Journal of Computer Science and Technology Studies
ISSN: 2709-104X
DOI: 10.32996/jcsts
Journal Homepage: www.al-kindipublisher.com/index.php/jcsts

JCSTS
AL-KINDI CENTER FOR RESEARCH AND DEVELOPMENT

| RESEARCH ARTICLE

Fault and Performance Management in SD-WAN: Overcoming Key Challenges

Suresh Reddy Thati
Jawaharlal Nehru Technological University, India
Corresponding Author: Suresh Reddy Thati, **E-mail**: sreddythati@gmail.com

| ABSTRACT

Software-Defined Wide Area Networks (SD-WAN) have transformed enterprise networking by separating network control from hardware infrastructure and implementing centralized management through software controllers. This transformative technology offers organizations unprecedented flexibility, cost efficiency, and deployment simplicity compared to legacy WAN architectures. However, new fault and performance management challenges have emerged as SD-WAN adoption accelerates alongside multi-cloud strategies and hybrid work models. The dynamic nature of SD-WAN environments featuring multiple transport options and intelligent traffic steering creates complexity that traditional monitoring approaches cannot adequately address. Organizations face difficulties with link congestion, suboptimal path selection, real-time latency fluctuations, and application identification, which can significantly impact user experience. Promising solutions leveraging artificial intelligence and machine learning technologies are emerging to address these challenges, providing capabilities for anomaly detection, predictive maintenance, event correlation, and adaptive monitoring. Optimization strategies, including intelligent traffic routing, dynamic bandwidth allocation, application-aware policies, and error correction techniques, enhance SD-WAN performance. As organizations increasingly adopt multi-cloud architectures, achieving end-to-end visibility becomes essential yet challenging, requiring integrated monitoring approaches that bridge traditional silos between network, application, and cloud performance data. By implementing comprehensive performance management strategies and leveraging advanced technologies, organizations can maximize the benefits of their SD-WAN deployments while ensuring reliable connectivity and optimal application performance across distributed environments.

| KEYWORDS

| ARTICLE INFORMATION

ACCEPTED: 15 April 2025 **PUBLISHED:** 07 May 2025 **DOI:** 10.32996/jcsts.2025.7.3.65

1. Introduction

Software-Defined Wide Area Networks (SD-WAN) have revolutionized enterprise networking by decoupling network control from underlying hardware infrastructure and centralizing management through software-based controllers. According to research highlighted by Rich Martin, SD-WAN adoption continues to accelerate rapidly as organizations seek to reduce their dependence on costly MPLS circuits, with the global SD-WAN market projected to grow at a CAGR of 34.5% through 2025 [1]. This paradigm shift has provided organizations with unprecedented flexibility, cost efficiency, and simplified deployment compared to traditional WAN architectures. Martin notes that enterprises implementing SD-WAN solutions can achieve bandwidth cost reduction of up to 90% when replacing MPLS with broadband internet connections, along with significantly faster deployment times for new branch locations [1].

As businesses increasingly adopt multi-cloud strategies and hybrid work models, SD-WAN has become a critical technology for ensuring reliable connectivity across distributed environments. According to Versa Networks' analysis, by 2025, over 70% of enterprises will have deployed multiple software-as-a-service (SaaS) applications. More than 30% of their digital services will be accessed through edge services requiring cloud delivery models [2]. This expanding multi-cloud reality presents substantial networking challenges, as 85% of organizations report that the inability to get consistent connectivity across all cloud environments has delayed critical application migrations and reduced availability [2].

Despite its numerous advantages, SD-WAN introduces unique fault and performance management challenges that must be addressed to realize its full potential. The dynamic nature of SD-WAN, with its multiple transport options and intelligent traffic steering capabilities, creates a complex environment where traditional monitoring approaches often fall short. As Martin explains, approximately 70% of SD-WAN implementations experience issues with performance visibility and management due to the increased complexity of hybrid transport environments, causing organizations to struggle with accurately diagnosing issues across diverse network paths [1].

