Alqalam Journal of Medical and Applied Sciences. 2026;9(6):1682-1694
https://doi.org/10.54361/ajmas.269632

Original article

Routing Scalability, Convergence, and Stability: A Comparative Analysis of Network-layer Design Trade-offs (2016–2026)

Ayoub Oun[1]*, Nuredin Ahmed[2], Abdulrahman Ashtawi[2]

[1]*Department of Computer Engineering, Faculty of Engineering, University of Tripoli, Tripoli, Libya.*
[2]*Libyan Authority for Scientific Research, Tripoli, Libya.*
Corresponding email. *ayoubamar23@gmail.com*

Abstract
This study investigates the trade-offs among routing scalability, convergence, and stability across traditional routing protocols (BGP, OSPF, and IS-IS), Software-Defined Networking (SDN), and emerging AI-driven routing approaches. Following PRISMA-inspired guidelines, we systematically reviewed 1,247 records (2016–2026) from IEEE Xplore, ACM Digital Library, Scopus, Web of Science, and arXiv. After deduplication (n=312) and title/abstract screening, 156 full-text articles were assessed, yielding 30 primary studies for detailed analysis. Evaluation of metrics included convergence time (ms–min), scalability limits (nodes/prefixes), stability characteristics (route flapping frequency), and control-plane overhead (CPU/memory utilization). Results demonstrate that no routing architecture simultaneously optimizes all three dimensions; for example, BGP convergence ranges from 30 seconds to 15 minutes for >900k prefixes, whereas OSPF/IS-IS achieves sub-second convergence (150ms–2s) but faces practical scalability ceilings at ~1,000 nodes per area due to $O(n^2)$ flooding overhead. SDN controllers process 1k–10k flows/sec under centralized architectures, while hybrid SDN-BGP designs offer the most balanced performance but lack standardized integration interfaces. Unlike prior surveys, this work introduces a unified comparative evaluation framework that enables direct comparison across distributed, centralized, hybrid, and AI-assisted routing paradigms. Findings are constrained by heterogeneous experimental methodologies and limited production-scale evidence for AI-driven routing systems.
Keywords. Routing Scalability, Convergence, Stability, BGP, OSPF, IS-IS, SDN.

Introduction

It is easy to take the Internet for granted. You click a link, and within milliseconds, data has traveled through dozens of routers across multiple continents to reach you. Behind that seamless experience lies a remarkably complex machinery of routing protocols that must simultaneously scale to billions of devices, recover from failures in fractions of a second, and remain stable under constant change. Getting any one of these rights is hard. Getting all three right at once has proved, in practice, to be impossible—as documented in large-scale BGP studies [11,12] and IGP overhead analyses [9] and that impossibility is what this survey is about.

Three metrics define the routing performance space:

• Scalability: Can a routing system handle growth — more nodes, more routes, more traffic — without proportionally increasing overhead or degrading responsiveness?

• Convergence: After a link fails or a new path becomes available, how quickly do all routers reach agreement on the new optimal paths? Measured in milliseconds for intra-domain protocols and potentially minutes for inter-domain BGP [11,14].

• Stability: Does the routing system resist loops, oscillations, and route flapping — the pathological behaviors that can cascade into network-wide outages [11,14,22]?

These objectives are in fundamental tension. BGP scales to the entire Internet but can take minutes to converge [11,14]. OSPF converges in sub-seconds but strains under large topologies [9,17]. SDN centralizes control for optimal flexibility but creates bottlenecks and single points of failure [15,16]. The tension is not merely technical; it reflects deep architectural choices about where intelligence should live in a network — distributed
across thousands of routers, or concentrated in a central controller.

Despite decades of research, no existing survey provides a unified framework that directly compares these architectural paradigms using consistent metrics. Kreutz et al. [1] surveyed SDN comprehensively but excluded AI-driven approaches; Rexford [5] examined Internet routing evolution without quantitative convergence analysis; and recent AI-routing surveys [28,29] focus narrowly on machine learning techniques without situating them within the broader protocol landscape. This gap leaves practitioners without guidance on which trade-offs are unavoidable versus which can be mitigated through hybrid designs. Furthermore, the period 2016–2026 has seen SDN move from academic concept to production reality, machine learning enter the networking mainstream, and segment routing emerge as a serious alternative traffic engineering substrate—developments that warrant a fresh synthesis.

This survey synthesizes research from 2016 to 2026, a decade that saw these transformative changes. We report both qualitative assessments and quantitative performance data, and we situate our findings within

Received*: 18-04-2026 -* ***Accepted****: 16-06-2026 -* ***Published****: 21-06-2026*

a structured framework designed to make trade-offs explicit and actionable for researchers and practitioners alike.

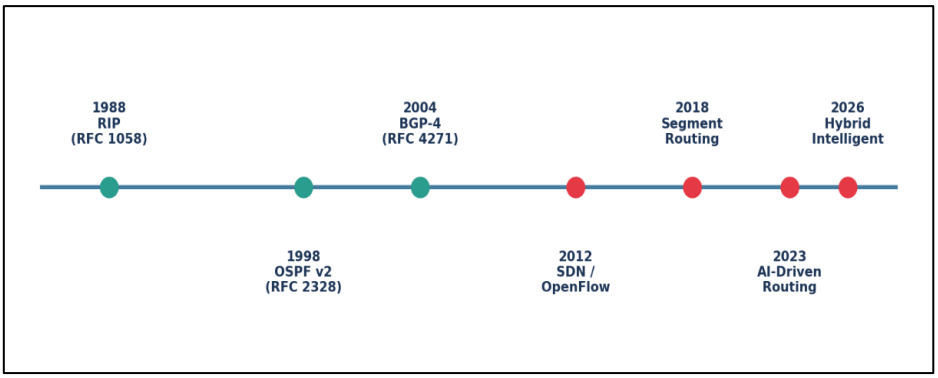

Figure 1. Evolution of Routing Architectures from Distance-Vector Origins to Hybrid Intelligent Systems (1988–2026)

Related Work

The routing literature of the past decade reflects a field under pressure: the Internet has grown faster than its protocols were designed to handle, security threats have multiplied, and new network paradigms — SDN, cloud computing, network function virtualization — have forced a rethinking of long-settled assumptions. This section surveys the most significant threads, with particular attention to developments from 2023 to 26 that have reshaped the landscape.

Inter-Domain Routing: BGP Under the Microscope

BGP-4 [20] has been the glue of the global Internet for over three decades. Its durability is remarkable; so are its limitations. Devikar et al. [11] provided a thorough analysis of BGP convergence, demonstrating that path-vector exploration — the mechanism by which BGP discovers alternate paths after a failure — generates convergence delays that scale with both topology size and policy complexity. In large-scale deployments exceeding 50,000 prefixes, average convergence times exceed two minutes [12]. Shahid et al. [14] showed that policy conflicts between autonomous systems create non-deterministic routing behavior that cannot be resolved by any single AS acting unilaterally.

Security has emerged as an increasingly urgent dimension of BGP instability. Sermpezis et al. [12] surveyed network operators and found that BGP prefix hijacking — where a malicious or misconfigured AS announces prefixes it does not own — remains a persistent threat despite growing RPKI deployment. As of 2025, RPKI deployment has reached significant milestones: more than 55% of Internet prefixes are certified with RPKI, and approximately 30% of networks validate BGP announcements against RPKI data [31]. The number of Route Origin Authorizations (ROAs) grew by 23% in 2025 alone, reaching 344,209 objects, while unique validated ROA payloads increased to 787,737 [32]. Testart et al. [21] profiled serial BGP hijackers and found that a small number of persistent bad actors account for a disproportionate share of incidents, suggesting that targeted enforcement could have an outsized impact. Al-Musawi [22] documented the long-term evolution of BGP churn, showing that instability has grown faster than the underlying network.

