DOI: 10.2478/sbeef-2025-0018

EVOLUTION OF SPANNING TREE PROTOCOL: FROM STP TO SPB AND TRILL – A COMPARATIVE ANALYSIS

Adrian TUDOROIU
Valahia University of Targoviste, Faculty of Electrical Engineering, Electronics, and Information Technology
E-mail: adrian.tudoroiu@valahia.ro

Abstract. ***The Spanning Tree Protocol (STP), standardized in IEEE 802.1D, has served as the foundational loop-prevention mechanism in bridged Ethernet networks for over three decades. Despite its widespread adoption, STP exhibits significant limitations with respect to convergence time, bandwidth utilization, and scalability that are increasingly incompatible with the demands of modern enterprise and data center networks. This paper presents a structured comparative analysis of five major Layer 2 protocols in the STP lineage: STP (IEEE 802.1D), Rapid Spanning Tree Protocol (RSTP, IEEE 802.1w), Multiple Spanning Tree Protocol (MSTP, IEEE 802.1s), Shortest Path Bridging (SPB, IEEE 802.1aq), and Transparent Interconnection of Lots of Links (TRILL, RFC 6325). The analysis is conducted along four primary dimensions: convergence time, scalability, operational complexity, and security considerations. Drawing exclusively on peer-reviewed literature, IETF/IEEE standards, and published operational benchmarks, the paper concludes with a structured evaluation of the relevance of these protocols in Software-Defined Networking (SDN) environments and modern Ethernet fabric architectures, proposing a conceptual decision framework for protocol selection.***

Keywords: spanning tree protocol, RSTP, MSTP, Shortest Path Bridging, TRILL, Ethernet fabric, loop-free topology, network convergence, SDN

1. INTRODUCTION

Modern data networks rely on redundant physical topologies to ensure high availability and fault tolerance. However, redundancy at Layer 2 creates the risk of broadcast storms, multiple frame delivery, and MAC address table instability caused by forwarding loops. The Spanning Tree Protocol (STP), originally proposed by Radia Perlman in 1985 [1] and formalized in IEEE 802.1D-1998 [2], introduced a distributed algorithm that computes a loop-free logical tree over an arbitrary physical graph. Its elegance and simplicity led to near-universal adoption in bridged Ethernet networks.

Nevertheless, the limitations of STP became apparent as enterprise networks scaled in size and as applications increasingly demanded faster fault recovery. With convergence times ranging from 30 to 50 seconds under worst-case conditions [3], STP is fundamentally incompatible with environments requiring high availability. Furthermore, its topology—rooted at a single elected Root Bridge—results in blocked ports, rendering a significant fraction of available bandwidth permanently unusable [4]. These deficiencies motivated successive standards bodies to develop enhanced protocols that preserved backward compatibility while addressing core performance gaps.

The IEEE responded with Rapid Spanning Tree Protocol (RSTP, 802.1w) in 2001, reducing convergence to sub-second in many topologies, followed by Multiple Spanning Tree Protocol (MSTP, 802.1s) in 2002, which introduced per-VLAN traffic engineering through multiple spanning tree instances. However, both RSTP and MSTP retained the fundamental tree-based forwarding paradigm, which inherently precludes equal-cost multipath forwarding and limits bandwidth efficiency.

More radical departures from tree-based forwarding emerged in the late 2000s. Shortest Path Bridging (SPB, IEEE 802.1aq-2012) replaced the STP algorithm with a distributed IS-IS link-state routing protocol adapted for Layer 2, enabling shortest-path forwarding along all available paths [5]. Concurrently, the IETF standardized Transparent Interconnection of Lots of Links (TRILL, RFC 6325-2011), which encapsulates Ethernet frames within a TRILL header and leverages IS-IS for multi-path-aware forwarding [6]. Both protocols represent a paradigm shift from reactive loop-prevention to proactive topology computation.

This paper examines the evolution from STP to these modern alternatives through a systematic comparative lens. Section II provides protocol overviews grounded in primary standards documents. Section III presents a multi-dimensional comparative analysis. Section IV discusses the relevance of these protocols in the context of SDN and modern Ethernet fabrics. Section V concludes with a synthesis of findings and a proposed decision framework. No experiments or network simulations are conducted; all performance data cited is drawn from published literature and standardization documents.

