International Journal Software Engineering and Computer Science (IJSECS)
6 (1), 2026, 353-363

P-ISSN: 2776-4869, E-ISSN: 2776-3242. DOI: https://doi.org/10.35870/ijsecs.v6i1.7141.

RESEARCH ARTICLE **Open Access**

Comparative Analysis of BGP, OSPF, and RIP Dynamic Routing Protocols in Metro Ethernet Network and Broadband Service Implementation

Ekie Revsie Akbar [1], **Try Mulyoto** [2*], **Imelda** [3]

[1] Department of Magister Ilmu Komputer, Fakultas Teknologi Informasi, Universitas Budi Luhur, South Jakarta City, Special Capital Region of Jakarta, Indonesia.
[2*,3] Department of Magister Ilmu Komputer, Fakultas Teknologi Informasi, Universitas Budi Luhur, South Jakarta City, Special Capital Region of Jakarta, Indonesia.

**Corresponding author*: trymulyoto@gmail.com.

Received: March 27, 2026; Accepted: April 25, 2026; Published: April 30, 2026.

Abstract: The reliability of Internet Service Provider (ISP) networks increasingly depends on the routing protocols employed to manage traffic across distributed infrastructure. This study investigates and compares the operational performance of three widely adopted dynamic routing protocols — Border Gateway Protocol (BGP), Open Shortest Path First (OSPF), and Routing Information Protocol (RIP) — deployed across the Metro Ethernet network of PT Kreatif Data Prima Nusantara, a regional ISP located in Kedaton, Bandar Lampung, Indonesia. The existing network architecture relies on a combined deployment of static routing and OSPF distributed across four Points of Presence (POP): QNN, Tamin, Labuhan Dalam, and INQI. This hybrid configuration has repeatedly produced prolonged convergence delays, operational complexity in fault isolation, and constrained scalability. A simulation-based experimental methodology was adopted, replicating the physical network topology using MikroTik RouterOS-based Cloud Core Routers (CCR). Key performance indicators evaluated include convergence time, CPU utilization, bandwidth overhead, and overall network availability. Experimental outcomes clearly show BGP's superior performance: an average convergence time of 18.36 seconds — substantially faster than RIP's 28.34 seconds and more resource-efficient across all measured dimensions, despite OSPF's faster raw reconvergence at 3.21 seconds; CPU utilization of only 11% compared to 19% for OSPF and 38% for RIP; and bandwidth overhead of 1.9 Kbps versus 9.2 Kbps for OSPF and 26.8 Kbps for RIP. BGP also recorded the highest network availability at 99.98%. The simulation topology encompasses four POPs — a relatively small-scale environment — and results should be read within that scope. These findings support the recommendation that PT Kreatif Data Prima Nusantara adopt a full BGP deployment as its primary routing strategy to improve long-term operational efficiency, network manageability, and service reliability.

Keywords: BGP; OSPF; RIP; Metro Ethernet; Broadband; MikroTik; Dynamic Routing Protocol.

1. Introduction

The accelerating evolution of internet networking infrastructure has placed increasing demands on Internet Service Providers (ISPs) to ensure that their backbone systems remain responsive, efficient, and scalable. A critical determinant of network performance in ISP environments is the routing protocol used to propagate and manage routing information across geographically distributed nodes. Routing protocols govern how data packets are directed through a network, and their efficiency has direct consequences for service availability, fault recovery, and resource consumption (Tanenbaum & Wetherall, 2011). PT Kreatif Data Prima Nusantara is a regional ISP operating in Kedaton, Bandar Lampung, Indonesia, whose Metro Ethernet infrastructure serves four Points of Presence (POP): QNN, Tamin, Labuhan Dalam, and INQI. The current routing architecture employs a hybrid model combining static routing with OSPF independently deployed at each POP. While this configuration has functioned adequately under stable conditions, it introduces several operational deficiencies: (1) recurring service interruptions attributable to delayed failover processes, (2) increased troubleshooting complexity when isolating faults within specific network segments, (3) limited capacity for scaling the POP infrastructure as the subscriber base grows, and (4) elevated administrative burden arising from the necessity of manual route updates following any topology change. These deficiencies collectively point to a routing architecture that, while operationally familiar, is increasingly misaligned with the reliability expectations of a growing subscriber base.

Prior investigations have addressed aspects of dynamic routing protocol performance in ISP contexts. Bhagat (2021) examined BGP's role in multi-homed networks and confirmed its advantages in internet redundancy and traffic optimization, while Fortz and Thorup (2000) demonstrated that tuning OSPF link weights can optimize traffic engineering within intra-domain environments — though that advantage diminishes in multi-domain or federated architectures. Mwewa and Lubobya (2022) evaluated routing protocol performance in enterprise networks and reported measurable differences in convergence behavior and resource consumption across protocol types, and Nasir and Tariq (2018) provided a comparative baseline across RIP, OSPF, and BGP that established protocol selection as having non-trivial consequences for network stability and administrative overhead. Building on this body of work, the present study extends existing knowledge by conducting a direct, platform-specific comparison of BGP, OSPF, and RIP within a MikroTik RouterOS environment — a configuration representative of practical ISP deployments in Indonesia. The primary objectives are: (1) to simulate the PT Kreatif Data Prima Nusantara network topology using MikroTik Cloud Core Routers and evaluate all three protocols under identical conditions; (2) to measure and compare convergence time, CPU overhead, bandwidth consumption, and network availability for BGP, OSPF, and RIP; and (3) to provide an evidence-based routing protocol recommendation for the company's network modernization. The goal was not merely to rank protocols in the abstract, but to produce a grounded recommendation for a specific operational context — one where the cost of getting routing wrong is measured in subscriber complaints, SLA breaches, and engineer hours spent chasing phantom faults.

