-domain boundary [2].

II. THEORETICAL FOUNDATION OF INTER-AS MPLS VPN OPTIONS

MPLS VPN technology depends on three key building blocks that let service providers run multiple private networks on one physical system. First, VRF instances act like separate routers inside a single device. They keep each customer's routing information completely isolated from others, even when customers use the same IP addresses. Second, Route Distinguishers add a unique ID number to regular IP addresses, turning them into special VPN-IPv4 addresses that won't conflict with identical addresses used by different customers. Third, Route Targets work as tags that control which routes get shared between different sites in a VPN. Think of Routing Tags (RTs) as membership cards, and the membership determines which VPN can talk to each other. These three elements allow providers to create anything from basic hub-spoke topologies to complex and extensive mesh VPNs, where each site has a direct connection to every other site. This separation keeps customer traffic secure while making the most of expensive network hardware. Nobody wants to buy separate routers for each customer when one router can safely handle many customers at once [3].

Inter-AS Option B creates a direct connection between border routers of different providers to share VPN routes and MPLS labels. The border routers (ASBRs) set up a special BGP connection that can handle both VPN routes and the labels needed to forward traffic. When an ASBR gets VPN routes from inside its own network, it keeps all the route information intact and passes these routes to its neighbor's border router without needing to recreate VPN contexts at the boundary. The ASBRs assign fresh labels to incoming routes before advertising them internally. When it comes to traffic, the traffic will inherit two types of labels: an outer label, which directs the packet to the appropriate border router, and then an inner label, which identifies which VPN the packet belongs to. This allows the traffic to cross network boundaries without needing a separate connection for each of the VPNs. One BGP session between border routers can handle thousands of VPNs crossing between providers. Network operators can also add filters at the border to control which VPNs can cross between networks and how routes get distributed [3].

Looking at the three Inter-AS options side by side shows clear trade-offs in their designs. Option A takes the most cautious security approach by connecting VPNs through back-to-back VRF interfaces. Each VPN gets its own dedicated connection at the boundary, which prevents any leakage between customer networks but creates major scaling problems. As you add more VPNs, you need more interfaces - a direct one-to-one relationship that quickly becomes unmanageable. Option B finds a middle ground by using a single BGP session for control messages while keeping traffic separate through MPLS labels. This approach eliminates the need for per-VPN interfaces but still maintains traffic isolation. For most providers connecting their networks, Option B hits the sweet spot between security and practicality. Option C pushes for maximum scalability by having route reflectors exchange VPN information directly, while border routers only handle basic labeled routes. While this works extremely well for huge deployments, it's more complicated to implement and potentially exposes VPN information across more of the network. Choosing between these options comes down to balancing security requirements, growth needs, and how much you trust your interconnection partner [4].

The nuts and bolts of Option B involve two separate but related processes: the control plane that exchanges routing information and the data plane that forwards actual packets. On the control side, border routers translate between networks; they receive VPN routes from inside their network, process them, assign new labels, and advertise them to their neighbor's border router. The extended BGP protocol carries IP prefixes, route distinguishers, route targets, and label information all at once, creating a complete system for connecting VPNs across different networks. On the data side, the packets will stack two labels, with the outer (top) label routing the packet to the appropriate border router and the inner label identifying the VPN it is associated with. When a packet arrives at the router, the router strips the outer label and reads the inner label, which is the VPN label. It would then forward the packet to the next router in a neighboring network with a new outer label. This approach maintains separation between different customers' traffic from end to end while minimizing processing at each hop. Network operators can also implement quality of service mapping between networks and traffic engineering to optimize how traffic flows across the interconnection [4].