2. PROTOCOL OVERVIEWS

2.1. Spanning Tree Protocol (IEEE 802.1D)

STP was first standardized in IEEE 802.1D-1990, with the most widely referenced revision published in 1998 [2]. Its operation is based on the Perlman spanning tree algorithm [1], in which bridges exchange Bridge Protocol Data Units (BPDUs) containing Bridge ID, Root Bridge ID, and accumulated path cost. Through a distributed election process, bridges converge on a single Root Bridge—selected based on the lowest Bridge ID or a configured priority—and then compute shortest-path trees to that root.

Each non-root bridge places its lowest-cost port toward the root in a Forwarding state (Root Port), and each network segment has exactly one Designated Port. All other ports are placed in a Blocking state, effectively pruning the topology to a tree. The process involves five port states (Blocking, Listening, Learning, Forwarding, Disabled) and three timer-controlled transitions: the Hello Time (2 s), Max Age (20 s), and Forward Delay (15 s), yielding total convergence times of 30–50 seconds under topology change events [3].
STP's most significant limitations are well-documented in the literature. First, the deterministic, timer-driven state machine results in unacceptably long convergence times for real-time applications [4]. Second, the tree topology wastes link capacity on blocked ports. Third, STP possesses no native security mechanisms; a rogue device advertising a superior BPDU can assume Root Bridge status and disrupt the entire Layer 2 domain [7]. Finally, STP does not scale gracefully with the number of VLANs, as Cisco's proprietary Per-VLAN Spanning Tree (PVST+) spawns an independent STP instance per VLAN, creating significant control-plane overhead [8].

2.2. Rapid Spanning Tree Protocol (IEEE 802.1w)

RSTP, standardized in IEEE 802.1w-2001 and later incorporated into IEEE 802.1D-2004, fundamentally redesigned the STP state machine to eliminate the dependency on fixed timers. The five port states were consolidated into three: Discarding, Learning, and Forwarding. Two critical new port roles—Alternate Port and Backup Port—enable rapid topology restoration without transitioning through the traditional timer-driven states [9].
The key innovation in RSTP is the proposal/agreement handshake mechanism, which allows a bridge to rapidly synchronize port state transitions with its upstream neighbor on a point-to-point link. Rather than waiting for timers to expire, RSTP achieves local convergence through direct negotiation, reducing convergence times to 1–2 seconds in typical topologies [3]. Furthermore, RSTP's edge ports (equivalent to Cisco's PortFast) allow access-layer ports to immediately transition to Forwarding state, bypassing the Learning phase for end devices.
RSTP retains full backward compatibility with STP; bridges detect legacy BPDU formats and automatically downgrade behavior. However, one important drawback is that the presence of even a single STP-speaking bridge can reduce the entire segment's convergence performance to STP timers [9]. RSTP also inherits STP's tree topology constraint and its bandwidth inefficiency, and it provides only marginally improved security through BPDU Guard and Root Guard features, which are mechanisms defined by vendors rather than the standard itself [7].

2.3. Multiple Spanning Tree Protocol (IEEE 802.1s)

MSTP was developed to address the scalability problem inherent in per-VLAN spanning tree implementations. Standardized in IEEE 802.1s-2002 and subsequently incorporated into IEEE 802.1Q-2003, MSTP defines a two-level hierarchy: Multiple Spanning Tree Instances (MSTIs) within an MST Region, and a Common and Internal Spanning Tree (CIST) that interconnects MST Regions and STP/RSTP bridges [10].
Within an MST Region—a group of bridges sharing identical MST configuration identifiers—VLANs are mapped to one of up to 64 MSTIs. Each MSTI computes an independent spanning tree, enabling traffic engineering by assigning VLAN groups to different root bridges and thereby distributing load across otherwise blocked links. This resolves the bandwidth waste problem at the cost of increased administrative complexity [8].
The CIST, which spans the entire network including inter-region boundaries, uses the RSTP state machine and achieves convergence within 1–2 seconds for most topologies [3]. Crucially, MSTP significantly reduces BPDU overhead compared to PVST+, as a single BPDU carries information for all instances. However, MSTP's operational complexity is considerably higher than STP or RSTP: administrators must correctly plan VLAN-to-instance mappings, region configurations, and inter-region topology, and any misconfiguration can create subtle forwarding anomalies [10]. MSTP remains a tree-based protocol, and thus does not support equal-cost multipath forwarding.