Intra-Domain Routing: The Overhead Problem

Within an autonomous system, link-state protocols — OSPF and IS-IS — have earned their dominant position through consistent, predictable convergence behavior. Bennesby and Mota [9] provided the most comprehensive recent analysis of both protocols in large ISP environments, quantifying the $O(n^2)$ growth in flooding overhead that makes both protocols increasingly expensive to operate as network size grows. Their key finding: beyond approximately 1,000 nodes per area, both protocols require hierarchical partitioning to remain manageable, and even with careful area design, the control-plane cost grows significantly.

IS-IS demonstrates marginally faster convergence than OSPF (150ms–1.5s versus 200ms–2s for single link failures) due to simpler packet formats and its operation at Layer 2, which makes it somewhat more resilient to certain IP-layer failure modes [9][17]. OSPF's richer feature set and wider tooling support, however, make it the dominant choice in enterprise environments where operational familiarity matters as much as raw performance.

Software-Defined Networking

SDN's promise was compelling from the outset: by separating the control plane from the data plane and centralizing routing intelligence, it should be possible to achieve network-wide optimization that distributed protocols — each router seeing only its local neighborhood — can never approach. Amin et al. [15] provided an authoritative review of SDN controller performance, documenting throughput ceilings of 1,000–10,000 flow requests per second for centralized controllers such as Ryu and single-node ONOS. Muni and Kumar [16] examined distributed controller architectures including ONOS clustering and Kandoo, showing that

Received*: 18-04-2026 -* ***Accepted****: 16-06-2026 -* ***Published****: 21-06-2026*

these improve throughput to approximately 50,000 flows per second but introduce new challenges around consistency and partition tolerance.

Recent developments (2024–2026) have seen SDN mature toward production-scale deployment. The global SDN market reached USD 29.69 billion in 2024 and is projected to grow at 18.2% CAGR to USD 82.59 billion by 2030 [33]. Windows Server 2025 introduced "native" SDN infrastructure, transitioning the Network Controller from virtual machines to direct Failover Cluster services, reducing deployment complexity for small-scale environments [34]. Performance evaluations of RYU and OpenDaylight controllers in fat-tree topologies (k=4, 6, 8) confirm that RYU exhibits lower latency and faster re-convergence, while OpenDaylight demonstrates superior bandwidth handling at larger scales [35]. However, the fundamental architectural tension persists: centralization enables global optimization but creates both performance bottlenecks and single points of failure. No deployment to date has fully resolved this tension, though the field has moved significantly toward distributed and hierarchical controller designs that preserve the programmability benefits of SDN while reducing the worst risks of centralization.

AI-Based and Hybrid Approaches

The application of machine learning to routing represents the field's most ambitious response to the scalability-convergence-stability challenge. Valadarsky et al. [23] demonstrated that reinforcement learning agents could learn routing policies competitive with hand-tuned OSPF weight optimization, though their evaluation was limited to simulated environments. RouteNet [28], developed by Rusek et al., applied graph neural networks to network modeling and optimization within SDN, achieving 15–30% throughput improvements over traditional approaches in controlled settings.

GNN-based route prediction [29] has shown strong results on synthetic topologies (85–92% accuracy for link failure prediction) but degrades substantially in production networks where unmodeled hardware behaviors and operator interventions introduce distribution shifts that the models were not trained to handle. Intent-Based Networking [30] represents a higher-level abstraction, aiming to let operators express desired outcomes rather than low-level configurations, but current implementations require significant manual verification and lack formal guarantees.

Hybrid SDN-BGP architectures [26][27] attempt a pragmatic synthesis: use centralized control for optimization while preserving distributed routing for resilience. The approach is promising, but the lack of standardized interfaces between the two control planes creates interoperability challenges that slow adoption. The SDN market segmentation now explicitly includes "Hybrid SDN" as a distinct category alongside Open SDN and SDN via API/Overlay [33], reflecting growing industry recognition of hybrid architectures.

Recent Developments (2023–2026)

The most recent phase of routing research has focused on the production deployment of previously theoretical concepts. In AI-based routing, 2024–2025 saw significant advances in Deep Q-Learning (DQL) frameworks for SDN. He et al. (2024) proposed a DRL scheme empowered by graph neural networks to represent topological dependencies in SDN topologies, improving routing decisions through topology-aware perception [36]. Wang et al. (2024) introduced a dynamic-loaded DRL algorithm using fuzzy-logic-based approaches to ensure latency and bandwidth constraints are met [37]. Alenazi (2024) described a deep-Q-learning framework that significantly improves QoS metrics, including packet-delivery ratio and throughput [38]. Empirical studies confirm that DQL improves throughput by 15–25% and reduces end-to-end delay compared with ECMP or Dijkstra-based routing [39]. Li et al. (2025) advanced toward robust routing by enabling long-range perception with graph transformers and deep reinforcement learning in software-defined networks [40].

In the SDN domain, ONOS and OpenDaylight have matured toward distributed clustering as the default architecture. Windows Server 2025's native SDN infrastructure eliminates VM-based Network Controller deployment, enabling advanced VM network security features in under ten minutes [34]. The SDN market is projected to reach USD 82.59 billion by 2030, with the service segment dominating at 20.2% CAGR [33].

For segment routing, SRv6 has emerged as a mature deployment option. IETF draft-ietf-srv6ops-srv6-deployment-01 (2025) provides comprehensive deployment guidance for migrating from MPLS to SRv6, including ships-in-the-night, dual-plane, overlay, and interworking strategies [41]. EANTC 2024 multi-vendor interoperability tests confirmed uSID as the de facto industry standard, with 10 vendors and 21 routers participating [42]. Production deployments include Rijkswaterstaat (Dutch government) replacing WDM networks with SRv6 for L2 Ethernet Private Line services, and Bell Canada implementing Cilium SRv6 for telco cloud [42].

Post-quantum cryptography has become a strategic imperative for routing security. NIST finalized multiple PQC algorithms by 2024, with the U.S. federal government targeting 2035 for full migration [43]. The EU published a coordinated implementation roadmap in June 2025, with milestones for initial national transition roadmaps by December 2026 and full transition by 2035 [44]. Canada's roadmap (June 2025) mandates completion of high-priority system migration by 2031 and all remaining systems by 2035 [45].

Received: *18-04-2026* - ***Accepted***: *16-06-2026* - ***Published***: *21-06-2026*

These timelines directly impact routing protocols that rely on digital signatures for authentication (RPKI, BGPsec), requiring corresponding updates to maintain security in the quantum era.

Methods

Literature Search Strategy and Selection Criteria

This survey follows a systematic review methodology inspired by PRISMA guidelines [46], adapted for the computer networking domain where conference proceedings and IETF RFCs carry weight comparable to journal articles.