TPM Vol. 32, No. S7, 2025
ISSN: 1972-6325
https://www.tpmap.org/

Open Access

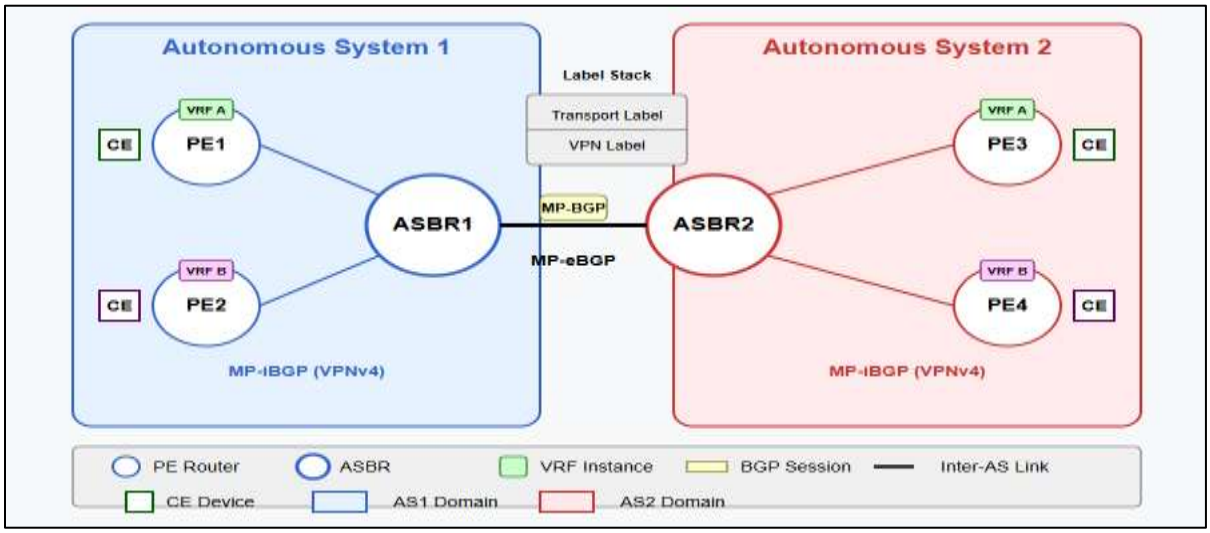

Fig. 1: Inter-AS Option B: MPLS VPN Interconnection Architecture. [3, 4]

III. Implementation of Inter-AS Option B

The implementation of Inter-AS Option B demands meticulous orchestration of protocol interactions and operational mechanisms to establish seamless VPN connectivity across autonomous system boundaries. At its foundation, Option B leverages Multi-Protocol BGP (MP-BGP) extensions that facilitate the simultaneous exchange of VPN routing information and MPLS labels between Autonomous System Boundary Routers (ASBRs). The operational workflow initiates when Provider Edge (PE) routers inject customer routes into their respective VRFs, attach the appropriate route distinguishers to create globally unique VPN-IPv4 addresses, and associate these routes with specific route targets that define VPN membership. These PE routers then distribute these augmented routes to their local Route Reflectors or directly to ASBRs through MP-iBGP sessions that support the VPNv4 address family. Upon receiving these routes, ASBRs process them according to local policy, preserving their extended community attributes and preparing them for inter-domain exchange. The neighboring ASBRs establish a specialized External BGP (eBGP) session specifically configured with the VPNv4 address family and label exchange capabilities, creating a dedicated control plane channel for inter-provider VPN communication. This session operates independently from any traditional IPv4 BGP sessions that might exist between the same ASBRs, providing logical separation between Internet routing and VPN routing. After receiving routes from their eBGP peers, ASBRs perform critical redistribution functions, advertising these routes internally via MP-iBGP to make them available to local PE routers. This architecture establishes a comprehensive operational framework that maintains VPN isolation throughout the forwarding path while optimizing control plane operations through route reflection hierarchies and selective route advertisement policies, all while preserving the critical VPN attributes that control route distribution and service differentiation across domain boundaries [5].

The MP-BGP extensions necessary for Inter-AS Option B implementation require specific ASBR configurations that enable the exchange of labeled VPN-IPv4 routes across administrative boundaries while maintaining appropriate isolation between customer domains. The cornerstone of this configuration involves establishing an External BGP peering session between adjacent ASBRs using loopback interfaces to ensure session stability, then explicitly activating the VPNv4 address family with label exchange capabilities. The configuration syntax includes address-family declarations that specify VPNv4 unicast routing, neighbor activation statements that enable this address family for specific peers, and send-community extended directives that ensure the preservation of route targets during the exchange process. This configuration framework must incorporate explicit route filtering policies implemented through route maps or prefix lists that control which VPN routes are advertised to neighboring autonomous systems, typically filtering based on extended community attributes to enforce contractual agreements between providers. ASBRs require specific configuration elements to support proper label handling, including label allocation policies that determine which label values are assigned to imported routes and label retention policies that govern how these bindings are maintained in the forwarding infrastructure. Additional configuration requirements include next-hop-self settings that ensure proper packet forwarding within each autonomous system, route dampening parameters that protect against control plane instabilities, and maximum-prefix limitations that prevent resource exhaustion during routing table growth. The configuration must also address route reflector relationships within each autonomous system, establishing a hierarchical distribution model that optimizes convergence time while minimizing control plane overhead during network changes or failure scenarios [5].