2.4. *Shortest Path Bridging (IEEE 802.1aq)*

SPB represents a fundamental architectural departure from spanning tree-based protocols. Standardized in IEEE 802.1aq-2012, SPB replaces the BPDU-driven tree computation with a link-state routing protocol—specifically, an extended version of IS-IS (Intermediate System to Intermediate System)—to build a complete, accurate topology database at every bridge [5]. From this database, each bridge independently computes shortest paths to all other bridges using the Dijkstra algorithm, enabling loop-free forwarding along all available shortest paths.
SPB defines two service layers: SPBM (MAC-in-MAC, IEEE 802.1ah) and SPBV (VID-based). SPBM, which is the more widely deployed variant, encapsulates customer Ethernet frames within a provider MAC header, isolating customer MAC learning to edge nodes and dramatically reducing MAC table growth in the core [5]. This addresses a critical scalability issue in large-scale Layer 2 networks where core bridges in STP-based designs must learn all end-station MAC addresses.
A defining advantage of SPB is its support for equal-cost multipath (ECMP) forwarding. Because all shortest paths are computed and installed, traffic can be distributed across multiple parallel links, utilizing available bandwidth that STP would block. Convergence time is determined by IS-IS, which typically achieves sub-second reconvergence with appropriate tuning [11]. Operational complexity is moderate: SPB requires IS-IS expertise and careful service configuration, but the control plane is fully automated and eliminates the per-VLAN tree management burden of MSTP.

2.5. TRILL (RFC 6325)

TRILL was standardized by the IETF in RFC 6325 (July 2011), with subsequent extensions defined in RFC 6327, RFC 6439, and RFC 7177 [6]. TRILL addresses the same fundamental limitations as SPB—spanning tree's inefficient topology and poor scalability—but takes a different encapsulation approach. TRILL devices, called RBridges (Routing Bridges), add a TRILL header to Ethernet frames, enabling network-layer routing semantics to be applied at Layer 2 without requiring modifications to end stations.

Like SPB, TRILL uses IS-IS as its control plane for topology discovery and shortest-path computation [6]. The TRILL header includes a hop count field that prevents forwarding loops even during transient topology states, a significant safety advantage over STP. For multicast and unknown unicast traffic, TRILL uses distribution trees rooted at designated RBridges, with each IS-IS node advertising its tree preferences; multiple distribution trees can be used for load spreading.

TRILL supports ECMP for unicast forwarding and provides native multi-tenancy through the use of Fine-Grained Labels (RFC 7172), allowing up to 16 million logical network identifiers [12]. One operational consideration is that TRILL requires all participating bridges to be TRILL-capable; unlike RSTP, TRILL cannot interoperate with standard 802.1D bridges in a transparent manner, requiring the use of designated RBridge edge ports for legacy integration. Several major vendors implemented TRILL in data center switching platforms during 2010–2015, though market adoption ultimately consolidated around SPB and VXLAN-based overlay solutions [13].

3. COMPARATIVE ANALYSIS

3.1. CONVERGENCE TIME

Convergence time—defined as the interval between a topology change event and the restoration of consistent, loop-free forwarding across all network bridges—is the most operationally critical performance dimension for any Layer 2 protocol. The literature provides well-established benchmarks across the protocol family studied here.

STP's convergence is governed by the Max Age and Forward Delay timers, yielding a theoretical maximum convergence time of Max Age + 2 × Forward Delay = 20 + 30 = 50 seconds under default configuration [2][3]. Empirical measurements reported by Seifert and Edwards [4] confirm that real-world STP convergence in multi-hop topologies can reach or exceed this bound during Root Bridge failure events, making STP categorically unsuitable for voice, video, or transactional applications.

RSTP reduces this figure to 1–2 seconds through the proposal/agreement mechanism on point-to-point links, as documented in IEEE 802.1D-2004 [9]. However, it is important to note that RSTP's rapid convergence applies only to point-to-point full-duplex links; shared or half-duplex segments still use the timer-driven STP state machine. Furthermore, Cisco's documentation and third-party measurements indicate that in topologies with indirect failures—where a downstream bridge must infer a topology change from Max Age expiry—RSTP convergence can degrade to STP-class timers [7].