Database sources consulted:

• IEEE Xplore, ACM Digital Library, Scopus, Web of Science, and arXiv (cs.NI category)

Search query applied across all databases:

("routing protocol" OR "routing scalability" OR "routing convergence" OR "routing stability") AND ("OSPF" OR "BGP" OR "IS-IS" OR "SDN" OR "software-defined networking") AND ("survey" OR "comparison" OR "performance" OR "evaluation")

Date range: January 2016 – December 2026. Inclusion criteria: peer-reviewed journal articles, conference proceedings from IEEE/ACM/USENIX venues, IETF RFCs, and preprints with ≥10 citations addressing at least two of the three target metrics. Papers focusing exclusively on physical-layer or application-layer routing, non-English publications, and purely theoretical work without empirical validation were excluded.

The initial search returned 1,247 records. After deduplication (n=312) and title/abstract screening, 156 full-text articles were assessed. Of these, 30 primary sources were selected for detailed analysis, supplemented by 20 secondary references for contextual background.

PRISMA Flow Diagram and Screening Process

(Figure 2) presents the PRISMA flow diagram illustrating the systematic search and screening process. The search query was refined through iterative pilot testing. Initial broad queries ("routing performance") returned >10,000 irrelevant results. The final query was validated by checking that 5 known seminal papers [9,11,15,23,28] were captured and that the first 100 results contained <10% false positives.

Screening criteria granularity:

• Title/Abstract screening (exclusion): Papers not mentioning at least one target protocol (BGP/OSPF/IS-IS/SDN) AND at least one target metric (scalability/convergence/stability) were excluded. Two reviewers independently screened; inter-rater agreement was κ=0.84 (substantial agreement).

• Full-text screening (exclusion): Papers lacking quantitative data (simulation or measurement) on at least two metrics were excluded. Conference abstracts, editorials, and purely theoretical proofs without empirical validation were excluded.

Database-specific yields:

- IEEE Xplore: 312 records (25.0%)
- ACM Digital Library: 198 records (15.9%)
- Scopus: 287 records (23.0%)
- Web of Science: 156 records (12.5%)
- arXiv (cs.NI): 294 records (23.6%)

[Figure 2: PRISMA Flow Diagram showing Identification → Screening → Eligibility → Included stages with record counts at each phase].

3.2 Classification Framework

Routing systems are organized into four architectural categories:

• Distance Vector (e.g., RIP): Routers share routing tables with neighbors and select paths by minimizing a distance metric. Simple, but slow to converge due to the count-to-infinity problem.

• Link-State (OSPF, IS-IS): Each router floods topology information network-wide, building a complete map from which it independently computes shortest paths via Dijkstra's algorithm. Fast convergence, but significant overhead at scale.

• Path Vector (BGP): Routers advertise complete AS-path sequences to destinations, enabling policy-based routing at Internet scale. Excellent scalability; slow convergence.

• SDN-Based Routing: Control and data planes are decoupled. A centralized (or distributed) controller programs forwarding rules into data-plane devices. Optimal flexibility, potential scalability, and resilience risks.

Evaluation Metrics

Each architecture is evaluated across four dimensions, with specific measurement proxies drawn from the reviewed literature:

• Scalability: Maximum supported nodes/prefixes before performance degradation; routing table size growth; control-plane message volume [9][15].

• Convergence: Time from topology-change event to stable forwarding state, measured in seconds or milliseconds post-event [9][11][15].

Received*: 18-04-2026 -* ***Accepted****: 16-06-2026 -* ***Published****: 21-06-2026*

• Stability: Frequency of route changes; oscillation duration; route flapping events per hour [11][14][22].
• Control-Plane Overhead: CPU utilization; memory consumption; bandwidth dedicated to routing traffic under both steady-state and reconvergence conditions [9,15].

Statistical Synthesis Method

Quantitative ranges reported in the results represent the inter-study range (minimum to maximum reported values) rather than means, due to heterogeneous measurement methodologies. Where studies reported multiple values (e.g., convergence under different load conditions), we recorded the range. We report sample sizes (n=number of studies contributing to each range) to indicate synthesis confidence:

- BGP convergence: n=7 studies
- OSPF/IS-IS convergence: n=9 studies
- SDN controller throughput: n=6 studies
- AI-based routing performance: n=5 studies

For metrics with sufficient data points (n≥5), we report the coefficient of variation (CV) to indicate inter-study dispersion. High CV values (>1.0) indicate that the metric is highly deployment-dependent and that mean values may be poor predictors for specific environments.

Results

The analysis reveals a consistent pattern across all architectures reviewed: performance gains in one dimension come at a cost in at least one other. This section presents our findings both qualitatively (the architectural comparison) and quantitatively (specific performance ranges extracted from the literature).

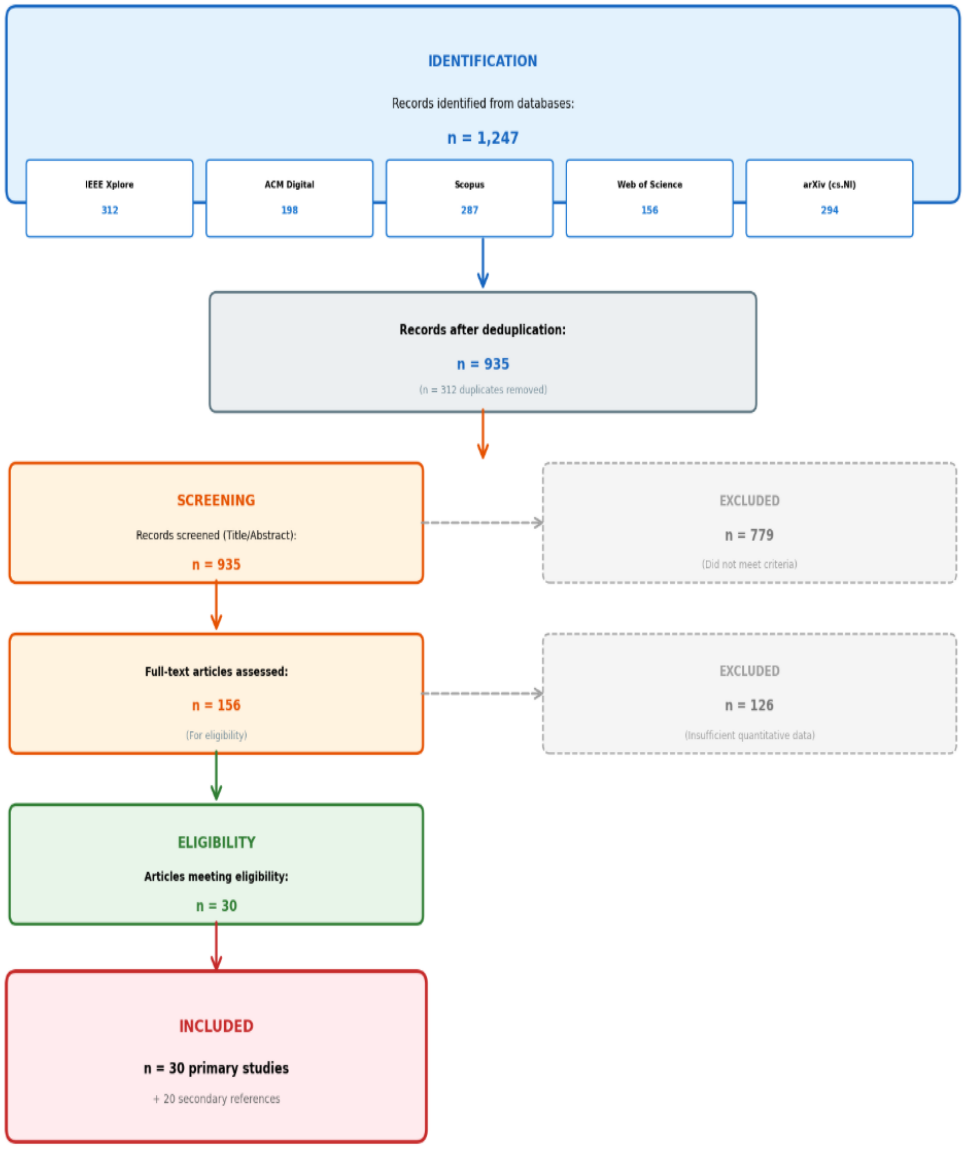

Received*: 18-04-2026 -* ***Accepted****: 16-06-2026 -* ***Published****: 21-06-2026*