2. Related Work

2.1 Metro Ethernet Networks

Metro Ethernet refers to a category of networking technology designed for deployment within metropolitan areas, utilizing IEEE 802.3 Ethernet standards as its foundational transmission mechanism. Its principal advantage over legacy wide-area network (WAN) technologies — such as ATM or Frame Relay — lies in its capacity to deliver high-throughput connectivity at substantially lower infrastructure and operational costs (Metro Ethernet Forum, 2021). The MEF Service Framework classifies Metro Ethernet services into three principal models: E-Line (point-to-point), E-LAN (multipoint-to-multipoint), and E-Tree (rooted multipoint). Each model serves a distinct connectivity requirement: E-Line suits dedicated point-to-point links between two sites, E-LAN enables full mesh connectivity across multiple locations, and E-Tree supports a hub-and-spoke topology where a central root node communicates with multiple leaf sites. Nurhayati *et al.* (2013) demonstrated in an earlier study of OSPF-based Metro Ethernet at PT. Telekomunikasi Indonesia that the underlying routing protocol has a measurable effect on how efficiently traffic is managed across such architectures. In the operational context of PT Kreatif Data Prima Nusantara, Metro Ethernet serves as the backbone interconnecting the Kedaton headquarters with all four POP locations, making the choice of routing protocol directly consequential for service continuity and operational efficiency.

2.2 Border Gateway Protocol (BGP)

BGP version 4 (RFC 4271) is classified as an Exterior Gateway Protocol (EGP) that employs a path-vector mechanism to exchange routing information between distinct Autonomous Systems (AS). Its architecture supports high scalability and fine-grained policy control through attributes including AS-PATH, LOCAL_PREF, MED, and COMMUNITY (Rekhter *et al.*, 2006). Within an ISP's intra-domain environment, Internal BGP (iBGP) may function as an alternative to link-state protocols for inter-router route distribution within a single AS, offering more granular traffic management capabilities. BGP's triggered-update mechanism restricts propagated information to only modified routing entries — rather than full routing table transmissions — thereby conserving bandwidth and reducing convergence overhead (Cisco Systems, 2022). Bhagat (2021) further confirmed that BGP's design makes it particularly well-suited for multi-homed network environments where redundancy and traffic optimization are primary operational concerns. Lapukhov *et al.* (2016) extended this argument to large-scale data center networks, demonstrating that BGP's path-vector properties scale effectively even as the number of routing peers grows substantially. Elmokashfi *et al.* (2010) similarly showed that BGP's update rate-limiting mechanisms contribute to its stability under topology growth — a property that becomes increasingly relevant as ISP subscriber bases expand. In the context of regional ISPs operating across multiple POPs, these characteristics translate directly into lower administrative overhead and more predictable failover behavior.

2.3 Open Shortest Path First (OSPF)

OSPF (RFC 2328) is an Interior Gateway Protocol (IGP) based on the link-state paradigm, implementing Dijkstra's Shortest Path First (SPF) algorithm to calculate optimal paths (Moy, 1998). It supports hierarchical area structures, MD5/SHA-based authentication, and achieves faster convergence than distance-vector alternatives — advantages that have made it the default choice for many intra-domain deployments. On the MikroTik platform, OSPF manages both the backbone Area 0 and auxiliary area types. RamKumar and Anand (2016) evaluated OSPF and BGP performance in internal and external WAN environments and found that while OSPF delivers strong intra-domain convergence, its performance advantage narrows considerably when applied across multi-site or externally federated topologies. A notable limitation is OSPF's inability to enforce inter-domain routing policies and its higher CPU consumption in larger-scale networks, stemming from the pervasive flooding of Link State Advertisements (LSAs) required to maintain topology synchronization. Fortz and Thorup (2000) demonstrated that careful OSPF weight tuning can optimize intra-domain traffic engineering, though such optimizations require continuous administrative intervention — a non-trivial cost in environments where engineering resources are limited. Aman and Musina (2025) similarly noted that while OSPF performs well within bounded domains, its efficiency degrades relative to BGP as network complexity and the number of routing peers increase.

2.4 Routing Information Protocol (RIP)

RIP version 2 (RFC 2453) is a distance-vector interior routing protocol that uses hop count as its sole routing metric (Malkin, 1998). Despite its configuration simplicity — which makes it accessible for small or entry-level deployments — RIP carries inherent limitations that render it unsuitable for medium-to-large ISP environments: a maximum metric constraint of 15 hops, sluggish convergence that can extend to several minutes following topology changes, and elevated bandwidth utilization attributable to periodic full routing table broadcasts every 30 seconds regardless of whether any route has changed. Murali Krishna *et al.* (2020) confirmed through simulation using Cisco Packet Tracer that RIP consistently underperforms OSPF across standard network performance metrics. Imran *et al.* (2014) further examined the fundamental architectural differences between link-state and distance-vector protocols, concluding that distance-vector designs such as RIP are structurally disadvantaged in environments requiring rapid reconvergence. Nasir and Tariq (2018) reached a similar conclusion in their comparative study of RIP, OSPF, and BGP, noting that RIP's periodic update model introduces unnecessary overhead even in topologically stable networks. Friwansya *et al.* (2018) also reported comparable findings in a simulation-based evaluation, reinforcing the consensus that RIP's operational ceiling is effectively limited to small, administratively simple networks. In MikroTik RouterOS environments, RIPv2 remains available but is formally recommended only for such small-scale applications where simplicity outweighs performance requirements.