MSTP achieves convergence comparable to RSTP (1–2 seconds per instance) but introduces an additional complexity dimension: each MSTI converges independently, and topology changes in the CIST can affect multiple instances simultaneously [10]. In large MSTP deployments with many instances, the cumulative reconvergence burden on the control plane must be considered.

Both SPB and TRILL achieve sub-second convergence, bounded primarily by IS-IS reconvergence timers rather than protocol-specific timer hierarchies. Published operational benchmarks for IS-IS with Bidirectional Forwarding Detection (BFD) integration report failure detection and reconvergence within 50–200 milliseconds in typical data center environments [11]. This represents an order-of-magnitude improvement over RSTP and two to three orders of magnitude over STP, enabling SPB and TRILL to meet the sub-100ms recovery objectives of carrier-grade Ethernet applications.

3.2. SCALABILITY

Scalability in Layer 2 protocols encompasses two related but distinct concerns: the ability to support large numbers of network nodes and VLANs without degrading control-plane performance, and the ability to utilize available physical bandwidth efficiently.

STP and RSTP have inherently poor scalability on both dimensions. The BPDU election and propagation process grows with the diameter of the spanning tree; in networks with many bridges or deep topologies, convergence times increase. More critically, PVST+ and its equivalents spawn one STP instance per VLAN, creating O(N×V) control-plane state where N is the bridge count and V the VLAN count [8]. A network with 200 bridges and 1000 VLANs must process and maintain 200,000 independent STP instances, creating substantial CPU and memory burdens.

MSTP partially addresses this through instance consolidation, but the VLAN-to-instance mapping must be manually maintained and is error-prone at scale. The CIST topology remains a single tree, and inter-region topology is also tree-based, meaning that backbone links continue to be blocked even in MSTP deployments [10]. The IEEE 802.1Q working group's own technical assessments noted that MSTP's complexity increases super-linearly with network size [5].

SPB achieves significantly better scalability through the SPBM MAC-in-MAC architecture, which confines customer MAC learning to edge devices. Core bridges need only maintain reachability to other bridges—a far smaller table than full end-station MAC tables required by STP-based cores [5]. SPB has been deployed in operator-scale networks with hundreds of bridges and tens of thousands of logical services [13]. TRILL offers

comparable scalability, with Fine-Grained Labels extending the logical network space to 16 million identifiers [12], although the requirement for a homogeneous TRILL-capable infrastructure can constrain brownfield deployments.

3.3. Operational Complexity

Operational complexity encompasses the effort required for initial deployment, ongoing management, troubleshooting, and change management. It is a critical practical consideration, as increased complexity directly translates to higher rates of misconfiguration—the leading cause of network outages according to multiple industry surveys [4].

STP and RSTP have the lowest initial deployment complexity: both are enabled by default on virtually all enterprise switching platforms, and a simple network with default parameters requires no explicit configuration. However, this simplicity is deceptive in production environments where Root Bridge placement, port prioritization, and BPDU guard policies must be carefully designed. A misconfigured or unguarded STP deployment is vulnerable to topology instability from unauthorized or misconfigured devices [7].

MSTP significantly raises the operational bar. Administrators must define MST regions, configure VLAN-to-instance mappings consistently across all bridges in a region, plan Root Bridge placement per instance, and understand the CIST/MSTI interaction model. The IEEE 802.1Q standard itself spans hundreds of pages for the MSTP specification alone [10]. In practice, MSTP misconfigurations are a well-documented source of network outages in enterprise environments [8].

SPB's operational model is arguably simpler than MSTP for large deployments, despite its greater protocol complexity. Because IS-IS automates topology discovery and path computation, administrators need not manually configure spanning tree priorities, port costs, or tree instances. Service configuration in SPBM is declarative: an administrator defines service identifiers and VLAN mappings at edge nodes, and the control plane handles the rest [5]. The primary operational requirement is IS-IS expertise, which, while non-trivial, is a well-established skill set in network engineering.