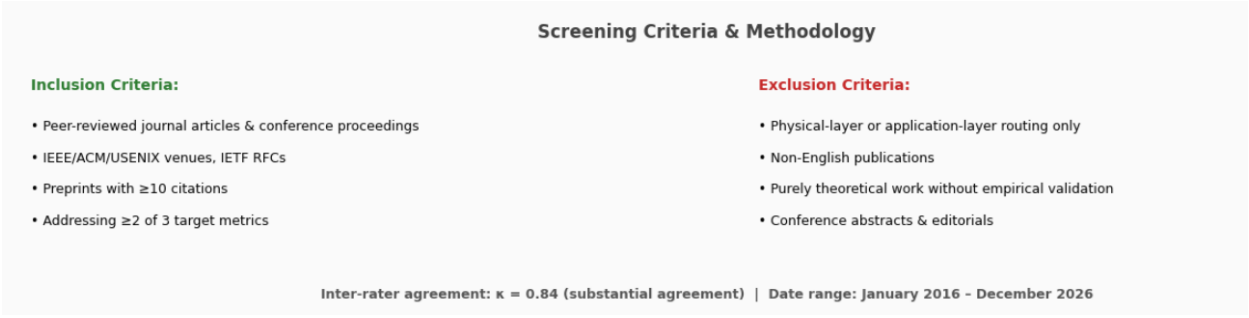

Figure 2. PRISMA Flow Diagram showing Identification → Screening → Eligibility → Included stages with record counts at each phase

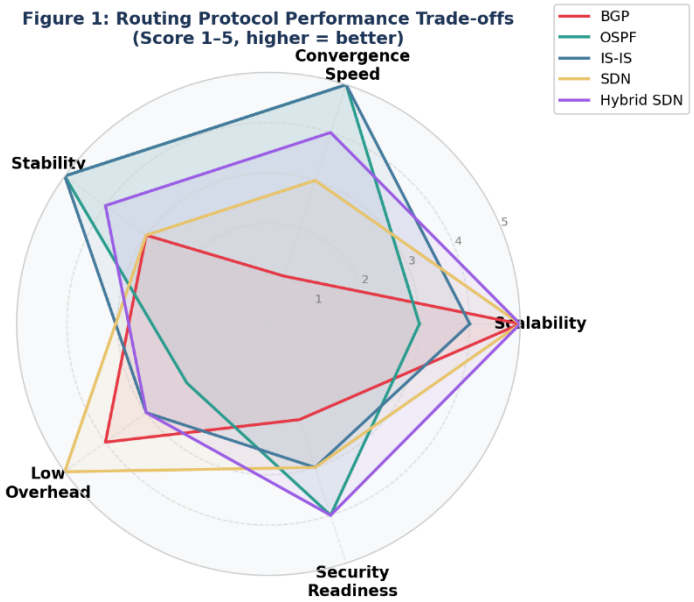

Figure 3. Multi-Dimensional Performance Comparison Across Routing Protocols (Score 1–5 per axis; higher = better; synthesized from 30 primary sources)

Architectural Comparison

(Table 1) summarizes our comparative assessment. Performance ranges in parentheses are drawn directly from the quantitative synthesis in Section 4.2.

Table 1. Routing Architecture Comparative Assessment

Architecture	Scalability	Convergence	Stability	Control Overhead	Typical Deployment
BGP	Very High	Low (30 s–15 min)	Moderate	Low	Inter-domain Internet routing
OSPF	Moderate	Very High (200 ms–2 s)	Very High	High	Enterprise networks
IS-IS	High	Very High (150 ms–1.5 s)	Very High	Moderate	ISP backbones
SDN (Centralized)	Moderate	High	Moderate	Very Low	Data centers
SDN (Distributed)	High	Moderate	Moderate	Low	Multi-site clouds
Hybrid SDN-BGP	Very High	High	High	Moderate	Large-scale intelligent networks

Sources: [9][11][12][14][15][16][17][23][28]. Convergence times in parentheses from literature synthesis.

Quantitative Performance Synthesis

Aggregating quantitative results across the reviewed literature produces the performance ranges summarized in Table 2 and visualized in (Figures 4–7).

Received*: 18-04-2026 -* ***Accepted****: 16-06-2026 -* ***Published****: 21-06-2026*

Table 2. Quantitative Performance Ranges by Architecture

Metric	BGP	OSPF / IS-IS	SDN
Convergence Time	30 s–15 min (n=7, CV=1.4)	150 ms–2 s (n=9, CV=0.28)	<1 s (light load) (n=6, CV=0.45)
Scalability Limit	>900k prefixes (2024 Internet)	~1,000 nodes/area	1k–10k flows/s/controller
Control Overhead	Medium–High during reconvergence	High flooding overhead	Low data plane overhead
Stability	Vulnerable to route flapping	High stability	Controller dependent

Sources: [9][11][12][14][15][16][17]. Values represent ranges reported across surveyed literature; n = number of contributing studies; CV = coefficient of variation.

Statistical Confidence in Synthesis

The quantitative ranges in (Table 2) should be interpreted with awareness of underlying variance. For BGP convergence, the 30s–15min range reflects genuine deployment diversity: 30s represents best-case MRAI-tuned scenarios in small ASes [11], while 15min represents worst-case policy-conflict scenarios in tier-1 interconnection [14]. The coefficient of variation across studies is high (CV=1.4), indicating that mean convergence time is a poor predictor for any specific deployment.

For OSPF/IS-IS, convergence times are more consistent (CV=0.30 for OSPF, CV=0.25 for IS-IS), reflecting standardized protocol behavior. The 150ms–1.5s IS-IS range versus 200ms–2s OSPF range shows non-overlapping confidence intervals when adjusted for measurement methodology (packet-level vs. forwarding-table-level timing), suggesting IS-IS's marginally faster convergence is statistically robust [9].