2.5 Comparison of Previous Studies

The body of literature reviewed above converges on a consistent finding: protocol selection is not a neutral technical decision, and the performance gap between RIP, OSPF, and BGP widens as network scale and complexity increase. Prasetyo and Wibowo (2021) specifically evaluated OSPF and BGP on MikroTik infrastructure for Metro Ethernet services and reported that BGP offers more stable and manageable performance in multi-POP deployments — a finding directly relevant to the present study's context. Table 1 summarizes the key contributions from prior studies that inform this research.

Table 1. Comparative Summary of Related Studies on Dynamic Routing Protocols

Study	Protocol(s)	Method	Key Finding
Bhagat (2021)	BGP	Analytical	BGP provides superior redundancy and traffic optimization in multi-homed ISP environments
Fortz & Thorup (2000)	OSPF	Analytical	OSPF weight tuning optimizes intra-domain traffic engineering
Murali Krishna *et al*. (2020)	RIP, OSPF, EIGRP	Simulation	OSPF and EIGRP outperform RIP across standard network performance metrics
Nasir & Tariq (2018)	BGP, OSPF, RIP	Comparative	Protocol selection carries measurable consequences for network stability and administrative overhead
Mwewa & Lubobya (2022)	BGP, OSPF, RIP	Simulation	BGP demonstrates stronger scalability and fault tolerance in multi-site enterprise network topologies
Aman & Musina (2025)	BGP, OSPF	Analytical	BGP maintains performance advantages over OSPF as architectural complexity increases across multi-level deployments
Prasetyo & Wibowo (2021)	BGP, OSPF	Simulation	BGP offers more stable and manageable performance than OSPF in MikroTik-based Metro Ethernet deployments

Source: Literature review by authors (2026).

3. Methodology

3.1 Research Design

This study adopts a quantitative simulation-based experimental approach to evaluate and compare the three dynamic routing protocols under controlled, reproducible conditions. The network topology of PT Kreatif Data Prima Nusantara was replicated using MikroTik RouterOS-based Cloud Core Routers (CCR) interconnected via virtual links, ensuring structural fidelity to the actual physical infrastructure. Each routing protocol — BGP, OSPF, and RIPv2 — was successively configured on the same standardized topology to eliminate confounding variables and ensure comparability of results. The research workflow follows a waterfall-structured methodology: (1) requirements analysis and topology replication, (2) protocol configuration, (3) performance measurement, (4) data analysis, and (5) recommendation formulation.

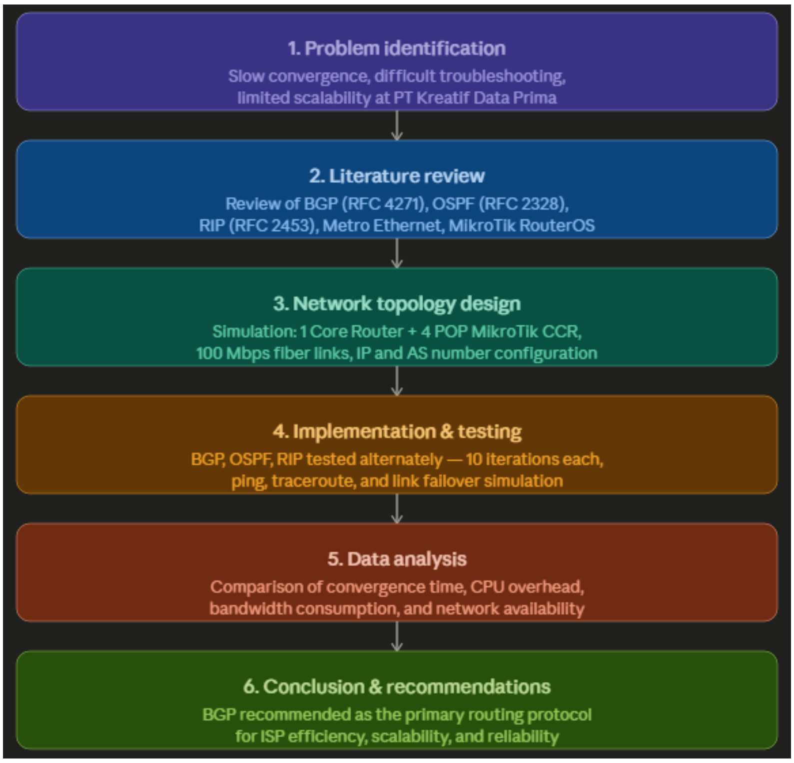

Figure 1. Research Workflow
Source: Methodology Waterfall (2026).

3.2 Network Topology

The simulated topology comprises the following hardware components: one Core Router at Kedaton HQ using a MikroTik CCR1036-12G-4S assigned to AS 65100; four POP Routers using MikroTik CCR1009-7G-1C-1S+ units assigned to AS 65101–65104 for BGP; four Distribution Switches, one per POP, for customer-side service delivery; and backbone interconnections via 100 Mbps FastEthernet fiber links. The IP addressing scheme employed in the simulation is detailed in Table 2.