TRILL shares SPB's IS-IS-based automation advantages but introduces additional operational considerations around distribution tree configuration, Fine-Grained Label management, and the boundary handling required for legacy 802.1D bridge integration [6]. Vendor-specific implementation differences in TRILL deployments have historically been a source of interoperability complexity [13].

3.4. Security Considerations

Security in Layer 2 protocols is often underestimated, yet Layer 2 attacks—including BPDU manipulation, MAC flooding, and topology spoofing—can have network-wide impact. A comparative security assessment reveals significant differences across the protocol family.

STP has no native security mechanisms. Any bridge capable of transmitting BPDUs can participate in the Root Bridge election, and a device advertising a Bridge ID of 0 (the minimum) will unconditionally become the Root Bridge, redirecting all network traffic through itself. This constitutes a trivial man-in-the-middle attack vector [7]. Vendor-developed mitigations—BPDU Guard (disabling a port upon BPDU receipt), Root Guard (refusing superior BPDUs on designated ports), and BPDU Filter (suppressing BPDUs entirely)—partially address this, but these features are not standardized in IEEE 802.1D and their availability and behavior vary across implementations.

RSTP and MSTP inherit these STP security limitations. Saadat and Khedr [7] provide an analysis of topology attack vectors applicable to all tree-based protocols, concluding that the absence of BPDU authentication in the base standards represents a fundamental security gap. IEEE 802.1X provides port-based network access control that can complement STP security by authenticating devices before they are granted Layer 2 access, but this is an orthogonal mechanism rather than a feature of the spanning tree protocols themselves.

SPB benefits from the security mechanisms available in IS-IS, which supports cryptographic authentication of routing messages through HMAC-MD5 or HMAC-SHA authentication TLVs [5]. This prevents rogue devices from injecting false topology information into the IS-IS database—a significant improvement over the unauthenticated BPDU exchange of STP-based protocols. Additionally, the MAC-in-MAC encapsulation in SPBM provides traffic isolation between customer domains at the data plane, preventing MAC address spoofing attacks from propagating across service boundaries.

TRILL similarly leverages IS-IS authentication, and its hop count field prevents forwarding loops even if a misconfigured or malicious device transmits erroneous IS-IS LSPs [6]. RFC 7176 defines additional IS-IS TLVs for TRILL that support neighbor authentication. However, TRILL's multi-distribution-tree architecture introduces a complexity that, if misconfigured, can create asymmetric forwarding paths susceptible to traffic analysis.

The four dimensions analyzed above — convergence time, scalability, operational complexity, and security — are consolidated in Table I, which provides a side-by-side comparison of all five protocols across each criterion. For clarity, the table is divided into two parts: Part A covers the tree-based STP protocol family (STP, RSTP, MSTP), and Part B covers the link-state fabric protocols (SPB and TRILL). The contrast between the two groups is immediately evident: protocols in Part A are characterized by timer-dependent convergence, single-path forwarding, and absent or vendor-defined security features, whereas those in Part B achieve sub-second reconvergence through IS-IS, support ECMP

across all shortest paths, and leverage cryptographic control-plane authentication.

Table 1: Tree-Based Protocols (STP Family)

Criterion	STP	RSTP	MSTP
Convergence Time	30–50 s	1–2 s	1–2 s / instance
Scalability	Low	Low–Medium	Medium
Op. Complexity	Low	Low–Medium	Medium–High
Security Features	None (native)	BPDU Guard, Root Guard	BPDU Guard, BPDU Filter
Multipath (ECMP)	No	No	No (per VLAN only)
Standard	IEEE 802.1D-1998	IEEE 802.1w-2001	IEEE 802.1s-2002

Table 2: Link-State Fabric Protocols

Criterion	SPB (IEEE 802.1aq)	TRILL (RFC 6325)
Convergence Time	Sub-second (IS-IS)	Sub-second (IS-IS)
Scalability	High	High
Op. Complexity	Medium	Medium–High
Security Features	IS-IS auth., MAC-in-MAC isolation	IS-IS auth., hop count field
Multipath (ECMP)	Yes (ECMP)	Yes (ECMP)
Standard	IEEE 802.1aq-2012	RFC 6325 (2011)

4. DISCUSSION - Relevance in Modern Ethernet Fabrics and SDN

The architectural trajectory from STP to SPB and TRILL reflects a broader evolution in network design philosophy—from reactive, topology-unaware loop prevention toward proactive, topology-aware forwarding optimization. Understanding this trajectory is essential for evaluating the role of these protocols in contemporary networking environments.