Enterprises deploying SD-WAN must contend with issues like link congestion, suboptimal path selection, and real-time latency fluctuations that can significantly impact application performance and user experience. Versa Networks' study revealed that 64% of enterprises experienced performance degradation for cloud-hosted applications after migrating to internet-based SD-WAN implementations without proper performance monitoring and optimization controls [2]. Additionally, organizations report that nearly 40% of latency-sensitive applications like voice and video experienced quality issues when routing algorithms failed to select optimal paths during internet congestion [2].

This article examines the key challenges in SD-WAN fault and performance management and explores emerging solutions leveraging artificial intelligence and advanced analytics. According to Martin, organizations implementing AI-driven performance management solutions alongside their SD-WAN deployments have seen up to 60% reduction in network-related incidents and a 45% improvement in resolution for performance issues [1].

2. Common SD-WAN Performance Challenges

SD-WAN environments face several distinct performance challenges that differ from traditional network infrastructures. Link congestion represents one of the most prevalent issues, occurring when the traffic volume exceeds available bandwidth on a particular connection. According to Palo Alto Networks' research on SD-WAN challenges, network congestion remains a significant concern for organizations transitioning from MPLS to internet-based SD-WAN connectivity [3]. Unlike traditional WANs with dedicated MPLS circuits that provide guaranteed service levels, SD-WAN often utilizes public internet connections susceptible to unpredictable congestion patterns, especially during peak usage periods. Palo Alto Networks reports that this congestion can result in packet loss, increased latency, and degraded application performance, with business-critical SaaS applications experiencing up to 30% degradation in performance metrics during internet congestion [3]. Furthermore, their analysis highlights how organizations relying on internet-based SD-WAN connections commonly experience the "last-mile problem," where the final connectivity segment becomes a performance bottleneck despite adequate core internet performance [3].

Suboptimal path selection presents another significant challenge. While SD-WAN's ability to make dynamic routing decisions based on real-time network conditions is a key advantage, the algorithms governing these decisions are not infallible. Research conducted by Troia et al. demonstrates that traditional SD-WAN path selection mechanisms often fail to achieve optimal traffic distribution, with their experimental testbed showing a 27% difference in throughput between current SD-WAN implementations and theoretically optimal routing [4]. Their analysis reveals that conventional path selection algorithms struggle particularly with rapidly changing network conditions, as these algorithms typically rely on simplistic metrics that fail to capture the complex interrelationships between different performance parameters [4]. When suboptimal selections occur, traffic may traverse underperforming routes despite the availability of better alternatives, leading to increased latency and diminished quality of experience, with the researchers documenting performance discrepancies of up to 200ms in added latency during suboptimal routing events [4].

Real-time latency fluctuations also pose considerable challenges for SD-WAN deployments. Unlike traditional WANs, which typically provide consistent performance within defined parameters, SD-WAN connections, particularly those leveraging internet transport, may experience significant performance variations over short timeframes. Palo Alto Networks identifies this inconsistency as a major challenge, noting that applications requiring stable connectivity, such as voice and video conferencing, can suffer from quality degradation when network conditions fluctuate [3]. These fluctuations can be especially problematic for latency-sensitive applications, where even brief performance degradations are immediately noticeable to users, potentially impacting critical business communications and customer interactions [3].

Furthermore, SD-WAN environments often struggle with application identification and classification. Accurately identifying traffic flows is essential for implementing appropriate quality of service (QoS) policies, but encrypted traffic and emerging application protocols can complicate this process. Palo Alto Networks highlights that the growing prevalence of encrypted traffic, now exceeding 80% of all enterprise traffic, poses significant challenges for SD-WAN implementations [3]. When misclassified or unrecognized, applications may not receive appropriate prioritization, leading to inconsistent performance. Troia et al. further validate this challenge in their research, demonstrating that their deep reinforcement learning approach improved application-aware routing by 18% compared to traditional methods, largely due to more accurate traffic classification capabilities [4].