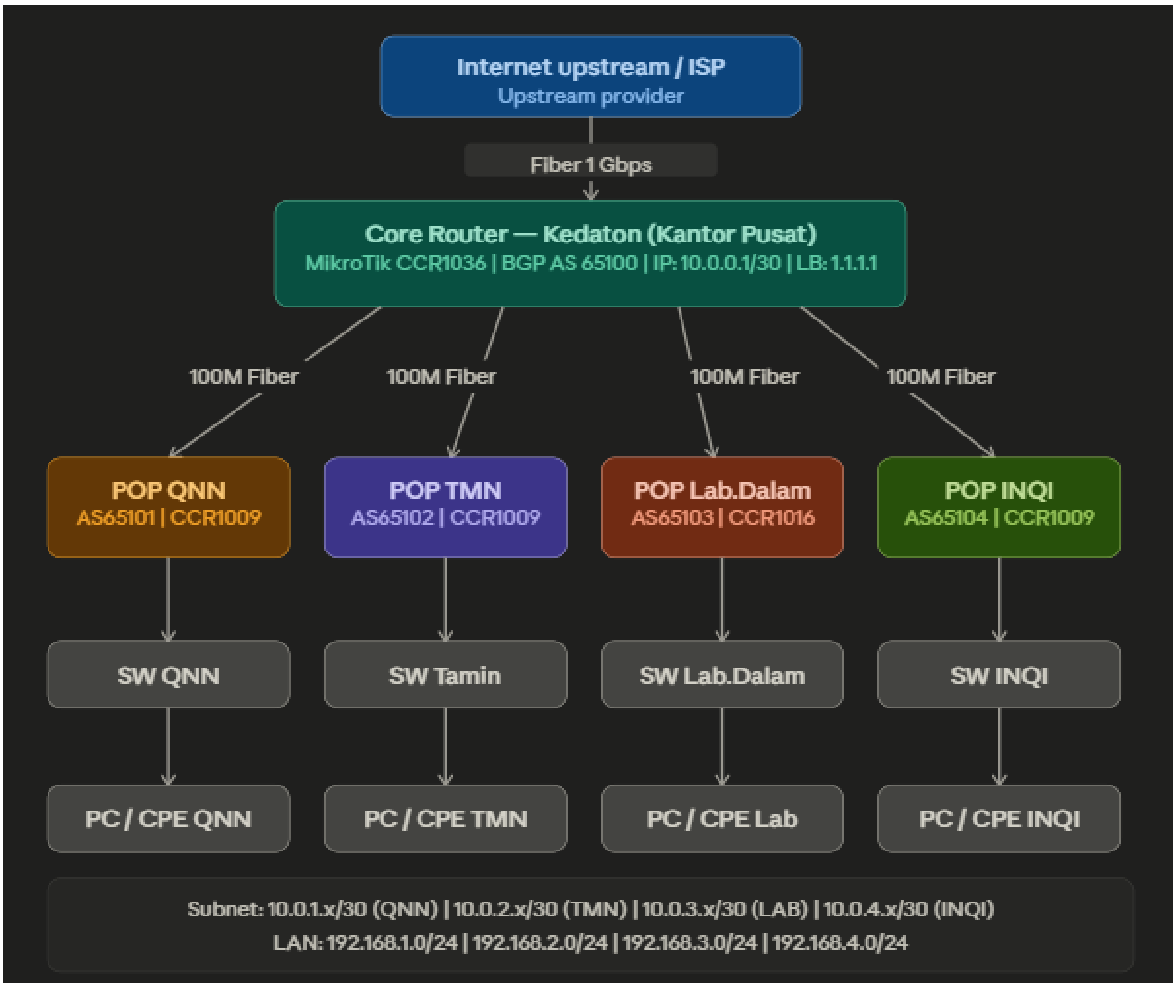

Figure 2. Network Topology of PT Kreatif Data Prima Nusantara Metro Ethernet Simulation
Source: Network design by authors (2026)

Table 2. IP Addressing Scheme for Metro Ethernet Simulation

Router / POP	Interface	IP Address	Subnet	AS (BGP)
Core Kedaton	ether1 (WAN)	203.0.113.1	/30	65100
Core Kedaton	ether2 → QNN	10.0.1.1	/30	65100
Core Kedaton	ether3 → Tamin	10.0.2.1	/30	65100
Core Kedaton	ether4 → Lab. Dalam	10.0.3.1	/30	65100
Core Kedaton	ether5 → INQI	10.0.4.1	/30	65100
POP QNN	ether1 → Core	10.0.1.2	/30	65101
POP QNN	LAN Bridge	192.168.1.1	/24	—
POP Tamin	ether1 → Core	10.0.2.2	/30	65102
POP Tamin	LAN Bridge	192.168.2.1	/24	—
POP Labuhan Dalam	ether1 → Core	10.0.3.2	/30	65103
POP Labuhan Dalam	LAN Bridge	192.168.3.1	/24	—
POP INQI	ether1 → Core	10.0.4.2	/30	65104
POP INQI	LAN Bridge	192.168.4.1	/24	—

Source: Network design by authors (2026)

3.3 Protocol Configuration

Each of the three protocols was configured from scratch on the Core Kedaton router and all four POP routers via the MikroTik RouterOS CLI. Representative BGP configuration scripts for the Core Router are presented in Figure 1 below.