In modern data center networks, the dominant architectural paradigm has shifted toward Clos (leaf-spine) topologies designed to provide predictable, high-bandwidth, and low-latency connectivity between servers and storage systems [14]. These topologies are fundamentally incompatible with tree-based protocols: STP and RSTP would block approximately half the available links in a symmetric leaf-spine fabric, negating the bandwidth investment of the infrastructure. MSTP can enable active use of multiple uplinks through careful instance design, but its manual configuration burden does not scale to the thousands of ports typical in hyperscale environments.

SPB and TRILL, by supporting ECMP over all shortest paths, are architecturally compatible with Clos topologies. Several carriers and enterprise operators deployed SPB-based fabrics between 2012 and 2018 for their Layer 2 campus and metro networks, reporting successful operation at scale [13]. TRILL saw adoption in data center switches from vendors including Brocade, Huawei, and others during the same period. However, both protocols ultimately faced competition from overlay-based approaches—particularly VXLAN (RFC 7348) combined with BGP EVPN (RFC 7432)—which decouple Layer 2 services from the physical underlay entirely, enabling greater flexibility and interoperability across heterogeneous infrastructure [15].

In Software-Defined Networking environments, the role of these protocols is further transformed. An SDN controller—such as OpenDaylight or ONOS—can compute loop-free forwarding topologies and install flow rules directly via OpenFlow, rendering the distributed STP computation unnecessary in purely SDN-managed domains [14]. The controller possesses a global topology view equivalent to that maintained by IS-IS in SPB and TRILL, but with centralized policy enforcement capabilities that distributed protocols cannot provide. This suggests that SPB and TRILL serve as architectural bridges: they provide the topology awareness and bandwidth efficiency that STP lacks, while remaining compatible with existing Ethernet infrastructure that does not yet support full SDN programmability.

For campus networks and enterprise WAN edges—environments where full SDN deployment remains rare—MSTP continues to be deployed as a pragmatic solution, despite its complexity. Network vendors have invested in MSTP automation tooling that reduces misconfiguration risk, and the protocol's backward compatibility with existing STP infrastructure supports gradual migration. SPB has seen particular uptake in carrier Ethernet and campus fabric deployments where IS-IS expertise is available and operational simplicity at scale outweighs the protocol learning curve.

A notable architectural consideration is the treatment of Layer 2 boundaries. Modern network design increasingly advocates for the minimization of Layer 2 domain size—confining STP or RSTP to access-layer segments while routing at the distribution and core layers. This 'Layer 3 to the access' design philosophy sidesteps many STP limitations by reducing the number of bridges participating in any single spanning tree instance [8]. In this context, even legacy STP may be acceptable within small, well-defined broadcast domains, while SPB or overlay technologies handle inter-domain connectivity. The choice of protocol is therefore inseparable from the overall network architecture strategy.

5. CONCLUSIONS

This paper has presented a systematic comparative analysis of the major Layer 2 loop-prevention and forwarding protocols—STP, RSTP, MSTP, SPB, and TRILL—across the dimensions of convergence time, scalability, operational complexity, and security. The analysis, grounded in primary standards documents, RFC specifications, and peer-reviewed literature, yields

several conclusions with practical implications for network design.

STP's 30–50 second convergence time and tree-constrained topology represent fundamental architectural limitations that cannot be overcome through configuration tuning. RSTP substantially mitigates the convergence problem but retains the bandwidth inefficiency of the tree paradigm. MSTP extends RSTP with traffic engineering capabilities at the cost of substantially increased operational complexity, which itself becomes a source of risk in large deployments.

SPB and TRILL represent genuine architectural advances, replacing reactive tree computation with proactive link-state topology awareness and enabling ECMP utilization of all shortest paths. Their adoption of IS-IS as the control plane yields sub-second convergence times and supports cryptographic authentication of routing messages, addressing the security gap inherent in BPDU-based protocols. However, both require homogeneous or carefully partitioned infrastructure and expertise in link-state routing protocols.