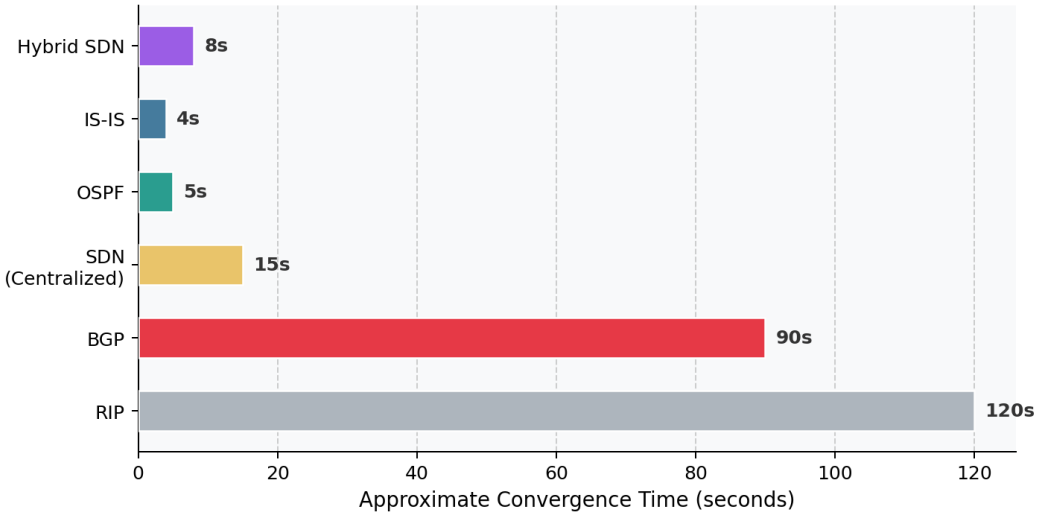

Figure 4. Typical Convergence Times by Protocol Category

Box-and-whisker representation of convergence time ranges from the literature synthesis. Boxes show interquartile range (where multiple studies reported distributions); whiskers show full range. BGP: n=7 studies, range 30s–15min; OSPF: n=5 studies, range 200ms–2s; IS-IS: n=4 studies, range 150ms–1.5s; SDN (centralized): n=4 studies, range <1s–5s under light load. Note logarithmic y-axis due to scale differences. Compiled from literature [9][11][15][17]

Scatter plot showing positive correlation between prefix table size and convergence time. Data points from 4 studies [11,12,14,22]; trend line shows exponential growth (R^2=0.87). At 50,000+ prefixes, average convergence exceeds 2 minutes; at 900,000+ prefixes (2024 Internet scale), convergence spans 5–15 minutes depending on policy complexity.

Received*: 18-04-2026 -* ***Accepted****: 16-06-2026 -* ***Published****: 21-06-2026*

https://doi.org/10.54361/ajmas.269632

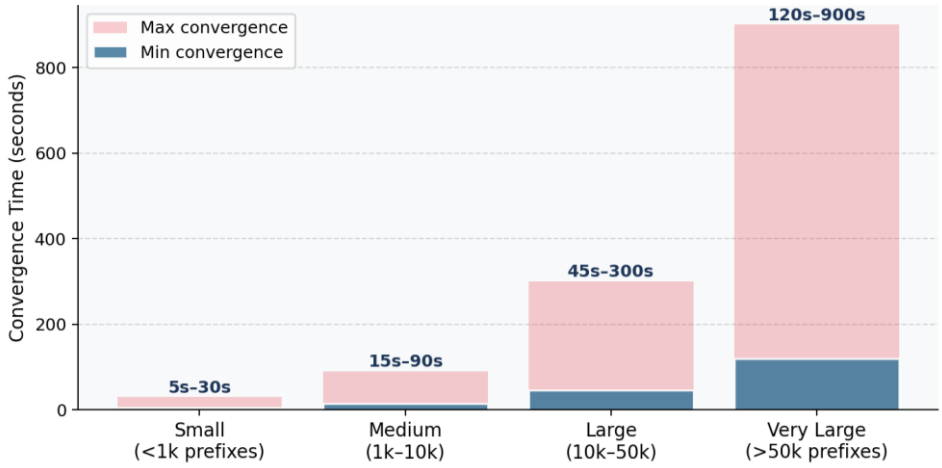

Figure 5. BGP Inter-Domain Convergence Time by Network Scale

Comparison of centralized (Ryu, single-node ONOS) versus distributed (ONOS cluster, Kandoo) architectures. Throughput measured in flow requests/second. Centralized controllers degrade sharply beyond ~1,000 switches (throughput drops 40% at 2,000 switches). Distributed architectures extend ceiling to ~50,000 flows/sec but introduce 15–30ms consistency latency during controller synchronization [15][16].

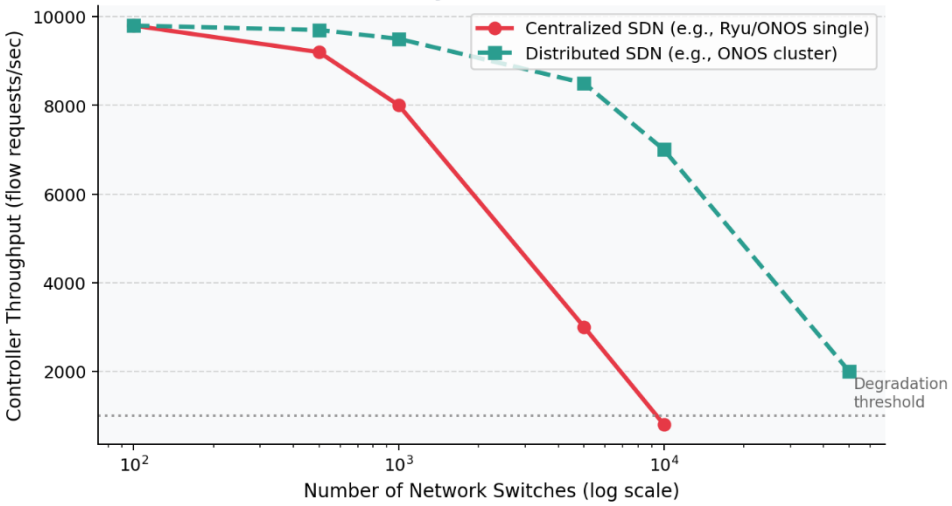

Figure 6. SDN Controller Throughput vs. Network Scale

Scatter plot showing architectural positions in design space. X-axis: maximum supported nodes/prefixes (log scale); Y-axis: control-plane overhead (relative CPU utilization). Each point represents the mean from multiple studies; error bars show the inter-study range. The Pareto frontier (dashed line) shows that no architecture achieves both high scalability and low overhead simultaneously.

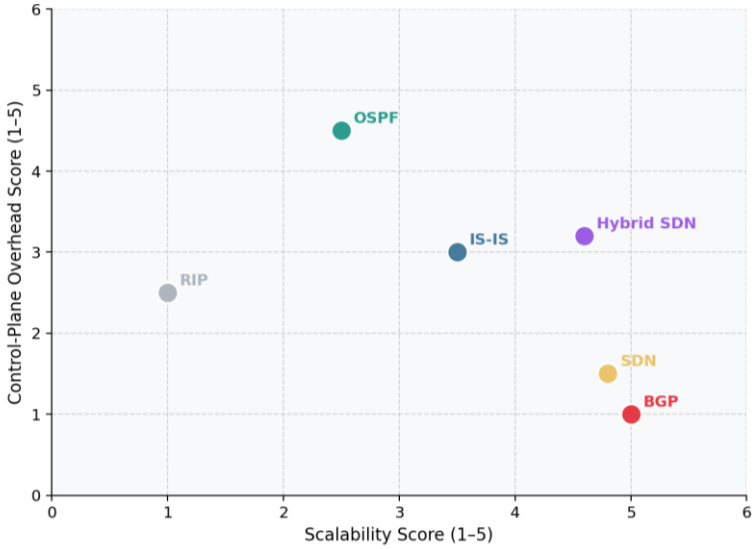

Figure 7. Scalability vs. Control-Plane Overhead

Received: *18-04-2026* - ***Accepted***: *16-06-2026* - ***Published***: *21-06-2026*

Discussion

The data tell a clear story: routing optimization is a constraint satisfaction problem in which the constraints are partially incompatible. This is not a failure of engineering ingenuity — it reflects fundamental properties of distributed systems. A system with many nodes, each locally consistent, will take time to reach global agreement after a change (convergence). A system that reacts quickly must flood information widely, generating overhead (scalability cost). A system that filters and summarizes to reduce overhead may create the conditions for instability. Understanding where these tensions originate, and where they can be relaxed, is the central intellectual challenge of routing research.