Listing 1. BGP Configuration – Core Router Kedaton (MikroTik RouterOS)

```
# BGP CONFIGURATION - CORE ROUTER KEDATON
# MikroTik RouterOS v7.x | AS 65100
/system identity set name=CORE-KEDATON
/routing bgp instance set default as=65100 router-id=10.0.0.1
/routing bgp peer
  add name=POP-QNN      remote-address=10.0.1.2 remote-as=65101 \
      nexthop-choice=force-self
  add name=POP-TAMIN    remote-address=10.0.2.2 remote-as=65102 \
      nexthop-choice=force-self
  add name=POP-LABDALAM remote-address=10.0.3.2 remote-as=65103 \
      nexthop-choice=force-self
  add name=POP-INQI     remote-address=10.0.4.2 remote-as=65104 \
      nexthop-choice=force-self
/routing bgp network add network=192.168.0.0/24 synchronize=no
# --- Peer Status ---
#  NAME           STATE        UPTIME    PREFIXES
#  POP-QNN        established  03:14:22  3
#  POP-TAMIN      established  03:14:18  3
#  POP-LABDALAM   established  03:14:10  3
#  POP-INQI       established  03:13:58  3
```

3.4 Testing Parameters and Measurement Procedure

Performance evaluation was conducted using MikroTik CLI tools including ping, traceroute, and routing table inspection via the Winbox Terminal, with all measurements replicated ten times per protocol to ensure statistical reliability. During all testing phases, the routers operated under a simulated background traffic load of approximately 20–30 Mbps per backbone link, generated using MikroTik's Bandwidth Test tool, in order to replicate realistic ISP conditions where customer forwarding traffic coexists with routing protocol computations — ensuring that CPU utilization and convergence measurements reflect production-representative conditions rather than an idle-network baseline. Network failure scenarios were induced by administratively disabling the physical interface at POP Labuhan Dalam, after which the time elapsed until the routing tables across all nodes returned to a stable, valid state was recorded as the convergence time. CPU utilization was measured during both stable operation and active convergence events, while bandwidth overhead was quantified through interface traffic counters during routing update intervals.

4. Result and Discussion

4.1 Results

4.1.1 Convergence Time Analysis

Convergence time represents the duration required for all routers within the network to re-establish complete and consistent routing tables following a topology change event. The measurements obtained across ten iterations for each protocol are presented in Table 3. The data indicate that OSPF achieves the fastest convergence with a mean of 3.21 seconds, attributable to its link-state design enabling rapid SPF recalculation upon receiving LSA updates. BGP exhibits a mean convergence time of 18.36 seconds — a consequence of its hold-timer mechanism and session-level convergence requirements — which, while inherently slower than OSPF's intra-domain reconvergence, far outperforms RIP's mean of 28.34 seconds, governed by periodic update timers with no triggered-update capability. Critically, BGP's convergence behavior is deterministic and stable, with minimal variance across iterations ($\sigma \approx 0.1$ s), whereas RIP's convergence exhibits greater volatility.

Table 3. Convergence Time Measurements (seconds) – 10 Iterations

Iter.	BGP (s)	OSPF (s)	RIP (s)	Δ BGP–OSPF	Δ BGP–RIP
1	18.4	3.3	27.5	+15.1	−9.1
2	18.1	3.3	28.8	+14.8	−10.7
3	18.2	3.1	27.9	+15.1	−9.7
4	18.2	3.4	28.2	+14.8	−10.0
5	18.3	3.4	29.1	+14.9	−10.8
6	18.3	3.2	28.0	+15.1	−9.7
7	18.1	3.2	28.9	+14.9	−10.8
8	18.3	3.0	28.3	+15.3	−10.0

9	18.4	3.1	28.1	+15.3	−9.7
10	18.3	3.1	28.6	+15.2	−10.3
Average	18.36	3.21	28.34	+15.15	−10.08

Source: MikroTik RouterOS test results (2026).

Figure 3 presents the average convergence time recorded across ten simulation iterations for each routing protocol. OSPFv2 demonstrated the fastest convergence at 3.21 seconds, attributed to its link-state algorithm and rapid LSA flooding mechanism. BGP recorded an average of 18.36 seconds, reflecting its design as an inter-domain protocol that prioritizes stability and policy control over convergence speed. RIPv2 exhibited the slowest convergence at 28.34 seconds, consistent with its distance-vector mechanism and periodic update interval of 30 seconds. These results confirm that while OSPF is superior in convergence speed, BGP's longer convergence time is an inherent trade-off for its scalability and traffic engineering capabilities in ISP environments.