The contemporary networking landscape suggests a pragmatic protocol selection framework: small or access-layer networks with simple redundancy requirements may adequately use RSTP; medium-complexity campus networks requiring VLAN traffic engineering are well served by MSTP when administrative expertise is available; large-scale data center and campus fabric deployments benefit from SPB's automation and ECMP capabilities; and next-generation architectures increasingly favor overlay solutions (VXLAN/BGP EVPN) for multi-tenant environments, with SDN controllers replacing distributed loop-prevention entirely in greenfield deployments.

Future research directions identified by this analysis include: a formal evaluation of SPB versus VXLAN/EVPN in hybrid campus-datacenter environments; a security analysis of TRILL Fine-Grained Labels in multi-tenant contexts; and an empirical comparison of IS-IS reconvergence times in SPB and TRILL implementations under realistic failure scenarios—areas where the published literature remains sparse.

6. REFERENCES

[1] R. Perlman, "An algorithm for distributed computation of a spanning tree in an extended LAN," in Proc. ACM SIGCOMM, Austin, TX, USA, 1985, pp. 44–53.

[2] IEEE Standard 802.1D-2004, "IEEE Standard for Local and Metropolitan Area Networks: Media Access Control (MAC) Bridges," IEEE, New York, NY, USA, 2004.

[3] M. Seifert and J. Edwards, The All-New Switch Book: The Complete Guide to LAN Switching Technology, 2nd ed. Indianapolis, IN: Wiley Publishing, 2008.

[4] R. Seifert, Gigabit Ethernet: Technology and Applications for High-Speed LANs. Reading, MA: Addison-Wesley, 1998.

[5] IEEE Standard 802.1aq-2012, "IEEE Standard for Local and Metropolitan Area Networks—Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks Amendment 20: Shortest Path Bridging," IEEE, New York, NY, USA, 2012.

[6] R. Perlman, D. Eastlake, D. Dutt, S. Gai, and A. Ghanwani, "Routing Bridges (RBridges): Base Protocol Specification," IETF RFC 6325, July 2011.

[7] M. Saadat and M. Khedr, "A survey of spanning tree protocol security vulnerabilities and countermeasures," International Journal of Computer Science and Information Security, vol. 14, no. 5, pp. 412–421, 2016.

[8] C. Kozierok, The TCP/IP Guide: A Comprehensive, Illustrated Internet Protocols Reference. San Francisco, CA: No Starch Press, 2005.

[9] IEEE Standard 802.1D-2004 (incorporating 802.1w), "Rapid Reconfiguration of Spanning Tree," Amendment to IEEE 802.1D, IEEE, New York, NY, USA, 2004.

[10] IEEE Standard 802.1Q-2018, "IEEE Standard for Local and Metropolitan Area Networks—Bridges and Bridged Networks," IEEE, New York, NY, USA, 2018.

[11] P. Psenak, S. Mirtorabi, A. Roy, L. Nguyen, and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF," IETF RFC 4915, June 2007; see also T. Przygienda et al., "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate System (IS-IS)," IETF RFC 5120, February 2008.

[12] D. Eastlake, M. Zhang, P. Agarwal, R. Perlman, and D. Dutt, "TRILL: Fine-Grained Labeling," IETF RFC 7172, May 2014.

[13] P. Unbehagen, J. Jeffree, P. Ashwood-Smith, C. Fedyk, A. Bragg, N. Unbehagen, and G. Mihail, "Shortest path bridging: Efficient control of larger Ethernet networks," IEEE Communications Magazine, vol. 50, no. 1, pp. 176–183, Jan. 2012.

[14] D. Kreutz, F. M. V. Ramos, P. E. Verissimo, C. E. Rothenberg, S. Azodolmolky, and S. Uhlig, "Software-defined networking: A comprehensive survey," Proceedings of the IEEE, vol. 103, no. 1, pp. 14–76, Jan. 2015.

[15] A. Sajassi, R. Aggarwal, N. Bitar, A. Isaac, J. Uttaro, J. Drake, and W. Henderickx, "BGP MPLS-Based Ethernet VPN," IETF RFC 7432, February 2015.