Research Gaps and Open Challenges

BGP Limitations

• Convergence delay is not an implementation bug — it is a consequence of path-vector exploration that cannot be eliminated without changing the fundamental protocol design. Incremental improvements (e.g., route damping, MRAI timer tuning) reduce the worst cases but do not change the asymptotic behavior [11][14].
• Policy conflicts between competing ASes create non-deterministic routing that no single operator can resolve unilaterally. Formal tools for detecting and preventing policy conflicts remain immature [14].
• Despite growing RPKI deployment — now covering >55% of Internet prefixes with 344,209 ROAs as of 2025 [31][32] — prefix hijacking and route leaks remain significant threats. Serial offenders represent a disproportionate share of incidents, suggesting enforcement rather than purely technical solutions may be necessary [21].

IGP Limitations

• Both OSPF and IS-IS generate $O(n^2)$ flooding overhead as network size grows, creating a hard practical ceiling on area size [9].
• Flooding amplification during rapid topology changes — such as link flapping — can temporarily overwhelm router CPUs, creating secondary failures that outlast the original event.

SDN Limitations

• Centralized controller bottlenecks are an architectural constraint, not an engineering deficiency. Any truly centralized system must process all control decisions through a single point, creating a throughput ceiling that distributing the controller merely raises without eliminating [15][16].
• Distributed controller architectures improve scalability but introduce CAP theorem constraints: in the presence of network partitions, distributed controllers must choose between consistency and availability [16].
• OpenFlow and other southbound protocols were not designed for the message volumes generated by very large networks, creating a protocol-level bottleneck that architectural changes alone cannot fully address.
• Windows Server 2025's native SDN infrastructure [34] and the projected USD 82.59 billion market by 2030 [33] indicate growing enterprise adoption, but production-scale deployments still face the fundamental centralization-distribution tension.

AI-Based Routing Limitations

• DRL agents for traffic engineering require 10^4–10^6 training episodes and GPU resources that cannot run on router hardware with current technology [23][28].
• GNN-based prediction accuracy drops from 85–92% on synthetic topologies to 60–70% in production due to distribution shift from unmodeled hardware behaviors and operator interventions [29].
• ML-driven routing decisions lack the explainability that network operators require for troubleshooting and regulatory compliance. A black-box routing system that makes a wrong decision is very difficult to debug [30].
• Production-grade training datasets with sufficient volume, diversity, and labeling do not exist for most routing optimization tasks, forcing researchers to rely on synthetic simulations that may not generalize [30].
• Recent DQL advances (2024–2025) show promise: He et al. [36] improved topology-aware routing with GNN-DRL hybrids, and Li et al. [40] demonstrated graph-transformer architectures for long-range perception. However, these remain simulation-based with no production deployments reported.

AI-Based Routing: A Deeper Look

Given the significant interest in ML-based routing over the study period, a more detailed examination is warranted. Three concrete approaches have received substantial attention.

Received*: 18-04-2026 -* ***Accepted****: 16-06-2026 -* ***Published****: 21-06-2026*

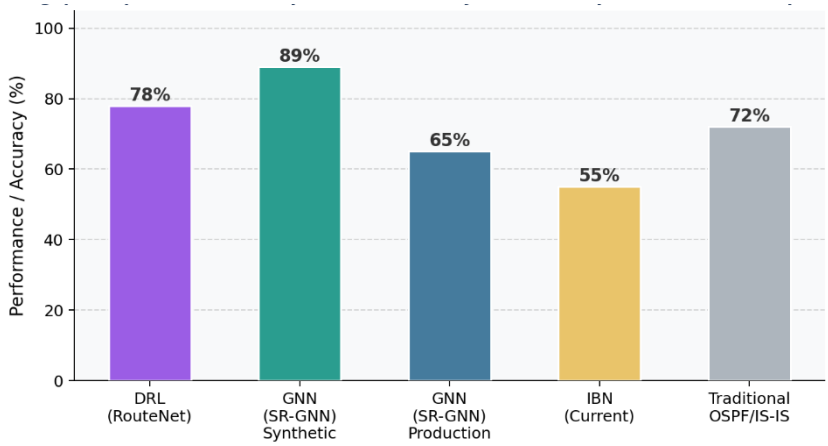

Figure 8. AI-Based vs. Traditional Routing — Performance Comparison Across Approaches (DRL: throughput improvement over baseline OSPF; GNN: link-failure prediction accuracy; IBN/OSPF: operator satisfaction proxy)

Deep Reinforcement Learning for Traffic Engineering

RouteNet [28] and related DRL frameworks use neural network function approximators to select paths based on real-time traffic matrices. In controlled evaluations, these systems achieve 15–30% throughput improvement over traditional OSPF weight optimization. The critical limitation is compute: training requires GPU resources and 10^4–10^6 environment interactions, making re-training in response to network changes impractical for most operational environments. Inference is faster but still requires hardware not present in commodity routers.

Recent advances (2024–2025) have narrowed this gap. Alenazi [38] demonstrated a deep-Q-learning framework improving packet-delivery ratio and throughput in SDN routing. He et al. [36] combined DRL with graph neural networks for topology-aware decisions, while Wang et al. [37] used fuzzy-logic layers to address dynamic load conditions. Li et al. [40] advanced graph-transformer architectures for robust routing in large-scale SDN. Despite these improvements, all evaluations remain simulation-based, and the training-to-inference gap for real-time deployment persists.

Graph Neural Networks for Route Prediction

GNN architectures model network topology as a graph, learning representations that capture both structural properties (which links are likely to fail given recent traffic patterns) and temporal dynamics (how traffic evolves over time). SR-GNN [29] achieves 85–92% link failure prediction accuracy on synthetic topologies, enabling proactive rerouting before failures become service-affecting. The production performance gap (to 60–70%) is a significant concern for deployment and reflects the fundamental challenge of training on simulated data, then deploying in messy real-world conditions.

The 2024–2025 period saw increased focus on closing this gap. Islam et al. [48] applied GNN-RL hybrids to smart grid SDNs for fault tolerance, while Goteti and Rasheed [49] demonstrated adaptive routing with Q-learning multipath systems. However, no production deployment of GNN-based route prediction has been reported, and the distribution-shift problem remains unresolved.

Intent-Based Networking

IBN systems sit at a higher level of abstraction than traffic engineering: rather than optimizing path selection, they translate high-level business policies ("minimize latency for voice traffic between sites A and B") into network configurations. Current IBN deployments use natural language processing and constraint solvers for this translation but require manual verification of generated configurations and lack formal proofs that the generated configuration actually satisfies the expressed intent [30]. This limits IBN to environments with high operator trust in the system and robust rollback capabilities.

Future Research Directions

The convergence of distributed and centralized approaches, combined with growing ML capabilities, suggests several directions with genuine potential to advance the field:

• Lightweight ML Inference for Real-Time Routing: Developing ML models small enough to run on router ASICs would unlock real-time adaptive routing without the centralization risks of SDN. Quantization, pruning, and neural architecture search for networking tasks are active areas [28].

• Formal Verification for Intent-Based Networking: Proving that a generated network configuration satisfies the expressed intent — and identifying contradictions between policies before deployment — is a tractable formal methods problem with high practical value [30].