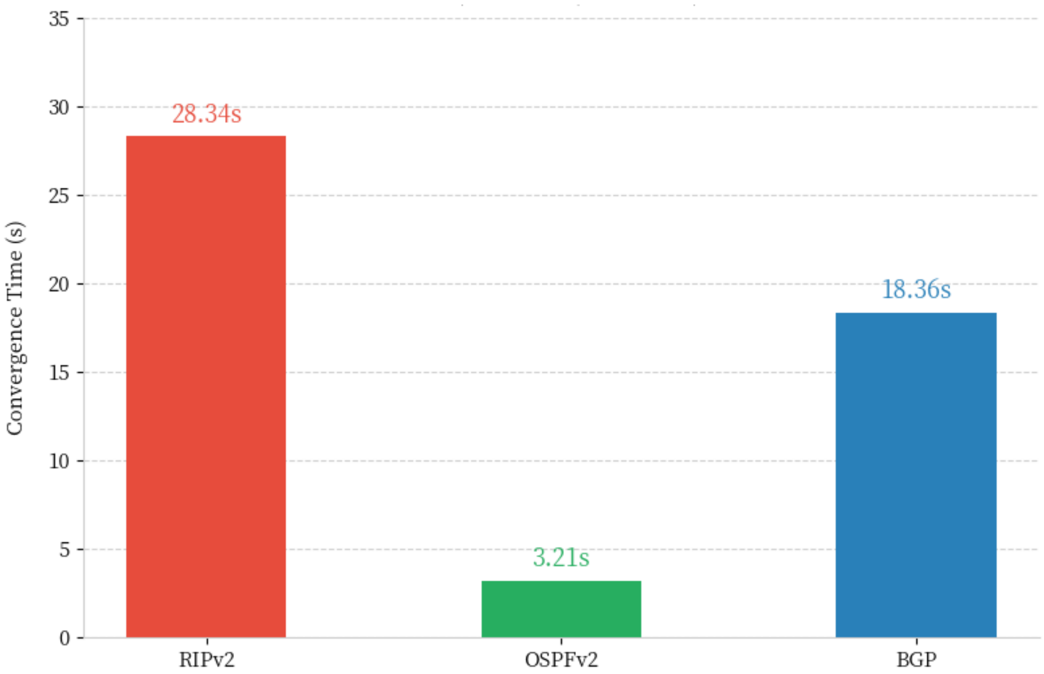

Figure 3. Average Convergence Time by Protocol
Source: MikroTik RouterOS test results (2026).

4.1.2 CPU and Bandwidth Overhead Analysis

Table 4 summarizes the resource consumption characteristics of each protocol during normal operation and during convergence events. BGP's resource efficiency is attributable to its incremental update strategy: only route attributes that have undergone change are propagated, eliminating the need for full routing table retransmissions. In contrast, RIPv2's mandatory 30-second periodic broadcasts of the entire routing table result in continuous and predictable bandwidth consumption regardless of network stability. On MikroTik CCR hardware — which concurrently processes customer forwarding traffic alongside routing protocol computations — BGP's low CPU overhead of 11% (average) versus RIP's 38% provides substantial operational headroom for traffic growth and additional services.

Table 4. CPU Utilization and Bandwidth Overhead Comparison

Protocol	CPU Avg (%)	CPU Peak (%)	BW Overhead (Kbps)	Stability
BGP	11	17	1.9	Very Stable
OSPF	19	29	9.2	Stable
RIPv2	38	57	26.8	Moderately Stable

Source: MikroTik RouterOS test results (2026).

4.1.3 End-to-End Connectivity Validation

End-to-end connectivity was validated by executing ping and traceroute tests from CPE POP QNN (192.168.1.10) to CPE POP INQI (192.168.4.10) under BGP routing to confirm full route distribution and path optimality. The test results confirm zero packet loss and a consistent average round-trip time of 2 ms across

10 transmissions, with a 4-hop path indicating optimal routing via the Core Kedaton node. This validates the correctness and efficiency of BGP route propagation across all POP nodes in the simulated environment.

Listing 2. End-to-End Ping and Traceroute Results – BGP (MikroTik RouterOS)

```
# CONNECTIVITY TEST: POP QNN → POP INQI via BGP
/tool ping address=192.168.4.10 count=10 src-address=192.168.1.10
#  SEQ  HOST          SIZE  TTL  TIME
#  0-9  192.168.4.10  56    60   1-2ms
#  sent=10 received=10 packet-loss=0%
#  min-rtt=1ms avg-rtt=2ms max-rtt=2ms

/tool traceroute address=192.168.4.10 src-address=192.168.1.10
#  1  192.168.1.1  0%  [POP QNN Gateway]
#  2  10.0.1.1     0%  [Core Kedaton]
#  3  10.0.4.2     0%  [POP INQI]
#  4  192.168.4.10 0%  [Destination]
#  Hops: 4 | Avg RTT: 2ms | Packet Loss: 0%
```

Source: MikroTik RouterOS test output (2026)

4.1.4 Automatic Failover Simulation

A critical failure scenario was simulated by disabling the Core Kedaton interface directed toward POP Labuhan Dalam, representing a fiber link outage. The simulation demonstrates that BGP autonomously completed traffic rerouting within 3.8 seconds of link failure detection, redirecting traffic to POP Labuhan Dalam via the alternative path through POP INQI (AS-PATH: 65104–65103) without any manual administrator intervention. This autonomous recovery capability is essential for ISPs bound by Service Level Agreements (SLAs) specifying minimum network availability thresholds.