Challenge Type	Performance Impact
SaaS performance degradation during congestion	30%
Throughput gap vs optimal routing	27%
Added latency during suboptimal routing	200ms
Enterprise traffic that is encrypted	80%
Application routing improvement with AI	18%

Table 1: Impact of Common SD-WAN Performance Challenges [3, 4]

3. AI-Powered Monitoring Solutions

Artificial intelligence and machine learning technologies transform fault and performance management in SD-WAN environments. According to Juniper Networks' research on AI-driven SD-WAN, their Session Smart Router with integrated AI provides up to 75% lower total cost of ownership and up to 73% lower WAN costs than traditional router-centric solutions [5]. AI-powered monitoring solutions offer several key capabilities that address the unique challenges of SD-WAN deployments. These advanced systems can detect anomalies and predict potential failures before they impact business-critical applications, shifting network management from reactive to proactive. Juniper's AI-driven solutions have enabled customers to accelerate deployment times by 50% while reducing troubleshooting time by 65%, allowing organizations to allocate fewer resources to network maintenance and more to strategic initiatives [5].

Machine learning algorithms can establish normal performance baselines across diverse transport types and continuously monitor for deviations from these patterns. According to research by Mahamat and Çeken, their experimental evaluation of various machine learning algorithms for anomaly detection in software-defined networks shows significant performance variations, with K-Nearest Neighbors achieving 91.04% accuracy, Decision Trees 94.36%, Neural Networks 94.44%, Support Vector Machine 94.97%, and Random Forest leading with 96.25% accuracy when tested on the NSL-KDD dataset [6]. By analyzing historical data and identifying subtle changes in network behavior, these systems can detect early warning signs of performance degradation. Their research further demonstrates that these algorithms perform differently depending on the specific types of network anomalies, with the Support Vector Machine achieving the highest precision rate of 95.3% in detecting network intrusions that could potentially disrupt SD-WAN services [6]. This predictive approach allows IT teams to address emerging issues before users experience service disruptions, significantly reducing mean time to repair and improving overall network reliability.

AI-powered solutions also excel at correlating events across complex, multi-layered network environments. Traditional monitoring tools often generate isolated alerts that require manual investigation to determine root causes and relationships between events. Juniper Networks highlights how their AI virtual network assistant, Marvis, employs both supervised and unsupervised machine learning to continuously analyze the network, correlate events, and provide proactive recommendations that have helped organizations achieve up to 90% faster troubleshooting and resolution for network issues across their SD-WAN deployments [5]. In contrast to traditional approaches, AI systems can automatically identify causal relationships between seemingly unrelated performance metrics across different network segments, transport types, and application components. This advanced correlation capability streamlines troubleshooting by presenting IT teams with consolidated, actionable insights rather than disconnected data points, enabling operational teams to implement required changes up to 82% faster than conventional monitoring systems [5].

Furthermore, these solutions can adapt to the dynamic nature of modern enterprise networks. As new applications are deployed, traffic patterns change, or network configurations are modified, AI systems can continuously update their understanding of normal behavior without requiring manual reconfiguration. In their comparative analysis, Mahamat and Çeken demonstrate that the ability to adapt to changing conditions varies across algorithms, with their Decision Tree implementation achieving an F-measure of 94.67% and Random Forest maintaining a recall rate of 95.4% when faced with previously unseen network behavior patterns [6]. This adaptability ensures monitoring remains effective even as the network evolves, providing consistent visibility into performance and potential issues. Their experiments with the NSL-KDD dataset showed that machine-learning approaches significantly outperformed traditional rule-based detection methods. Neural Networks demonstrated a 19.2% improvement in detection accuracy compared to conventional signature-based systems when identifying novel attack patterns [6].