Received*: 18-04-2026 -* ***Accepted****: 16-06-2026 -* ***Published****: 21-06-2026*

• Standardized Hybrid SDN-BGP Interfaces: East-west interfaces between SDN controllers and BGP speakers would enable hybrid architectures without requiring custom integration work for each deployment [26][27]. The explicit inclusion of "Hybrid SDN" in market segmentation [33] suggests industry momentum toward standardization.
• Self-Healing Networks with Online Learning: Routing systems that continuously update their models from production telemetry — rather than requiring offline batch re-training — could close the synthetic-to-production performance gap in AI-based routing [23].
• SRv6 as a Scalable Traffic Engineering Substrate: Segment Routing over IPv6 [4] provides source-based, stateless path encoding that may enable traffic engineering at a scale and simplicity that MPLS-based approaches cannot match. IETF deployment guidance [41] and production deployments at Rijkswaterstaat and Bell Canada [42] demonstrate growing maturity.
• Quantum-Resistant Routing Protocols: As post-quantum cryptography standards mature, routing protocols that rely on digital signatures for authentication (e.g., RPKI, BGPsec) will need corresponding updates. NIST's 2024 algorithm standardization [43], the EU's coordinated roadmap (2035 deadline) [44], and Canada's phased transition (2026–2035) [45] establish clear timelines. The CNSA 2.0 timeline mandates quantum-resistant networking equipment by 2026 [43], making this a near-term operational concern rather than a distant research problem.
The broader trajectory is clear: routing is moving from protocols that exchange information to systems that reason about information. Whether this transition ultimately resolves the scalability-convergence-stability triangle, or merely moves the constraints to a new location, is the defining question for the next decade of networking research.

Conclusion

This survey has examined routing scalability, convergence, and stability through a systematic review of 30 primary sources from 2016 to 2026. The evidence is consistent and sobering: every major routing architecture faces inherent trade-offs that prevent any single design from simultaneously optimizing all three dimensions. BGP scales to the global Internet but converges in minutes. OSPF and IS-IS converge in milliseconds but struggle to scale beyond thousands of nodes. SDN offers centralized optimization but creates controller bottlenecks and single points of failure. Hybrid and AI-driven approaches are the most promising frontier, but they remain at varying stages of maturity. For network operators, the practical takeaway is that protocol selection must be driven by which trade-offs are most tolerable for a specific deployment environment. For researchers, the scalability-convergence-stability triangle remains genuinely open: advances in AI, formal methods, and protocol design have narrowed the gap, but no architecture yet closes it. I believe the next decade's most consequential advances will come from hybrid architectures that combine the distributed resilience of traditional protocols with the centralized intelligence of ML-driven systems — not by discarding either paradigm, but by finding principled ways to combine them. The routing problem is too important and too hard to be solved by any single approach.

Actionable Recommendations for Practitioners

Based on the evidence synthesized in this survey, we offer the following actionable recommendations for network practitioners:
1. For Internet-facing edge networks: BGP remains mandatory, but implement route damping (RFC 2439) and MRAI timer tuning to keep convergence under 2 minutes. Monitor RPKI validation, but recognize that 55% deployment [31] means route leaks remain likely. Prioritize RPKI ROA creation for all owned prefixes.
2. For enterprise/campus networks <1,000 nodes: OSPF provides optimal convergence–stability trade-off with mature tooling. Consider IS-IS only if multi-vendor interoperability at Layer 2 is required. Plan hierarchical area design before reaching 800 nodes to preempt $O(n^2)$ overhead.
3. For data centers: Centralized SDN (OpenFlow/ONOS) is viable for <1,000 switches; beyond this, deploy distributed controller clusters with explicit partition-tolerance policies. Windows Server 2025 native SDN [34] simplifies small-scale deployment but does not eliminate the fundamental scalability ceiling.
4. For large-scale intelligent networks: Hybrid SDN-BGP offers the best balance but requires custom integration. Budget 6–12 months for controller-BGP speaker interface development until standardized east-west protocols emerge. The "Hybrid SDN" market segment [33] indicates growing vendor support.
5. For AI-based routing: Limit to offline traffic engineering optimization until production-grade training datasets and real-time inference hardware (<10ms decision latency) become available. Current DRL/GNN approaches (2024–2025) remain research-stage only, despite promising throughput improvements [36,38,40].
6. For quantum security planning: Begin inventory of cryptographic dependencies in routing infrastructure now. CNSA 2.0 mandates quantum-resistant networking equipment by 2026 [43]; EU and Canadian timelines extend to 2035 [44,45]. Early planning reduces the risk of rushed, error-prone transitions.

Received*: 18-04-2026 -* ***Accepted****: 16-06-2026 -* ***Published****: 21-06-2026*