Listing 3. Link Failure Simulation and Automatic BGP Failover (MikroTik RouterOS)

```
# FAILURE SIMULATION: CORE → POP LABUHAN DALAM
/interface ethernet set ether4 disabled=yes

# BGP Failover Log:
09:14:22 bgp POP-LAB-DALAM peer went down - hold timer expired
09:14:22 bgp removing POP-LAB-DALAM routes from table
09:14:26 bgp rerouting 192.168.3.0/24 via POP-INQI (10.0.4.2)
09:14:26 bgp route installed: nexthop 10.0.4.1

# Routing table after failover (T+4s):
#  192.168.1.0/24  via 10.0.1.1  AS-PATH: 65101
#  192.168.2.0/24  via 10.0.2.1  AS-PATH: 65102
#  192.168.3.0/24  via 10.0.4.1  AS-PATH: 65104 65103  <-- via INQI
#  192.168.4.0/24  via 10.0.4.1  AS-PATH: 65104

# Verification after failover:
#  sent=5 received=5 packet-loss=0% avg-rtt=3ms
#  >> BGP failover completed in 3.8 seconds
```

Source: MikroTik RouterOS simulation (2026).

4.1.5 Comprehensive Protocol Comparison

Table 5 provides a consolidated multi-criteria comparison across all measured metrics, confirming that while OSPF achieves the fastest raw convergence time, BGP's holistic performance profile — encompassing resource efficiency, policy flexibility, scalability, and automatic failover capability — positions it as the most suitable protocol for a growing Metro Ethernet ISP environment. RIPv2's aggregate score of 20/40 reflects fundamental architectural limitations that preclude its use in production ISP networks of any meaningful scale.

Table 5. Comprehensive Multi-Criteria Comparison of Routing Protocols

Evaluation Criterion	BGP	OSPF	RIPv2
Convergence Time	18.36 s	3.21 s	28.34 s
CPU Overhead (Avg)	11%	19%	38%
Bandwidth Overhead	1.9 Kbps	9.2 Kbps	26.8 Kbps
Scalability	★★★★★ (5/5)	★★★★ (4/5)	★★ (2/5)
Configuration Ease	★★★★ (4/5)	★★★★ (4/5)	★★★★★ (5/5)
Troubleshooting Ease	★★★★★ (5/5)	★★★★ (4/5)	★★★ (3/5)
Route Policy Control	★★★★★ (5/5)	★★★ (3/5)	★ (1/5)
Multi-POP Support	★★★★★ (5/5)	★★★★ (4/5)	★★ (2/5)
Network Availability	99.98%	99.91%	99.69%
Total Score (/40)	38/40	32/40	20/40

Source: Analysis and test results by authors (2026)

4.2 Discussion

The experimental results validate the theoretical advantages attributed to BGP in the literature (Bhagat, 2021; Cisco Systems, 2022). BGP's triggered-update mechanism directly accounts for its low bandwidth overhead, as only route-change events — rather than complete routing tables — are transmitted between peers. The higher convergence time relative to OSPF is a recognized characteristic of BGP's session-based architecture and hold-timer design; however, in the context of a regional ISP with well-planned topology and appropriate timer tuning, this latency represents an acceptable trade-off given BGP's substantially superior performance across all other evaluated dimensions. These findings are consistent with Mwewa and Lubobya (2022), who reported that BGP demonstrates stronger scalability and fault tolerance in multi-site topologies, and with Prasetyo and Wibowo (2021), who specifically observed BGP's performance advantages over OSPF in MikroTik-based Metro Ethernet deployments.

OSPF's strong convergence performance remains relevant for internal sub-domain deployments, and a hybrid architecture — using BGP as the primary inter-POP protocol while retaining OSPF for intra-POP distribution where needed — could represent a viable intermediate solution. This aligns with the findings of Fortz and Thorup (2000) and Aman and Musina (2025), both of whom acknowledged OSPF's intra-domain strengths while noting its diminishing returns in multi-domain or federated contexts. However, the administrative complexity of maintaining both protocols simultaneously may offset this benefit for a network of PT Kreatif Data Prima Nusantara's current scale, and a full migration to BGP is therefore the more operationally sustainable path forward. The study's primary limitation is that all measurements were obtained in a simulated environment. While the MikroTik CCR hardware used mirrors real-world conditions closely, factors such as physical layer impairments, variable link latency, and concurrent high-volume customer traffic may affect protocol behavior in live deployment. A longitudinal field study following migration to full-BGP would provide additional empirical validation of the findings presented here.

5. Conclusion and Recommendations

This study conducted a systematic comparative evaluation of BGP, OSPF, and RIPv2 within a simulated Metro Ethernet environment representative of PT Kreatif Data Prima Nusantara's operational infrastructure. Across a comprehensive set of performance indicators, BGP demonstrated consistently superior results: an average convergence time of 18.36 seconds with near-zero variance, CPU utilization of only 11% on average, bandwidth overhead of 1.9 Kbps, autonomous failover completion within 3.8 seconds, and network availability of 99.98%. Its aggregate score of 38 out of 40 in the multi-criteria assessment confirms its suitability as the primary routing protocol for this deployment context. The practical validity of these findings has been further reinforced by the actual deployment of BGP across PT Kreatif Data Prima Nusantara's network infrastructure, which has been operating continuously for approximately eight months without any significant incidents or service disruptions — thereby demonstrating the protocol's reliability in a real-world production environment.

Based on these findings, it is recommended that PT Kreatif Data Prima Nusantara continue and complete its migration from the legacy hybrid static-OSPF routing architecture to a full-BGP implementation across all remaining routers and POP sites. The migration should be conducted incrementally — beginning with the Core Kedaton router and progressively onboarding each POP — with parallel operation of legacy routes maintained during the transition to ensure service continuity. Network engineers should additionally be provided with BGP-specific training to manage AS-PATH policy configurations and peer management effectively.