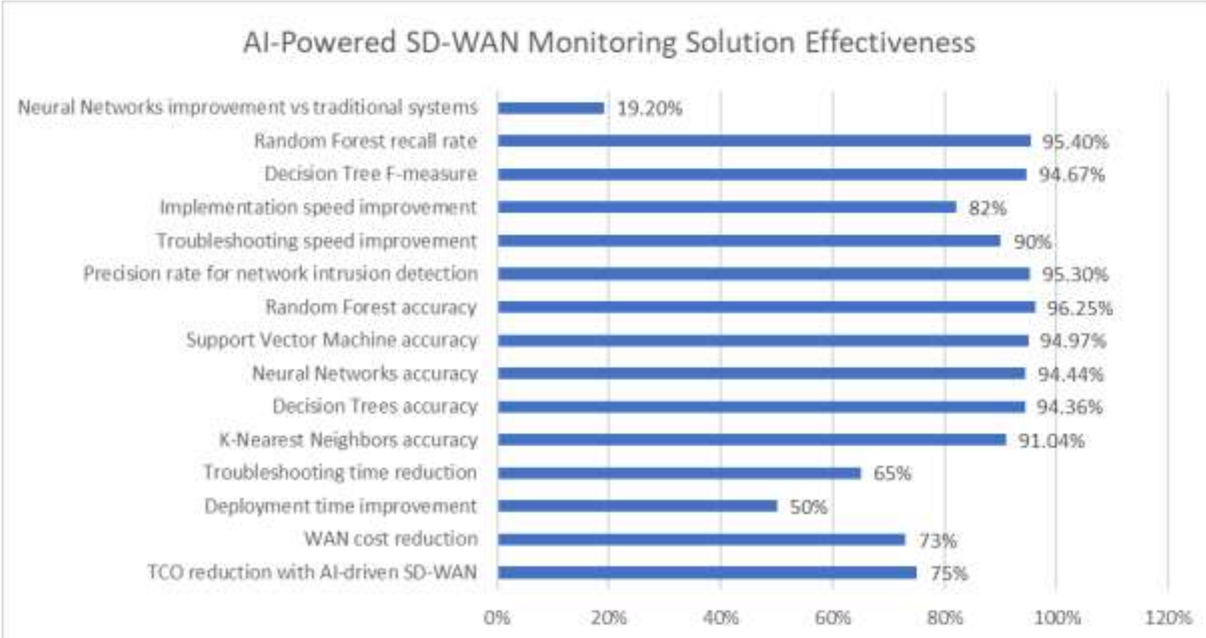

Fig. 1: Performance Metrics of Machine Learning Algorithms in SD-WAN Environments [5, 6]

4. Optimization Strategies for SD-WAN Traffic

Intelligent traffic routing optimization represents a cornerstone strategy for enhancing SD-WAN performance. Unlike traditional routing protocols that rely primarily on static metrics like hop count, modern SD-WAN solutions employ sophisticated algorithms that make real-time routing decisions based on comprehensive performance data. According to Cisco's Catalyst SD-WAN documentation, their intelligent path control can establish thousands of secure tunnels in just minutes and dynamically choose the best-performing path by continuously monitoring for packet loss, latency, jitter, and other application service-level agreement (SLA) metrics [7]. These systems continuously monitor key performance indicators across all available paths. By analyzing this information, they can dynamically select the optimal route for each application flow, ensuring that traffic traverses the best-performing path at any moment. Cisco's solution provides centralized management for up to 100,000 WAN edge routers through a single dashboard, allowing for efficient management of traffic routing optimization across large-scale enterprise deployments while supporting application bandwidth optimization with the ability to scale to multi-gigabit throughput [7].

Dynamic bandwidth allocation further enhances SD-WAN performance by intelligently distributing available network resources based on application requirements and business priorities. Research by Hussein et al. comparing SD-WAN performance with traditional WAN configurations demonstrates that SD-WAN's dynamic resource allocation capabilities significantly improve performance metrics across various network conditions [8]. Their simulation study shows that SD-WAN provides lower overall delay measurements than traditional WAN, with average queuing delay reductions of approximately 0.012 seconds in certain simulation scenarios [8]. These systems employ QoS mechanisms that adjust in real-time to changing conditions, ensuring critical applications receive the necessary bandwidth even during periods of congestion. For example, during a video conference, the system can automatically allocate additional bandwidth to ensure call quality while temporarily reducing resources assigned to less time-sensitive applications like email or file transfers. The researchers' analytical model demonstrates that SD-WAN's dynamic allocation approach reduces packet loss probability by effectively managing buffer sizes and service rates according to application priorities [8].