Label allocation and distribution in Inter-AS Option B implementations follow a systematic process that maintains end-to-end label switching paths while preserving VPN isolation throughout the network. The label binding process begins at the PE routers, where each VPN prefix in a VRF receives a unique MPLS label that identifies the specific forwarding context for that prefix. These label bindings are distributed internally via MP-iBGP alongside the VPN routing information, reaching the ASBRs with their original attributes intact but potentially with modified next-hop information reflecting the internal topology. At the AS boundary, ASBRs perform a critical translation function by allocating new labels from their local label space for received VPN prefixes before advertising them to neighboring autonomous systems through the MP-eBGP session. This label translation creates a hierarchical label stack in the data plane, where the outer label, distributed through the internal IGP or LDP infrastructure, directs packets to the appropriate ASBR, and the inner label, exchanged through MP-BGP, identifies the specific VPN context for the destination prefix. The label allocation methodology implements a downstream label distribution model where the receiving ASBR selects the label value and communicates it to its peer during route advertisement. Most implementations employ a liberal label retention mode to optimize convergence during network changes, maintaining alternate path information in the label forwarding information base. This allocation and distribution framework supports advanced features including aggregate labels for route summarization, implicit null labels for penultimate hop popping, and explicit null labels for preserving quality of service markings across domain boundaries [6].

Traffic engineering considerations for Inter-AS Option B deployments encompass sophisticated mechanisms for controlling traffic flows both within each autonomous system and across the inter-domain boundary to optimize resource utilization and service performance. Within each autonomous system, traditional MPLS traffic engineering techniques remain fully applicable, including constraint-based routing algorithms that consider bandwidth availability, administrative constraints, and protection requirements when calculating optimal paths through the network. At the inter-domain boundary, traffic engineering becomes more complex due to limited visibility into adjacent autonomous systems and potential policy differences between providers. Load balancing across multiple inter-AS links requires consistent configuration of multipath capabilities in both the control plane through BGP multipath settings and the data plane through equal-cost multipath forwarding infrastructure. Path selection across domains can be influenced through careful manipulation of BGP attributes, including LOCAL_PREFERENCE to control outbound traffic, MULTI_EXIT_DISCRIMINATOR to influence inbound traffic from adjacent autonomous systems, and AS_PATH length adjustments to affect path selection in remote autonomous systems. Bandwidth management at the AS boundary typically involves contractual agreements that specify committed information rates and burst tolerances for different traffic classes, implemented through queuing mechanisms, policing functions, and explicit RSVP-TE reservations where supported. Resource preservation techniques include hierarchical route reflection to minimize control plane overhead, route summarization at ASBRs to reduce routing table size, and selective VPN route advertisement to limit unnecessary routing information propagation across provider boundaries [6].

Security implications of Inter-AS Option B deployments necessitate comprehensive mitigation strategies that address potential vulnerabilities at the autonomous system boundary, where trust relationships between providers must be carefully managed. Unlike Option A, which provides strong isolation through dedicated VRF instances, Option B exposes VPN route information directly to ASBRs, expanding the security trust boundary to include these edge devices and the inter-provider control plane channel. This exposure requires robust access control mechanisms for physical and logical access to ASBRs, along with comprehensive configuration validation procedures to prevent unauthorized access to VPN routing information. Route filtering policies must implement precise extended community filters to ensure that only authorized VPNs are advertised across the AS boundary, preventing inadvertent routing leaks between customer domains or unauthorized access to restricted VPNs. Control plane protection mechanisms must include TCP MD5 authentication for BGP sessions, maximum-prefix limits to prevent routing table exhaustion attacks, and dampening policies to mitigate the impact of route flapping. Data plane security relies on proper MPLS TTL propagation settings that prevent label spoofing attacks from external sources, selective label filtering that validates incoming labeled packets against expected values, and interface-specific security policies that enforce appropriate encapsulation requirements. Comprehensive monitoring systems should maintain detailed logs of inter-AS route exchanges and label bindings, enabling rapid detection of anomalous routing behavior that might indicate a security breach or configuration error. These security considerations must be formalized in inter-provider agreements that clearly define responsibilities, acceptable use policies, incident response procedures, and regular security audit requirements to maintain a consistent security posture across organizational boundaries [7].