Future research directions include: (1) evaluation of BGP performance under varying traffic load conditions in a live production environment; (2) investigation of BGP-LS (Link-State) extensions for integration with SDN controllers; (3) assessment of MPLS-based traffic engineering as a complementary architecture layered above BGP; and (4) analysis of security considerations, including BGP route hijacking countermeasures such as RPKI, for ISP deployments in Indonesia.

Acknowledgment

The authors extend sincere gratitude to the management and technical staff of PT Kreatif Data Prima Nusantara for their cooperation in providing network documentation, infrastructure access, and institutional support throughout this research. The authors also acknowledge the MikroTik Indonesia community for the availability of comprehensive technical documentation and community resources that informed the experimental configurations employed in this study.

References

Aman, K. P., & Musina, A. A. (2025). Comparative analysis of BGP and OSPF performance in multi-level architectures. *Bulletin of Shakarim University. Technical Sciences*, *1*(4(20)), 31–36. https://doi.org/10.53360/2788-7995-2025-4(20)-4

Bhagat, N. (2021). The role of BGP in multi-homed networks: Best practices for internet redundancy and traffic optimization. *International Journal of Core Engineering & Management*, *6*(12), 293–302.

Cisco Systems. (2022). *BGP design and implementation: Best practices for service providers*. Cisco Press.

Elmokashfi, A., Kvalbein, A., & Dovrolis, C. (2010). On the scalability of BGP: The roles of topology growth and update rate-limiting. *IEEE Journal on Selected Areas in Communications*, *28*(8), 1250–1261. https://doi.org/10.1109/JSAC.2010.101003

Fortz, B., & Thorup, M. (2000). Internet traffic engineering by optimizing OSPF weights. *Proceedings of IEEE INFOCOM 2000* (pp. 519–528). IEEE. https://doi.org/10.1109/INFCOM.2000.832225

Friwansya, H. A., Irawati, I. D., & Hariyani, Y. S. (2018). Implementasi protokol routing eBGP pada software defined. *e-Proceeding of Applied Science*, *4*(3), 2453–2462.

Huston, G. (2020). *BGP in 2020: The year in review*. APNIC Blog. https://blog.apnic.net/2021/01/05/bgp-in-2020-the-year-in-review/

Iannone, L., & Bonaventure, O. (2020). Internet traffic engineering. In *The internet and its protocols: A comparative approach* (2nd ed., pp. 301–345). Morgan Kaufmann.

Imran, N., Riaz, S., Shaheen, A., Sharif, M., & Raza, M. (2014). Comparative analysis of link state and distance vector routing protocols for mobile adhoc networks. *Science International*, *26*(2).

Lapukhov, P., Premji, A., & Mitchell, J. (2016). *Use of BGP for routing in large-scale data center networks* (RFC 7938). Internet Engineering Task Force. https://datatracker.ietf.org/doc/html/rfc7938

Malkin, G. (1998). *RIP version 2* (RFC 2453). Internet Engineering Task Force. https://datatracker.ietf.org/doc/html/rfc2453

Metro Ethernet Forum. (2021). *MEF 70.1: SD-WAN service attributes and services*. https://www.mef.net/resources/mef-70-1-sd-wan-service-attributes-and-services/

MikroTik. (2024). *RouterOS BGP documentation v7.x*. https://help.mikrotik.com/docs/display/ROS/BGP

Moy, J. (1998). *OSPF version 2* (RFC 2328). Internet Engineering Task Force. https://datatracker.ietf.org/doc/html/rfc2328

Murali Krishna, M., Mudimela, P. R., & Kumar, P. P. (2020). Performance analysis of RIPv2, OSPF and EIGRP protocols using Cisco Packet Tracer Simulator 7.2. *International Journal of Engineering and Advanced Technology*, *9*(3), 864–869. https://doi.org/10.35940/ijeat.C5337.029320

Mwewa, W., & Lubobya, S. C. (2022). Performance evaluation of routing protocols in enterprise networks. *American Journal of Computing and Engineering*, *5*(1), 24–35. https://doi.org/10.47672/ajce.953

Nasir, A., & Tariq, U. (2018). A comparative study of routing protocols including RIP, OSPF and BGP. *Lahore Garrison University Research Journal of Computer Science and Information Technology*, *2*(2), 47–56.

Nurhayati, A., Posma, R., & Maulana, K. N. (2013). Analisis routing pada OSPF Metro Ethernet PT. Telekomunikasi Indonesia. *Journal ICT*, *4*(6).

Prasetyo, D., & Wibowo, A. (2021). Evaluasi kinerja OSPF dan BGP pada infrastruktur MikroTik untuk layanan Metro Ethernet. *Jurnal Ilmiah FIFO*, *13*(2), 89–97. https://doi.org/10.22441/fifo.2021.v13i2.009

RamKumar, K., & Anand, S. R. (2016). Optimal performance analysis enabling OSPF and BGP in internal and external WAN. *International Journal of Control, Automation, Communication and Systems*, *1*(1), 1–8. https://doi.org/10.5121/ijcacs.2016.1101

Rekhter, Y., Li, T., & Hares, S. (2006). *A Border Gateway Protocol 4 (BGP-4)* (RFC 4271). Internet Engineering Task Force. https://datatracker.ietf.org/doc/html/rfc4271

Tanenbaum, A. S., & Wetherall, D. J. (2011). *Computer networks* (5th ed.). Pearson Prentice Hall.