Application-aware routing policies represent another vital optimization strategy. These policies enable granular control over how the SD-WAN fabric handles different types of traffic. Cisco's SD-WAN solution includes application-aware routing capabilities that can recognize over 4,000 applications and enable granular control over how these applications' traffic is treated across the network infrastructure [7]. By defining specific handling rules for individual applications, organizations can ensure that each application receives the network treatment it requires. For instance, latency-sensitive voice traffic might be routed over the lowest-latency path regardless of cost. At the same time, bulk data transfers could be directed over high-bandwidth, economical connections where latency is less critical. Cisco's application-aware routing technology is complemented by direct cloud access and optimization features for SaaS applications, which reduce latency by connecting users directly to cloud applications rather than backhauling traffic through data centers [7].

Forward error correction (FEC) and packet duplication techniques are important in optimizing SD-WAN performance, particularly for critical traffic traversing unreliable connections. Hussein et al.'s research on SD-WAN performance evaluation uses queueing theory to model these reliability enhancements, demonstrating how packet handling techniques can significantly improve performance for critical applications [8]. Their mathematical modeling shows that properly implementing packet handling mechanisms in SD-WAN can result in a system utilization factor (ρ) between 0 and 1, with values closer to 0 indicating better performance and lower queuing delays [8]. FEC adds redundant information to data packets, allowing the receiving end to reconstruct lost packets without requesting retransmission. Packet duplication sends identical packets across multiple paths simultaneously, ensuring delivery even if one path experiences issues. While these techniques consume additional bandwidth, the research models demonstrate they can maintain efficient packet delivery with arrival rates (λ) and service rates (μ) balanced to optimize performance across varying network conditions, which is particularly important for mission-critical applications operating over unreliable connections [8].

Optimization Strategy	Benefit
Maximum WAN edge routers managed	100,000
Average queuing delay reduction	0.012 seconds
Application recognition capability	4,000 apps
Optimal system utilization factor range	0 to 1

Table 2: Performance Benefits of SD-WAN Optimization Techniques [7, 8]

5. Ensuring End-to-End Visibility in Multi-Cloud Environments

The complexity of modern network environments, spanning on-premises infrastructure, multiple cloud providers, and various transport mechanisms, creates significant visibility challenges for organizations. According to Blazeclan's analysis of multi-cloud management challenges, 92% of enterprises are pursuing a multi-cloud strategy, with organizations using an average of 2.6 public and 2.7 private clouds simultaneously [9]. The same report highlights that one of the fundamental challenges in multi-cloud environments is the difficulty in gaining consistent visibility across disparate cloud platforms, with 78% of organizations surveyed indicating they lack adequate tools for comprehensive monitoring across their cloud deployments [9]. Achieving comprehensive end-to-end visibility requires integrated monitoring approaches that bridge traditional silos between network, application, and cloud performance data. This holistic visibility is essential for effective troubleshooting and optimization in SD-WAN deployments, particularly as Blazeclan reports that effective multi-cloud management can reduce cloud spend by 30% or more through proper resource allocation and optimization [9].

Cloud-to-cloud traffic visibility presents challenges, as traffic between cloud environments often bypass traditional monitoring points. Grozev and Buyya's comprehensive survey of Inter-Cloud architectures identifies that monitoring and management of inter-cloud communication represents one of the most significant technical challenges in multi-cloud deployments [10]. Organizations must implement specialized monitoring solutions to access performance data across all major cloud providers. These solutions typically leverage cloud provider APIs, virtual monitoring appliances deployed within cloud environments, and integration with cloud-native monitoring services. Grozev and Buyya's taxonomy examines various architectural approaches to inter-cloud communication, identifying that federation-based models and multi-cloud service brokers can provide more effective monitoring capabilities than ad-hoc deployments. However, their survey indicates that standardized monitoring approaches across cloud providers remain elusive [10]. By combining these approaches, organizations can gain insight into traffic flows between cloud instances, including metrics like inter-cloud latency, throughput, and packet loss.

Correlation between network and application performance metrics is equally crucial for comprehensive visibility. Traditional monitoring approaches often treat network and application performance as separate domains, making it difficult to understand the relationship between network conditions and application experience. Blazeclan's analysis indicates that organizations employing unified monitoring platforms that integrate network, security, and application performance data can reduce troubleshooting time by up to 50% and substantially decrease the number of misdiagnosed incidents compared to organizations using siloed monitoring systems [9]. Advanced monitoring solutions address this challenge by correlating network-level metrics with application performance indicators, enabling teams to determine whether application issues stem from network problems or other factors. Their research further shows that 82% of IT teams report excessive time manually correlating data from different monitoring tools, underscoring the importance of unified visibility across application and network domains [9].