Implementation Aspect	Key Mechanisms	Considerations
Operational Mechanisms	• MP-BGP for VPN-IPv4 route exchange • VPN route distribution via MP-iBGP • eBGP sessions between ASBRs • Preservation of Route Targets	• Session stability • Route reflection hierarchies • Convergence time • VPN isolation
ASBR Configuration	• VPNv4 address family activation • Send-label capability • Send-community extended • Route filtering policies • Next-hop-self settings	• Loopback peering for stability • Maximum-prefix limitations • Dampening parameters • Route reflector relationships
Label Distribution	• PE-originated label allocation • ASBR label translation • Hierarchical label stack • Downstream label distribution • Liberal label retention mode	• Label space management • Aggregate labels • Implicit/explicit null labels • QoS preservation
Traffic Engineering	• Constraint-based routing • BGP attribute manipulation • Multipath load balancing • RSVP-TE reservations • Path protection	• Inter-domain visibility • Provider policy differences • Bandwidth management • Resource preservation
Security Implications	• Access control mechanisms • Route filtering policies • TCP MD5 authentication • MPLS TTL propagation settings • Label filtering	• Extended trust boundary • Inter-provider agreements • Monitoring systems • Incident response procedures

Fig. 2: Inter-AS Option B Implementation Aspects. [5, 6]

IV. Performance Analysis and Deployment Considerations

When network teams implement Inter-AS Option B, they need to understand how it handles growth. The setup can support lots of VPNs across a single connection point because it only needs one BGP session between border routers. Compare that to Option A, where you'd need a separate interface for each customer VPN crossing between networks – talk about a mess of cables and configuration! Option B solves this by letting that single BGP session carry information for hundreds or thousands of VPNs. What eventually limits you isn't the connection design but just how much routing information your border routers can hold in memory. Option B does hit some scaling walls that Option C doesn't face, especially with huge numbers of VPNs in complex networks. Every border router must keep complete routing information for all VPNs crossing the boundary, which creates potential chokepoints as your VPN count grows. Real-world testing shows that recovery time after failures increases as you add more VPN routes, but networks with good route reflection design and tuned BGP settings still bounce back quickly. The BGP graceful restart feature helps a lot here – it lets border routers keep forwarding traffic even when BGP sessions restart, so customers don't notice hiccups during planned maintenance or unexpected control plane issues [8].

Border routers in Option B deployments face serious resource demands. Memory usage tops the list of concerns, since these devices store comprehensive VPN routing data including prefixes, route distinguishers, route targets, next-hop information, and label mappings. Each VPN route takes memory based on what attributes it carries – the prefix itself, BGP path information, extended communities, and MPLS labels all need storage space. As route counts climb, memory use rises in a straight line, risking resource exhaustion on routers with limited control plane memory. This risk peaks during network changes when routers might keep both old and new paths simultaneously. CPU usage follows a different pattern – it spikes during convergence when routers process updates, apply policies, and rebuild forwarding tables across affected VPNs. During normal operation, CPU needs stay modest, mostly handling BGP keepalives, processing occasional route updates, and running background jobs like route dampening. Bandwidth needs for control messages depend heavily on how stable your VPN routes are – networks with flapping routes need more bandwidth between border routers to handle all the updates. Graceful restart helps reduce control traffic during temporary BGP disconnections by letting forwarding continue based on previous information while the control plane recovers, cutting the bandwidth impact of reconvergence and keeping networks stable during maintenance windows or equipment problems [8].

Setting up Option B in the real world starts with getting your MPLS foundation right. First, enable the proper features, set MTU values high enough for label stacks, and establish your internal label distribution using LDP or RSVP-TE. Next come VRF definitions – you'll specify route distinguishers to make customer addresses globally unique, plus import and export route targets that control which VPNs can talk to each other, along with address family settings for

IPv4 and maybe IPv6. The setup's core is MP-BGP configuration, where you'll turn on VPNv4 address families and label exchange for internal connections, as well as the link to other independent systems.