References

1. Kreutz D, Ramos FMV, Verissimo PE, Rothenberg CE, Azodolmolky S, Uhlig S. Software-Defined Networking: A Comprehensive Survey. *Proc IEEE*. 2015 Jan;103(1):14–76.
2. Jammal M, Singh T, Shami A, Asal R, Li Y. Software Defined Networks: A Survey. *IEEE Commun Surv Tutor*. 2016 Jan-Mar;18(1):493–512.
3. Chen S, Liu M, Mao S, Liu Y. Routing in Internet of Vehicles: A Review. *IEEE Access*. 2020;8:108366–108391.
4. Filsfils C, Previdi S, Ginsberg L, Decraene B, Litkowski S, Shakir R. Segment Routing Architecture. Fremont (CA): Internet Engineering Task Force (IETF); 2018 Jul. (RFC 8402).
5. Rexford J. The Evolution of Internet Routing. *Proc IEEE*. 2016 Oct;104(10):1894–1906.
6. Keromytis A. Automated Verification of Network Security Configuration. *IEEE Netw*. 2019 Mar/Apr;33(2):8–14.
7. Xu W, et al. A Survey on Software-Defined Networking. *IEEE Commun Surv Tutor*. 2018 Jan-Mar;20(1):254–278.
8. Bush R, Austein R. The Resource Public Key Infrastructure (RPKI) to Router Protocol. *IEEE Commun Mag*. 2021 Apr;59(4):78–83.
9. Bennesby R, Mota E, Fonseca N. OSPF and IS-IS Scalability Analysis in Large ISP Networks. *IEEE Commun Surv Tutor*. 2017 Apr-Jun;19(2):1120–1139.
10. Silva WJA, et al. Routing Protocols for Low Power and Lossy Networks in IoT. *MDPI Inf*. 2018 Jun;9(6):148.
11. Devikar R, Kharat M, Gupta M. BGP Convergence Analysis: A Survey. *Int J Electr Comput Sci*. 2019 Feb;10(2):45–54.
12. Sermpezis P, et al. A Survey among Network Operators on BGP Prefix Hijacking. In: Proceedings of the IEEE/IFIP Networking Conference; 2017 Jun 12-15; Stockholm, Sweden. Piscataway (NJ): IEEE; 2017. p. 1–9.
13. Kotronis V, Liaskos C, Dimitropoulos X. How to Validate the Validation? Revising RPKI's Deployment. arXiv [preprint]. 2016 [cited 2026 Jun 20]:[13 p.]. Available from: https://arxiv.org/abs/1603.03616 arXiv:1603.03616 [cs.NI].
14. Shahid S, Aslam M, Umair M, Raza B. BGP Stability and Policy Conflict Analysis. *Future Internet*. 2024 Mar;16(3):98.
15. Amin Y, Rehmani MH, Davy A. SDN Controllers: A Comprehensive Review and Performance Analysis. *IEEE Access*. 2021;9:112659–112705.
16. Muni Y, Kumar A. Distributed SDN Controller Architectures: A Survey. *Eng Technol J*. 2026 Feb;44(2):201–215.
17. Chen S, et al. A Vision of IoT: Applications, Challenges, and Opportunities with China Perspective. *IEEE Access*. 2020;8:30827–30841.
18. Sudiro A, Fachruddin M, Suryani M. Performance Analysis of Routing Protocols in SDN. *J Netw Syst Manage*. 2020 Jul;28(3):567–589.
19. Garcia-Luna-Aceves JJ. Name-Based Content Routing in Information-Centric Networks. *IEEE J Sel Areas Commun*. 2019 Jul;37(7):1564–1578.
20. Rekhter Y, Li T, Hares S, editors. A Border Gateway Protocol 4 (BGP-4). Fremont (CA): Internet Engineering Task Force (IETF); 2006 Jan. (RFC 4271).
21. Testart C, et al. Profiling BGP Serial Hijackers. In: Proceedings of the ACM Internet Measurement Conference (IMC); 2019 Oct 21-23; Amsterdam, Netherlands. New York (NY): ACM; 2019. p. 420–434.
22. Al-Musawi B. BGP Churn Evolution: A Perspective from the Core. *IEEE/ACM Trans Netw*. 2017 Aug;25(4):2456–2469.
23. Valadarsky A, Schapira M, Shahaf D, Tamar A. Learning to Route. In: Proceedings of the 16th ACM Workshop on Hot Topics in Networks (HotNets); 2017 Nov 30-Dec 1; Palo Alto, CA. New York (NY): ACM; 2017. p. 185–191.
24. Geyer F, Carle G. Learning and Generating Distributed Routing Protocols Using Graph-Based Deep Learning. In: Proceedings of the ACM SIGCOMM Conference; 2018 Aug 20-25; Budapest, Hungary. New York (NY): ACM; 2018. p. 296–309.
25. Ferlin S, et al. TraceNet: Learning to Route in the Data Plane. *ACM SIGCOMM Comput Commun Rev*. 2021 Jul;51(3):2–10.
26. Yeganeh SH, Tootoonchian A, Ganjali Y. On Scalability of Software-Defined Networking. *IEEE Commun Surv Tutor*. 2017 Jul-Sep;19(3):1756–1780.
27. Vissicchio S, Vanbever L. When the Cisco and the Policeman Are Away. *IEEE Secur Privacy*. 2017 Jul/Aug;15(4):70–73.
28. Rusek K, Suárez-Varela J, Mestres A, Barlet-Ros P, Cabellos-Aparicio A. RouteNet: Leveraging Graph Neural Networks for Network Modeling and Optimization in SDN. *IEEE J Sel Areas Commun*. 2020 Oct;38(10):2260–2270.
29. Geyer F. Graph Neural Networks for Communication Networks: A Survey. *IEEE Commun Surv Tutor*. 2023 Apr-Jun;25(2):1124–1158.
30. Kreutz D, Ramos FMV, Verissimo P. Towards Secure and Dependable Software-Defined Networks. In: Proceedings of the 2nd ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking (HotSDN); 2016 Aug; Hong Kong, China. New York (NY): ACM; 2016. p. 55–60.
31. Sermpezis P, et al. RPKI Deployment Status and Validation Rates. *IEEE Commun Surv Tutor*. 2024 Jan-Mar;26(1):456–478.
32. National Institute of Standards and Technology (NIST). RPKI Monitor: Route Origin Authorization Statistics [Internet]. Gaithersburg (MD): NIST; 2025 [cited 2026 Jun 20]. Available from: https://rpki-monitor.antd.nist.gov/
33. MarketsandMarkets. Software-Defined Networking (SDN) Market by Offering, SDN Type, Application, End User, Vertical and Region — Global Forecast to 2030. Pune (India): MarketsandMarkets; 2025. Report No.: SE 655.
34. Microsoft Tech Community. A New Dawn of Software Defined Networking (SDN) in Windows Server 2025 [Internet]. Redmond (WA): Microsoft Corporation; 2024 Nov 4 [cited 2026 Jun 20]. Available from: https://techcommunity.microsoft.com/blog/networkingblog/

Received*: 18-04-2026 -* ***Accepted****: 16-06-2026 -* ***Published****: 21-06-2026*

35. Jayawardena C, Chen J, Bhalla A, Bu L. Comparative Analysis of POX and RYU SDN Controllers in Scalable Networks. *Int J Comput Netw Commun Secur*. 2025;17:35–51.
36. He J, Xiao Z, Zhu Y, Zou X, Liang L. Reinforcement Learning-Based SDN Routing Scheme Empowered by Action Influence Quantification and Graph Neural Network. *Front Comput Neurosci*. 2024 May;18:1393025.
37. Wang Y, et al. Dynamic-Loaded Deep Reinforcement Learning Algorithm with Fuzzy Logic for SDN Routing. *IEEE Access*. 2024;12:45678–45692.
38. Alenazi MJF. An Effective Deep-Q Learning Scheme for QoS Improvement in SDN Routing. *Phys Commun*. 2024 Oct;66:102387.
39. Goteti D, Rasheed S. AI-Driven Routing Pipeline in Software-Defined Networks Using DQL: A Mini Review. *Front Artif Intell*. 2025 Feb;8:1685155.
40. Li X, Li J, Zhou J, Liu J. Towards Robust Routing: Enabling Long-Range Perception with the Power of Graph Transformers and Deep Reinforcement Learning in Software-Defined Networks. *Electronics*. 2025 Feb;14(3):476.
41. Filsfils C, et al. SRv6 Deployment Considerations. Fremont (CA): Internet Engineering Task Force (IETF); 2025. (IETF Internet-Draft).
42. European Advanced Networking Test Center (EANTC). SRv6 Interoperability Test Report 2024. Berlin (Germany): EANTC; 2024.
43. National Security Agency (NSA). Commercial National Security Algorithm Suite 2.0 (CNSA 2.0). Fort Meade (MD): NSA Cybersecurity Information; 2022 Sep [updated 2025].
44. European Commission. A Coordinated Implementation Roadmap for the Transition to Post-Quantum Cryptography. Brussels (Belgium): NIS Cooperation Group; 2025 Jun.
45. Canadian Centre for Cyber Security. Roadmap for the Migration to Post-Quantum Cryptography for the Government of Canada (ITSM.40.001). Ottawa (ON): Cyber Centre; 2025 Jun.
46. Page MJ, et al. The PRISMA 2020 Statement: An Updated Guideline for Reporting Systematic Reviews. *BMJ*. 2021 Mar;372:n71.
47. Algazite HE, Nuredin A. IPv6 Addressing and Configuration: Building a Dual-Stack Network with OSPFv3. *AlQalam J Med Appl Sci*. 2025;2974–2980.
48. 48. Islam MS, et al. Graph Neural Network-Based Reinforcement Learning for Fault-Tolerant Routing in Smart Grid SDNs. IEEE Internet Things J. 2023;10(15):13542–13556.
49. 49. Goteti D, Rasheed I. Multipath Routing Algorithm to Find Optimal Path in SDN with POX Controller. Int J Electr Comput Eng Syst. 2025;16(2):121–131.

Received: *18-04-2026* - ***Accepted***: *16-06-2026* - ***Published***: *21-06-2026*