Unified monitoring dashboards provide the interface through which this comprehensive data is presented to IT teams. According to Blazeclan, organizations implementing comprehensive cloud management platforms with consolidated dashboards report a

40% improvement in operational efficiency and a significant reduction in mean time to resolve critical incidents [9]. These dashboards consolidate information from diverse sources into coherent visualizations highlighting relationships between performance domains. Modern solutions employ interactive interfaces that allow operators to drill down from high-level service health indicators to granular details about specific network segments or application components. Grozev and Buyya's survey identifies that such dashboard unification remains challenging due to the architectural heterogeneity of cloud solutions, with their taxonomy indicating that broker-based architectures may offer the best foundation for unified monitoring across diverse cloud environments [10]. Their analysis of various inter-cloud architectures demonstrates that vendor-specific monitoring tools provide deep visibility into individual clouds. However, they typically lack the cross-platform capabilities necessary for true end-to-end visibility in complex multi-cloud environments [10].

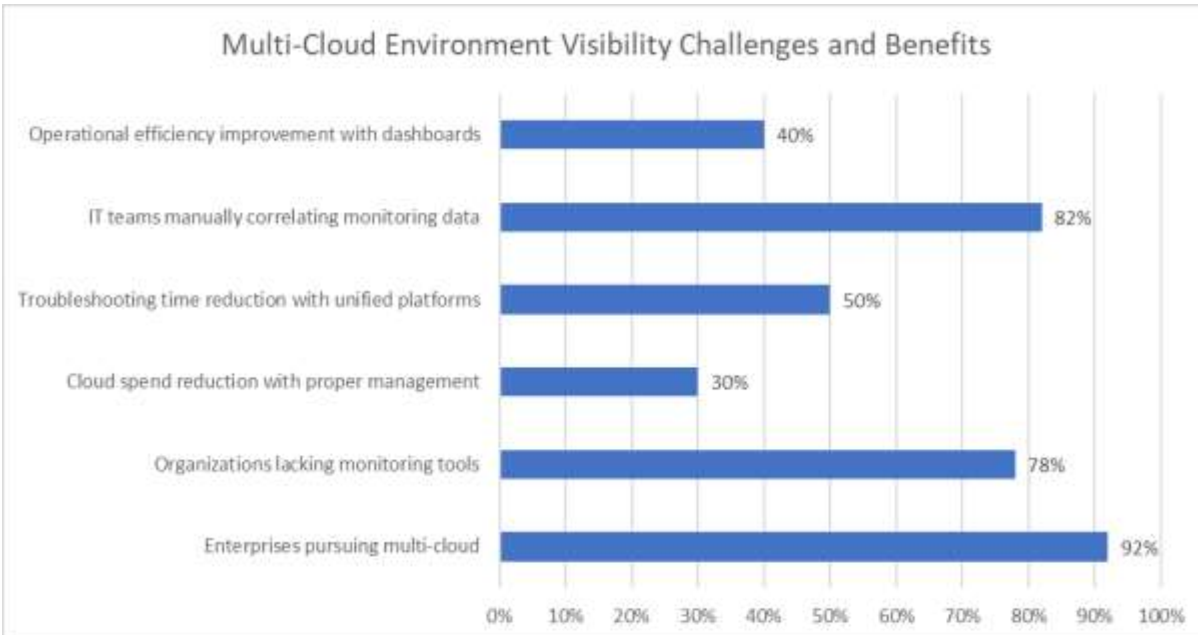

Fig. 2: Cloud Management Metrics in SD-WAN Multi-Cloud Deployments [9, 10]