For the cross-domain connection, your border routers need specific eBGP settings: the send-label capability for exchanging MPLS labels and send-community extended commands to preserve route targets during exchange. Route filtering through prefix lists and route-maps controls which VPN routes cross between networks, typically filtering based on extended communities to enforce business agreements. Many deployments include redistribution rules that manage route flow between different protocols or between global and VRF tables. QoS settings ensure traffic gets appropriate treatment across network boundaries by mapping internal service markings to values agreed with neighboring providers. Don't forget operational tools like logging, SNMP monitoring for key interfaces and protocols, and possibly BFD for quick failure detection to speed up recovery during link or node outages [9].

Option B works in many different scenarios across provider and enterprise networks. The most common case involves service providers connecting their MPLS networks to extend VPN coverage across regions or countries, letting business customers maintain seamless connectivity globally. These providers typically implement full Option B with direct VPN-IPv4 route exchange between border routers, using detailed filtering to strictly control which VPNs can cross boundaries based on contracts. Corporate mergers create another common use case, where companies need to join previously separate MPLS networks without service disruptions. These projects often use a phased approach – starting with limited VPN connectivity for critical services, then gradually expanding as teams align their operational practices. Managed service providers create a third scenario when they sell white-labeled VPN services to a smaller regional provider, which can be seen as having a hierarchical model where the regional provider manages customer-facing services and the core provider oversees the routing architecture. These scenarios also demonstrate that option B provides flexibility to satisfy a wide variety of business requirements, from strict regulatory requirements in financial and healthcare fields to media distribution and real-time patient collaboration that support high-bandwidth and low-latency usage cases. Each deployment adapts the configuration details to match specific technical and business constraints [9].

Fixing problems in Option B setups requires understanding both control and data plane operations across network boundaries. The most common headaches involve route propagation failures – VPN routes don't show up in remote PE routers despite seemingly correct configuration. This usually happens because of route target mismatches between networks, BGP session setup errors preventing VPNv4 route exchange, or filters accidentally blocking specific routes based on prefix properties or extended communities. Start troubleshooting by checking the BGP session status between border routers, making sure neighbors exchange VPNv4 routing information with label capabilities turned on. Then, examine route target configurations for consistency between import/export policies across domains, and trace route propagation step-by-step from the source PE through local and remote border routers to the destination PE. Label problems form another common issue category, showing up as "label not available" errors or black-holed traffic, where packets disappear without any error notification. These typically stem from incompatible label allocation between providers, MPLS MTU mismatches preventing proper label forwarding, or incorrect TTL propagation settings causing packets to get dropped by devices along the path. If you are testing the data plane using a special-purpose tool (e.g., ping or traceroute), you can determine whether you have a label forwarding problem by testing the entire label-switch path on both networks and identifying where the connectivity is failing. If you are following a structured testing process, test physical connectivity first, followed by verifying control plane status (e.g., through BGP sessions), verifying route propagation, and then testing data plane delivery end-to-end [10].

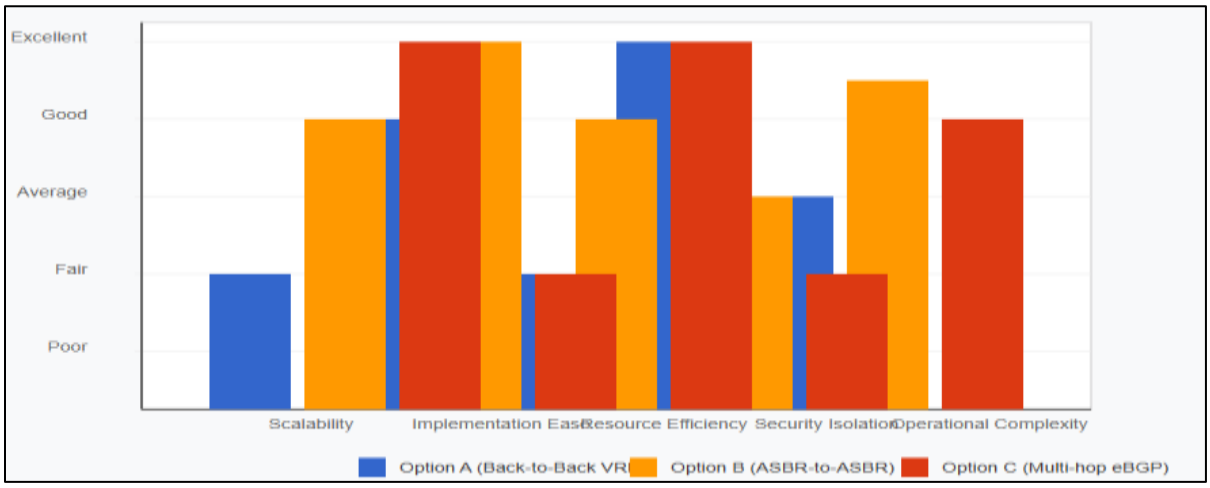

Fig. 3: Inter-AS Option B Performance Analysis. [9, 10]

CONCLUSION

Inter-AS Option B represents an optimal middle ground for service providers seeking to interconnect MPLS VPNs across autonomous system boundaries. By facilitating direct exchange of VPN-IPv4 routes between ASBRs while maintaining label-based forwarding separation, Option B delivers a balanced combination of scalability, security, and operational simplicity. The architecture effectively addresses key challenges of cross-domain VPN connectivity without the excessive interface requirements of Option A or the expanded trust boundaries of Option C. Implementation experience demonstrates that Option B provides sufficient isolation for most inter-provider scenarios while efficiently supporting numerous VPNs across a single interconnection point. As networks continue evolving toward greater automation and programmability, Option B's straightforward control plane model positions it well for integration with emerging software-defined networking architectures and intent-based systems. For service providers managing multi-domain environments, Option B remains the most practical interconnection strategy, particularly when combined with robust security practices, including comprehensive route filtering, proper MPLS TTL handling, and formalized inter-provider agreements.

REFERENCES

[1] Yakov Rekhter, Eric C. Rosen, "BGP/MPLS IP Virtual Private Networks (VPNs) RFC 4364," Data Tracker, 2006. [Online]. Available: https://datatracker.ietf.org/doc/rfc4364/
[2] Huawei"S1720, S2700, S5700, and S6720 V200R011C10 Configuration Guide - VPN," 2022. [Online]. Available: https://support.huawei.com/enterprise/en/doc/EDOC1000178179/ac916d6d/overview-of-bgp-mpls-ipv6-vpn
[3] Cisco, "BGP MPLS VPNs," 2019. [Online]. Available: https://www.cisco.com/c/en/us/td/docs/wireless/asr_5000/21-25/VPC-DI-Sys-Admin/21-25-vpc-di-sys-admin/21-17-VPC-DI-Sys-Admin_chapter_011001.pdf
[4] Cisco, "MPLS VPN Inter-Autonomous System Support," 2012. [Online]. Available: https://www.cisco.com/c/en/us/td/docs/ios-xml/ios/mp_ias_and_csc/configuration/15-sy/mp-vpn-inter-sys-sup.html
[5] E. Rosen et al., "BGP/MPLS IP Virtual Private Networks (VPNs)," Network Working Group, 2006. [Online]. Available: https://datatracker.ietf.org/doc/html/rfc4364
[6] Huawei, "ME60 V800R022C00SPC600 Feature Description," 2022. [Online]. Available: https://support.huawei.com/enterprise/en/doc/EDOC1100278569/e68f24ef/understanding-bgp-mpls-ip-vpn
[7] K. Kompella, Y. Rekhter, "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling," 2007. [Online]. Available: https://datatracker.ietf.org/doc/html/rfc4761
[8] Yakov Rekhter, Rahul Aggarwal, "Graceful Restart Mechanism for BGP with MPLS RFC 4781," Data Tracker, 2018. [Online]. Available: https://datatracker.ietf.org/doc/rfc4781/
[9] Noction, "BGP/MPLS Layer 3 VPNs Practical Configuration", 2018. [Online]. Available: https://www.noction.com/blog/bgp-mpls-layer3-vpn-practical-configuration
[10] Networking with FISH, "Troubleshooting Basic MPLS L3VPN – Part 1 – BGP," 2021. [Online]. Available: https://www.networkingwithfish.com/troubleshooting_basic_l3vpn/