6. Conclusion

Software-Defined Wide Area Networks have revolutionized enterprise networking by providing unprecedented flexibility and cost efficiency compared to traditional architectures. Decoupling network control from underlying hardware infrastructure enables centralized management and simplified deployment, making SD-WAN an essential technology for organizations adopting multi-cloud strategies and hybrid work models. Despite these advantages, SD-WAN environments introduce complex performance challenges that require sophisticated management approaches. Link congestion, suboptimal path selection, latency fluctuations, and application identification difficulties can significantly degrade user experience when not properly addressed. Artificial intelligence and machine learning technologies are invaluable in transforming SD-WAN fault and performance management from reactive to proactive approaches. These technologies establish performance baselines, detect anomalies before they impact service, correlate events across complex environments, and adapt to changing network conditions—all critical capabilities for modern distributed networks. Optimization strategies enhance SD-WAN deployments through intelligent traffic routing, dynamic bandwidth allocation based on application requirements, granular application-aware policies, and reliability-enhancing techniques like forward error correction. As organizations increasingly implement multi-cloud architectures, achieving comprehensive end-to-end visibility becomes more essential and challenging. Specialized monitoring solutions that access performance data across cloud providers, correlate network and application metrics, and present information through unified dashboards are key to effective troubleshooting and optimization. The future of SD-WAN lies in increasingly intelligent, automated systems that can predict, identify, and remediate performance issues before they impact business operations, ultimately delivering on the promise of flexible, reliable connectivity across distributed environments while minimizing management complexity and operational costs.

Funding: This research received no external funding.
Conflicts of Interest: The authors declare no conflict of interest.
Publisher's Note: All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers.

References

[1] Rich Martin, "The Hidden Challenges of SD-WAN & How to Overcome Them," Itential, 2025. [Online]. Available: https://www.itential.com/blog/company/hybrid-cloud-automation/the-hidden-challenges-of-sd-wan-how-to-overcome-them/

[2] Versa Networks, "The Role of a Secure SD-WAN in Multi-Cloud Transformation." [Online]. Available: https://versa-networks.com/documents/white-papers/secure-sd-wan-and-multi-cloud-transformation.pdf

[3] Palo Alto Networks, "Top 7 SD-WAN Challenges: SD-WAN Risks, Issues, & Solutions." [Online]. Available: https://www.paloaltonetworks.in/cyberpedia/sd-wan-challenges

[4] Sebastian Troia et al., "On Deep Reinforcement Learning for Traffic Engineering in SD-WAN," ResearchGate, December 2020. [Online]. Available: https://www.researchgate.net/publication/346647694_On_Deep_Reinforcement_Learning_for_Traffic_Engineering_in_SD-WAN

[5] Juniper Networks, "AI-Driven SD-WAN Accelerates Digital Transformation with an Agile, Efficient, and Resilient Network." [Online]. Available: https://www.juniper.net/content/dam/www/assets/solution-briefs/us/en/routers/ai-driven-sd-wan-accelerates-digital-transformation-with-an-agile-efficient-and-resilient-network.pdf

[6] Soumaine Bouba Mahamat, Celal Çeken, "Anomaly Detection in Software-Defined Networking Using Machine Learning," Düzce University Journal of Science & Technology, 7 (2019) 748-756, 2018. [Online]. Available: https://dergipark.org.tr/tr/download/article-file/648648

[7] Cisco, "Cisco Catalyst SD-WAN Solution Overview," November 22, 2023. [Online]. Available: https://www.cisco.com/c/en/us/solutions/collateral/enterprise-networks/sd-wan/nb-06-sd-wan-sol-overview-cte-en.html

[8] Samaa Adel Ibrahim Hussein et al., "Performance evaluation of software-defined wide area network based on queueing theory," The Institute of Engineering and Technology, 05 July 2022. [Online]. Available: https://ietresearch.onlinelibrary.wiley.com/doi/10.1049/ntw2.12039

[9] Team Blazeclan, "Multi-Cloud Management: Challenges & Best Practices," Blazeclan, 14 July 2023. [Online]. Available: https://blazeclan.com/india/blog/multi-cloud-management-challenges-and-best-practices-to-overcome-them/

[10] [10] Nikolay Grozev, Rajkumar Buyya, "Inter-Cloud architectures and application brokering: taxonomy and survey," Wiley, 12 December 2012. [Online]. Available: https://onlinelibrary.wiley.com/doi/10.1002/spe.2168